

Article

Comparing Markov Chain Samplers for Molecular Simulation

Robert D. Skeel ^{*,†}  and Youhan Fang 

Department of Computer Science, Purdue University, West Lafayette, IN 47907, USA; yfang@purdue.edu

* Correspondence: rskeel@purdue.edu

† Current address: School of Mathematical and Statistical Sciences, Arizona State University, Phoenix, Tempe, AZ 85287, USA.

Received: 6 September 2017; Accepted: 19 October 2017; Published: 21 October 2017

Abstract: Markov chain Monte Carlo sampling propagators, including numerical integrators for stochastic dynamics, are central to the calculation of thermodynamic quantities and determination of structure for molecular systems. Efficiency is paramount, and to a great extent, this is determined by the integrated autocorrelation time (IAcT). This quantity varies depending on the observable that is being estimated. It is suggested that it is the maximum of the IAcT over all observables that is the relevant metric. Reviewed here is a method for estimating this quantity. For reversible propagators (which are those that satisfy detailed balance), the maximum IAcT is determined by the spectral gap in the forward transfer operator, but for irreversible propagators, the maximum IAcT can be far less than or greater than what might be inferred from the spectral gap. This is consistent with recent theoretical results (not to mention past practical experience) suggesting that irreversible propagators generally perform better if not much better than reversible ones. Typical irreversible propagators have a parameter controlling the mix of ballistic and diffusive movement. To gain insight into the effect of the damping parameter for Langevin dynamics, its optimal value is obtained here for a multidimensional quadratic potential energy function.

Keywords: Markov chain Monte Carlo; stochastic dynamics integrators; decorrelation time; integrated autocorrelation time

1. Summary

Thermodynamic and structural properties of a molecular system can be precisely defined as ensemble-dependent expectations of functions of its configuration. The calculation of such expectations is feasible only with the use of Markov chain Monte Carlo (MCMC) methods or approximations thereof. Considered here are sampling propagators that do not compromise the Markov property. Included are unbiased samplers that conditionally accept proposed moves and biased samplers that unconditionally accept such moves, in particular, discretized stochastic dynamics. Many such sampling propagators are proposed in the literature, and, in virtually all cases, experiments are conducted to substantiate claims of superiority. Too often though, a good metric is not used to measure the computational cost of a propagator. The aim of this article is threefold: first, to explore some practicalities related to measuring the efficiency of a propagator; second, to highlight the superior efficiency of irreversible propagators, namely, those that do not satisfy detailed balance; and third, to provide some insight into the optimal mix of diffusive and ballistic movement for Langevin dynamics.

Let $\rho_{\mathbf{Q}}(\mathbf{q})$, $\mathbf{q} = [q_1, q_2, \dots, q_v]^T$ denote the probability density function corresponding to the ensemble of interest. This function is assumed to be known only up to a multiplicative factor. An observable is an expectation $\mathbb{E}[u(\mathbf{Q})] = \int u(\mathbf{q})\rho_{\mathbf{Q}}(\mathbf{q}) d\mathbf{q}$ for some “preobservable” $u(\mathbf{q})$. This can be estimated by the mean $\bar{U}_N = (1/N) \sum_{n=0}^{N-1} u(\mathbf{Q}_n)$, where the random values \mathbf{Q}_n are obtained from

a Markov chain $\mathbf{Q}_0 \rightarrow \mathbf{Q}_1 \rightarrow \dots \rightarrow \mathbf{Q}_{N-1}$ that samples from $\rho_{\mathbf{Q}}(\mathbf{q})$. Note the use here of upper case to denote random variables.

In practice, sampling performance is improved if the configuration variables \mathbf{q} are augmented with auxiliary variables \mathbf{p} , e.g., momenta, yielding phase space variables $\mathbf{z} = (\mathbf{q}, \mathbf{p})$. The probability density is extended to $\rho(\mathbf{q}, \mathbf{p})$ so that

$$\int \rho(\mathbf{q}, \mathbf{p}) d\mathbf{p} = \rho_{\mathbf{Q}}(\mathbf{q}),$$

and an MCMC scheme is constructed to produce a chain $\mathbf{Z}_0 \rightarrow \mathbf{Z}_1 \rightarrow \dots \rightarrow \mathbf{Z}_{N-1}$, where $\mathbf{Z} = (\mathbf{Q}, \mathbf{P})$.

The samples from the chain tend to be highly correlated, and this greatly reduces the convergence rate as $N \rightarrow +\infty$. As explained in Section 2, the variance of an estimate for $\mathbb{E}[u(\mathbf{Z})]$ is

$$\text{Var}[\bar{U}_N] = \frac{\tau}{N} \text{Var}[u(\mathbf{Z})] + \mathcal{O}\left(\frac{1}{N^2}\right) \quad (1)$$

where τ is the *integrated autocorrelation time* (IAcT) for $u(\mathbf{z})$. The *effective sample size* is therefore N/τ , and the appropriate metric for evaluating a propagator is the effective sample size divided by the computing time.

In a great many practical simulations, the effective sample size is probably close to zero. One can disagree on the significance of such simulations [1]. In any case, for the comparison of sampling algorithms, it is possible to choose molecular systems, restrained if necessary, for which it is feasible to attain a decent effective sample size.

