

On the utility of graphics cards to perform massively parallel simulation with advanced Monte Carlo methods

Anthony Lee ^{*} Christopher Yau [†] Michael B. Giles [‡] Arnaud Doucet [§]
Christopher C. Holmes [¶]

Abstract

We present a case-study on the utility of graphics cards to perform massively parallel simulation with advanced Monte Carlo methods. Graphics cards, containing multiple Graphics Processing Units (GPUs), are self-contained parallel computational devices that can be housed in conventional desktop and laptop computers. For certain classes of Monte Carlo algorithms they offer massively parallel simulation, with the added advantage over conventional distributed multi-core processors that they are cheap, easily accessible, easy to maintain, easy to code, dedicated local devices with low power consumption. On a canonical set of stochastic simulation examples including population-based Markov chain Monte Carlo methods and Sequential Monte Carlo methods, we find speedups from 35 to 500 fold over conventional single-threaded computer code. Our findings suggest that GPUs have the potential to facilitate the growth of statistical modelling into complex data rich domains through the availability of cheap and accessible many-core computation. We believe the speedup we observe should motivate wider use of parallelizable simulation methods and greater methodological attention to their design.

Keywords: Sequential Monte Carlo, Population-Based Markov Chain Monte Carlo, General Purpose Computation on Graphics Processing Units, Many-Core Architecture, Stochastic Simulation, Parallel Processing

1 Introduction

We describe a case-study in the utility of graphics cards involving Graphics Processing Units (GPUs) to perform local, dedicated, massively parallel stochastic simulation. GPUs were originally developed as dedicated devices to aid in real-time graphics rendering. However recently there has been

^{*}University of Oxford, Department of Statistics, 1 South Parks Road, Oxford OX1 3TG, UK, and Oxford-Man Institute, 9 Alfred Street, Oxford OX1 4EH, UK. Email: lee@stats.ox.ac.uk

[†]University of Oxford, Department of Statistics, 1 South Parks Road, Oxford OX1 3TG, UK. Email: yau@stats.ox.ac.uk

[‡]University of Oxford, Mathematical Institute, 24-29 St Giles, Oxford OX1 3LB, UK, and Oxford-Man Institute, 9 Alfred Street, Oxford OX1 4EH, UK. Email: mike.giles@maths.ox.ac.uk

[§]Institute of Statistical Mathematics, 4-6-7 Minami-Azabu, Minato-ku, Tokyo 106-8569, Japan. Email: arnaud@cs.ubc.ca

[¶]University of Oxford, Department of Statistics, 1 South Parks Road, Oxford OX1 3TG, UK, and Oxford-Man Institute, 9 Alfred Street, Oxford OX1 4EH, UK. Email: cholmes@stats.ox.ac.uk

an emerging literature on their use for scientific computing as they house multi-core processors. Examples include Stone et al. (2007) and Friedrichs et al. (2009), which discuss their use in molecular modelling and dynamics.

To gain an understanding of the potential benefits to statisticians we have investigated speedups on a canonical set of examples taken from the advanced Monte Carlo literature. These include Bayesian inference for a Gaussian mixture model computed using a population-based Markov chain Monte Carlo (MCMC) method and a sequential Monte Carlo (SMC) sampler and sequential Bayesian inference for a multivariate stochastic volatility model implemented using a standard SMC method, also known as a particle filter in this context. In these examples we report substantial speedups from the use of GPUs over conventional CPUs.

The potential of parallel processing to aid in statistical computing is well documented (see, for e.g. Kontogiorgos (2006)). However, previous studies have relied on distributed multi-core clusters of CPUs for implementation. In contrast, graphics cards for certain generic types of computation offer parallel processing speedups with advantages on a number of fronts, including:

- Cost: graphics cards are relatively cheap. At time of writing the cards used in our study retail at around 200 US dollars.
- Accessibility: graphics cards are readily obtainable from high street computer stores or over the internet.
- Maintenance: the devices are self-contained and can be hosted on conventional desktop and laptop computers.
- Speed: in line with multi-core CPU clusters, graphics cards offer significant speedup, albeit for a restricted set of class of scientific computing algorithms.
- Power: GPUs are low energy consumption devices compared to clusters of traditional computers, with a graphics card requiring around 200 Watts. While improvements in energy efficiency are application-specific, it is reasonable in many situations to expect a GPU to use around 10 per cent of the energy to that of an equivalent CPU cluster.
- Dedicated and local: the graphics cards slot into conventional computers offering the user ownership without the need to transport data externally.

The idea of splitting the computational effort of parallelizable algorithms amongst processors is certainly not new to statisticians. In fact, distributed systems and clusters of computers have been around for decades. Previous work on parallelization of MCMC methods on a group of networked computers include, among others, Rosenthal (2000) and Brockwell (2006). Rosenthal (2000) discusses how to deal with computers running at different speeds and potential computer failure while Brockwell (2006) discusses the parallel implementation of a standard single chain MCMC algorithm. The latency and bandwidth of communication in these systems make them suitable only in cases where communication between threads of computation is infrequent and low in volume. In other words, while many algorithms involve computation that could theoretically be distributed amongst processors, the overhead associated with distributing the work erases any speedup. In contrast, many-core processor communication has very low latency and very high

bandwidth. For many algorithms, this makes parallelization viable where it previously was not. In addition, the energy efficiency of a many-core computation compared to a single-core or distributed computation can be improved. This is because the computation can both take less time and require less overhead.

We choose to investigate the speed up for the simulation of random variates from complex distributions, a common computational task when performing inference using Monte Carlo. In particular, we focus on population-based MCMC methods and SMC methods for producing random variates as these are not algorithms that typically see significant speedup on clusters due to the need for frequent, high-volume communication between computing nodes. The algorithms are implemented for the Compute Unified Device Architecture (CUDA) and make use of graphics processing units (GPUs) which support this architecture. CUDA offers a fairly mature development environment via an extension to the C programming language. We estimate that a programmer proficient in C should be able to code effectively in CUDA within a few weeks of dedicated study. For our applications we use CUDA version 2.1 with an NVIDIA GTX 280 as well as an NVIDIA 8800 GT. The GTX 280 has 30 multiprocessors, each of which can run up to 1024 threads in parallel while the 8800 GT has 14 multiprocessors, each of which can support 768 active threads. At present, the retail price of the GTX 280 is just over double that of the 8800GT and it requires just over twice the power.

