

Efficient Monte Carlo sampling by parallel marginalization

Jonathan Weare[†]

Department of Mathematics, University of California, Berkeley, CA 94720

Communicated by Alexandre J. Chorin, University of California, Berkeley, CA, June 8, 2007 (received for review January 12, 2007)

Markov chain Monte Carlo sampling methods often suffer from long correlation times. Consequently, these methods must be run for many steps to generate an independent sample. In this paper, a method is proposed to overcome this difficulty. The method utilizes information from rapidly equilibrating coarse Markov chains that sample marginal distributions of the full system. This is accomplished through exchanges between the full chain and the auxiliary coarse chains. Results of numerical tests on the bridge sampling and filtering/smoothing problems for a stochastic differential equation are presented.

Markov chain Monte Carlo | renormalization | multi-grid | filtering | parameter estimation

To understand the behavior of a physical system, it is often necessary to generate samples from complicated high dimensional distributions. The usual tools for sampling from these distributions are Markov chain Monte Carlo (MCMC) methods by which one constructs a Markov chain whose trajectory averages converge to averages with respect to the distribution of interest. For some simple systems, it is possible to construct Markov chains with independent values at each step. In general, however, spatial correlations in the system of interest result in long correlation times in the Markov chain and hence slow convergence of the chain's trajectory averages. Here, a method is proposed to alleviate the difficulties caused by spatial correlations in high dimensional systems. The method, parallel marginalization, is tested on two stochastic differential equation conditional path sampling problems.

Parallel marginalization takes advantage of the shorter correlation lengths present in marginal distributions of the target density. Auxiliary Markov chains that sample approximate marginal distributions are evolved simultaneously with the Markov chain that samples the distribution of interest. By swapping their configurations, these auxiliary chains pass information between themselves and with the chain sampling the original distribution. As shown below, these swaps are made in a manner consistent with both the original distributions and the approximate marginal distributions. The numerical examples indicate that improvement in efficiency of parallel marginalization over standard MCMC techniques can be significant.

The design of efficient methods to approximate marginal distributions was addressed by Chorin (1) and Stinis (2). The use of Monte Carlo updates on coarse subsets of variables is not a new concept (see ref. 3 and the references therein). The method presented in ref. 3 does not use marginal distributions. However, attempts have been made previously to use marginal distributions to accelerate the convergence of MCMC (see refs. 4 and 5). In contrast to parallel marginalization, the methods proposed in refs. 4 and 5 do not preserve the distribution of the full system and therefore are not guaranteed to converge. The parallel construction used here is motivated by the parallel tempering method (see ref. 6), and allows efficient comparison of the auxiliary chains and the original chain. See refs. 6 and 7 for expositions of standard MCMC methods.

Parallel marginalization for problems in Euclidean state spaces is described in detail below. In the final sections, the

conditional path sampling problem is described and numerical results are presented for the bridge sampling and smoothing/filtering problems.

Parallel Marginalization

For the purposes of the discussion in this section, we assume that appropriate approximate marginal distributions are available. As discussed in a later section, they may be provided by coarse models of the physical problem as in the examples below, or they may be calculated via the methods in refs. 1 and 2.

Assume that the d_0 dimensional system of interest has a probability density, $\pi_0(x_0)$, where $x_0 \in \mathbb{R}^{d_0}$. Suppose further that, by the Metropolis–Hastings or any other method (see ref. 6), we can construct a Markov chain, $Y_0^n \in \mathbb{R}^{d_0}$, which has π_0 as its stationary measure. That is, for two points $x_0, y_0 \in \mathbb{R}^{d_0}$

$$\int T_0(x_0 \rightarrow y_0) \pi_0(x_0) dx_0 = \pi_0(y_0),$$

where $T_0(x_0 \rightarrow y_0)$ is the probability density of a move to $\{Y_0^{n+1} = y_0\}$ given that $\{Y_0^n = x_0\}$. Here, n is the algorithmic step. Under appropriate conditions (see ref. 6), averages over a trajectory of $\{Y_0^n\}$ will converge to averages over π_0 , i.e. for an objective function $g(x_0)$

$$\frac{1}{N} \sum_{n=0}^{N-1} g(Y_0^n) \rightarrow \mathbf{E}[g(X_0)].$$

The size of the error in the above limit decreases as the rate of decay of the time autocorrelation

$\text{Corr}[g(Y_0^n), g(Y_0^0)]$

$$= \frac{\mathbf{E}[(g(Y_0^n) - \mathbf{E}[g(X_0)])(g(Y_0^0) - \mathbf{E}[g(X_0)])]}{\text{Var}[g(X_0)]}$$

increases. In this formula, Y_0^0 is assumed to be drawn from π_0 .

