Iid sampling with self-avoiding particle filters: the pinball sampler

Christian P. Robert†

CEREMADE, Université Paris Dauphine and CREST, Insee, Paris, France

Kerrie L. Mengersen‡

School of Mathematical and Physical Sciences, University of Newcastle, Australia

Summary. Particle filters are usually implemented in sequential settings or for processing and analysis of large datasets. We show they can also produce iid samples for a given reference distribution π. The proposal chosen for the moves of the particle system is based on a self-avoiding random walk with a bouncing relaunch mechanism and a delayed rejection acceptance rule that avoids the immediate vicinity of the other particles, while preserving the stationary distribution π.

Keywords: Adaptive scheme, delayed rejection, MCMC algorithms, particle filter, pseudo-reference distribution, self-avoiding random walk

1. Introduction

Particle filters, that is simultaneous processing of a vector of random variables $\left(\theta_1^{(t)}, \ldots, \theta_M^{(t)}\right)$ towards simulation from a reference distribution π and computation of related integrals,

$$\int h(\theta)\pi(d\theta) \tag{1}$$

are currently used in sequential settings such as tracking a moving target (see, e.g., Gilks and Berzuini, 2000, or Doucet *et al.*, 2001) or in static settings with large datasets as exploratory tools (see, e.g., Berzuini and Gilks, 2001, or Chopin, 2001). Various improvements have occurred in the past years with the introduction of additional moves based on MCMC techniques, as detailed in Doucet *et al.* (2001).

The current paper shows that systems of particles also have an appeal from a purely Monte Carlo point of view in the sense that they can produce iid samples from the reference distribution π, a feature that is only achieved with difficulty via the complex perfect sampling technique in regular MCMC methodology. We present an updating system for the particle system that is based on a standard random walk with corrections to avoid the immediate vicinity of other particles. It therefore classifies as a self-avoiding random walk for the whole chain $\left(\theta_1^{(t)}, \ldots, \theta_M^{(t)}\right)_t$.

A major difference with the current work on particle filters is that we move away from the importance sampling justification and produce an iid and unbiased representation of

†Research supported by EU TMR network ERB–FMRX–CT96–0095 on *"Computational and Statistical Methods for the Analysis of Spatial Data"*.

‡Supported by Centre National de la Recherche Scientifique during a visit at CREST in April 2001 and by ARC Large Grant.

the reference distribution. Paradoxically, the corresponding importance resampling particle system based on the same proposal enjoys poor properties like high degeneracy and low mixing. The method we propose thus pertains more closely to the theme of adaptive MCMC algorithms in that it takes advantage of the vector $\left(\theta_1^{(t)},\ldots,\theta_M^{(t)}\right)$ to speed up the exploration of the distribution of interest. It thus relates to earlier adaptive schemes such as Gilks and Roberts (1995), Haario, Sacksman and Tamminen (2000) and Warnes (2000). In particular, it is similar to the multistate sampler of Warnes (2000) in that it simulates simultaneously a set of values in a dependent manner, two major differences being that, first, we do not aim at approximating the posterior distribution through a kernel estimator, as in Oudjane (2000) and Warnes (2000), and, second, the particles are connected via a "repulsive" proposal. The difference with adaptive proposals as in Haario, Sacksman and Tamminen (2000) and Andrieu and Robert (2001) is that the chain we simulate is homogeneous and, hence, standard ergodic theorems apply.

The paper is organised as follows: We establish the general properties of our particle filter in Section 2, with the description of the general particle filter algorithm and a focus on the iid nature of the resultant samples. The bouncing relaunch mechanism is then detailed in Section 3 and its justification is embedded in current concepts of delayed rejection. Some illustrations of the performance of the new algorithm are considered in Section 4.

2. Particle filters

A system of particles $\left(\theta_1^{(t)},\ldots,\theta_M^{(t)}\right)_t$ is a random vector evolving over time t according to a Markovian updating scheme. It is customarily associated with a vector of weights, $\left(\omega_1^{(t)},\ldots,\omega_M^{(t)}\right)$, $\sum_{i=1}^M \omega^{(t)} = 1$, that is used to approximate integrals such as (1) through importance sampling approximations

$$\sum_{k=1}^M \omega_k^{(t)} h(\theta_k^{(t)}) .$$

Therefore, the original particle filter algorithm does not directly simulate the $\theta_k^{(t)}$'s, but rather updates the weights $\omega_k^{(t)}$ until the discrepancy between those weights is too high, in which case a resampling move based on the $\omega_k^{(t)}$'s is launched (and possibly supplemented with additional MCMC moves, to avoid degeneracy, as in Berzuini and Gilks, 2001, or Crisan and Doucet, 2001).

In this paper we follow a completely different approach in that (a) we do not use importance sampling schemes nor weights $\omega_k^{(t)}$, and (b) we resample the *whole* vector $\left(\theta_1^{(t)},\ldots,\theta_M^{(t)}\right)$ at each iteration of the algorithm. In this respect, we are closer to regular Monte Carlo simulation than to standard particle filtering, principally because we simulate a whole vector instead of a single variable θ. We can also draw a link with parallel MCMC algorithms where a thread of MCMC chains is generated, usually to assess convergence to the stationary distribution and mixing properties of the chain, as in Gelman and Rubin (1992) and subsequent literature. The difficulty with these early implementations is that the convergence and the dynamics of the various chains $\left(\theta_k^{(t)}\right)_t$ were considered individually, rather than as a whole: the distribution of $\theta_k^{(t)}$ is thus of the form $K_t\left(\theta_k^{(0)},\theta_k^{(t)}\right)$ and