For the examples we consider, we find that computation time can be significantly lowered for all applications, and drastically lowered in some cases. This means that we can obtain the samples we want in seconds instead of hours and minutes instead of days. The types of speedup observed are dependent on the ability of the methods to be parallelized. In particular, speedup increases with the number of auxiliary distributions in population-based MCMC and the number of particles in SMC until some device-specific capacity is reached.

2 GPUs for parallel processing




Figure 1: Link between host and graphics card. The thicker lines represent higher data bandwidth while the squares represent processor cores.

GPUs have evolved into many-core processing units, currently with up to 30 multiprocessors per card, in response to commercial demand for real-time graphics rendering, independently of demand

for many-core processors in the scientific computing community. As such, the architecture of GPUs is very different to that of conventional central processing units (CPUs). An important difference is that GPUs devote proportionally more transistors to arithmetic logic units and less to caches and flow control in comparison to CPUs. This makes them less general purpose but highly effective for data-parallel computation with high arithmetic intensity, ie. computations where the same instructions are executed on different data elements and where the ratio of arithmetic operations to memory operations is high. This single instruction, multiple data (SIMD) architecture puts a heavy restriction on the types of computation that optimally utilize the GPU but in cases where the architecture is suitable it reduces overhead. Figure 1 gives a visualization of the link between a host machine and the graphics card, emphasizing the data bandwidth characteristics of the links and the number of processing cores.

CUDA provides the interface to compliant GPUs by extending the C programming language to allow users to define special functions, called kernels, which are executed in parallel on the GPU by a user-specified number of threads. Loosely speaking, each thread executes identical code but can use a unique identifier to do operations on different parts of the data. In addition, the framework allows threads to branch conditionally and end up executing different instructions but this can severely impact the time required for the threads to return as different instructions cannot be executed in parallel.

2.1 Implementation Details

```
__global__ void iw(int N, float* d_array, float* d_warray) {
    // get thread identifier
    const int tid = blockDim.x * blockIdx.x + threadIdx.x;
    // get total number of threads
    const int tt = blockDim.x * gridDim.x;
    int i;
    float p, q, x;
    for (i = tid; i < N; i += tt) {
        x = d_array[i];
        p = target_pdf(x, args_p);
        q = proposal_pdf(x, args_q);
        d_warray[i] = p / q;
    }
}
```

Figure 2: Kernel to evaluate an importance weight

Implementing code in CUDA requires programmers to define kernels that are run by blocks of threads on a GPU. All of the blocks together can be thought of as a grid. Example code for a kernel that computes an importance weight for an Importance Sampler is given in Figure 2 where we observe the C like syntax. In the code, `target_pdf` and `proposal_pdf` are device functions that evaluate the target and proposal densities respectively while `args_p` and `args_q` are arrays in constant device memory. It is clear that each thread will be executing the same code simultaneously if `N` is a multiple of `tt`, and so this code takes full advantage of the SIMD architecture of the GPU. The kernel is called on the host via the code

```
iw<<<nb,nt>>>(N, d_array, d_warray);
```

where `nb` and `nt` are the number of blocks and number of threads respectively. CUDA allows threads within the same block to synchronize and share dynamically allocated, fast memory. However, a block of threads cannot be split across multiprocessors on the graphics card. In most cases, the number of blocks and the number of threads per block is a detail for the application programmer and not a concern for those designing advanced Monte Carlo methods. Nevertheless, the thread hierarchy makes designers acutely aware of the benefit of writing kernels with minimal interaction between threads since this is relatively computationally expensive. In addition, the SIMD execution architecture makes certain mechanisms less attractive, an example being rejection sampling since all concurrent threads will need to wait for the slowest thread to complete. A final note is that due to the specialized nature of the hardware, kernel code must be written in a subset of the C programming language. For example, functions on the device cannot use function pointers, make recursive calls or have a variable number of arguments. In most cases, the restrictions are a direct consequence of the fact that device code is compiled into one large function since the device's control flow architecture does not support function calls.

2.2 Single precision issues

The current generation of GPUs is 4-8 times faster at single precision arithmetic than double precision. Although this ratio will decrease in the future, there will probably remain a factor 2 difference in speed, the same as for Intel CPUs when using SSE instructions. This raises the question of whether single precision arithmetic is adequate for statistical applications.

There are two particular areas in which care must be taken. The first concerns the much more limited range of single precision floating point numbers. Because of their 8-bit exponent, their magnitude must lie in the approximate range $[10^{-38}, 10^{+38}]$, whereas the magnitude of double precision variables is in the approximate range $[10^{-308}, 10^{+308}]$. Consequently, when working in single precision it is often necessary to work with log-likelihoods, rather than the likelihoods themselves though this is rarely a restriction for statisticians.

The second area of potential issues concerns the averaging of N floating point values, for $N \gg 1$. The simplest implementation uses an accumulator, to which the values are added one at a time. However, this may lead to a large increase in the error due to finite machine precision. When all of the values are of the same sign, the relative error is amplified by factor $O(N)$ in the worst case, and $O(\sqrt{N})$ in the more typical case where the rounding error at each step can be modelled as a random variable with zero mean. This behaviour is well understood Higham (1993, 2002) and the growth can be reduced to $O(\log N)$ by using a binary tree summation algorithm in which the values are summed in pairs, and then those new values are summed in pairs, and the process is repeated until a single value is obtained. This is the natural approach for the parallel implementation of a reduction operation. Example code is provided by NVIDIA on their CUDA website, and the implementation in our code is based on this.

Despite these concerns, single precision seems perfectly sufficient for the applications in this paper. The statistical variability due to the use of random numbers within the algorithms exceeds the perturbations due to finite machine precision.

2.3 Parallel random number generation

One important aspect of any Monte Carlo simulation is the generation of pseudorandom numbers. Fortunately, many uniform pseudorandom number generators can be implemented efficiently in parallel. The key idea is that each thread computes a contiguous block of numbers within a single overall stream. The thread can jump to the start of its block of numbers using a “skip-ahead” algorithm which enables it to skip n places in $O(\log n)$ operations (e.g. see L’Ecuyer et al. (2002)). The uniform pseudorandom numbers can then be transformed to match various different output distributions as needed. In our applications we use a parallelized version of the multiple recursive generator MRG32k3a presented in L’Ecuyer (1999) as well as a parallelized version of a xorshift random number generator Marsaglia (2003). In the case of the xorshift random number generator, more time must be spent to compute the seeds for each thread before any computation is done but the random number generation itself is faster and the initialization can be done offline.