It is well known that judicious elimination of variables by renormalization can reduce long range spatial correlations (for example, see refs. 8 and 9). The variables are removed by averaging out their effects on the full distribution. If the original density is $\pi(\hat{x}, \bar{x})$ and we wish to remove the \bar{x} variables, the distribution of the remaining \hat{x} variables is given by the marginal density (see refs. 1 and 6)

$$\bar{\pi}(\hat{x}) = \int \pi(\hat{x}, \bar{x}) d\bar{x}. \quad [1]$$

Author contributions: J.W. designed research, performed research, contributed new reagents/analytic tools, analyzed data, and wrote the paper.

The author declares no conflict of interest.

Abbreviation: MCMC, Markov chain Monte Carlo.

[†]E-mail: weare@math.berkeley.edu.

© 2007 by The National Academy of Sciences of the USA

The full distribution can be factored as

$$\pi(\hat{x}, \bar{x}) = \bar{\pi}(\hat{x})\pi(\bar{x}|\hat{x}),$$

where $\pi(\bar{x}|\hat{x})$ is the conditional density of \bar{x} given \hat{x} . Because they exhibit shorter correlation lengths, the marginal distributions are useful in the acceleration of MCMC methods.

With this in mind, we consider a collection of lower dimensional Markov chains $Y_i^n \in \mathbb{R}^{d_i}$, which have stationary distributions $\pi_i(x_i)$ where $d_0 > \dots > d_i$. For each $i \leq L$, let T_i be the transition probability density of Y_i^n , i.e. $T_i(x_i \rightarrow y_i)$ is the probability density of $\{Y_i^{n+1} = y_i\}$ given that $\{Y_i^n = x_i\}$. The $\{\pi_i\}$ are approximate marginal distributions. For example, divide the x_i variables into two subsets, $\hat{x}_i \in \mathbb{R}^{d_{i+1}}$ and $\bar{x}_i \in \mathbb{R}^{d_i - d_{i+1}}$, so that $x_i = (\hat{x}_i, \bar{x}_i)$. The \bar{x}_i variables represent the variables of x_i that are removed by marginalization, i.e.,

$$\pi_{i+1}(\hat{x}_i) \approx \int \pi_i(\hat{x}_i, \bar{x}_i) d\bar{x}_i.$$

After arranging these chains in parallel, we have the larger process

$$Y^n = (Y_0^n, \dots, Y_L^n) \in \mathbb{R}^{d_0} \times \dots \times \mathbb{R}^{d_L}.$$

The probability density of a move to $\{Y^{n+1} = y\}$ given that $\{Y^n = x\}$ for $x, y \in \mathbb{R}^{d_0} \times \dots \times \mathbb{R}^{d_L}$ is given by

$$T(x \rightarrow y) = \prod_{i=0}^L T_i(x_i \rightarrow y_i).$$

Because

$$\int \left(T(x \rightarrow y) \prod_{i=0}^L \pi_i(x_i) \right) dx_0 \dots dx_L = \prod_{i=0}^L \pi_i(y_i),$$

the stationary distribution of Y^n is

$$\prod(x_0, \dots, x_L) = \pi_0(x_0) \dots \pi_L(x_L).$$

The next step in the construction is to allow interactions between the chains $\{Y_i^n\}$ and to thereby pass information from the rapidly equilibrating chains on the lower dimensional spaces (large i) down to the chain on the original space ($i = 0$). This is accomplished by swap moves. In a swap move between levels i and $i + 1$, we take a d_{i+1} dimensional subset, \hat{x}_i , of the x_i variables and exchange them with the x_{i+1} variables. The remaining $d_i - d_{i+1}$ \bar{x}_i variables are resampled from the conditional distribution $\pi_i(\bar{x}_i|x_{i+1})$. For the full chain, this swap takes the form of a move from $\{Y^n = x\}$ to $\{Y^{n+1} = y\}$ where

$$x = (\dots, \hat{x}_i, \bar{x}_i, x_{i+1}, \dots)$$

and

$$y = (\dots, x_{i+1}, \bar{y}_i, \hat{x}_i, \dots).$$

The ellipses represent components of Y^n that remain unchanged in the transition and \bar{y}_i is drawn from $\pi_i(\bar{x}_i|x_{i+1})$.

