



Comments on “Particle Markov Chain
Monte Carlo” by
C. Andrieu, A. Doucet and R. Holtenstein

J. Cornebise and G.W. Peters

Technical Report #2009-7
November 18, 2009

This material was based upon work partially supported by the National Science Foundation under Grant DMS-0635449 to the Statistical and Applied Mathematical Sciences Institute. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Statistical and Applied Mathematical Sciences Institute
PO Box 14006
Research Triangle Park, NC 27709-4006
www.samsi.info

Comments on "Particle Markov Chain Monte Carlo" by C. Andrieu, A. Doucet and R. Holenstein

J. Cornebise* and G.W. Peters†‡

November 18, 2009

Abstract

We merge in this note our two discussions about the Read Paper “Particle Markov chain Monte Carlo” ([Andrieu, Doucet, and Holenstein, 2010](#)) presented on October 16th 2009 at the Royal Statistical Society, appearing in the Journal of the Royal Statistical Society Series B. We also present a more detailed version of the ABC extension.

1 Introduction

The article [Andrieu et al. \(2010\)](#) is clearly going to have significant impact on scientific disciplines with a strong interface with computational statistics and non-linear state space models. Our comments are based on practical experience with PMCMC implementation in latent process multifactor SDE models for commodities ([Peters et al., 2009](#)), wireless communications ([Nevat et al., 2009](#)) and population dynamics ([Hayes et al., 2009](#)), using Rao-Blackwellised particle filters ([Doucet et al., 2000](#)) and adaptive MCMC ([Roberts and Rosenthal, 2009](#)).

2 Generic comments

- From our implementations, ideal use cases consist of highly non-linear dynamic equations for a small dimension d_x of the state-space, large dimension d_θ of the static parameter, and potentially large length T of the time series. In our cases d_x was 2 or 3, d_θ up to 20, and T between 100 and 400.

*Statistical and Applied Mathematical Sciences Institute, P.O. Box 14006, Research Triangle Park, NC 27709-4006, USA, jcornebise@samsi.info

†School of Mathematics and Statistics, University of New South Wales, Sydney, NSW, 2052, Australia, garethpeters@unsw.edu.au

‡This material was based upon work supported by the National Science Foundation under Agreement No. DMS-0635449. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

- In PMH, non-adaptive MCMC proposals for θ (e.g. tuned according to pre-simulation chains or burn-in iterations) would be costly for large T , and requires to keep N fixed over the whole run of the Markov chain. Adaptive MCMC proposals such as the Adaptive Metropolis sampler (Roberts and Rosenthal, 2009), avoid such issues and proved particularly relevant for large d_θ and T , as can be seen in Figure 2.
- The Particle Gibbs (PG) could potentially stay frozen on a state $x_{1:T}(i)$. Consider a state space model with state transition function almost linear in x_n for some range of θ , from which $y_{1:T}$ is considered to result, and strongly non-linear elsewhere. If the PG samples $\theta(i)$ in those regions of strong non-linearity, the particle tree would likely coalesce on the trajectory preserved by the conditional SMC, leaving it with a high importance weight, maintaining $(\theta(i+1), x_{1:T}(i+1)) = (\theta(i), x_{1:T}(i))$ over several iterations. Using PMH within PG would help escape this region, especially using PRC and adaptive SMC kernels, outlined in another comment, to fight the degeneracy of the filter and the high variance of $\hat{p}_\theta(y_{1:T})$.

3 Adaptive Sequential Monte Carlo

Our comments on adaptive SMC relate to Particle marginal Metropolis-Hastings (PMMH) which has acceptance probability given in Equation (13) of the read paper for proposed state $(\theta^*, X_{1:T}^*)$, relying on the estimate $\hat{p}_{\theta^*}(y_{1:T}) = \prod_{n=1}^T \frac{1}{N} \sum_{k=1}^N w_n(x_{1:n}^{*,k})$. Although a small N suffices to approximate the mode of joint path space distribution, producing a reasonable proposal for $x_{1:T}$, it results in high variance estimates of $\hat{p}_{\theta^*}(y_{1:T})$. We study a population dynamics example from (Hayes et al., 2009, Model 3 excerpt), involving a log-transformed theta-logistic state space model, see (Wang, 2007, Equation 3(a), 3(b)) for parameter settings. PMCMC performance depends on the trade-off between degeneracy of the filter, N , and design of the SMC mutation kernel. Regarding the latter:

- A Rao-Blackwellised filter (Doucet et al., 2000) can improve acceptance rates, Nevat et al. (see 2009).
- Adaptive mutation kernels, which in PMCMC, can be considered as adaptive SMC proposals, can reduce degeneracy on the path space, allowing for higher dimensional state vectors x_n . Adaption can be local (within filter) or global (sampled Markov chain history). Though currently particularly designed for ABC methods, the work of Peters et al. (2008) incorporates into the mutation kernel of SMC Samplers (Del Moral et al., 2006) the Partial Rejection Control (PRC) mechanism of Liu (2001), which is also beneficial for PMCMC. PRC adaption reduces degeneracy by rejecting a particle mutation when its incremental importance weight is below a threshold c_n . The PRC mutation kernel

$$q_\theta^*(x_n|y_n, x_{n-1}) = r(c_n, x_{n-1})^{-1} \min \left[1, W_{n-1}(x_{n-1}) \frac{w_n(x_{n-1}, x_n)}{c_n} \right] q_\theta(x_n|y_n, x_{n-1}), \quad (1)$$

can also be used in PMH, where $q_\theta(x_n|y_n, x_{n-1})$ is the standard SMC proposal, and

$$r(c_n, x_{n-1}) = \int \min \left[1, W_{n-1}(x_{n-1}) \frac{w_n(x_{n-1}, x_n)}{c_n} \right] q_\theta(x_n|y_n, x_{n-1}) dx_n. \quad (2)$$

As presented in [Peters et al. \(2008\)](#), algorithmic choices for $q_\theta^*(x_n|y_n, x_{n-1})$ can avoid evaluation of $r(c_n, x_{n-1})$. [Cornebise \(2009b\)](#) extend this work, developing PRC for Auxiliary SMC samplers, also useful in PMH. Threshold c_n can be set adaptively: locally either at each SMC mutation or Markov chain iteration; or globally based on chain acceptance rates. Additionally, c_n can be set adaptively via quantile estimates of pre-PRC incremental weights, see [Peters et al. \(2009\)](#).

- [Cornebise et al. \(2008\)](#) state that adaptive SMC proposals can be designed by minimizing function-free risk theoretic criteria such as Kullback-Leibler divergence between a joint proposal in a parametric family and a joint target. [Cornebise \(2009a, Chapter 5\)](#) and [Cornebise et al. \(2009\)](#) use a mixture of experts, adapting kernels of a mixture on distinct regions of the state-space separated by a softmax partition. These results extend to PMCMC settings.

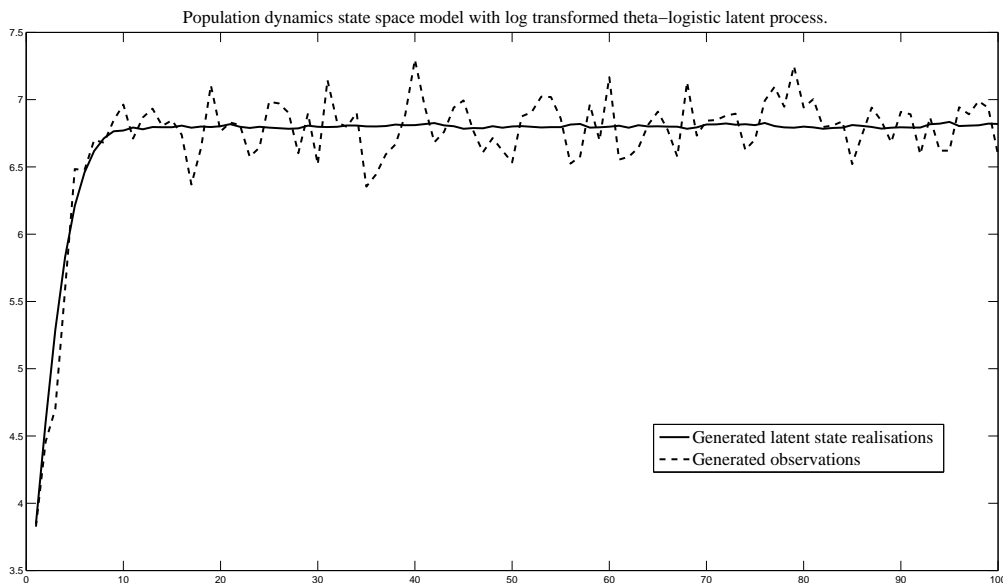


Figure 1: Sequence of simulated states and observations for the population dynamic log-transformed theta logistic model from [Wang \(2007\)](#), with static parameter $\theta = (r, \zeta, K)$ under constraints $K > 0$, $r < 2.69$, $\zeta \in \mathbb{R}$. State transition is $f_\theta(x_n|x_{n-1}) = \mathcal{N}(x_n; x_{n-1} + r(1 - (\exp(x_{t-1})/K)^\zeta), 0.01)$, and local likelihood is $g_\theta(y_n|x_n) = \mathcal{N}(y_n; x_n, 0.04)$, for $T = 100$ timesteps.