3 Parallelizable Sampling Methods

There are an abundance of statistical problems that are essentially computational in nature, especially in Bayesian inference. In many such cases, the problem can be distilled into one of sampling from a probability distribution whose density π we can compute pointwise and up to a normalizing constant, ie. we can compute $\pi^*(\cdot)$ where $\pi(\mathbf{x}) = \pi^*(\mathbf{x})/Z$. A common motivation for wanting samples from π is so we can compute expectations of certain functions. If we denote such a function by ϕ , the expectation of interest is

$$I \triangleq \int_{\mathbf{x} \in \mathcal{X}} \phi(\mathbf{x}) \pi(\mathbf{x}) d\mathbf{x}$$

The Monte Carlo estimate of this quantity is given by

$$\hat{I}_{MC} \triangleq \frac{1}{N} \sum_{i=1}^N \phi(\mathbf{x}^{(i)})$$

where $\{\mathbf{x}^{(i)}\}_{i=1}^N$ are samples from π .

Clearly, we need samples from π in order to compute this estimate. In practice, however, we often cannot sample from π directly. There are two general classes of methods for dealing with this. The first are importance sampling methods, where we generate weighted samples from π by generating N samples according to some importance density γ proportional to γ^* and then estimating I via

$$\hat{I}_{IS} \triangleq \sum_{i=1}^N W^{(i)} \phi(\mathbf{x}^{(i)})$$

where $W^{(i)}$ are normalized importance weights

$$W^{(i)} = \frac{w(\mathbf{x}^{(i)})}{\sum_{j=1}^N w(\mathbf{x}^{(j)})} \text{ and } w(\mathbf{x}^{(i)}) = \frac{\pi^*(\mathbf{x}^{(i)})}{\gamma^*(\mathbf{x}^{(i)})}$$

The asymptotic variance of this estimate is given by $C(\phi, \pi, \gamma)/N$, ie. a constant over N . For many problems, however, it is difficult to come up with an importance density γ such that $C(\phi, \pi, \gamma)$ is small enough for us to attain reasonable variance with practical values of N .

The second general class of methods are MCMC methods, in which we construct an ergodic π -stationary Markov chain sequentially. Once the chain has converged, we can use all the dependent samples to estimate I . The major issue with MCMC methods is that their convergence rate can be prohibitively slow in some applications.

There are many ways to parallelize sampling methods that are not the focus of this work. For example, naive importance sampling, like classical Monte Carlo, is intrinsically parallel. Therefore, in applications where we have access to a good importance density γ we can get linear speedup with the number of processors available. Similarly, in cases where MCMC converges rapidly we can parallelize the estimation of I by running separate chains on each processor. While these situations are hoped for, they are not particularly interesting from a parallel architecture standpoint because they can run equally well in a distributed system. Finally, this paper is not concerned with problems for which the computation of individual kernel moves or importance weights are very expensive but themselves parallelizable. While the increased availability of parallel architectures will almost certainly be of help in such cases, the focus here is on potential speedups by parallelizing general sampling methods.

Much work in recent years has gone into dealing with the large constants in the variance of importance sampling estimates and slow convergence rates in MCMC and it is in these ‘advanced’ Monte Carlo methods that we direct our interest. This is mainly because while they are parallelizable, they are not trivially so and stand to benefit enormously from many-core architectures. In the remainder of this section we briefly review three such methods: population-based MCMC, SMC and SMC samplers.

3.1 Population-Based MCMC

A common technique in facilitating sampling from a complex distribution π with support in \mathcal{X} is to introduce an auxiliary variable $\mathbf{a} \in \mathcal{A}$ and sample from a higher-dimensional distribution $\bar{\pi}$ with support in the joint space $\mathcal{A} \times \mathcal{X}$, such that $\bar{\pi}$ admits π as a marginal distribution. With such samples, one can discard the auxiliary variables and be left with samples from π .

This idea is utilized in population-based MCMC, which attempts to speed up convergence of an MCMC chain for π by instead constructing a Markov chain on a joint space \mathcal{X}^M using $M - 1$ auxiliary variables each in \mathcal{X} . In general, we have M parallel ‘subchains’ each with stationary distribution $\pi_i, i \in \mathcal{M} \triangleq \{1, \dots, M\}$ and $\pi_M = \pi$. Associated with each subchain i is an MCMC kernel L_i that leaves π_i invariant, and which we run at every time step. Of course, without any further moves, the stationary distribution of the joint chain is

$$\bar{\pi}(\mathbf{x}_{1:M}) \triangleq \prod_{i=1}^M \pi_i(\mathbf{x}_i)$$

and so $\mathbf{x}_M \sim \pi$. This scheme does not affect the convergence rate of the independent chain M . However, since we can cycle mixtures of $\bar{\pi}$ -stationary MCMC kernels without affecting the

stationary distribution of the joint chain Tierney (1994), we can allow certain types of interaction between the subchains which can speed up convergence (Geyer 1991; Hukushima and Nemoto 1996). In general, we apply a series of kernels that act on subsets of the variables. For the sake of clarity, let us denote the number of second-stage kernels by R and the kernels themselves as K_1, \dots, K_R , where kernel K_j operates on variables with indices in $\mathcal{I}_j \subset \mathcal{M}$. The idea is that the R kernels are executed sequentially and it is required that each K_j leave $\prod_{i \in \mathcal{I}_j} \pi_i$ invariant.

Given π , there are a wide variety of possible choices for M , $\pi_{1:M-1}$, $L_{1:M}$, R , $\mathcal{I}_{1:R}$ and $K_{1:R}$ which will affect the convergence rate of the joint chain. For those interested, Jasra et al. (2007) gives a review of some of these. It is clear that the first stage of moves involving $L_{1:M}$ is trivially parallelizable. However, the second stage is sequential in nature. For a parallel implementation, it is beneficial for the \mathcal{I}_j 's to be disjoint as this allows the sequence of exchange kernels to be run in parallel. Of course, this implies that $\mathcal{I}_{1:R}$ should vary with time since otherwise there will be no interaction between the disjoint subsets of chains. Furthermore, if the parallel architecture used is SIMD (Single Instruction Multiple Data) in nature, it is desirable to have the K_j 's be nearly identical algorithmically. The last consideration for parallelization is that while speedup is generally larger when more computational threads can be run in parallel, it is not always helpful to increase M arbitrarily as this can affect the convergence rate of the chain. However, in situations where a suitable choice of M is dwarfed by the number of computational threads available, one can always increase the number of chains with target π to produce more samples.