If these swaps are undertaken unconditionally, the resulting chain will equilibrate rapidly, but will not, in general, preserve the product distribution Π . To remedy this, we introduce the swap acceptance probability

$$A_i = \min \left\{ 1, \frac{\bar{\pi}_i(x_{i+1}) \pi_{i+1}(\hat{x}_i)}{\bar{\pi}_i(\hat{x}_i) \pi_{i+1}(x_{i+1})} \right\}. \quad [3]$$

In this formula, $\bar{\pi}_i$ is the function on $\mathbb{R}^{d_{i+1}}$ resulting from marginalization of π_i as in Eq. 1. Given that $\{Y^n = x\}$, the probability density of $\{Y^{n+1} = y\}$, after the proposal and either acceptance with probability A_i or rejection with probability $1 - A_i$, of a swap move, is given by

$$S_i(x \rightarrow y) = (1 - A_i)\delta_{\{y=x\}} + A_i \pi_i(\bar{y}_i|x_{i+1}) \delta_{\{(\bar{y}_i, y_{i+1})=(\bar{x}_i, x_{i+1})\}} \prod_{j \neq \{i, i+1\}} \delta_{\{y_j=x_j\}}$$

for $x, y \in \mathbb{R}^{d_0} \times \dots \times \mathbb{R}^{d_L}$. δ is the Dirac delta function.

We have the following lemma.

Lemma 1. *The transition probabilities S_i satisfy the detailed balance condition for the measure Π , i.e.,*

$$\prod(x)S_i(x \rightarrow y) = \prod(y)S_i(y \rightarrow x)$$

where $x, y \in \mathbb{R}^{d_0} \times \dots \times \mathbb{R}^{d_L}$.

The detailed balance condition stipulates that the probability of observing a transition $x \rightarrow y$ is equal to that of observing a transition $y \rightarrow x$ and guarantees that the resulting Markov chain preserves the distribution Π . Therefore, under general conditions, averages over a trajectory of $\{Y^n\}$ will converge to averages over Π . Because

$$\pi_0(x_0) = \int \prod(x_0, \dots, x_L) dx_1 \dots dx_L,$$

we can calculate averages over π_0 by taking averages over the trajectories of the first d_0 components of Y^n .

“Exact” Approximation of Acceptance Probability

Note that formula 3 for A_i requires the evaluation of $\bar{\pi}_i$ at the points $\hat{x}_i, x_{i+1} \in \mathbb{R}^{d_{i+1}}$. Although the approximation of $\bar{\pi}_i$ by functions on $\mathbb{R}^{d_{i+1}}$ is in general a very difficult problem, its evaluation at a single point is often not terribly demanding. In fact, in many cases, including the examples in this paper, the \hat{x}_i variables can be chosen so that the remaining \bar{x}_i variables are conditionally independent given \hat{x}_i .

Despite these mitigating factors, the requirement that we evaluate $\bar{\pi}_i$ before we accept any swap is a little onerous. Fortunately, and somewhat surprisingly, this requirement is not necessary. In fact, standard strategies for approximating the point values of the marginals yield Markov chains that themselves preserve the target measure. Thus, even a poor estimate of the ratio appearing in Eq. 3 can give rise to a method that is exact in the sense that the resulting Markov chain will asymptotically sample the target measure.

To illustrate this point, we consider the following example of a swap move. Assume that the current position of the chain is $\{Y^n = x\}$ where

$$x = (\dots, \hat{x}_i, \bar{x}_i, x_{i+1}, \dots).$$

The following steps will result in either $\{Y^{n+1} = x\}$ or $\{Y^{n+1} = y\}$ where

$$y = (\dots, x_{i+1}, \bar{y}_i, \hat{x}_i, \dots)$$

and $\bar{y}_i \in \mathbb{R}^{d_i - d_{i+1}}$.

1. Let $v^0 = \bar{x}_i$ and let $v^j \in \mathbb{R}^{d_i - d_{i+1}}$ for $j = 1, \dots, M - 1$ be independent samples from $p_i(\cdot|\hat{x}_i)$, where $p_i(\cdot|\hat{x}_i)$ is a reference density conditioned by \hat{x}_i . For example, $p_i(\cdot|\hat{x}_i)$ could be a Gaussian approximation of $\pi_i(\bar{x}_i|\hat{x}_i)$. How p_i is chosen depends on the problem at hand (see numerical examples below). In general, $p_i(\cdot|\hat{x}_i)$ should be easily evaluated and independently

sampled, and it should “cover” $\pi_i(\cdot|\hat{x}_i)$ in the sense that areas of \mathbb{R}^{d_i} where $\pi_i(\cdot|\hat{x}_i)$ is not negligible should be contained in areas where $p_i(\cdot|\hat{x}_i)$ is not negligible.

2. Let $u^j \in \mathbb{R}^{d_i-d_{i+1}}$ for $j = 0, \dots, M-1$ be independent random variables sampled from $p_i(\cdot|x_{i+1})$ (recall that we are considering a swap of \hat{x}_i and x_{i+1} , which live in the same space). Note, that the $\{u^j\}$ variables depend on x_{i+1} , whereas the $\{v^j\}$ variables depend on \hat{x}_i .
3. Define the weights

$$w_v^j = \frac{\pi_i(\hat{x}_i, v^j)}{p_i(v^j|\hat{x}_i)} \text{ and } w_u^j = \frac{\pi_i(x_{i+1}, u^j)}{p_i(u^j|x_{i+1})}.$$

The choice of p_i made above affects the variance of these weights, and therefore the variance of the acceptance probability below.