Figure 2: Path of three sampled latent states x_2 , x_{37} , x_{93} , and of the sampled parameters $\theta = (r, \zeta, L)$, over 100,000 PMH iterations based on $N = 200$ particles using a simple SIR filter – the one dimensional state did not call for Rao-Blackwellisation. Note also the effect the Adaptive MCMC proposal for θ , set-up to start at iteration 5,000, particularly visible on the mixing of parameter K . The most noticeable property of the algorithm is the remarkable mixing of the chain, in spite of the high total dimension of the sampled state: each iteration involves a proposal of $(X_{1:T}, \theta)$ of dimension 103.

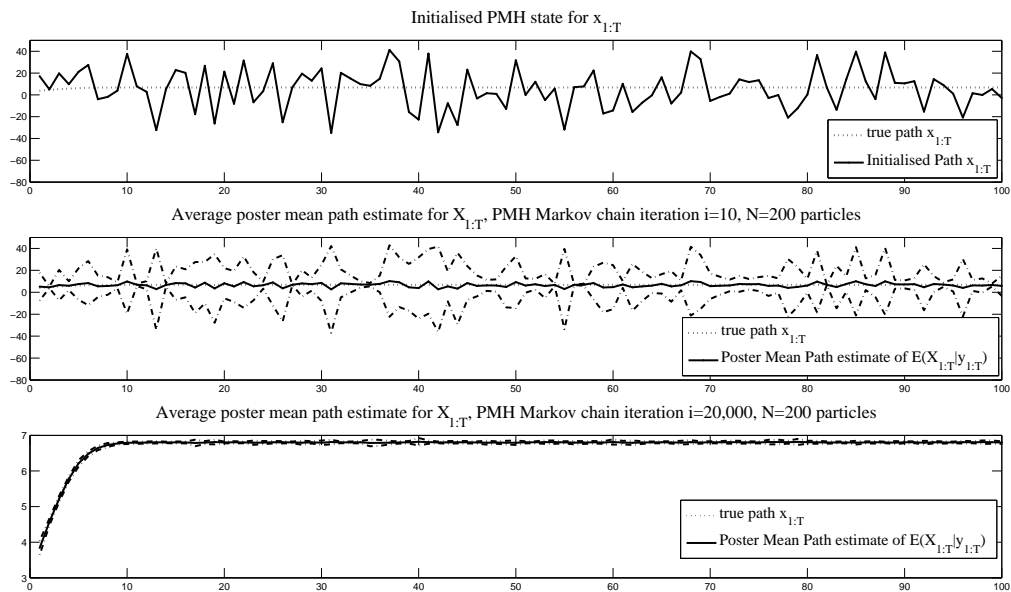


Figure 3: Convergence of the distribution of the path of latent states $x_{1:T}$. Note the change of vertical scale. Initializing PMH on a very unlikely initial path does not prevent the MMSE estimate of the latent states converging: as few as 10 PMH iterations already begins to concentrate the sampled paths around the true path – assumed here to be close to the mode of the posterior distribution thanks to the small observation noise –, with very satisfactory results after 20,000 iterations.

4 Approximate Bayesian Computation and PMCMC

For intractable joint likelihood $p_\theta(y_{1:T}|x_{1:T})$, we could design a SMC-ABC algorithm (see e.g. [Peters et al., 2008](#); [Ratmann, 2010](#), Chapter 1) for a fixed ABC tolerance ϵ , using the approximations

$$\hat{p}_\theta^{ABC}(y_{1:T}) := \frac{1}{N} \sum_{k=1}^N \frac{\frac{1}{S} \sum_{s=1}^S \mathbb{I}(\rho(y_1^k(s), y_1) < \epsilon) \mu_\theta(x_1^k)}{q_\theta(x_1^k|y_1)} \times \prod_{n=2}^T \left(\frac{1}{N} \sum_{k=1}^N \frac{\frac{1}{S} \sum_{s=1}^S \mathbb{I}(\rho(y_n^k(s), y_n) < \epsilon) f_\theta(x_n^k|x_{n-1}^{A_{n-1}^k})}{q_\theta(x_n^k|y_n, x_{n-1}^{A_{n-1}^k})} \right)$$