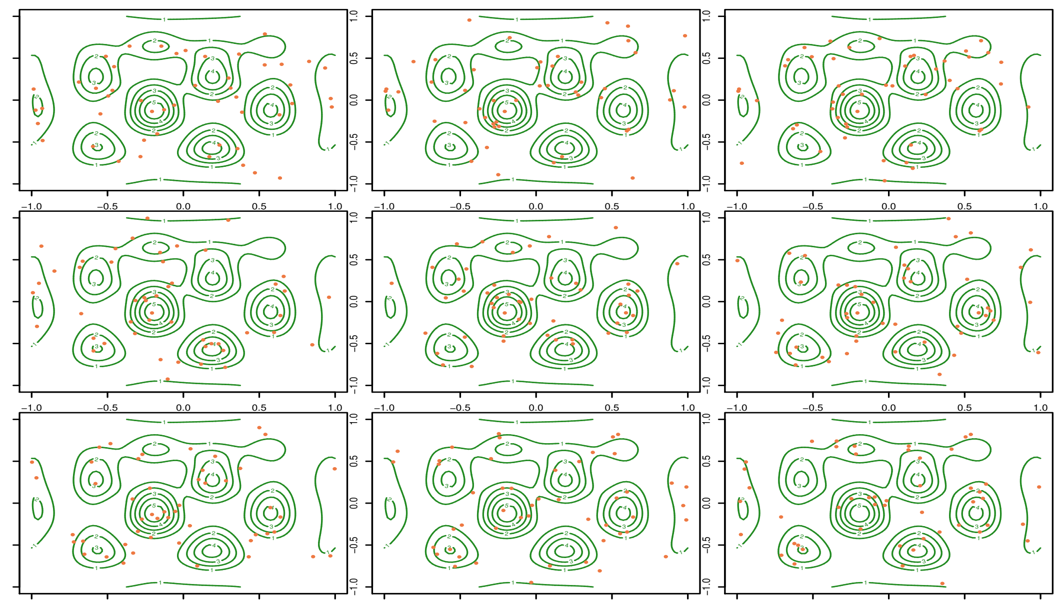

Fig. 1. Successive iterations of 50 particles over the contours of a two-dimensional target distribution

convergence characteristics are dependent on the starting value $\theta_k^{(0)}$.

The algorithm developed in this paper introduces a kernel defined over the vector of particles; the corresponding stationary distribution is π^M, and we prove that convergence occurs for this vector. This approach also has its limitations in that the number of particles is necessarily fixed over the iterations to achieve convergence (unless some branching process extension applies or can be devised). Figure 1 shows 9 successive iterations of 50 particles from a particle system as the particles explore the surface of a target distribution described in Section 4. This snapshot clearly indicates how the particles interact over the target surface, visiting the different modes in proportion to their amplitude.

The move from the system at time t to the system at time $t+1$ is Gibbs-like, in that the update of each particle is sequential and is conditional on the current value of the whole particle system:

[**Algorithm GPF**]

(a) Construct a grid $\left(\theta_1^{(0)}, \ldots, \theta_M^{(0)}\right)$ of starting values over the support of π

(b) For $t = 0, \ldots, T$,
  for $k = 1, \ldots, M$,
    simulate

$$\theta_k^{(t+1)} \sim K_k\left(\theta \,\middle|\, \theta_1^{(t+1)}, \ldots, \theta_{k-1}^{(t+1)}, \theta_k^{(t)}, \ldots, \theta_M^{(t)}\right)$$

---

with proposals $K_k(\cdot|\theta_1,\ldots,\theta_M)$ that are associated with the stationary distribution π, e.g. that satisfy the detailed balance condition (see, e.g., Robert and Casella, 1999)

$$\begin{aligned}\pi\left(\theta_k^{(t)}\right) K_k\left(\theta_k^{(t+1)}\left|\theta_1^{(t+1)},\ldots,\,\theta_{k-1}^{(t+1)},\theta_k^{(t)},\ldots,\theta_M^{(t)}\right.\right) &=\\ \pi\left(\theta_k^{(t+1)}\right) K_k\left(\theta_k^{(t)}\left|\theta_1^{(t+1)},\ldots,\,\theta_{k-1}^{(t+1)},\theta_k^{(t+1)},\ldots,\theta_M^{(t)}\right.\right)&\,. \end{aligned}\tag{2}$$

(A particular choice for K_k that pushes particles farther apart from one another will be detailed in the next section.) When the support of π is unbounded, the grid in the initial step of [Algorithm GPF] can be constructed from a starting distribution μ or via a reparameterisation of the distribution onto a compact support like $[-1,1]^p$ (see, e.g., Robert, 2001) and a uniform grid.

A crucial feature of this sequential update is that it approximates iid sampling from π:

THEOREM 1. *The stationary distribution associated with* [Algorithm GPF] *is*

$$(\theta_1,\ldots,\theta_M) \sim \pi(\theta_1)\times\ldots\times\pi(\theta_M) = \pi^M(\theta_1,\ldots,\theta_M)$$