4. Choose \tilde{y}_i from among the $\{u^j\}$ according to the multinomial distribution with probabilities

$$P(\tilde{y}_i = u^j) = \frac{w_u^j}{\sum_{l=0}^{M-1} w_u^l}.$$

Note that \tilde{y}_i is an approximate sample from $\pi_i(\cdot|x_{i+1})$.

5. Set

$$Y^{n+1} = (\dots, x_{i+1}, \tilde{y}_i, \hat{x}_i, \dots)$$

with probability

$$A_i^M = \min \left\{ 1, \frac{\pi_{i+1}(\hat{x}_i) \sum_{j=0}^{M-1} w_u^j}{\pi_{i+1}(x_{i+1}) \sum_{j=0}^{M-1} w_v^j} \right\} \quad [4]$$

and

$$Y^{n+1} = Y^n = (\dots, \hat{x}_i, \tilde{x}_i, x_{x+1}, \dots)$$

with probability $1 - A_i^M$.

The transition probability density for the above swap move from $x \rightarrow y$ for $x, y \in \mathbb{R}^{d_0} \times \dots \times \mathbb{R}^{d_L}$ is given by

$$S_i^M(x \rightarrow y) = (1 - R)\delta_{\{y=x\}} + R\delta_{\{(\tilde{y}_i, y_{i+1})=(x_{i+1}, \hat{x}_i)\}} \prod_{j \notin \{i, i+1\}} \delta_{\{y_j=x_j\}},$$

where

$$R = M \int p_i(u^0|x_{i+1}) \frac{w_u^0}{\sum_{j=0}^{M-1} w_u^j} A_i^M \times \prod_{j=1}^{M-1} p_i(v^j|\hat{x}_i) p_i(u^j|x_{i+1}) dv^j du^j$$

and δ is again the Dirac delta function. In other words, S_i^M dictates that the Markov chain accepts the swap with probability R and rejects it with probability $1 - R$.

Although the preceding swap move corresponds to a method for approximating the ratio

$$\frac{\bar{\pi}_i(x_{i+1})}{\bar{\pi}_i(\hat{x}_i)}$$

appearing in the formula for A_i above, it also has some similarities with the multiple-try Metropolis method presented in ref. 10, which uses multiple suggestion samples to improve acceptance rates of standard MCMC methods. The following lemma is suggested by results in ref. 10.

Lemma 2. *The transition probabilities S_i^M satisfy the detailed balance condition for the measure Π .*

As before, the detailed balance condition guarantees that averages over trajectories of the first d_0 dimensions of Y^n will converge to averages over π_0 .

The A_i^M contain an approximation to the ratio of marginals in 3

$$\frac{\sum_{j=0}^{M-1} w_u^j}{\sum_{j=0}^{M-1} w_v^j} = \frac{\frac{1}{M} \sum_{j=0}^{M-1} \frac{\pi_i(x_{i+1}, u^j)}{p_i(u^j|x_{i+1})}}{\frac{1}{M} \sum_{j=0}^{M-1} \frac{\pi_i(\hat{x}_i, v^j)}{p_i(v^j|\hat{x}_i)}} \xrightarrow{M \rightarrow \infty} \frac{\mathbf{E}_{p_i} \left[\frac{\pi_i(x_{i+1}, \tilde{X}_i)}{p_i(\tilde{X}_i|x_{i+1})} \middle| \{\tilde{X}_i = x_{i+1}\} \right]}{\mathbf{E}_{p_i} \left[\frac{\pi_i(\hat{x}_i, \tilde{X}_i)}{p_i(\tilde{X}_i|\hat{x}_i)} \middle| \{\tilde{X}_i = \hat{x}_i\} \right]} = \frac{\bar{\pi}_i(x_{i+1})}{\bar{\pi}_i(\hat{x}_i)},$$

where \mathbf{E}_{p_i} denotes expectation with respect to the density p_i . When $0 < \mathbf{E}_{p_i} [w_v^j | \{\tilde{X}_i = \hat{x}_i\}] < \infty$, the convergence above follows from the strong law of large numbers and the fact that

$$\mathbf{E}_{p_i} \left[\frac{\pi_i(\tilde{X}_i, \tilde{X}_i)}{p_i(\tilde{X}_i|\tilde{X}_i)} \middle| \{\tilde{X}_i = \hat{x}_i\} \right] = \int \frac{\pi_i(\hat{x}_i, \tilde{x}_i)}{p_i(\tilde{x}_i|\hat{x}_i)} p_i(\tilde{x}_i|\hat{x}_i) d\tilde{x}_i = \int \pi_i(\hat{x}_i, \tilde{x}_i) d\tilde{x}_i = \bar{\pi}_i(\hat{x}_i).$$