Often, the *spectral gap* is cited as the relevant quantity. To understand its role, it is helpful to express ideas in a direct way as in Refs. [2,3]. As detailed in Section 2, introduce a forward transfer operator \mathcal{F} to express the ratio ρ_n/ρ in terms of ρ_{n-1}/ρ , where $\rho_n(\mathbf{z})$ is the probability density for \mathbf{Z}_n . Let $\mathcal{F}_0 = \mathcal{F} - \mathcal{E}$ where $\mathcal{E}u$ denote the *function* with constant value $\mathbb{E}[u(\mathbf{Z})]$. Assume that the operator \mathcal{F}_0 has its spectrum strictly inside the unit circle, as it does in practice. The error in $(1/N) \sum_{n=0}^{N-1} \rho_n(\mathbf{z})$ can be shown [1] to be “proportional to” $(1 - \mathcal{F}_0)^{-1}$ and therefore to the reciprocal of the spectral gap $|1 - \lambda_2|$, where λ_2 is a nonunit eigenvalue of \mathcal{F} nearest 1.

Estimates of the IAcT are obtained by summing the terms of the autocorrelation function, which is constructed from autocovariances of increasing time lags normalized by the (lag zero) covariance. Each term contributes a roughly equal statistical error but a signal that decays as the lag time increases. Therefore, in practice, the terms are weighted by a function called a lag window. The lag window must be tailored to the autocorrelation function, and choosing a suitable lag window is very difficult, as mentioned in Section 2.1.

Reliable estimates of the IAcT are impractical in general. Section 3 reviews the concept of quasi-reliability, which aims to enforce a necessary condition for reliable estimates. Informally, the goal is to ensure good coverage of those “parts” of phase space that has been explored, to reduce the possibility of missing an opening to an unexplored part of phase space. More precisely, for an arbitrary subset of phase space, we ask that the proportion of samples in that subset differ from its expectation by no more than some tolerance *tol*, with, say, 95% confidence. This is shown to be true if the IAcT τ for any preobservable $u(\mathbf{z})$ satisfies $\tau \leq \text{tol}^2 N$. The maximum IAcT τ_{\max} is the greatest eigenvalue of

$$\mathcal{G} = 1 - \mathcal{E} + \mathcal{F}_0(1 - \mathcal{F}_0)^{-1} + \mathcal{F}_0^\dagger(1 - \mathcal{F}_0^\dagger)^{-1}, \quad (2)$$

where † denotes the adjoint with respect to the inner product

$$\langle v, u \rangle = \int \overline{v(\mathbf{z})} u(\mathbf{z}) \rho(\mathbf{z}) d\mathbf{z}.$$

For a reversible propagator, where $\mathcal{F}^\dagger = \mathcal{F}$, τ_{\max} is 1 less than twice the reciprocal of the spectral gap. However, for an irreversible propagator, it can be much smaller, as demonstrated by a simple

example in Section 4.1, or larger as in Section 5. As a practical algorithm, it is suggested to estimate the maximum IAcT by discretizing the space of preobservables. Consideration of a quadratic energy potential suggests using a linear combination of phase space variables (and possibly quadratic terms in these variables).

The idea of seeking the preobservable that maximizes the IAcT is suggested already in Ref. [4], which considers a set of indicator functions as preobservables and uses the greatest IAcT of these to assess sampling thoroughness. In general, maximizing over a linear combination of preobservables can lead to a much larger result than taking the maximum of them individually, due to correlations that might exist between different preobservables. This does not necessarily apply, however, to those considered in Refs. [1,4].

Section 3.1 notes that that typical irreversible propagators, termed “quasi-reversible” here, have a forward transfer operator $\mathcal{F} = \mathcal{R}\bar{\mathcal{F}}$, where each of $\bar{\mathcal{F}}$ and \mathcal{R} is reversible and $\mathcal{R}^2 = 1$. For such propagators, the estimation of τ_{\max} simplifies somewhat.

Theoretical results [5] indicate that adding irreversibility reduces the autocorrelation times of observables. Section 4 gives a couple of very elementary examples illustrating the dramatic increase in τ_{\max} if an irreversible propagator \mathcal{F} is replaced by its reversible “counterpart” $\frac{1}{2}(\mathcal{F} + \mathcal{F}^\dagger)$.

Discretized Langevin dynamics is a particularly effective general-purpose propagator. Unfortunately, one must specify a value for the damping coefficient γ . Section 5 analyzes τ_{\max} for a quadratic potential and obtains an optimal value for the coefficient, namely, the value $(3/8)^{1/2} = 0.612 \dots$ times the critical damping coefficient for lowest frequency ω_1 . This value is such that the lowest frequency mode is moderately underdamped, with higher frequencies increasingly underdamped.

2. Preliminaries

Assuming that \mathbf{Z}_0 has probability density $\rho(\mathbf{z})$, the variance of the estimate \bar{U}_N is exactly

$$\text{Var}[\bar{U}_N] = \frac{1}{N} \text{Var}[u(\mathbf{Z})] \left(1 + 2 \sum_{k=1}^{N-1} \left(1 - \frac{k}{N} \right) \frac{C(k)}{C(0)} \right)$$

where the autocovariances

$$C(k) = \mathbb{E} [(u(\mathbf{Z}_0) - \mathbb{E}[u(\mathbf{Z}_0)]) (u(\mathbf{Z}_k) - \mathbb{E}[u(\mathbf{Z}_k)])].$$