Proof. Since the kernels $K_k(\cdot|\theta_1,\ldots,\theta_M)$ satisfy the (conditional) balance condition (2), the fact that $\left(\theta_1^{(t+1)},\ldots,\theta_M^{(t+1)}\right)\sim\pi^M$ if $\left(\theta_1^{(t)},\ldots,\theta_M^{(t)}\right)\sim\pi^M$ then follows by a chain rule:

$$\begin{aligned}\left(\theta_1^{(t+1)},\ldots,\theta_M^{(t+1)}\right) &\sim \int \pi^M\left(d\theta_1^{(t)},\ldots,d\theta_M^{(t)}\right) K_1\left(\theta_1^{(t+1)}\left|\theta_1^{(t)},\ldots,\theta_M^{(t)}\right.\right)\\ &\qquad \cdots K_M\left(\theta_M^{(t+1)}\left|\theta_1^{(t+1)},\ldots,\theta_M^{(t)}\right.\right)\\ &= \pi\left(\theta_1^{(t+1)}\right)\int\pi^{M-1}\left(d\theta_2^{(t)},\ldots,d\theta_M^{(t)}\right) K_2\left(\theta_2^{(t+1)}\left|\theta_1^{(t+1)},\ldots,\theta_M^{(t)}\right.\right)\\ &\qquad \cdots K_M\left(\theta_M^{(t+1)}\left|\theta_1^{(t+1)},\ldots,\theta_M^{(t)}\right.\right)\\ &\;\;\vdots\\ &= \pi^M\left(\theta_1^{(t+1)},\ldots,\theta_M^{(t+1)}\right)\end{aligned}$$

□

Note that, like the Gibbs sampler, [Algorithm GPF] does not satisfy the detailed balance condition on the whole vector $(\theta_1,\ldots,\theta_M)$ since it depends on the order in which the particles are updated. It could be transformed into a reversible algorithm by introducing a random permutation of the indices but we fail to see the gain in this additional step.

Provided the standard irreducibility conditions (see Robert and Casella, 1999) are satisfied by the componentwise transition kernels $K_k(\cdot|\theta_1,\ldots,\theta_M)$, the k-fold Markov chain $\left(\theta_1^{(t)},\ldots,\theta_M^{(t)}\right)_t$ is thus ergodic and positive reccurent with stationary distribution π^M. Therefore, once the particle system $\left(\theta_1^{(t)},\ldots,\theta_M^{(t)}\right)$ has eliminated its dependence on the initial value $\left(\theta_1^{(0)},\ldots,\theta_M^{(0)}\right)$, it can be considered as an iid sample from π at any given time t, rather than in the long run as in regular MCMC sampling. The result is therefore more interesting that usual ergodicity theorems in that a slice *(in time)* of the process provides an unbiased evaluation of the reference distribution π, once stationarity is achieved. An evaluation of the simulation output as in regular Monte Carlo experiments such as normal approximation confidence intervals can then be undertaken. The central effort in implementing our method bears on devising proposals that speed up mixing, as detailed in the following section.

3. Bouncing relaunch

The interest in using a particle sample $\left(\theta_1^{(t)},\ldots,\theta_M^{(t)}\right)_t$, rather than a single chain $\left(\theta_1^{(t)}\right)_t$, is primarily to increase the possibilities in the exploration of the state space and to avoid trapping setups such as the usual concentration around a *single* mode of the reference distribution π that depends on the starting point for the chain. While the choice of the kernels K_k is quite open [for convergence to π^M to occur], we propose in this section a particular Metropolis-within-Metropolis set of kernels K_k that is designed to address these problems. The basic ideas behind our choice are

(a) to introduce a pseudo-reference distribution π^R that pushes particles farther apart from one another;
(b) to implement a Metropolis move calibrated for π^R;
(c) to increase the efficiency of this proposal via a (currently deterministic) delayed rejection mechanism;
(d) to validate the move with a final Metropolis acceptance probability calibrated for the true reference distribution π.

3.1. *Repulsive proposal distribution*

First, we introduce the *pseudo reference distribution*

$$\pi_k^R(\theta) \propto \pi(\theta) \prod_{j\neq k} e^{-\xi/\pi(\theta_j)||\theta-\theta_j||^2}\,, \tag{3}$$

which is derived from the distribution of interest π by adding exponential terms that create holes around the other particles θ_j, since the j-th term $e^{-\xi/\pi(\theta_j)||\theta-\theta_j||^2}$ converges to 0 when θ goes to θ_j. These terms thus induce a repulsive effect (hence the R) around the other particles. The factor $\pi(\theta_j)$ in the exponential is there to moderate the repulsive effect in zones of high probability and to enhance it in zones of low probability. (This effect can be further attenuated by replacing $\pi(\theta_j)$ with $\max(\zeta,\pi(\theta_j))$ so that points θ_j with almost zero density $\pi(\theta_j)$ do not induce a trapping effect of their own by excluding almost the entire state space.) The parameter ξ is a *tempering factor*, since the repulsive part in π_k^R can be seen as a term raised to the power ξ; this factor can be calibrated during the simulation against the number of particles M and the acceptance rate. Note also that the construction of π_k^R does not depend on the normalisation constant of π, since such a dependence would jeopardize the implementation of the method in most MCMC setups: indeed, when the normalisation constant in π is unknown, ξ can be rescaled as $\xi\pi(\theta_0)$, where θ_0 is an arbitrary value in the parameter space with $\pi(\theta_0)>0$, or as the average of the $\pi(\theta_j)$'s over the other particles.

The move related to π_k^R maps $\theta_k^{(t)}$ to $\tilde\theta_k^{(t)}$ via a Metropolis-Hastings kernel

$$\tilde K_k\left(\theta_k \left|\theta_1^{(t+1)},\ldots,\ \theta_{k-1}^{(t+1)},\theta_k^{(t)},\ldots,\theta_M^{(t)}\right.\right)$$

tuned for the pseudo reference distribution π_k^R. (Note that $\tilde{K}_k$ includes a Dirac mass at $\theta_k^{(t)}$). Then the value $\tilde{\theta}_k^{(t)}$ is evaluated again π, using the Metropolis-Hastings acceptance probability

$$1 \wedge \frac{\pi\left(\tilde{\theta}_k^{(t)}\right)}{\pi\left(\theta_k^{(t)}\right)} \frac{\pi_k^R\left(\theta_k^{(t)}\right)}{\pi_k^R\left(\tilde{\theta}_k^{(t)}\right)}$$

as justified by the following lemma:

Lemma 1. *If θ' given θ is generated by a Metropolis-Hastings mechanism with stationary distribution π', the acceptance probability*

$$\alpha(\theta, \theta') = 1 \wedge \frac{\pi(\theta')}{\pi(\theta)} \frac{\pi'(\theta)}{\pi'(\theta')}$$

makes the move a Metropolis-Hastings mechanism with stationary distribution π.