or

$$\hat{p}_\theta^{ABC}(y_{1:T}) := \frac{1}{N} \sum_{k=1}^N \frac{\frac{1}{S} \sum_{s=1}^S \mathcal{N}(y_1^k(s); y_1, \epsilon^2) \mu_\theta(x_1^k)}{q_\theta(x_1^k|y_1)} \times \prod_{n=2}^T \left(\frac{1}{N} \sum_{k=1}^N \frac{\frac{1}{S} \sum_{s=1}^S \mathcal{N}(y_n^k(s); y_n, \epsilon^2) f_\theta(x_n^k|x_{n-1}^{A_{n-1}^k})}{q_\theta(x_n^k|y_n, x_{n-1}^{A_{n-1}^k})} \right)$$

with ρ a distance on the observation space and $y_n^k(s) \sim g_\theta(\cdot|x_n^k)$ simulated observations. Additional degeneracy on the path space induced by ABC approximation should be controlled, e.g. with PRC ([Peters et al., 2008](#)), see Equation (1). More details on this algorithm are available in Appendix A, which is not contained in our comment to JRSSB due to space restrictions.

A Algorithmic details of ABC filtering within PMCMC

Here we expand on the comment we made above in which we approximated the local likelihood $g_\theta(y_n|x_n)$ of the SMC-based filtering part of the PMCMC algorithm. by the ABC approximation

$$g_\theta^{ABC}(y_n|x_n, y_n(1:S), \epsilon) := \frac{1}{S} \sum_{s=1}^S \pi_\theta(y_n(s)|x_n, y_n, \epsilon) \quad (3)$$

where possible choices for π_θ are

$$\pi_\theta^{\mathbb{I}}(y_n(s)|x_n, y_n, \epsilon) := \mathbb{I}(\rho(y_n(s), y_n) < \epsilon) \text{ or } \pi_\theta^{\mathcal{N}}(y_n(s)|x_n, y_n, \epsilon) := \mathcal{N}(y_n(s); y_n, \epsilon^2)$$

with ρ a distance on the observation space and $y_n(s) \sim g_\theta(\cdot|x_n)$ simulated observations – assumed here to be univariate for sake of brevity, but generalisation to multivariate setting and summary statistics is straightforward.

We note it is critical in the filtering context to ensure that the particle system under approximation does not collapse into uniformly null incremental weights $w_n(X_{1:n}^k) = 0$ which

Algorithm 1 SMC-ABC-PRC filtering algorithm targeting $p_\theta(x_{1:T}|y_{1:T})$ as required in Step 2(b) of the PMMH of (Andrieu et al., 2010, Section 2.4.2). Replaces the SMC algorithm presented in (Andrieu et al., 2010, Section 2.2.1). Approximation g_θ^{ABC} is defined in Equation (3) and function r in Equation (2).

Step 1: Initialize ϵ and c_1

Step 2: At time $n = 1$,

(a) for $k = 1, \dots, N$

(i) sample $X_1^k \sim q_\theta(\cdot|y_1)$

(ii) sample $Y_1^k(s) \sim g_\theta(\cdot|X_1^k)$ for $s = 1, \dots, S$

(iii) compute the incremental weight

$$\tilde{w}_1(X_1^k) := \frac{\mu_\theta(X_1^k)g_\theta^{ABC}(y_1|X_1^k, Y_1^k(1:S), \epsilon)}{q_\theta(X_1^k|y_1)},$$

(iv) with probability $1 - p_1^k = 1 - \min\{1, \tilde{w}_1(X_1^k)/c_1\}$, reject X_1^k and go to (i)

(v) otherwise, accept X_1^k and set

$$w_1(X_1^k) = \tilde{w}_1(X_1^k)r(c_1)/p_1^k$$

(b) normalise the weights $W_1^k := w_1(X_1^k) / \sum_{m=1}^N w_1(X_1^m)$.

Step 3: At times $n = 2, \dots, T$,

(a) possibly adapt c_n online

(b) for $k = 1, \dots, N$

(i) sample $A_{n-1}^k \sim \mathcal{F}(\cdot|\mathbf{W}_{n-1})$,

(ii) sample $X_n^k \sim q_\theta(\cdot|y_n, X_{n-1}^{A_{n-1}^k})$ and set $X_{1:n}^k = (X_{n-1}^{A_{n-1}^k}, X_n^k)$, and

(iii) sample $Y_n^k(s) \sim g_\theta(\cdot|X_n^k)$ for $s = 1, \dots, S$

(iv) compute the incremental weight

$$\tilde{w}_n(X_{1:n}^k) := \frac{f_\theta(X_n^k|X_{n-1}^{A_{n-1}^k})g_\theta^{ABC}(y_n|X_n^k, Y_n^k(1:S), \epsilon)}{q_\theta(X_n^k|y_n, X_{n-1}^{A_{n-1}^k})},$$