The limit $N \rightarrow +\infty$ gives Equation (1) where the integrated autocorrelation time

$$\tau = 1 + 2 \sum_{k=1}^{+\infty} \frac{C(k)}{C(0)}. \tag{3}$$

As an example of augmenting configuration space, consider $\rho(\mathbf{q}, \mathbf{p}) \propto \exp(-\beta(V(\mathbf{q}) + \frac{1}{2}\mathbf{p}^\top M^{-1}\mathbf{p}))$, where $\mathbf{p} = [p_1, p_2, \dots, p_\nu]^\top$. A good propagator for this is the BAOAB integrator [6] for Langevin dynamics, whose equations are

$$d\mathbf{Q}_t = M^{-1}\mathbf{P}_t dt, \quad d\mathbf{P}_t = \mathbf{F}(\mathbf{Q}_t) dt - \gamma\mathbf{P}_t dt + \sqrt{\frac{2\gamma}{\beta}} M_{1/2} d\mathbf{W}_t, \tag{4}$$

where M is a matrix chosen to compress the range of vibrational frequencies, $\mathbf{F}(\mathbf{q}) = -\nabla V(\mathbf{q})$, $M_{1/2}M_{1/2}^\top = M$, and \mathbf{W}_t is a vector of independent standard Wiener processes. Each step of the integrator consists of the following substeps:

- B: $\mathbf{P}'_n = \mathbf{P}_{n-1} + \frac{1}{2}\Delta t \mathbf{F}(\mathbf{Q}_{n-1})$,
- A: $\mathbf{Q}'_n = \mathbf{Q}_{n-1} + \frac{1}{2}\Delta t M^{-1}\mathbf{P}'_n$,
- O: $\mathbf{P}''_n = \exp(-\gamma\Delta t)\mathbf{P}'_n + \sqrt{1 - \exp(-2\gamma\Delta t)}\beta^{-1/2}M_{1/2}\mathbf{R}_n$,

$$\begin{aligned} \text{A: } \mathbf{Q}_n &= \mathbf{Q}'_n + \frac{1}{2}\Delta t M^{-1} \mathbf{P}''_n, \\ \text{B: } \mathbf{P}_n &= \mathbf{P}''_n + \frac{1}{2}\Delta t \mathbf{F}(\mathbf{Q}_n), \end{aligned}$$

where \mathbf{R}_n is a vector of independent standard Gaussian random variables. The samples generated from this process are shown [7] to be those from a distribution that differs from the correct one by only $\mathcal{O}(\Delta t^2)$. The special choice $\gamma = 1/(2\Delta t)$ is the Euler–Leimkuhler–Matthews integrator [6] for Brownian dynamics with step size $\delta t = \Delta t^2/2$. Remarkably, the invariant density of this integrator differs from the correct one by only $\mathcal{O}(\delta t^2)$. This integrator can be expressed as a Markov chain $\mathbf{Q}'_1 \rightarrow \mathbf{Q}'_2 \rightarrow \dots \rightarrow \mathbf{Q}'_{N-1}$ in configuration space, with propagator

$$\begin{aligned} \mathbf{Q}_n &= \mathbf{Q}'_n + \frac{1}{2}\sqrt{2\delta t}\beta^{-1/2}M_{1/2}^{-T}\mathbf{R}_n, \\ \mathbf{Q}'_{n+1} &= \mathbf{Q}_n + \delta t M^{-1}\mathbf{F}(\mathbf{Q}_n) + \frac{1}{2}\sqrt{2\delta t}\beta^{-1/2}M_{1/2}^{-T}\mathbf{R}_n, \end{aligned}$$

which is a discretization of Brownian dynamics

$$d\mathbf{Q}_t = M^{-1}\mathbf{F}(\mathbf{Q}_t) dt + \sqrt{\frac{2}{\beta}}M_{1/2}^{-T}d\mathbf{W}_t. \tag{5}$$

The desired samples $\{\mathbf{Q}_n\}$ are available as part of the process (and, as a theoretical observation, they can be recovered from the Markov chain $\{\mathbf{Q}'_n\}$ alone, by eliminating \mathbf{R}_n in the two equations above and solving for \mathbf{Q}_n).

For any MCMC propagator, the forward transfer operator is defined so that

$$u_n = \mathcal{F}u_{n-1}$$

where $u_n(\mathbf{z}) = \rho_n(\mathbf{z})/\rho(\mathbf{z})$ and $\rho_n(\mathbf{z})$ is the probability density for \mathbf{Z}_n . In particular,

$$\mathcal{F}u_{n-1}(\mathbf{z}) = \frac{1}{\rho(\mathbf{z})} \int \rho(\mathbf{z}|\mathbf{z}')u_{n-1}(\mathbf{z}')\rho(\mathbf{z}')d\mathbf{z}'$$

where $\rho(\mathbf{z}|\mathbf{z}')$ is the transition probability for the chain. The operator \mathcal{F} has an eigenfunction $\varphi_1(\mathbf{z}) \equiv 1$ for eigenvalue $\lambda_1 = 1$.

A reversible propagator is one that satisfies detailed balance, which means that

$$\rho(\mathbf{z}'|\mathbf{z})\rho(\mathbf{z}) = \rho(\mathbf{z}|\mathbf{z}')\rho(\mathbf{z}').$$

Detailed balance is equivalent to $\mathcal{F}^\dagger = \mathcal{F}$, where the adjoint \mathcal{F}^\dagger is defined by the condition that $\langle \mathcal{F}v, u \rangle = \langle v, \mathcal{F}^\dagger u \rangle$ for arbitrary $u(\mathbf{z})$ and $v(\mathbf{z})$. The BAOAB integrator is not reversible as a sampling propagator, except for the special case $\gamma = 1/(2\Delta t)$; however, a scheme consisting of a fixed number of BAOAB steps followed by a momenta flip is reversible. The unmodified BAOAB integrator is, however, in a class of “quasi-reversible” integrators introduced in Section 3.1.