Proof. Once again, the proof is based on the balance condition. If K' is the kernel associated with the original move,

$$\pi'(\theta)K'(\theta,\theta') = \pi'(\theta')K'(\theta',\theta)$$

and the overall kernel is

$$K^\star(\theta,\theta') = \alpha(\theta,\theta')K'(\theta,\theta') + \int (1-\alpha(\theta,\xi))K'(\theta,d\xi)\delta_\theta(\theta'),$$

where $\delta_\theta(\theta')$ denotes the Dirac mass at θ. It therefore satisfies

$$\begin{aligned}
\pi(\theta)K^\star(\theta,\theta') &= \{\pi(\theta)/\pi'(\theta) \wedge \pi(\theta')/\pi'(\theta')\}\pi'(\theta)K'(\theta,\theta') \\
&\quad + \int (1-\alpha(\theta,\xi))K'(\theta,d\xi)\pi(\theta)\delta_\theta(\theta') \\
&= \{\pi(\theta)/\pi'(\theta) \wedge \pi(\theta')/\pi'(\theta')\}\pi'(\theta')K'(\theta',\theta) \\
&\quad + \int (1-\alpha(\theta',\xi))K'(\theta',d\xi)\pi(\theta')\delta_{\theta'}(\theta) \\
&= \pi(\theta')K^\star(\theta',\theta).
\end{aligned}$$

□

Notice that this correction for using the wrong target distribution can also be found in Neal (1996) and Celeux *at al.* (2000), in tempered moves.

3.2. Avoiding the other particles: a delayed rejection mechanism

We now concentrate on the construction of the kernel $\tilde{K}_k(\cdot|\theta_1,\ldots,\theta_M)$ which is designed to approximate simulation from the repulsive distribution π_k^R.

To this end, we will make use of the delayed rejection method of Tierney and Mira (1999): when a move brings a particle θ_k too close to other particles, it is likely to be rejected because of the repulsive part in π_k^R. Rather than leaving θ_k in its current state, we propose a new move for θ_k which deflects it further away from those particles. This new move can be evaluated in its turn with a corrected acceptance probability, which takes into account the fact that the first move has been rejected. Continuing, if this second move is also rejected, further moves can be considered in the same way. As in Tierney and Mira (1999), these moves can depend on the previous [rejected] proposed values. The new feature is that, given we are working with a system of particles, the moves can also depend on the other particles.

In order to achieve reversibility, a critical feature that ensures detailed balance, we use a random walk proposal, followed by iterated *deterministic* reflections if the current proposal is rejected. More precisely, given $\left(\theta_1^{(t+1)}, \ldots, \theta_k^{(t)}, \ldots, \theta_M^{(t)}\right)$, we generate the first proposal θ_{k0} from a symmetric distribution $q_\tau\left(||\theta - \theta_k^{(t)}||\right)$, with scale factor τ calibrated for the reference distribution π. This proposal is accepted with probability

$$\alpha_{k0}\left(\theta_k^{(t)}, \theta_{k0}\right) = 1 \wedge \frac{\pi_k^R\left(\theta_{k0}\right)}{\pi_k^R\left(\theta_k^{(t)}\right)}$$

If θ_{k0} is rejected, this may be for two reasons: either the new value θ_{k0} has fallen in a region of low density, and the rejection is governed by the first part of π_k^R, or θ_{k0} is too close to one or to several particles, and the rejection is governed by the second part of π_k^R. We then devise the delayed rejection for this second possibility, by introducing a *pinball effect*, i.e. by taking the particle away from the particle closest to θ_{k0}, θ_j say, to θ_{k1} through a symmetry of $\theta_k^{(t)}$ with respect to the (θ_j, θ_{k0}) line. Obviously, if the rejection is due to a low value for $\pi(\theta)$, the additional proposal may send the particle into a region of even lower dimension. This danger, however, can be partially alleviated in a compact space by devising reflection rules that bring back the particle to the center of the space after a series of moves toward the boundary of the space.

If $\sigma(\theta_0, \theta_1, \theta_2)$ denotes the density of the second move, the delayed rejection acceptance probability of Tierney and Mira (1999) is

$$\alpha_{k1}\left(\theta_k^{(t)}, \theta_{k0}, \theta_{k1}\right) = 1 \wedge \frac{\pi_k^R(\theta_{k1}) q_\tau(||\theta_{k0} - \theta_{k1}||)\sigma\left(\theta_{k1}, \theta_{k0}, \theta_k^{(t)}\right)\left[1 - \alpha_{k0}(\theta_{k1}, \theta_{k0})\right]}{\pi_k^R\left(\theta_k^{(t)}\right) q_\tau\left(||\theta_{k0} - \theta_k^{(t)}||\right) \sigma\left(\theta_k^{(t)}, \theta_{k0}, \theta_{k1}\right)\left[1 - \alpha_{k0}\left(\theta_k^{(t)}, \theta_{k0}\right)\right]}.$$