3.2 Sequential Monte Carlo

SMC methods are a powerful extension of importance sampling methodology that are particularly popular for sampling from a sequence of probability distributions. In the context of state-space models, these methods are known as particle filtering methods; Doucet and Johansen (2008) gives a recent survey of the field. In this context, let $\{\mathbf{x}_t\}_{t \geq 0}$ be an unobserved Markov process of initial density $\mathbf{x}_0 \sim p_0(\cdot)$ and transition density $\mathbf{x}_t \sim f(\cdot | \mathbf{x}_{t-1})$ for $t \geq 1$. We only have access to an observation process $\{\mathbf{y}_t\}_{t \geq 0}$; the observations are conditionally independent conditional upon $\{\mathbf{x}_t\}_{t \geq 0}$ of marginal density $\mathbf{y}_t \sim g(\mathbf{y}_t | \mathbf{x}_t)$ for $t \geq 1$. SMC methods are used to approximate recursively in time the filtering densities $p(\mathbf{x}_{0:t} | \mathbf{y}_{0:t})$ which are proportional to $p(\mathbf{x}_{0:t}, \mathbf{y}_{0:t}) \triangleq p_0(\mathbf{x}_0) \prod_{i=1}^t f(\mathbf{x}_i | \mathbf{x}_{i-1}) \prod_{i=1}^t g(\mathbf{y}_i | \mathbf{x}_i)$ for $t = 1, \dots, T$. These distributions are approximated with a set of random samples called particles through use of a sequential version of importance sampling and a special particle-interaction step known as resampling.

Parallelization of SMC methods is reasonably straightforward. The importance sampling step used at each time step is trivially parallelizable as it involves only the local state of a particle. The resampling step, in which some particles are replicated and others destroyed depending on their normalized importance weights, comprises the construction of an empirical cumulative distribution function for the particles based on their importance weights followed by sampling from this N times, where N is the fixed number of particles used throughout the computation. While neither of these tasks is trivially parallelizable, they can benefit moderately from parallelization. However, the bulk of the speedup will generally come from the parallelization of the evolution and weighting steps. Therefore, using criteria like effective sample size (Liu 2008) to avoid resampling at every time step can also improve speedup.

3.3 Sequential Monte Carlo Samplers

SMC samplers (Del Moral et al. 2006) are a more general class of methods that utilize a sequence of auxiliary distributions π_0, \dots, π_T , much like population-based MCMC as discussed out in Jasra et al. (2007). However, in contrast to population-based MCMC, SMC samplers start from an auxiliary distribution π_0 and recursively approximate each intermediate distribution in turn until finally $\pi_T = \pi$ is approximated. The algorithm has the same general structure as classical SMC, with differences only in the types of proposal distributions, target distributions and weighting functions used in the algorithm. As such, parallelization of SMC samplers closely follows that of SMC.

4 Canonical Examples

To demonstrate the types of speed increase one can attain by utilizing GPUs, we apply each method to a representative statistical problem. We use Bayesian inference for a Gaussian mixture model as an application of the population-based MCMC and SMC samplers, while we use a factor stochastic volatility state-space model to gauge the speedup of our parallel SMC method. We ran our parallel code on a computer equipped with an NVIDIA 8800 GT GPU, a computer equipped with an NVIDIA GTX 280 GPU and we ran reference single-threaded code on a Xeon E5420 / 2.5 GHz processor. The resulting processing times and speedups are given in Tables 1 - 3.

The applications we discuss here are representative of the types of problems that these methods are commonly used to solve. In particular, while the distribution of mixture means given observations is only one example of a multimodal distribution, it can be thought of as a canonical distribution with multiple well-separated modes. Therefore, the ability to sample points from this distribution is indicative of the ability to sample points from a wide range of multimodal distributions. Similarly, performance of a latent variable sampler in dealing with observations from a factor stochastic volatility model is indicative of performance on observations from reasonably well-behaved but non-linear and non-Gaussian continuous state-space models.

4.1 Mixture Modelling

Let $\mathbf{y} = y_{1:m}$ denote observations where $y_j \in \mathbb{R}$ for $j \in \{1, \dots, m\}$. A univariate Gaussian mixture model with k components states that each observation is i.i.d. with density

$$p(y_j | \mu_{1:k}, \sigma_{1:k}, w_{1:k-1}) = \sum_{i=1}^k w_i f(y_j | \mu_i, \sigma_i)$$

where f denotes the density of the univariate normal distribution. The density of \mathbf{y} is then proportional to $\prod_{j=1}^m p(y_j | \mu_{1:k}, \sigma_{1:k}, w_{1:k-1})$.

For simplicity, we assume that k , $w_{1:k-1}$ and $\sigma_{1:k}$ are known and that the prior distribution on μ is uniform on the k -dimensional hypercube $[-10, 10]^k$. We set $k = 4$, $\sigma_i = \sigma = 0.55$, $w_i = w = 1/k$

for $i \in \{1, \dots, k\}$. We simulate $m = 100$ observations for $\boldsymbol{\mu} = \mu_{1:4} = (-3, 0, 3, 6)$. The resulting posterior distribution for $\boldsymbol{\mu}$ is given by

$$p(\boldsymbol{\mu}|\mathbf{y}) \propto p(\mathbf{y}|\boldsymbol{\mu})\mathbb{I}(\boldsymbol{\mu} \in [-10, 10]^4)$$

Generating N samples from this posterior is a popular method for determining the ability of samplers to explore a high-dimensional space with multiple well-separated modes. Invariance of the posterior to permutation of the labels of the parameters gives it $k! = 24$ symmetric modes, which should all be represented in the samples. It is worth noting that basic random-walk MCMC and importance sampling methods fail for practical values of N . It should be noted that while it might not be necessary to sample from all the symmetric modes in the case of a mixture model, the successful traversal of all the modes suggests that the sampler would succeed in traversing non-symmetric modes in other distributions, so long as symmetry is not exploited by the sampler.