As a model problem for Brownian dynamics, given by Equation (5), consider $F(q) = -m\omega^2q$. (The absence of boldface indicates scalar quantities.) Changing variables $Q_t = (m\beta)^{-1/2}Q'_t$ and dropping the prime gives the simple stochastic differential equation

$$dQ_t = -\omega^2Q_t dt + \sqrt{2}dW_t, \tag{6}$$

where $W(t)$ denotes a standard Wiener process. A perfect realization of $Q(t)$ at discrete points is obtained by the reversible propagator

$$Q_n = \exp(-\omega^2\delta t)Q_{n-1} + \sqrt{1 - \exp(-2\omega^2\delta t)}\frac{1}{\omega}R_n. \tag{7}$$

The probability density $\rho(q, t)$ for $Q(t)$ satisfies the Fokker–Planck equation $(\partial/\partial t)\rho = (\partial/\partial q)(\omega^2 q \rho) + (\partial^2/\partial q^2)\rho$. Writing $\rho(q, t) = u(q, t)\rho(q)$ gives $(\partial/\partial t)u = -\omega^2 q(\partial/\partial q)u + (\partial^2/\partial q^2)u$. The operator on the right-hand side has eigenfunctions $u(q) = He_k(\omega q)$ with eigenvalues $-k\omega^2$ for $k = 0, 1, \dots$. The modified version of the Hermite polynomial of degree k is given by $He_k(x) = (-1)^k \exp(x^2/2)(d^k/dx^k) \exp(-x^2/2)$. The forward transfer operator \mathcal{F} for the propagator defined by Equation (7) has these same eigenfunctions and has eigenvalues $\lambda_{k+1} = \exp(-k\omega^2 \delta t)$ for $k = 1, 2, \dots$. The spectral gap is $1 - \exp(-\omega^2 \delta t)$. In the multidimensional case with normal mode frequencies $0 < \omega_1 \leq \omega_2 \leq \dots \leq \omega_\nu$, the time step δt is some fraction, say $\frac{1}{2}$, of $1/\omega_\nu^2$ and the spectral gap is *very nearly* $\frac{1}{2}\omega_1^2/\omega_\nu^2$. To see the applicability to practical numerical integrators, consider the Euler–Leimkuhler–Matthews discretization of Equation (6):

$$Q'_n = (1 - \omega^2 \delta t)Q'_{n-1} + (1 - \frac{1}{2}\omega^2 \delta t)\sqrt{2\omega^2 \delta t}\frac{1}{\omega}R_n. \tag{8}$$

Comparing Equations (8) and (7) and equating the coefficients of the two terms on their right-hand sides enables Equation (8) to be written in the form of Equation (7) with modified values for ω and δt . In particular, the modified value of $\exp(-\omega^2 \delta t)$ is $1 - \omega^2 \delta t$, so spectral gap is exactly $\omega^2 \delta t$. In the multidimensional case, where $\omega_1^2 \delta t \ll 1$, the spectral gap for the discrete stochastics is *very nearly* that of the continuous stochastics.

2.1. Estimating Integrated Autocorrelation Time

Suggested [8] as a covariance estimate is the quantity

$$C_N(k) = \frac{1}{N} \sum_{n=0}^{N-k-1} (u(\mathbf{z}_n) - \bar{U}_N)(u(\mathbf{z}_{n+k}) - \bar{U}_N),$$

with justification in Ref. [9] (pp. 323–324).

The use of all possible terms $C_N(k)/C_N(0)$ in Equation (3) to estimate the IAcT does not converge in the limit $N \rightarrow +\infty$, so, in practice, a *lag window* $w(k)$ is used to increasingly damp terms as the noise to signal ratio increases:

$$\tau \approx 1 + \sum_{k=1}^{N-1} w(k) \frac{C_N(k)}{C_N(0)}.$$

An interesting algorithm called *acor* for estimating the IAcT is available on the web [10]. Estimating the IAcT can be quite difficult, and *acor* can give unsatisfactory results. An attempt to improve it [1] is at best marginally successful. For reversible methods, there are properties of the autocorrelation function that may be useful for improving estimates of it [8].

3. Quasi-Reliable Estimates of Sampling Thoroughness

The idea of quasi-reliability is to require that the sample size N be large enough that, with say 95% confidence, the *estimated* probability of any subset Ω of phase space differs by no more than *tol* from its correct value. More specifically, for any subset Ω of phase space, an estimate $\bar{1}_{\Omega, N}$ of $\mathbb{E}[1_\Omega(\mathbf{Z})] = \Pr(\mathbf{Z} \in \Omega)$ must satisfy

$$\text{Var}[\bar{1}_{\Omega, N}] \leq \frac{1}{4} \text{tol}^2.$$

Because

$$\text{Var}[\bar{1}_{\Omega, N}] \approx \tau_\Omega \frac{1}{N} \text{Var}[1_\Omega(\mathbf{Z})] \leq \frac{1}{4N} \tau_\Omega,$$

where τ_Ω is the IAcT for 1_Ω , it is enough that

$$\frac{1}{4N} \tau_\Omega \leq \frac{1}{4} \text{tol}^2 \quad \text{for all } \Omega. \tag{9}$$

For molecular simulations, this requires that only those conformations or clusters of conformations having a probability greater than tol be sampled.

In practice, molecular systems have many symmetries, which dramatically reduces the amount of sampling needed. For example, water molecules are generally considered interchangeable as are many subsets of atoms on a given molecule, e.g., the 2 hydrogen atoms of any water molecule. More formally, there are permutations P of the variables \mathbf{z} such that $\rho(P\mathbf{z}) = \rho(\mathbf{z})$ and that $\mathcal{P}^{-1}\mathcal{F}\mathcal{P} = \mathcal{F}$ where $(\mathcal{P}u)(\mathbf{z}) = u(P\mathbf{z})$. For such symmetries P , the quasi-reliability requirement considers only those Ω for which $1_\Omega(P\mathbf{z}) = 1_\Omega(\mathbf{z})$.