For small values of M in 4, calculation of the swap acceptance probabilities is very cheap. However, higher values of M may improve the acceptance rates. For example, if the $\{\pi_i\}_{i>0}$ are exact marginals of π_0 , then $A_i \equiv 1$, whereas $A_i^M \leq 1$. Results similar to Lemma 2 hold when more general approximations replace the one given above; for example, when the $\{u^j\}$ and $\{v^j\}$ are generated by a Metropolis–Hastings rule. In practice, one has to balance the speed of evaluating A_i^M for small M with the possible higher acceptance rates for M large.

It is easy to see that a Markov chain, which evolves only by swap moves, cannot sample all configurations. These swap moves must therefore be used in conjunction with a transition rule that can reach any region of space, such as T from expression 2. More precisely, T should be Π -irreducible and aperiodic (see ref. 11). The transition rule for parallel marginalization is

$$P(x \rightarrow y) = (1 - \alpha)T(x \rightarrow y) + \alpha \int T(x \rightarrow z)S(z \rightarrow y)dz,$$

where

$$S(x \rightarrow y) = \sum_{k=0}^{L-1} \frac{1}{L} S_i^M(x \rightarrow y)$$

and $\alpha \in [0, 1)$ is the probability that a swap move occurs. P dictates that, with probability α , the chain attempts a swap move between levels I and $I + 1$, where I is a random variable chosen uniformly from $\{0, \dots, L - 1\}$. Next, each level of the chain evolves independently according to the $\{T_i\}$. With probability $1 - \alpha$, the chain does not attempt a swap move, but does evolve each level. The next result follows trivially from Lemma 2 and guarantees the invariance of Π under evolution by P .

Theorem 1. *The transition probability P satisfies the detailed balance condition for the measure Π , i.e.,*

$$\prod (x)P(x \rightarrow y) = \prod (y)P(y \rightarrow x)$$

where $x, y \in \mathbb{R}^{d_0} \times \dots \times \mathbb{R}^{d_N}$.

Thus, by combining standard MCMC steps on each component governed by the transition probability T , with swap steps between the components governed by S , an MCMC method results that not only uses information from rapidly equilibrating lower dimensional chains, but is also convergent.

Numerical Example 1: Bridge Path Sampling

In the bridge path sampling problem, we wish to approximate conditional expectations of the form

$$\mathbb{E}[g(Z^s) | \{Z^0 = z^-, \{Z^T = z^+\}],$$

where $s \in (0, T)$ and $\{Z^t\}$ is the real valued processes given by the solution of the stochastic differential equation

$$dZ^t = f(Z^t)dt + \sigma(Z^t)dW^t. \tag{5}$$

g, f , and σ are real valued functions of \mathbb{R} . Of course, we can also consider functions g of more than one time. This problem arises, for example, in financial volatility estimation. Because, in general, we cannot sample paths of **5**, we must first approximate $\{Z^t\}$ by a discrete process for which the path density is readily available. Let $t_0 = 0, t_1 = T/K, \dots, t_K = T$ be a mesh on which we wish to calculate path averages. One such approximate process is given by the linearly implicit Euler scheme (a balanced implicit method, see ref. 12),

$$\begin{aligned} X^{t_{k+1}} &= X^{t_k} + f(X^{t_k})\Delta \\ &+ (X^{t_{k+1}} - X^{t_k})f'(X^{t_k})\Delta + \sigma(X^{t_k})\sqrt{\Delta}\xi^k, \tag{6} \\ X^0 &= Z^0 \quad X^T = Z^T. \end{aligned}$$

The $\{\xi^k\}$ are independent Gaussian random variables with mean 0 and variance 1, and $\Delta = T/K$. K is assumed to be a power of 2. The choice of this scheme over the Euler scheme (see ref. 13) is due to its favorable stability properties, as explained later. Without the condition $X^{t_k} = Z^{t_k}$ above, generating samples of (X^0, \dots, X^{t_K}) is a relatively straightforward endeavor. One simply generates a sample of Z^0 , then evolves the system with this initial condition. However, the presence of information about $\{Z^t\}_{t>0}$ complicates the task. In general, a sampling method that requires only knowledge of a function proportional to conditional density of $(X^{t_1}, \dots, X^{t_{K-1}})$ must be applied. The approximate path density associated with discretization **6** is

$$\pi_0(x_0^{t_1}, \dots, x_0^{t_{K-1}} | x_0^0, x_0^{t_K}) \propto \exp\left(-\sum_{k=0}^{K-1} V(x_0^k, x_0^{k+1}, \Delta)\right) \tag{7}$$

where

$$V(x, y, \Delta) = \frac{[(1 - \Delta f'(x))(y - x) + \Delta f(x)]^2}{2\sigma^2(x)\Delta}.$$