(v) with probability $1 - p_n^k = 1 - \min\{1, \tilde{w}_n(X_{1:n}^k)/c_n\}$, reject X_n^k and go to (ii)

(vi) otherwise, accept X_n^k and set

$$w_n(X_{1:n}^k) = \tilde{w}_n(X_{1:n}^k)r(c_n, X_{n-1}^{A_{n-1}^k})/p_n^k$$

(c) normalise the weights $W_n^k := w_n(X_{1:n}^k) / \sum_{m=1}^N w_n(X_{1:n}^m)$.

may occur for π_{θ}^{\parallel} at any stage of the filtering during each PMCMC iteration, especially for small tolerances ϵ . The PRC mutation kernel $q_{\theta}^*(x_n|y_n, x_{n-1})$ defined in Equation (1) is critical to overcome both this collapse and the additional degeneracy on the path space introduced by the ABC approximation. The algorithm presented in McKinley et al. (2009, Sections 3.4 and 3.5.1) is a special case of the SMC samplers PRC-ABC algorithm of Peters et al. (2008) in which the PRC rejection threshold $c_n = 0$, the mutation kernel is global and resampling is performed at each stage of the filter, which avoids the computation of the normalizing constant $r(c_n, x_{n-1})$ defined in Equation (2). We further note that the work of Cornebise (2009a) casts the SMC sampler PRC algorithm (Peters et al., 2008) with rejection of the ancestor index A_{n-1}^k – here in Step 3.(a).(v) – into an Auxiliary SMC sampler framework. The combination of these two concepts recovers a generalized version of McKinley et al. (2009), see Algorithm 1, which has advantage in the PMCMC setting of allowing for adaptation of the threshold c_n .

References

- Andrieu, C., A. Doucet, and R. Holenstein (2010). Particle Markov chain Monte Carlo methods. *J. R. Statist. Soc. B* 72(2), 1–33.
- Cornebise, J. (2009a). *Adaptive Sequential Monte Carlo Methods*. Ph. D. thesis, University Pierre and Marie Curie – Paris 6.
- Cornebise, J. (2009b). Auxiliary SMC samplers with applications to PRC and ABC. Working paper, Statistical and Applied Mathematical Sciences Institute.
- Cornebise, J., E. Moulines, and J. Olsson (2008). Adaptive methods for sequential importance sampling with application to state space models. *Statistics and Computing* 18(4), 461–480.
- Cornebise, J., E. Moulines, and J. Olsson (2009). Adaptive sequential Monte Carlo by means of mixture of experts. Working paper, Telecom ParisTech.
- Del Moral, P., A. Doucet, and A. Jasra (2006). Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 68(3), 411–436.
- Doucet, A., N. de Freitas, K. Murphy, and S. Russell (2000). Rao-Blackwellised particle filtering for dynamic Bayesian networks. In *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, pp. 176–183.
- Hayes, K., G. Hosack, and G. Peters (2009). Searching for the allee effect in latent state space models via adaptive particle Markov chain Monte Carlo. Working paper, Department of Mathematics and Statistics, University of New South Wales.
- Liu, J. (2001). *Monte Carlo Strategies in Scientific Computing*. New York: Springer.
- McKinley, T., A. Cook, and R. Deardon (2009). Inference in epidemic models without likelihoods. *The International Journal of Biostatistics* 5(1), 24.

- Nevat, I., G. Peters, and A. Doucet (2009). Channel tracking for relay networks via adaptive particle MCMC. Working paper, University of New South Wales.
- Peters, G., M. Briers, P. Shevchenko, and A. Doucet (2009). Online calibration and filtering for multi factor commodity models with seasonality: incorporating futures and options contracts. Working paper, Department of Mathematics and Statistics, University of New South Wales.
- Peters, G., Y. Fan, and S. Sisson (2008). On sequential Monte Carlo, partial rejection control and approximate Bayesian computation. Technical report, Department of Mathematics and Statistics, University of New South Wales.
- Peters, G., Y. Fan, and S. Sisson (2009). Likelihood-free Bayesian inference for α -stable models. Technical report, Department of Mathematics and Statistics, University of New South Wales.
- Ratmann, O. (2010). *Approximate Bayesian Computation under model uncertainty, with an application to stochastic processes of network evolution*. Ph. D. Thesis, to appear, Imperial College.
- Roberts, G. and J. Rosenthal (2009). Examples of adaptive MCMC. *Journal of Computational and Graphical Statistics* 18(2), 349–367.
- Wang, G. (2007). On the latent state estimation of nonlinear population dynamics using bayesian and non-Bayesian state-space models. *Ecological Modelling* 200(3-4), 521–528.