Since the first proposal is symmetric,

$$q_\tau\left(||\theta_{k0} - \theta_k^{(t)}||\right) = q_\tau(||\theta_{k0} - \theta_{k1}||),$$

since $\sigma(\theta_0, \theta_1, \theta_2)$ reduces here to a Dirac mass, and since the first proposal was rejected, implying

$$\pi_k^R(\theta_{k0}) < \pi_k^R\left(\theta_k^{(t)}\right),$$

the delayed rejection acceptance probability simplifies into

$$\begin{aligned}
\alpha_{k1}\left(\theta_k^{(t)}, \theta_{k0}, \theta_{k1}\right) &= 1 \wedge \frac{\pi_k^R(\theta_{k1})\left[1 - 1 \wedge \frac{\pi_k^R(\theta_{k0})}{\pi_k^R(\theta_{k1})}\right]}{\pi_k^R\left(\theta_k^{(t)}\right)\left[1 - \frac{\pi_k^R(\theta_{k0})}{\pi_k^R\left(\theta_k^{(t)}\right)}\right]} \\
&= 1 \wedge \frac{0 \vee \left[\pi_k^R(\theta_{k1}) - \pi_k^R(\theta_{k0})\right]}{\pi_k^R\left(\theta_k^{(t)}\right) - \pi_k^R(\theta_{k0})}
\end{aligned}$$

as in Mira (1998, Section 5.3).

If θ_{k1} is rejected, the delayed rejection mechanism described above can be repeated, as a series of symmetries with respect to a line defined by the closest of the other particles. We can then generalise Mira's (1998) formula (5.13), obtained in the symmetric case, to our setting to establish the following result:

LEMMA 2. *Delayed rejection associated with the acceptance probabilities* $(1 \leq m \leq H)$

$$\alpha_{km}(\theta_k^{(t)}, \ldots, \theta_{km}) = 1 \wedge \frac{0 \vee [\pi_k^R(\theta_m) - \max_{0 \leq \ell < m} \pi_k^R(\theta_\ell)]}{\pi_k^R(\theta_k^{(t)}) - \max_{0 \leq \ell < m} \pi_k^R(\theta_\ell)} \tag{4}$$

preserves π_k^R *as its stationary distribution.*

Proof. It is sufficient, as in Mira (1998, Section 5.3), to consider the third candidate, with acceptance probability

$$\begin{aligned}
\alpha_{k2}(\theta_k^{(t)}, \theta_{k0}, \theta_{k1}, \theta_{k2}) &= 1 \wedge \frac{\pi_k^R(\theta_{k2})}{\pi_k^R(\theta_k^{(t)})} \frac{1 - \alpha_{k0}(\theta_{k2}, \theta_{k1})}{1 - \alpha_{k0}(\theta_k^{(t)}, \theta_{k0})} \frac{1 - \alpha_{k1}(\theta_{k2}, \theta_{k1}, \theta_{k0})}{1 - \alpha_{k1}(\theta_k^{(t)}, \theta_{k0}, \theta_{k1})} \\
&= 1 \wedge \frac{\pi_k^R(\theta_{k2})}{\pi_k^R(\theta_k^{(t)})} \frac{1 - 1 \wedge \frac{\pi_k^R(\theta_{k2})}{\pi_k^R(\theta_{k1})}}{1 - \frac{\pi_k^R(\theta_{k0})}{\pi_k^R(\theta_k^{(t)})}} \frac{1 - 1 \wedge \frac{[\pi_k^R(\theta_{k0}) - \pi_k^R(\theta_{k1})] \vee 0}{[\pi_k^R(\theta_{k2}) - \pi_k^R(\theta_{k1})] \vee 0}}{1 - 1 \wedge \frac{[\pi_k^R(\theta_{k1}) - \pi_k^R(\theta_{k0})] \vee 0}{\pi_k^R(\theta_k^{(t)}) - \pi_k^R(\theta_{k0})}}.
\end{aligned}$$

This ratio is equal to 0 if $\pi_k^R(\theta_{k2}) < \pi_k^R(\theta_{k1})$. If $\pi_k^R(\theta_{k2}) > \pi_k^R(\theta_{k1})$,

$$\begin{aligned}
\alpha_{k2}(\theta_k^{(t)}, \theta_{k0}, \theta_{k1}, \theta_{k2}) &= 1 \wedge \frac{\pi_k^R(\theta_{k2}) - \pi_k^R(\theta_{k1})}{\pi_k^R(\theta_k^{(t)}) - \pi_k^R(\theta_{k0})} \frac{1 - 1 \wedge \frac{[\pi_k^R(\theta_{k0}) - \pi_k^R(\theta_{k1})] \vee 0}{\pi_k^R(\theta_{k2}) - \pi_k^R(\theta_{k1})}}{1 - 1 \wedge \frac{[\pi_k^R(\theta_{k1}) - \pi_k^R(\theta_{k0})] \vee 0}{\pi_k^R(\theta_k^{(t)}) - \pi_k^R(\theta_{k0})}} \\
&= 1 \wedge \frac{\pi_k^R(\theta_{k2}) - \pi_k^R(\theta_{k1}) - [\pi_k^R(\theta_{k0}) - \pi_k^R(\theta_{k1})] \vee 0}{\pi_k^R(\theta_k^{(t)}) - \pi_k^R(\theta_{k0}) - [\pi_k^R(\theta_{k1}) - \pi_k^R(\theta_{k0})] \vee 0}
\end{aligned}$$