4.1.1 Population-Based MCMC




Figure 3: Estimated marginal posterior density $p(\mu_{1:2}|\mathbf{y})$ from MCMC samples

We select the auxiliary distributions $\pi_{1:M-1}$ following the parallel tempering methodology, ie. $\pi_i(\mathbf{x}) \propto \pi(\mathbf{x})^{\beta_i}$ with $0 < \beta_1 < \dots < \beta_M = 1$ and use $M = 200$. This class of auxiliary distributions is motivated by the fact that MCMC converges more rapidly when the target distribution is flatter. For this problem, we use the cooling schedule $\beta_i = (i/M)^2$ and a standard $\mathcal{N}(\mathbf{0}, I_k)$ random walk Metropolis-Hastings kernel for the first stage moves.

For the second stage moves, we use only the basic exchange move (Geyer 1991; Hukushima and Nemoto 1996): chains i and j swap their values with probability $\min\{1, \alpha_{ij}\}$ where

$$\alpha_{ij} = \frac{\pi_i(\mathbf{x}_j)\pi_j(\mathbf{x}_i)}{\pi_i(\mathbf{x}_i)\pi_j(\mathbf{x}_j)}$$

Further, we allow exchanges to take place only between adjacent chains so that all moves can be done in parallel. We use $R = M/2$ and $\mathcal{I}_{1:R}$ is either $\{\{1, 2\}, \{3, 4\}, \dots, \{M-1, M\}\}$ or

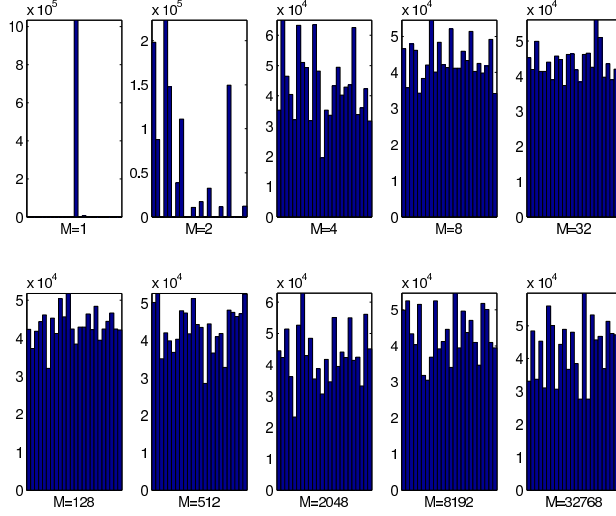


Figure 4: Number of MCMC samples from each mode

$\{\{2, 3\}, \{4, 5\}, \dots, \{M-2, M-1\}, \{M, 1\}\}$, each with probability half. While use of permutation or crossover moves would be appropriate for this particular model, we felt that they would detract from the ability to generalize our results to the case where the likelihood is not invariant to permutations of the labels.

Table 1: Running times for the Population-Based MCMC Sampler for various numbers of chains M .

M	CPU (mins)	8800GT (secs)	Speedup	GTX280 (secs)	Speedup
8	0.0166	0.887	1.1	1.083	0.9
32	0.0656	0.904	4	1.098	4
128	0.262	0.923	17	1.100	14
512	1.04	1.041	60	1.235	51
2048	4.16	1.485	168	1.427	175
8192	16.64	4.325	230	2.323	430
32768	66.7	14.957	268	7.729	527
131072	270.3	58.226	279	28.349	572

To test the computational time required by our algorithms we allow the number of chains to vary but fix the number of points we wish to sample from the marginal density $\pi_M = \pi$ at 8192. As such, an increase in the number of chains leads to a proportional increase in the total number of points sampled. Processing times for our code are given in Table 1, in which one can see that using 131072 chains is impractical on the CPU but entirely reasonable using the GPU. Figure 3 shows the estimated posterior density $p(\mu_{1:2}|\mathbf{y})$ from a set of 2^{20} MCMC samples from π_M with $M = 32768$, which is nearly identical to the estimated marginal posterior densities of any other pair of components of $\boldsymbol{\mu}$. This marginal density has 12 well-separated modes in \mathbb{R}^2 but it is worth noting that the joint density $p(\mu_{1:4}|\mathbf{y})$ has 24 well-separated modes in \mathbb{R}^4 . Figure 4 shows the number of points from each mode for various values of M . We also computed the average number of iterations taken for the samplers to traverse all modes for the different values of M . For $M = 1$ and $M = 2$,

the sampler did not traverse all the modes at all, while for values of M between 4 and 32 the traversal time decreased from 80000 to 10000, after which it was unchanged with increases in M . These numbers should be compared to $24 \times H_{24} \approx 91$ - the expected number of samples required to cover every mode if one could sample independently from π - where H_i is the i th harmonic number.

4.1.2 SMC Sampler




Figure 5: Estimated marginal posterior density $p(\mu_{1:2}|\mathbf{y})$ from SMC samples

As with population-based MCMC, we use a tempering approach and the same cooling schedule, ie. $\pi_t(\mathbf{x}) \propto \pi(\mathbf{x})^{\beta_t}$ with $\beta_t = (t/M)^2$ and $M = 200$. We use the uniform prior to generate the samples $\{\mathbf{x}_0^{(1:N)}\}$ and perform 10 MCMC steps with the standard $\mathcal{N}(\mathbf{0}, I_k)$ random walk Metropolis-Hastings kernel at every time step. We use the generic backwards kernel suggested in Crooks (1998), Neal (2001) and Del Moral et al. (2006) for the case where each kernel is π_t -stationary so that the unnormalized incremental importance weights are of the form $\pi_t(\mathbf{x}_{t-1})/\pi_{t-1}(\mathbf{x}_{t-1})$.

Table 2: Running times for the Sequential Monte Carlo Sampler for various values of N .

N	CPU (mins)	8800GT (secs)	Speedup	GTX280 (secs)	Speedup
8192	4.44	1.192	223.5	0.597	446
16384	8.82	2.127	249	1.114	475
32768	17.7	3.995	266	2.114	502
65536	35.3	7.889	268	4.270	496
131072	70.6	15.671	270	8.075	525
262144	141	31.218	271	16.219	522

We also ran the SMC sampler with no resampling at all, which for these settings corresponds to the Annealed importance sampling (AIS) method proposed in Neal (2001); see also Crooks (1998) for a similar method in physics. The resulting samples were less successful at characterizing the full multi-modality of the posterior distribution. This is consistent with the numerical findings and




Figure 6: Effective number of SMC samples from each mode

theoretical results discussed in (Del Moral et al. 2006, Sections 3.5 and 4.2.3): when the MCMC kernels used at every time step mixes reasonably well, it is very beneficial to resample. In addition, when ESS is used as a threshold for resampling in the SMC sampler, the resampling step takes very little time compared to the evolution and weighting of the particles simply because it happens so infrequently. As such, the running time of the SMC sampler compared to AIS is practically identical. Processing times for our code are given in Table 2. Figure 5 shows the estimated posterior density $p(\mu_{1:2}|\mathbf{y})$ from the SMC sampler with $N = 65536$. Figure 6 shows the number of points from each mode for various values of N .