At this point, we wish to apply the parallel marginalization sampling procedure to the density π_0 . However, as indicated above, a prerequisite for the use of parallel marginalization is the ability to estimate marginal densities. In some important problems, homogeneities in the underlying system yield simplifications in the calculation of these densities by the methods in refs. 1 and 2. These calculations can be carried out before implementation of parallel marginalization, or they can be integrated into the sampling procedure.

In some cases, the numerical estimation of the $\{\pi_i\}_{i>0}$ can be completely avoided. The examples presented here are two such cases. Let $S_i = \{0, 2^i, 3(2^i), 4(2^i), \dots, K\}$. Decompose S_i as $\hat{S}_i \cup \tilde{S}_i$ where

Table 1. Swap acceptance rates for bridge sampling and filtering/smoothing problems

Levels*	0/1	1/2	2/3	3/4	4/5	5/6	6/7	7/8	8/9
BS†	0.86	0.83	0.75	0.69	0.54	0.45	0.30	0.22	0.26
FS‡	0.86	0.83	0.74	0.65	0.46	0.23	0.04	NA	NA

NA, not applicable.
 *Swaps between levels i and $i + 1$.
 †Bridge sampling problem.
 ‡Filtering/smoothing problem.

$$\hat{S}_i = \{0, 2(2^i), 4(2^i), 6(2^i), \dots, K\}$$

and

$$\tilde{S}_i = \{2^i, 3(2^i), 5(2^i), 7(2^i), \dots, K - 2^i\}.$$

In the notation of the previous sections, $x_i = (\hat{x}_i, \tilde{x}_i)$, where $\hat{x}_i = \{x_i^{t_k}\}_{k \in \hat{S}_i \setminus \{0, K\}}$ and $\tilde{x}_i = \{x_i^{t_k}\}_{k \in \tilde{S}_i}$. In words, the hat and tilde variables represent alternating time slices of the path. For all i , fix $x_i^0 = z^-$ and $x_i^{t_K} = z^+$. We choose the approximate marginal densities

$$\pi_i(\{x_i^{t_k}\}_{k \in \hat{S}_i \setminus \{0, K\}} | x_i^0, x_i^{t_K}) \propto q_i(\{x_i^{t_k}\}_{k \in \hat{S}_i}),$$

where for each i , q_i is defined by successive coarsenings of **6**. That is,

$$q_i(\{x_i^{t_k}\}_{k \in \hat{S}_i}) = \exp\left(-\sum_{k=0}^{K/2^i-1} V(x_i^{t_{2^k}}, x_i^{t_{2^{k+1}}}, 2^i\Delta)\right).$$

Because π_i will be sampled using a Metropolis–Hastings method with x^0 and x^{t_K} fixed, knowledge of the normalization constants

$$Z_i(x_i^0, x_i^{t_K}) = \int q_i \prod_{k \in \hat{S}_i \setminus \{0, K\}} dx_i^{t_k}$$

is unnecessary.

Note from **7** that, conditioned on the values of $x^{t_{k-1}}$ and $x^{t_{k+1}}$, the variance of x^{t_k} is of order Δ . Thus, any perturbation of x^{t_k} that leaves x^{t_j} fixed for $j \neq k$ and is compatible with joint distribution **7** must be of the order $\sqrt{\Delta}$. This finding suggests that distributions defined by coarser discretizations of **7** will allow larger perturbations, and consequently will be easier to sample. However, it is important to choose a discretization that remains stable for large values of Δ . For example, although the linearly implicit Euler method performs well in the experiments below, similar tests using the Euler method were less successful due to limitations on the largest allowable values of Δ .

In this numerical example, bridge paths are sampled between time 0 and time 10 for a diffusion in a double-well potential

$$f(x) = -4x(x^2 - 1) \text{ and } \sigma(x) = 1.$$

The left and right end points are chosen as $z^- = z^+ = 0$. $\Delta = 2^{-10}$. $Y_i^n \in \mathbb{R}^{10/(2^i\Delta)+1}$ is the i th level of the parallel marginalization Markov chain at algorithmic time n . There are 10 chains ($L = 9$ in expression **2**). The observed swap acceptance rates are reported in Table 1. Let $Y_{mid}^n \in \mathbb{R}$ denote the midpoint of the path defined by Y_0^n (i.e., an approximate sample of the path at time 5). In Fig. 1, the autocorrelation of Y_{mid}^n

$$\text{Corr}[Y_{mid}^n, Y_{mid}^0]$$

is compared to that of a standard Metropolis–Hastings rule. In Fig. 1, the time scale of the autocorrelation for the Metropolis–