Therefore,

$$\begin{aligned}
\alpha_{k2}(\theta_k^{(t)}, \theta_{k0}, \theta_{k1}, \theta_{k2}) &= \begin{cases}
0 & \text{if } \pi_k^R(\theta_{k2}) < \pi_k^R(\theta_{k1}), \\
1 \wedge \frac{\pi_k^R(\theta_{k2}) - \pi_k^R(\theta_{k1})}{\pi_k^R(\theta_k^{(t)}) - \pi_k^R(\theta_{k0}) - \pi_k^R(\theta_{k1}) + \pi_k^R(\theta_{k0})} & \text{if } \pi_k^R(\theta_{k2}) > \pi_k^R(\theta_{k1}), \\
& \text{and } \pi_k^R(\theta_{k1}) > \pi_k^R(\theta_{k0}), \\
1 \wedge \frac{\pi_k^R(\theta_{k2}) - \pi_k^R(\theta_{k1}) - \pi_k^R(\theta_{k1}) + \pi_k^R(\theta_{k0})}{\pi_k^R(\theta_k^{(t)}) - \pi_k^R(\theta_{k0})} & \text{if } \pi_k^R(\theta_{k2}) > \pi_k^R(\theta_{k1}), \\
& \text{and } \pi_k^R(\theta_{k1}) < \pi_k^R(\theta_{k0}),
\end{cases} \\
&= 1 \wedge \frac{[\pi_k^R(\theta_{k2}) - \max(\pi_k^R(\theta_{k1}), \pi_k^R(\theta_{k0}))] \vee 0}{\pi_k^R(\theta_k^{(t)}) - \max(\pi_k^R(\theta_{k1}), \pi_k^R(\theta_{k0}))}
\end{aligned}$$

as in Mira (1998). □

The delayed rejection algorithm thus reads as follows:

[Algorithm DRPF]

(a) Construct a grid $(\theta_1^{(0)}, \ldots, \theta_M^{(0)})$ of starting values over the support of π
(b) For $t = 0, \ldots, T$,
  for $i = 1, \ldots, M$,

  (i) Generate $\theta_{i0} \sim \mathcal{N}(\theta_i^{(t)}, \tau^2)$
  (ii) Determine the Metropolis–Hastings acceptance probability $\alpha_0(\theta_k^{(t)}, \theta_{i0})$
  (iii) While θ_{im} is rejected $(m = 0, \ldots, H-1)$,

    A. Determine $\theta_m^\star$ the closest particle to θ_{im}
    B. If $m = 0$, take θ_{i1} as the symmetric of θ_{i0} wrt the $(\theta_m^\star, \theta_k^{(t)})$ line
    Else, take $\theta_{i(m+1)}$ as the symmetric of $\theta_{i(m-1)}$ wrt the $(\theta_m^\star, \theta_{im})$ line
    C. Compute the delayed rejection acceptance probability

$$\alpha_{i(m+1)}\left(\theta_i^{(t)}, \ldots, \theta_{i(m+1)}\right)$$

  (iv) Take $\theta_i^{(t+1)} = \theta_{im}$

The deterministic move from θ_{im} to $\theta_{i(m+1)}$ can be replaced by a stochastic move centered at the symmetry, provided reversibility is preserved. Note that the overall scheme is dimension-free, and thus can be implemented in any dimension.

4. Illustrations

To illustrate how the particle system explores the whole space, we first chose a two-dimensional distribution on $[-1, 1]^2$ represented in Figure 2. The density of this distribution is proportional to

$$\begin{aligned} f(x,y) &= ((x-1)(x+1)\sin(7y) + (y-1)(y+1)\sin(8x))^2 \\ &\quad +((x-1)(x+1)\cos(6y) - (y-1)(y+1)\sin(8x))^2 , \end{aligned}$$

which was chosen for its multimodal features. Figure 3 represents the changes in the system after one iteration, while Figure 4 represents the particle filter after 50 iterations. The strong dispersion of the particles across the entire space is evident within and across iterations. A visually acceptable representation of the distribution is achieved as a result, with coverage of all modes according to the target density.

As a second illustration we consider a diffuse bivariate t_1 distribution with correlation of 0.4., based on the *a priori* belief that the DRPF algorithm will better describe tail behaviour than the traditional Metropolis algorithm. A bivariate distribution was again adopted for ease of representation, although the methodology accommodates problems of general dimension. Figure 5 depicts perspective plots of the target distribution, typical

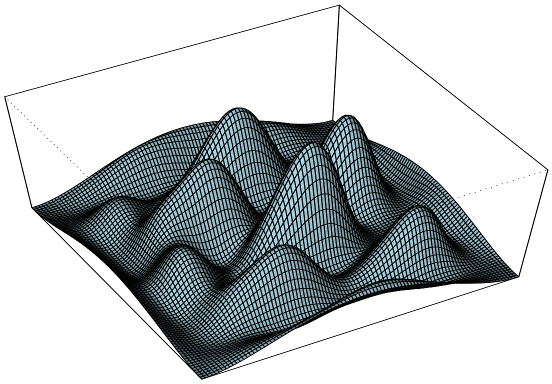

Fig. 2. Representation of the reference distribution $\pi(x,y) = ((x-1)(x+1)\sin(7y) + (y-1)(y+1)\sin(8x))^2 + ((x-1)(x+1)\cos(6y) - (y-1)(y+1)\sin(8x))^2$ on $[-1,1]^2$.