4.1.3 Comparison

While both methods are capable of exploring the posterior distribution for $\boldsymbol{\mu}$, there are important differences in how the methods make use of parallelization. In particular, the SMC sampler parallelizes across particles approximating the same auxiliary distribution whilst the MCMC sampler parallelizes across auxiliary distributions at the same iteration. As such, to make full use of the graphics card the SMC sampler requires many particles while the MCMC sampler requires many auxiliary distributions. In most cases, however, one will be happy to use in excess of 8192 particles for SMC but one may not want to use in excess of 32768 auxiliary distributions. Indeed, for the application described above there seems to be no benefit in increasing the number of chains beyond 128, although this might also be due to the choice of cooling schedule and random walk variances. Furthermore, we utilized only the simplest information exchange and proposal moves in our samplers so as not to trivialize the problem. It should be noted, however, that there are situations in which a large number of intermediate temperatures are required for exchange acceptance probabilities to be greater than some preset value, for example when the dimension of the distribution of interest increases Predescu et al. (2004). As mentioned in Section 4.1.1, we would like to emphasize that the use of very large numbers of chains is possible using these parallel methods with only a modest increase in computation time.

The SMC sampler appears to be more efficient than the MCMC sampler for this problem. Indeed, with only 8192 particles the SMC sampler gives a reasonable representation of the posterior, taking only 597ms. The MCMC sampler requires around 2^{20} samples to give a reasonably uniform number of samples per mode, and this takes just over 2 minutes. In addition, although we have not done so here, it is possible with both the SMC and MCMC approaches to use the samples from the auxiliary distributions to estimate integrals of interest by computing appropriate importance weights. An interesting recent proposal on how to effectively combine estimates using such samples can be found in Gramacy et al. (2008).

For Bayesian inference in mixture models, there are many ways of dealing with the identifiability of the mixture parameters; Jasra (2005) includes a review of these. It is worth mentioning that for this type of model, we can permute samples as a post-processing step or within an MCMC kernel so traversal of the modes can be achieved trivially. The speedup of both methods is unaffected by the use of such mechanisms. In addition, the speedup is unaffected by increases in the number of observations since this only increases the amount of computation that each thread must do by a constant factor. Increasing the number of observations also has little effect on the difficulty for the sampler to move between modes since the modes are already well separated. Similarly, the speedup observed is robust to changes in the number of mixture components. The computation of each likelihood requires memory that is linear in the number of components, while the memory required per thread dictates the number of threads that can be run in parallel. However, as the number of components increases we usually have more observations, providing an opportunity to parallelize the computation of a single likelihood across multiple threads. This allows the amount of memory per thread to be significantly lower than if the complete likelihood was calculated by each thread and thus still allows many threads to be run in parallel.

4.2 Factor Stochastic Volatility

We consider the factor stochastic volatility model most similar to that proposed in Liu and West (2000):

$$\begin{aligned}\mathbf{y}_t &\sim \mathcal{N}(\mathbf{B}\mathbf{f}_t, \mathbf{\Psi}) \\ \mathbf{f}_t &\sim \mathcal{N}(\mathbf{0}, \mathbf{H}_t) \\ \mathbf{x}_t &\sim \mathcal{N}(\mathbf{\Phi}\mathbf{x}_{t-1}, \mathbf{U})\end{aligned}$$

where

$$\begin{aligned}\mathbf{\Psi} &\triangleq \text{diag}(\psi_1, \dots, \psi_M) \\ \mathbf{H}_t &\triangleq \text{diag}(\exp(\mathbf{x}_t)) \\ \mathbf{\Phi} &\triangleq \text{diag}(\phi_1, \dots, \phi_K)\end{aligned}$$

Here, \mathbf{f}_t is K -dimensional, \mathbf{y}_t is M -dimensional and \mathbf{B} is an $M \times K$ factor loading matrix with zero entries above the diagonal for reasons of identifiability. The latent variable at each time step t is the K -dimensional vector \mathbf{x}_t . The likelihood of the data, \mathbf{y}_t , given \mathbf{x}_t is Gaussian with

$$\mathbf{y}_t | \mathbf{x}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{B}\mathbf{H}_t\mathbf{B}^T + \mathbf{\Psi})$$

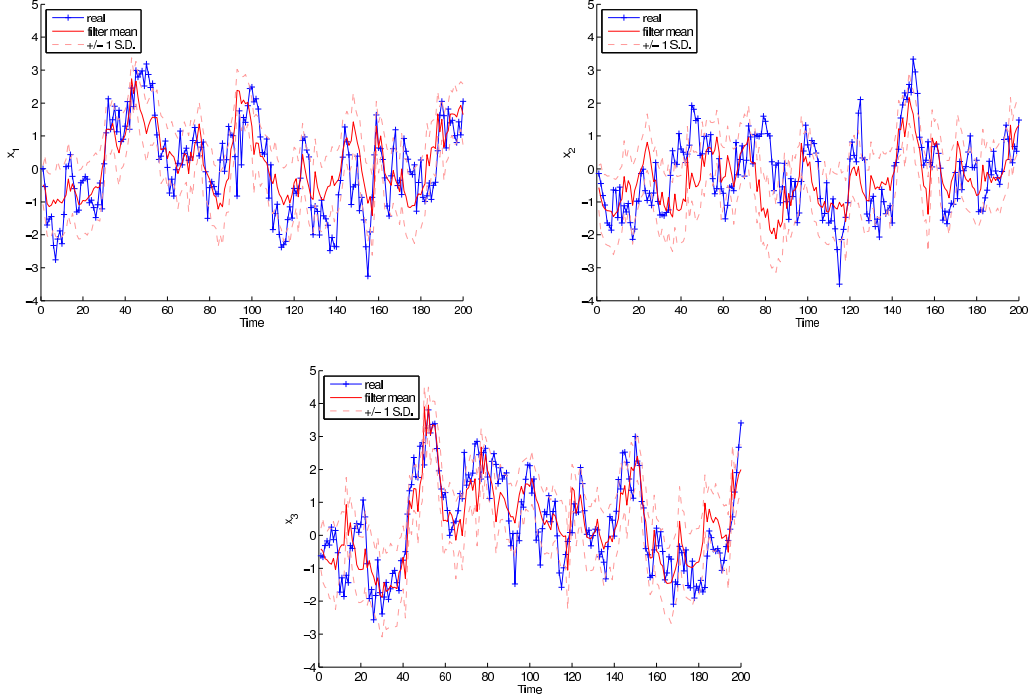


Figure 7: Estimated and Real Values of \mathbf{x}

We generate data for times $t = 1, \dots, T = 200$, $M = 5$, $K = 3$, $\mathbf{x}_0 = \mathbf{0}$, $\psi_i = 0.5, i \in \{1, \dots, M\}$, $\phi_i = 0.9, i \in \{1, \dots, K\}$,