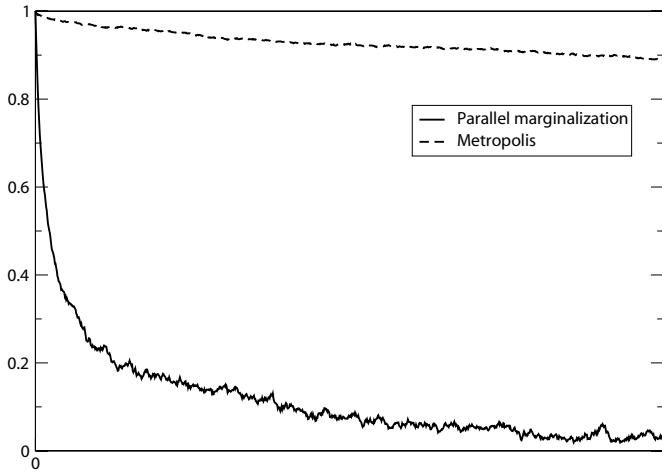


Fig. 1. Autocorrelation of parallel marginalization and standard Metropolis methods for bridge sampling problem.

Hastings method has been scaled by a factor of 1/10 to more than account for the extra computational time required per iteration of parallel marginalization. The relaxation time of the parallel chain is clearly reduced. In these numerical examples, the algorithm in the previous section is applied with a slight simplification. First generate M independent Gaussian random paths $\{\zeta^j(t_k)\}_{k \in \tilde{S}_i}$ with independent components $\zeta^j(t_k)$ of mean 0 and variance $2^{i-1}\Delta$. For each j and $k \in \tilde{S}_i$, let

$$u^j(t_k) = \zeta^j(t_k) + 0.5(x_{i+1}^{t_k-1} + x_{i+1}^{t_k+1}).$$

If in step 4, $\tilde{y}_i = u^{j^*}$, then in step 1 we set $v^0 = \tilde{x}_i$ and for each $k \in \tilde{S}_i$

$$\{v^j(t_k)\}_{j>0} = \{\zeta^j(t_k) + 0.5(\hat{x}_i^{t_k-1} + \hat{x}_i^{t_k+1})\}_{j \neq j^*}.$$

All other steps remain the same. This change yields a slightly faster though less generally applicable swap step that also preserves the density Π . Note that this modification implies that the reference density p_i is given by

$$p_i(\tilde{x}_i | \hat{x}_i) \propto \exp\left(\sum_{k \in \tilde{S}_i} -\frac{(\tilde{x}_i^{t_k} - 0.5(\hat{x}_i^{t_k-1} + \hat{x}_i^{t_k+1}))^2}{2^i \Delta}\right).$$

For this problem, the choice of M in 4, the number of samples of $\{u^j\}$ and $\{v^j\}$, seems to have little effect on the swap acceptance rates. In the numerical experiment $M = i + 1$ for swaps between levels i and $i + 1$.

Numerical Example 2: Nonlinear Smoothing/Filtering

In the nonlinear smoothing and filtering problem, we wish to approximate conditional expectations of the form

$$\mathbf{E}[g(Z^s) | \{H^j = h^j\}_0^s],$$

where $s \in (0, T)$ and the real valued processes $\{Z^t\}$ and $\{H^j\}$ are given by the system

$$\begin{aligned} dZ^t &= f(Z^t)dt + \sigma(Z^t)dW^{-t}, \\ H^j &= r(Z_j^S) + \chi^j, \\ Z^0 &\sim \rho, \quad X^n \sim i.i.d. \mu. \end{aligned}$$

$g, f, \sigma,$ and r are real valued functions of \mathbb{R} . The $\{X^j\}$ are real valued independent random variable drawn from the density μ




Fig. 2. Autocorrelation of parallel marginalization and standard Metropolis methods for filtering and smoothing problem.

and are independent of the Brownian motion $\{W^t\}$. $\{s_j\} \subset \{t_j\}$, and $0 = s_0 < s_1 < \dots < s_J = T$. The process Z^t is a hidden signal, and the $\{H^j\}$ are noisy observations.

Again, the system must first be discretized. The linearly implicit Euler scheme gives

$$\begin{aligned} X^{t_{k+1}} &= X^{t_k} + f(X^{t_k})\Delta + (X^{t_{k+1}} - X^{t_k})f'(X^{t_k})\Delta + \sigma(X^{t_k})\sqrt{\Delta}\xi^k, H^j \\ &= r(X^{s_j}) + X^j X^0 = Z^0 X^n \sim i.i.d. \mu. \end{aligned}$$

The $\{\xi^k\}$ are independent Gaussian random variables with mean 0 and variance 1, and $\Delta = T/K$. The $\{\xi^k\}$ are independent of the $\{X^j\}$. K is again assumed to be a power of 2.