It is helpful to express the IAcT in terms of the forward transfer operator. It can be shown that

$$\mathbb{E}[v(\mathbf{Z}_0)u(\mathbf{Z}_k)] = \langle \mathcal{F}^k v, u \rangle.$$

and, in particular,

$$\begin{aligned} C(k) &= \mathbb{E}[u(\mathbf{Q}_0)u(\mathbf{Q}_k)] - \mathbb{E}[u(\mathbf{Q}_0)]\mathbb{E}[u(\mathbf{Q}_k)] = \langle \mathcal{F}^k u, u \rangle - \langle \mathcal{E}u, u \rangle \\ &= \begin{cases} \langle (1 - \mathcal{E})u, u \rangle, & \text{for } k = 0, \\ \langle \mathcal{F}_0^k u, u \rangle, & \text{for } k \geq 1, \end{cases} \end{aligned} \tag{10}$$

using the fact that $\mathcal{E}\mathcal{F}_0 = \mathcal{F}_0\mathcal{E} = 0$. The integrated autocorrelation time can be rewritten as

$$\tau = \frac{C(0) + 2\sum_{k=1}^{+\infty} C(k)}{C(0)} = \frac{\langle (1 - \mathcal{E} + 2\sum_{k=1}^{+\infty} \mathcal{F}_0^k)u, u \rangle}{\langle (1 - \mathcal{E})u, u \rangle} = \frac{\langle \mathcal{G}u, u \rangle}{\langle (1 - \mathcal{E})u, u \rangle}$$

where

$$\mathcal{G} = 1 - \mathcal{E} + \mathcal{F}_0(1 - \mathcal{F}_0)^{-1} + (\mathcal{F}_0(1 - \mathcal{F}_0)^{-1})^\dagger,$$

which is a self-adjoint operator.

For a reversible propagator, for which \mathcal{F} and hence \mathcal{F}_0 is self adjoint, an arbitrary preobservable u is in many cases expressible as a linear combination of the eigenfunctions $\varphi_k(\mathbf{z})$ of \mathcal{F}_0 , corresponding to eigenvalues $1 > \lambda_2 \geq \lambda_3 \geq \dots > -1$. (For a more rigorous treatment, see Ref. [8] (Section 2).) The IAcT for u is then simply a weighted average, of the values

$$1 + \frac{2\lambda_k}{1 - \lambda_k} = \frac{1 + \lambda_k}{1 - \lambda_k}$$

with weights $\langle \varphi_k, u \rangle^2 / (\langle \varphi_k, \varphi_k \rangle \langle u, u \rangle)$. This is maximized for $u = \varphi_2$, since $(1 + \lambda_2)/(1 - \lambda_2)$ is the largest of these values. Note that, for λ_2 negative, τ could be much less than 1. Having $\tau < 1$ may appear paradoxical until it is recognized that Equation (1) is simply an asymptotic approximation for $N \rightarrow +\infty$.

For the simple example with $F(q) = -\omega^2 q$, the eigenfunction $\varphi_2(q) = \sqrt{2}\omega q$. The indicator function 1_Ω that is richest in $\varphi_2(q)$ is the one with $\Omega = [0, +\infty]$, for which the first weight is $2/\pi$. This means that the IAcT for 1_Ω is at least $2/\pi$ of the maximum IAcT. For a multimodal distribution, the eigenfunction $\varphi_2(q)$ corresponding to the subdominant eigenvalue λ_2 [2,3] is closer than is the function $\sqrt{2}\omega q$ to an indicator function modulo a constant. Therefore, little is lost and simplicity is gained, if we use the maximum IAcT over all observables satisfying the symmetries instead of just indicator functions:

$$\tau_{\max} = \sup_{u \in W} \frac{\langle \mathcal{G}u, u \rangle}{\langle (1 - \mathcal{E})u, u \rangle} \tag{11}$$

where

$$W = \{u = u(\mathbf{z}) \mid u(P\mathbf{z}) = u(\mathbf{z}) \text{ for symmetries } P\}$$

is our set of preobservables. This can be simplified to

$$\tau_{\max} = \sup_{u \in W} \frac{\langle \mathcal{G}u, u \rangle}{\langle u, u \rangle}. \tag{12}$$

To see this, note that the same supremum is obtained for both objective functions if the function u is restricted so that $\mathcal{E}u = 0$ and that for such a function u the two objective functions are equal.

For the simple example of Equation (8), the IAcT is maximized by $u(q) = q$. In the case of a multidimensional Gaussian distribution, the maximum occurs for linear combination $\mathbf{a}^T \mathbf{q}$ of \mathbf{q} where \mathbf{a} is the eigenvector of the Hessian of $V(\mathbf{q})$ corresponding to its smallest eigenvalue ω_1^2 . For multimodal distributions in one dimension, it appears that the maximizing preobservable $\varphi_2(q)$ is qualitatively similar to q in the sense that it is monotonic [2,3].

As is customary when seeking an unknown function, one considers a finite linear combination $u(\mathbf{q}) = \mathbf{a}^T \mathbf{u}(\mathbf{q})$ of basis functions $u_i \in W$ where the a_i are chosen to maximize the IAcT. The number of basis functions is limited by computational cost. The result of Section 5 hints that the ideal number might be $\mathcal{O}(\nu^2)$. From Equation (10), the autocovariance for such $u(\mathbf{q})$ is

$$C(k) = \mathbf{a}^T C_k \mathbf{a}$$

where

$$C_k = \begin{cases} \langle (1 - \mathcal{E})\mathbf{u}, \mathbf{u}^T \rangle, & \text{for } k = 0, \\ \langle \mathcal{F}_0^k \mathbf{u}, \mathbf{u}^T \rangle, & \text{for } k \geq 1. \end{cases} \tag{13}$$

Consequently,

$$\tau_{\max} \approx \max_{\mathbf{a}} \frac{\mathbf{a}^T K \mathbf{a}}{\mathbf{a}^T C_0 \mathbf{a}} \quad \text{where } K = C_0 + 2 \sum_{k=1}^{+\infty} C_k, \tag{14}$$

which is a solution of the generalized eigenvalue problem

$$\frac{1}{2}(K + K^T)\mathbf{a} = C_0 \mathbf{a} \tau. \tag{15}$$

The approximation of Equation (14) improves with the number of basis functions.