$$\mathbf{B} = \begin{pmatrix} 1 & 0 & 0 \\ 0.5 & 1 & 0 \\ 0.5 & 0.5 & 1 \\ 0.2 & 0.6 & 0.3 \\ 0.8 & 0.7 & 0.5 \end{pmatrix} \text{ and } \mathbf{U} = \begin{pmatrix} 0.5 & 0.2 & 0.1 \\ 0.2 & 0.5 & 0.2 \\ 0.1 & 0.2 & 0.5 \end{pmatrix}$$

This is a simple example of a multivariate, non-linear and non-Gaussian continuous state-space model for which particle filters are commonly employed to sample from the posterior $p(\mathbf{x}_{0:T}|\mathbf{y}_{1:T})$. Processing times for our code are given in Table 3. In Figure 7 we plot the filter means for each component of \mathbf{x} with ± 1 sample standard deviations alongside the true values of \mathbf{x} used to generate the observations.

Table 3: Running time (in seconds) for the Sequential Monte Carlo method for various values of N .

N	CPU	8800GT	Speedup	GTX280	Speedup
8192	2.167	0.263	8	0.082	26
16384	4.325	0.493	9	0.144	30
32768	8.543	0.921	9	0.249	34
65536	17.425	1.775	10	0.465	37
131072	34.8	3.486	10	.929	37

The speedups obtained in this application are considerably less than for mixture model inference problem. This can be explained by lower arithmetic intensity and higher space complexity in each thread. The mixture model likelihood calculation contains a compute-intensive product-sum operation involving 104 values whilst the factor stochastic volatility likelihood consists mainly of matrix operations. In the latter case, the speedup is independent of T but not the dimension of the observations since the amount of memory required per thread increases quadratically in the dimension of each observation. This type of performance, however, still provides considerable speedup and may be more representative of the type of speedup practitioners can expect in general.

4.3 Floating Point Precision

For all three algorithms discussed above, we ran identical algorithms with the same random numbers on the CPU using double precision floating point numbers and the resulting estimates of expectations of interest were affected only slightly. Furthermore, those differences were dominated by the Monte Carlo variance of the estimates. The actual paths sampled, of course, were different due to the sensitivity of all of the algorithms to perturbations but this did not affect the ability of the samples to approximate the target distribution.

5 Discussion

The speedup for the population-based MCMC algorithm and the SMC sampler is tremendous. In particular, the evaluation of $p(\mathbf{y}|\boldsymbol{\mu})$ for the mixture-modelling application has high arithmetic intensity since it consists of a product-sum operation with 400 Gaussian log-likelihood evaluations involving only 104 values. In fact, because of the low register and memory requirements, so many threads can be run concurrently that SIMD calculation of this likelihood can be sped up by 500 times on the 8800 GT and 800 times on the GTX 280. However, the speedup attained for the standard SMC algorithm may be more representative of the kinds of gains one can expect in most applications with only reasonable arithmetic intensity. Even so, speedups of 10 to 35 make many problems tractable that previously were not by reducing a week’s worth of computation to a few hours. For example, estimation of static parameters in continuous state-space models or the use of SMC proposals within MCMC can require thousands of runs, so a speedup of this scale can substantially reduce the computation time of such approaches (see for eg. Andrieu et al. (2009)). It is worth noting also that we can expect speedups in the vicinity of 500 with SMC if few resampling steps are required and each weighting step has small space complexity and moderate time complexity.

The GTX 280 GPU outperforms the 8800 GT GPU by a factor of about 2 in all situations in which the GPU is used to capacity. This is the case in all but the population-based MCMC algorithm, in which the number of threads is determined by the number of auxiliary distributions. The reason for this is simple: the algorithms presented are register-bound on the inputs given, in that the number of registers required by each thread is the critical quantity that bounds the number of threads that can be run concurrently. The GTX 280 has twice the number of registers per multiprocessor and more than twice the multiprocessors compared to the 8800 GT. Hence, one could expect more speedup on many-core chips with even more registers. In fact, further improvements could be made

using multiple cards with shared memory, configurations of which are now available in NVIDIA’s Tesla line. It is also possible to design algorithms that are memory-bound, though we have not encountered this in the context of Monte Carlo simulation. It is certainly possible that both register and memory limitations can affect the parallelizability of the mentioned algorithms when facing very high-dimensional problems. However, in such cases it is possible that alternative uses of many-core architecture can provide speedup in these situations.

The acceleration of the Monte Carlo methods discussed here have practical benefits not only to computation time but also to energy efficiency. A general purpose CPU allocates extra circuits and hence power to flow control and caching which is unnecessary for the types of computation described here. As such, reasonable decreases in power consumption can be realized by using specialized many-core architectures like SIMD instead.

It should be noted that while we have used CUDA to implement the parallel components of algorithms, the results are not necessarily specific to this framework or to GPUs. It is expected that the many-core processor market will grow and there will be a variety of different devices and architectures to take advantage of. However, the SIMD parallelization technique and the sacrifice of caching and flow control for arithmetic processing is unlikely to disappear, particularly because when it is well-suited to a problem it will nearly always deliver considerable speedup. In addition, GPUs are affordable, off-the-shelf components that can be easily installed on a personal computer. Of particular interest is an emerging framework, the Open Computing Language (OpenCL), which provides a uniform programming environment for developers that enables them to write portable code for a variety of parallel devices, including GPUs and CPUs.