The approximate path measure for this problem is

$$\begin{aligned} \pi_0(x_0^0, \dots, x_0^{t_k} | h^0, \dots, h^T) &\propto \exp\left(-\sum_{k=0}^{K-1} V(x_0^{t_k}, x_0^{t_{k+1}}, \Delta)\right) \\ &\times \rho(x_0^0) \prod_{n=0}^J \mu(x_0^{s_j} - r(h^{s_j})). \end{aligned}$$

The approximate marginals are chosen as

$$\pi_i(\{x_i^{t_k}\}_{k \in \tilde{S}_i} | h^0, \dots, h^T) \propto q_i(\{x_i^{t_k}\}_{k \in \tilde{S}_i}) \rho(x_i^{t_0}) \prod_{n=0}^J \mu(x_i^{s_j} - r(h^{s_j})),$$

where $V, q_i,$ and S_i are as defined in the previous section.

In this example, samples of the smoothed path are generated between time 0 and time 10 for the same diffusion in a double-well potential. The densities μ and ρ are chosen as

$$\mu = N(0, 0.01) \text{ and } \rho(x) \propto \exp(-(x^2 - 1)^2).$$

The observation times are $s_0 = 0, s_1 = 1, \dots, s_{10} = 10$ with $H^j = -1$ for $j = 0, \dots, 5$ and $H^j = 1$ for $j = 6, \dots, 10$. $\Delta = 2^{-10}$. There are eight chains ($L = 7$ in expression 2). The observed swap acceptance rates are reported in Table 1. Again, $Y_{mid}^n \in \mathbb{R}$ denotes the midpoint of the path defined by Y_0^n (i.e., an approximate sample of the path at time 5). In Fig. 2, the autocorrelation of Y_{mid}^n is compared to that of a standard Metropolis-Hastings rule. Fig. 2 has been adjusted as in the previous example. The relaxation time of the parallel chain is again clearly reduced. The algorithm is modified as in the previous example. For this problem, acceptable swap rates

require a higher choice of M in 4 than needed in the bridge sampling problem. In this numerical experiment $M = 2^i$ for swaps between levels i and $i + 1$.

Conclusion

A MCMC method has been proposed and applied to two conditional path sampling problems for stochastic differential equations. Numerical results indicate that this method, parallel marginalization, can have a dramatically reduced equilibration time when compared to standard MCMC methods.

Note that parallel marginalization should not be viewed as a stand-alone method. Other acceleration techniques, such as hybrid Monte Carlo, can and should be implemented at each

level within the parallel marginalization framework. As indicated by the smoothing problem, the acceptance probabilities at coarser levels can become small. The remedy for this is the development of more accurate approximate marginal distributions by, for example, the methods in refs. 1 and 2.

I thank Prof. A. Chorin for his guidance during this research, which was carried out during my Ph.D. studies at the University of California, Berkeley, and Dr. P. Okunev, Dr. P. Stinis, and the referees for their very helpful comments. This work was supported by the Director, Office of Science, Office of Advanced Scientific Computing Research, of the U.S. Department of Energy under Contract DE-AC03-76SF00098 and National Science Foundation Grant DMS0410110.

1. Chorin A (2003) *Multiscale Model Sim* 1:105–118.
2. Stinis P (2005) *J Comput Phys* 208:691–703.
3. Goodman J, Sokal A (1989) *Phys Rev D* 40:2035–2071.
4. Brandt A, Ron D (2001) *J Stat Phys* 102:163–186.
5. Okunev P (2005) Ph.D. thesis (Univ of California, Berkeley).
6. Liu J (2002) *Monte Carlo Strategies in Scientific Computing* (Springer, Berlin).
7. Binder K, Heermann D (2002) *Monte Carlo Simulation in Statistical Physics* (Springer, Berlin).
8. Binney J, Dowrick N, Fisher A, Newman M (1992) *The Theory of Critical Phenomena: An Introduction to the Renormalization Group* (Oxford Univ Press, New York).
9. Kadanoff L (1966) *Physics* 2:263.
10. Liu J, Liang F, Wong W (2000) *J Am Stat Assoc* 95:121–134.
11. Tierney L (1994) *Ann Stat* 22:1701–1728.
12. Milstein G, Platen E, Schurz H (1998) *SIAM J Numer Anal* 35:1010–1019.
13. Kloeden P, Platen E (1992) *Numerical Solution of Stochastic Differential Equations* (Springer, Berlin).