Without information about the actual distribution $\rho(\mathbf{z})$, a general choice for basis functions might be linear polynomials, which are the “subdominant” eigenfunctions for a Gaussian distribution. Simple examples in Ref. [1] (consisting of a mixture of two Gaussians, a one-node neural “network”, and logistic regression) demonstrate that the use of linear basis functions can yield an IAcT much greater than that of a preobservable of “interest”. For molecular simulation, it is clear that instead of atomic coordinates, a better choice of a basis function is the distance between two atoms, each of which is uniquely distinguishable. In particular, for a protein, one might choose α -carbons distributed along the backbone chain of a protein. For further suggestions, consult Ref. [11], which considers the automatic construction of indicator functions for estimating τ_{\max} , based on the dynamics of the propagator.

3.1. Quasi-Reversible Propagators

As stated above, the BAOAB integrator would be reversible if the final B substep were modified to include a momenta flip, i.e.,

$$\text{BR: } \mathbf{P}_n = -(\mathbf{P}'_n + \frac{1}{2}\Delta t \mathbf{F}(\mathbf{Q}_n)).$$

However, reversing direction runs counter to intuition. Additionally, empirical evidence [12] favors irreversible propagators. If the flipped BAOAB integrator is followed by another momenta flip, the result is the original irreversible BAOAB integrator. More generally, a sampler might be said to be “quasi-reversible” if its forward transfer operator $\mathcal{F} = \mathcal{R}\bar{\mathcal{F}}$ where each of $\bar{\mathcal{F}}$ and \mathcal{R} is reversible and $\mathcal{R}^2 = 1$. For a momenta flip, the operator \mathcal{R} is defined by $(\mathcal{R}u)(\mathbf{q}, \mathbf{p}) = u(\mathbf{q}, -\mathbf{p})$.

Quasi-reversible propagators are special in that their covariance matrices satisfy a special property if basis functions u_i are chosen so that

$$\mathcal{R}\mathbf{u} = D\mathbf{u} \quad \text{where } D = \text{diag}(I', -I'')$$

and I' and I'' are identity matrices of possibly different dimension. Such basis functions are easy to construct since any function u can be written as a sum of its “even part” $\frac{1}{2}(u + \mathcal{R}u)$ and its “odd part” $\frac{1}{2}(u - \mathcal{R}u)$. In particular, using the fact that $\mathcal{R}\mathcal{E} = \mathcal{E}\mathcal{R}$,

$$\begin{aligned} C_k &= \langle \mathcal{F}^k \mathbf{u}, \mathbf{u}^\top \rangle - \langle \mathcal{E} \mathbf{u}, \mathbf{u}^\top \rangle = \langle \mathbf{u}, (\bar{\mathcal{F}}\mathcal{R})^k \mathbf{u}^\top \rangle - \langle \mathbf{u}, \mathcal{E} \mathbf{u}^\top \rangle \\ &= \langle \mathcal{R} \mathbf{u}, (\mathcal{R}\bar{\mathcal{F}})^k \mathcal{R} \mathbf{u}^\top \rangle - \langle \mathcal{R} \mathbf{u}, \mathcal{R} \mathcal{E} \mathbf{u}^\top \rangle = \langle D \mathbf{u}, \mathcal{F}^k \mathbf{u}^\top D \rangle - \langle D \mathbf{u}, \mathcal{E} \mathbf{u}^\top D \rangle = DC_k^\top D. \end{aligned}$$

The matrix C_k thus partitions into four blocks. The symmetric part of C_k consists of the two diagonal blocks, and the skew-symmetric part consists of the two off-diagonal blocks. Empirical estimates of C_k lack these symmetry properties. The expected symmetries provides twice as many samples for estimating the sampling error in the off-diagonal elements. Additionally, since the IAcT requires only the symmetric part of C_k , it is unnecessary to compute the off-diagonal blocks, and the eigenvalue problem decouples into two smaller problems. In addition, it follows that the maximizing linear combination is either a linear combination of the even functions or of the odd functions. For molecular dynamics at least, velocities “relax” much more quickly than positions, so it seems that the long time scales are present only in the position coordinates, so only the (1, 1) block might be computed.

4. Irreversible Samplers and Their Superiority

Uniform sampling—if it could be used—converges like $\mathcal{O}(1/N)$, which is superior to the $\mathcal{O}(1/\sqrt{N})$ convergence of random sampling. It would be desirable to combine them by promoting near-uniform sampling within layers of near-constant energy and using diffusion to move among energy layers.