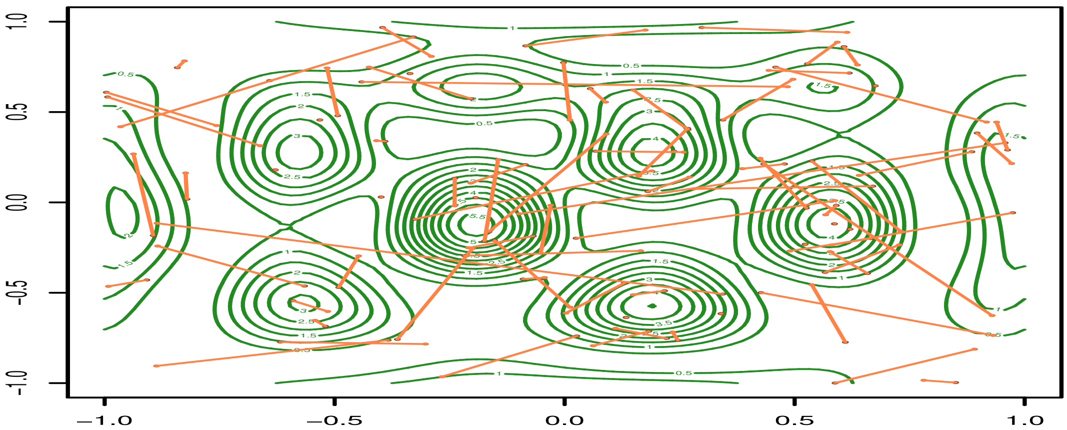

Fig. 3. One iteration of the particle system with moves indicated by lines

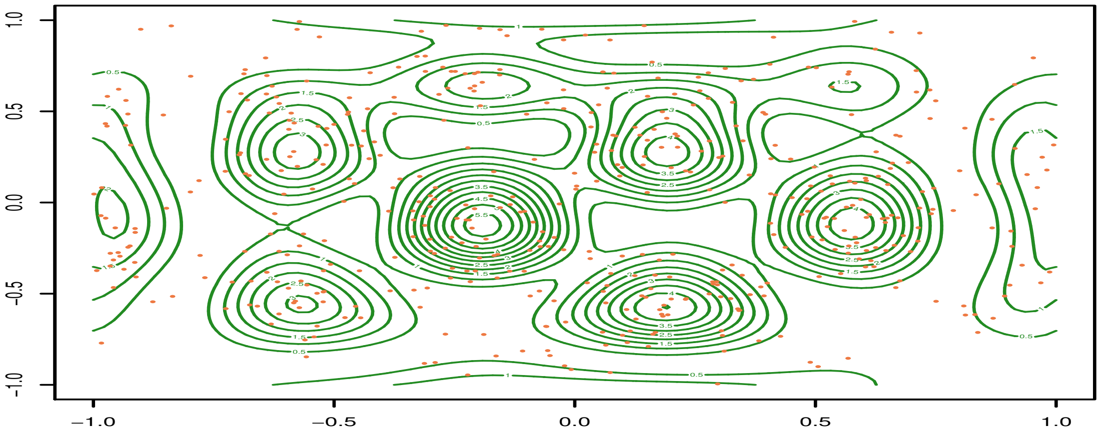

Fig. 4. The particle system after 50 iterations

samples arising from the Metropolis and DRPF algorithms and a corresponding quantile-quantile plot of the two sample outputs. Comparison of the ranges in the plots of the two sample outputs indicates that the DRPF algorithm may tend to push particles further, although the qq plot confirms that this is a tenuous conclusion. Indeed, despite the disperse nature of the distribution, the attraction to the mode overwhelms the repulsive effect of neighbouring particles. Interestingly, the dispersion is thus most evident in the upper tails rather than the outer tails of the distribution.

This latter tendency was further explored by comparing the proportion of the sample under each algorithm that lies outside defined "confidence bands", where the latter are defined in terms of distance from the centre (0,0) of the target distribution. Thus a 90distance to the mean. Under this comparison, not dissimilar proportions of the Metropolis and DRPF samples were observed in the 95% to 70% bands and again in the 20% to 5% bands, but greater proportions of DRPF samples were seen in the 60%-30% bands.

5. Conclusion

This paper has two main aims. First, particle sampling algorithms are demonstrated to produce iid samples from the target distribution. Second, a new algorithm is introduced that enhances exploration of a space and avoidance of trapping spaces by "bouncing" particles away from each other using a delayed rejection mechanism. The new approach is embedded in current particle filtering and delayed rejection literature and its performance is examined both theoretically and empirically.

The illustrations considered in Section 4 provide some indication of the performance of the proposed algorithm. While the first example validates the claim that the DRPF algorithm indeed facilitates exploration of multimodal spaces, the second example indicates that its performance is at least equivalent to that of the traditional Metropolis algorithm in standard but difficult cases, with the major gain being the generation of iid samples.