The speedups attainable with many-core architectures have broad implications in the design, analysis and application of SMC and population-based MCMC methods. With respect to SMC, it allows more particles to be used for the same or even less computation time, which can make these samplers viable where they previously were not. When faced with designing a population-based MCMC sampler, the results expectedly show that there is little cost associated with increasing the number of auxiliary distributions until the GPU reaches the critical limit of threads it can run concurrently. After this, there is a doubling in the computation time when the number of chains is doubled. In our application, this does not occur until we have around 4096 auxiliary distributions. Of course, we can expect this application-specific number to decrease when dealing with higher-dimensional distributions or those whose density evaluations require more registers or memory. Nevertheless, practitioners have more freedom to increase the number of auxiliary distributions to achieve a faster rate of convergence as the computation time associated with each step is not as closely tied to this value as it is on a single-core processor. In both SMC and MCMC, it is also clear from this case-study that it is beneficial for each thread to use as few registers as possible since this determines the number of threads that can be run simultaneously. This may be of interest to the methodology community since it creates a space-time tradeoff that might be exploited in some applications.

A consequence of the space-time tradeoff mentioned above is that methods which require large numbers of registers per thread are not necessarily suitable for parallelization using GPUs. For example, operations on large, dense matrices that are unique to each thread can restrict the number of threads that can run in parallel and hence dramatically affect potential speedup. In cases where data is shared across threads, however, this is not an issue. For example, a mixture model with large amounts of data does not affect the number of registers required whilst increasing the number

of components increases the number of registers required only linearly. In contrast, increasing the number of observed assets in a factor stochastic volatility model leads to a quadratic increase in the number of registers required, substantially affecting scalability in this regard. An increase in the number of observations itself has no effect on speedup. In principle, it is not the size of the data that matters but the space complexity of the algorithm in each thread that dictates how scalable the parallelization is.

The parallelization of the advanced Monte Carlo methods described here opens up challenges for both practitioners and for algorithm designers. There are already an abundance of statistical problems that are being solved computationally and technological advances, if taken advantage of by the community, can serve to make previously impractical solutions eminently reasonable and motivate the development of new methods.

Acknowledgments

The authors acknowledge support from the Oxford-Man Institute for Quantitative Finance and the Medical Research Council.

Appendix: Software

We have created a website resource at http://portal.stats.ox.ac.uk/userdata/cholmes/cuda_mc/ which will be available for the statistics community with the code used in these examples as well as useful information on GPU programming for scientific computing with links to tutorials and relevant papers.

References

- Andrieu, C., Doucet, A., and Holenstein, R. “Particle Markov Chain Monte Carlo (with discussion).” *J. Royal Statistical Soc. B* (2009). Under revision.
- Brockwell, A. E. “Parallel Processing in Markov chain Monte Carlo Simulation by Pre-Fetching.” *J. Comp. Graph. Statist.*, 15(1):246–261 (2006).
- Crooks, G. E. “Nonequilibrium measurements of free energy differences for microscopically reversible Markovian systems.” *J. Stat. Phys.*, 90:1481–1387 (1998).
- Del Moral, P., Doucet, A., and Jasra, A. “Sequential Monte Carlo Samplers.” *J. Royal Statistical Soc. B*, 68(3):411–436 (2006).
- Doucet, A. and Johansen, A. M. “A Tutorial on Particle Filtering and Smoothing: Fifteen years later.” In Crisan, D. and Rozovsky, B. (eds.), *Handbook of Nonlinear Filtering*. Oxford University Press (2008).

- Friedrichs, M. S., Eastman, P., Vaidyanathan, V., Houston, M., Legrand, S., Beberg, A. L., Ensign, D. L., Bruns, C. M., and Pande, V. S. “Accelerating molecular dynamic simulation on graphics processing units.” *Journal of Computational Chemistry*, 30(6):864–872 (2009).
- Geyer, C. J. “Markov chain Monte Carlo maximum likelihood.” In Kerigamas, E. (ed.), *Computing Science and Statistics: Proceedings of 23rd Symposium on the Interface Interface Foundation*, 156–163. Fairfax Station (1991).
- Gramacy, R. B., Samworth, R. J., and King, R. “Importance Tempering.” *Statistics and Computing* (2008). To appear.
- Higham, N. J. “The accuracy of floating point summation.” *SIAM Journal on Scientific Computing*, 4(4):783–799 (1993).
- . *Accuracy and stability of numerical algorithms (second edition)*. SIAM (2002).
- Hukushima, K. and Nemoto, K. “Exchange Monte Carlo Method and Application to Spin Glass Simulations.” *J. Phys. Soc. Jpn.*, 65:1604–1608 (1996).
- Jasra, A. “Bayesian Inference for Mixture Models via Monte Carlo Computation.” Ph.D. thesis, Imperial College London (2005).
- Jasra, A., Stephens, D. A., and Holmes, C. C. “On population-based simulation for static inference.” *Statistics and Computing*, 17(3):263–279 (2007).
- Kontoghiorghes, E. J. (ed.). *Handbook of Parallel Computing and Statistics*. Chapman & Hall (2006).
- L’Ecuyer, P. “Good parameter sets for combined multiple recursive random number generators.” *Operations Research*, 47(1):159–164 (1999).
- L’Ecuyer, P., Chen, E. J., and Kelton, W. D. “An object-oriented random-number package with many long streams and substreams.” *Operations Research*, 50(6):1073–1075 (2002).
- Liu, J. and West, M. “Combined parameter and state estimation in simulation-based filtering.” In Doucet, A., de Freitas, N., and Gordon, N. (eds.), *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York (2000).
- Liu, J. S. *Monte Carlo Strategies in Scientific Computing*. Springer (2008).
- Marsaglia, G. “Xorshift RNGs.” *Journal of Statistical Software*, 8(14):1–6 (2003).
- Neal, R. M. “Annealed Importance Sampling.” *Statistics and Computing*, 11(2):125–139 (2001).
- Predescu, C., Predescu, M., and Ciobanu, C. V. “The incomplete beta function law for parallel tempering sampling of classical canonical systems.” *Journal of Chemical Physics*, 120(9):4119–4128 (2004).
- Rosenthal, J. S. “Parallel computing and Monte Carlo algorithms.” *Far East J. Theor. Stat.*, 4:207–236 (2000).
- Stone, J. E., Phillips, J. C., Freddolino, P. L., Hardy, D. J., Trabuco, L. G., and Schulten, K. “Accelerating Molecular Modeling Applications with Graphics Processors.” *Journal of Computational Chemistry*, 28(16):2618–2640 (2007).

Tierney, L. “Markov Chains for Exploring Posterior Distributions.” *The Annals of Statistics*, 22(4):1701–1762 (1994).