4.1. A Very Simple Example

The following very simple example is a discrete analog of Example 2.8 of Ref. [5] and similar to an example of Ref. [13]. Assuming probabilities are represented as row vectors, let \mathcal{F} be an n by n probability transition matrix given by

$$\mathcal{F}_{ij} = \begin{cases} \theta, & \text{if } i = j, \\ 1 - \theta, & \text{if } i = j + 1 \text{ mod } n, \\ 0, & \text{otherwise,} \end{cases}$$

where $0 < \theta < 1$. The stationary probability vector is $(1/n)[1, 1, \dots, 1]$. The inner product for two row vectors u^\top, v^\top is $\langle v^\top, u^\top \rangle = v^\top u$, and the adjoint of a transition matrix is its transpose. Since \mathcal{F} is a circulant matrix, both it and its adjoint have as eigenvectors $u_k^\top = [1, \zeta^{k-1}, \dots, (\zeta^{k-1})^{n-1}]$, where ζ denotes the n th root of unity $\exp(2\pi i/n)$. A straightforward, though lengthy, calculation shows that

$$u_k^\top \mathcal{F}_0 = (\theta + (1 - \theta)\zeta^{k-1})u_k^\top \quad \text{and} \quad u_k^\top \mathcal{F}_0^\top = (\theta + (1 - \theta)\zeta^{1-k})u_k^\top,$$

for $k = 2, 3, \dots, n$, and

$$u_k^\top \mathcal{G} = \frac{\theta}{1 - \theta} u_k^\top.$$

Therefore,

$$2/|1 - \lambda_2| = 2/((1 - \theta)|1 - \zeta|) = 1/((1 - \theta) \sin(\pi/n)),$$

but

$$\tau_{\max} = \theta / (1 - \theta).$$

This is a most elementary illustration of the assertion [5] that “the asymptotic variance for every observable can be made as small as desired under large irreversible drift”. If $\theta \ll 1$, the IAcT is likewise much less than 1 despite the very strong correlation between successive values of the state j . The reason for this is that the IAcT is a measure of sample quality not independence. The measure τ simply says that N samples are as good as N/τ independent random samples. The \mathcal{F} operator encourages a uniform sampling, which is better than uncorrelated random samples. However, as shown in the example that follows, such a dramatic difference between (one less than twice) the reciprocal of the spectral gap and the IAcT does not hold when sampling is inhibited more by energy than by entropy barriers. A reversible propagator can be formed from this simple example by using instead the symmetric part of its propagator: $\mathcal{F}^{rv} = \frac{1}{2}(\mathcal{F} + \mathcal{F}^T)$. This time

$$u_k^T \mathcal{F}_0^{rv} = \left(\theta + \frac{1}{2}(1 - \theta)(\zeta^{k-1} + \zeta^{1-k}) \right) u_k^T, \quad \text{for } k = 2, 3, \dots, n,$$

and

$$u_k^T \mathcal{G}^{rv} = \left(\frac{4}{(1 - \theta)(2 - \zeta^{k-1} - \zeta^{1-k})} - 1 \right) u_k^T.$$

Therefore

$$2/|1 - \lambda_2| = 4 / ((1 - \theta)(2 - \zeta - \zeta^{-1})) = 1 / ((1 - \theta) \sin^2(\pi/n)),$$

and

$$\tau_{\max} = 1 / ((1 - \theta) \sin^2(\pi/n)) - 1.$$

Clearly, the irreversible propagator is much better for $n \gg 1$.

4.2. A Very Simple Example with a Barrier

Consider the $2n$ by $2n$ transition matrix given by

$$\mathcal{F} = \left[\begin{array}{ccc|cc} 0 & & 1 - \varepsilon & & \varepsilon \\ 1 & 0 & & & \\ & \ddots & \ddots & & \\ & & 1 & 0 & \\ \hline & & & \varepsilon & 0 & 1 - \varepsilon \\ & & & & 1 & 0 \\ & & & & & \ddots & \ddots \\ & & & & & & 1 & 0 \end{array} \right].$$

If ε were zero, the eigenvalue 1 would have multiplicity 2, with orthogonal eigenvectors given by the vector of all ones and by $f^T = [1, 1, \dots, 1, -1, -1, \dots, -1]$. For small $\varepsilon > 0$, it is expected that the eigenvector corresponding to the subdominant eigenvalue would be a perturbation of f^T . Indeed, it happens that $f^T \mathcal{F}^n = (1 - 2\varepsilon)f^T$. This also holds for $(\mathcal{F}^n)^T$, so the IAcT for \mathcal{F}^n is $1/\varepsilon - 1$. This suggests a value of about $n(1/\varepsilon - 1)$ for the IAcT of \mathcal{F} . This is corroborated by numerical computation, as shown in Table 1a.

The excess IAcT due to making the propagator reversible is given by Table 1b. It appears to grow quadratically with n , consistent with the diffusive nature of the propagator. Summing pairs of entries from the two tables shows that for this example the advantage of irreversibility depends on the relative importance of entropy barriers to energy barriers, which is consistent with intuition from molecular dynamics.

Table 1. τ_{\max} is the IAcT for \mathcal{F} and τ_{\max}^{rv} is that for $\mathcal{F}^{\text{rv}} = \frac{1}{2}(\mathcal{F} + \mathcal{F}^{\text{T}})$.

(a) τ_{\max}			
$n \setminus \varepsilon$	0.1	0.01	0.001
10	90	990	9990
100	900	9900	99,900
1000	9000	99,000	999,000
(b) $\tau_{\max}^{\text{rv}} - \tau_{\max}$			
$n \setminus \varepsilon$	0.1	0.01	0.001
10	35	33	33
100	3910	3511	3355
1000	403,612	389,921	351,047

5. Optimal Langevin Damping for a Model Problem

To analyze the effect of the damping parameter γ on the IAcT for Langevin dynamics, given by Equation (4), consider the standard model problem $F(q) = -m\omega^2q$. Changing variables $Q_t = (\beta m)^{-1/2}Q'_t$, $P_t = \beta^{-1/2}m^{1/2}P'_t$, and dropping the primes gives

$$dQ_t = P_t dt, \quad dP_t = -\omega^2Q_t dt - \gamma P_t dt + \sqrt{2\gamma} dW_t.$$

Assume exact integration with step size Δt .