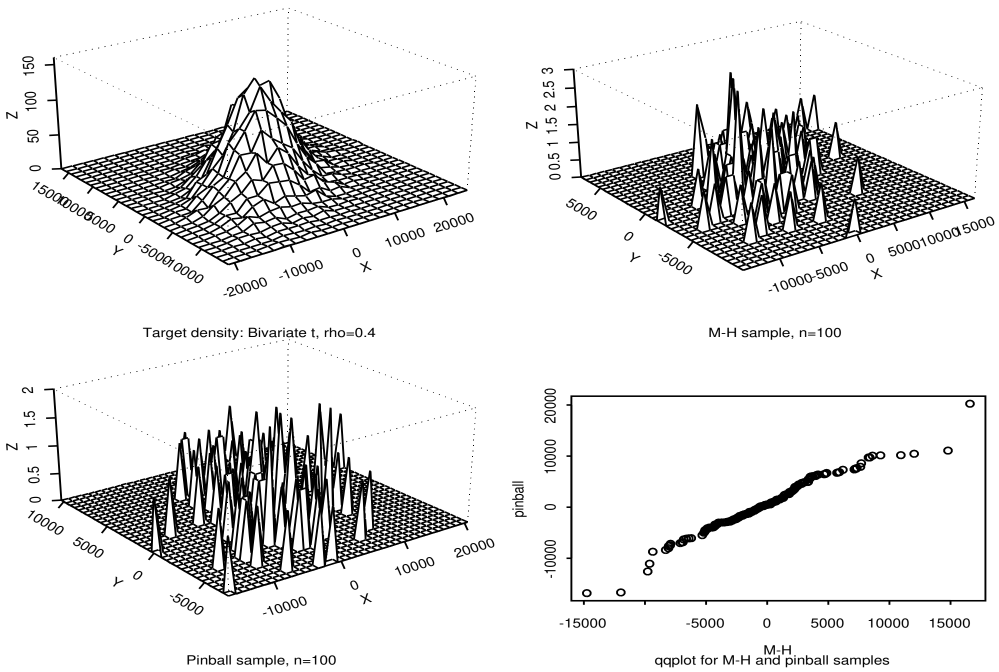

Fig. 5. Comparison of Metropolis and DRPF samples of a bivariate t distribution with one degree of freedom and correlation 0.4. The four plots represent the target density (top left), a typical sample of 100 iterations from a Metropolis algorithm (top right), a typical sample of 100 particles from a DRPF algorithma (bottom left) and a quantile-quantile plot of the Metropolis and pinball samples (bottom right).

These gains are obtained at minimal programming cost in that, provided the distribution π is reparametrised over a compact set like $[-1,1]^p$ (or relatively compact set like $(-1,1)^p$), the same program applies. Finer details like calibration of (ξ,τ) obviously need to be polished for new problems, but this can be achieved in a learning round where various values of (ξ,τ) are assessed against the acceptance rate (a) of the whole system, (b) of the various particles, with assessment focused on the maintenance of reasonable acceptance rates. In addition, the asymptotically iid nature of the particule system allows for simpler convergence assessment than in regular MCMC settings. For instance, non-parametric homogeneity tests can be run at various times.

Caveats with the method are acknowledged, including obvious issues such as the resurgence of the *curse of dimensionality*: the starting values are assumed to provide a fair coverage of the support of π. In a generic setting, this assumption is untenable in high dimensions. More fundamentally perhaps, the proposal and rejection moves all rely on Euclidean balls: as the dimension increases, the amount of space lost in a coverage by these balls increases as well. It may be that substituting balls based on the L_1 or maximum distances to the Euclidean balls brings better performance, but the implementation is likely to be more delicate.

The other difficulty is that the pseudo-distribution (3) has a complexity of order $\mathrm{O}(M)$ and induces a complexity $\mathrm{O}(M^2)$ for a step of [Algorithm GPF]. This constraint prevents the derivation of large samples when mixing is poor and convergence requires many iterations. It could be relaxed, however, by considering only certain neighbours in (3) at most steps and by updating the set of neighbours at fixed intervals. The attempts we made at implementing this method were, however, unsucessful in that the definition of the neighbourhoods cancels reversibility: using the same algorithm as above is then not valid, as shown by the behaviour of the corresponding particle systems.

Acknowledgements

A conversation with Robert L. Wolpert along rue Estienne d'Orves in 1993 acted as a generator for this work. Thanks also to C. Andrieu, N. Chopin, P. Green, and M. Hurn for helpful discussions about this paper.

References

Andrieu, C., and Robert, C.P. (2001) Controlled MCMC for Optimal Sampling. Tech. report 2001-25, CEREMADE, Paris.

Berzuini, C., and Gilks, W.R. (2001) Resample-move filtering with cross-model jumps. In Sequential Monte Carlo in Practice (eds. Doucet, A., de Freitas, J.F.G., and Gordon, N.J.). New York: Springer-Verlag.

Cappé, O., and Robert, C.P. (2000) Ten years and still running! *J. American Statistical Society* **95**(4), 1282–1286.

Celeux, G., Hurn, M., and Robert, C.P. (2000) Computational and inferential difficulties with mixtures posterior distribution *J. American Statistical Society* **95**(3), 957–979.

Chopin, N. (2001) A sequential particle filter for static models. *Biometrika* (to appear).

Crisan, D., and Doucet, A. (2000) Convergence of sequential Monte Carlo methods. Tech. report, U. of Cambridge.

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

Gilks, W.R., Best, N.G., and Tan, K.K.C. (1995) Adaptive rejection Metropolis sampling within Gibbs sampling. *J. Royal Statistical Society, Series C* **44**, 455–472.

Gilks, W.R. and Berzuini, C. (2001) Following a moving target: Monte Carlo inference for dynamic Bayesian models. *J. Royal Statistical Society, Series B* (to appear).

Haario, H., Saksman, E., and Tamminen, J. (1999) Adaptive Proposal Distribution for Random Walk Metropolis Algorithm. *Computational Statistics* **14**, 375-395.

Neal, R. (1996) Sampling from multimodal distributions using tempered transitions. *Statistics and Computing* **4**, 353–366

Oudjane, N. (2000) Stabilité et approximations particulaires en filtrage non-linéaire. Ph. D. thesis, INRIA, Rennes.

Robert, C.P. (2001) Discussion of G. Roberts' "Advances in MCMC". In *Highly Structured Stochastic Systems* (eds, Green, P.G., Hjört, N., and Richardson, S.). Chapman and Hall, London.

Robert, C.P., and Casella, G. (1999) *Monte Carlo Statistical Methods*. Springer-Verlag, New York.

Warnes, G.R. (2000) The Normal Kernel Coupler. Ph.D. thesis, U. Washington, Seattle.