The analysis is rather lengthy, so for the benefit of the reader who wishes to omit it, the discussion and conclusions are given here:

1. Reaching precise conclusions is difficult for most of the analysis unless $\omega\Delta t$ is bounded above by some constant of order 1, which seems to be satisfied in practice. This assumption underlies the statements that follow.
2. The spectral gap is an increasing function of ω , so for a multidimensional quadratic potential energy, the value of γ that maximizes the spectral gap depends on the lowest frequency ω_1 .
3. The spectral gap is maximized for $\gamma \leq 2\omega$, corresponding to an underdamped system, for which the spectral gap is $\omega\Delta t + \mathcal{O}(\Delta t^2)$.
4. The eigenfunctions of the operator \mathcal{G} can be partitioned into eigenspaces \mathbb{P}'_k , $k = 0, 1, \dots$, where \mathbb{P}'_k is a linear combination of $k + 1$ specific polynomials of degree k in ωq and p . The greatest eigenvalue of \mathcal{G} is $\tau_{\max} = \max_k \tau_{\max}^{(k)}$ where $\tau_{\max}^{(k)}$ is the maximum IAcT over \mathbb{P}'_k .
5. Figure 1 shows that, for fixed Δt and γ , the value $\tau_{\max}^{(k)}$ is an increasing function of ω , at least for $k = 1, 2, 3, 4$. Hence, as for the spectral gap, it is the lowest frequency ω_1 that dictates the maximum τ .
6. Figure 2 indicates that, for fixed Δt and ω , the maximizing τ is either $\tau_{\max}^{(1)}$ or $\tau_{\max}^{(2)}$, depending on the value of γ . The optimal damping coefficient is $\gamma = (\sqrt{6}/2)\omega$, for which $\tau_{\max}^{(1)} = \tau_{\max}^{(2)} = \sqrt{6}/(\omega\Delta t)$.
7. For the preobservable ωq , and, indeed, for any odd polynomial, the IAcT τ becomes arbitrary small as γ goes to zero. This does not mean, however that the variance goes to 0, because Equation (1) is an asymptotic result, and the IAcT is a prefactor for the leading order $1/N$ term in the expression for the variance. The next order $1/N^2$ term would dominate if γ were very small. An order $1/N^2$ error is characteristic of uniform sampling, which would be the consequence of nearly ballistic movement. In addition, odd polynomials are special in that their expectation is independent of total energy, so it matters not that barely diffusive movement samples different values of total energy only at a very slow rate.

8. The eigenfunction for $\tau_{\max}^{(2)}$ is a specific linear combination of $\omega^2 q^2 - 1$ and $p^2 - 1$. For quadratic polynomials, the total energy does affect its expected value, which is why it is necessary that γ be large enough to sample different energies on a reasonable time scale.

5.1. The Forward Transfer Operator

The forward transfer operator is

$$\mathcal{F}u = \rho^{-1}e^{\Delta t \mathcal{L}_K}(\rho u)$$

where

$$\rho(q, p) \propto \exp\left(-\frac{1}{2}p^2 - \frac{1}{2}\omega^2 q^2\right).$$

and \mathcal{L}_K is the Fokker–Planck operator [14] (Equations (10.1)–(10.3), (10.9))

$$\mathcal{L}_K f = -p \frac{\partial}{\partial q} f + \omega^2 q \frac{\partial}{\partial p} f + \gamma \frac{\partial}{\partial p} (p f) + \gamma \frac{\partial^2}{\partial p^2} f.$$

Therefore,

$$\mathcal{F} = e^{\Delta t \mathcal{L}}$$

where

$$\mathcal{L} f = \frac{1}{\rho} \mathcal{L}_K(\rho f) = -p \frac{\partial}{\partial q} f + \omega^2 q \frac{\partial}{\partial p} f - \gamma p \frac{\partial}{\partial p} f + \gamma \frac{\partial^2}{\partial p^2} f.$$

Using the relation $\mathcal{L} f = \frac{1}{\rho} \mathcal{L}_K(\rho f)$, it is easy to show that the adjoint

$$\mathcal{L}^+ f = p \frac{\partial}{\partial q} f - \omega^2 q \frac{\partial}{\partial p} f - \gamma p \frac{\partial}{\partial p} f + \gamma \frac{\partial^2}{\partial p^2} f.$$

5.2. Gamma for Maximum Spectral Gap

Ref. [14] (Chapter 10, Equations (1), (2), (3b), (9), (22), (52), (60), (71), (72), (74), (77), (78), (82) and (83)) gives the eigenvalues of \mathcal{L}_K as

$$\mu_{k,l} = -\frac{1}{2}(k+l)\gamma - \frac{1}{2}(k-l)\delta, \quad k, l = 0, 1, \dots,$$

with eigenfunctions

$$\psi_{k,l} = \rho(q, p)^{1/2} (k! l! \delta^{k+l})^{-1/2} (\sqrt{\gamma_+} \mathcal{B} - \sqrt{\gamma_-} \mathcal{A})^k (-\sqrt{\gamma_-} \mathcal{B} + \sqrt{\gamma_+} \mathcal{A})^l \rho(q, p)^{1/2},$$

where

$$\gamma_{\pm} = \frac{1}{2}(\gamma \pm \delta) \quad \text{and} \quad \delta = \sqrt{\gamma^2 - 4\omega^2},$$

with the radical sign denoting the principal square root, and operators \mathcal{A}, \mathcal{B} defined by

$$\mathcal{A} f = -\omega^{-1} \frac{\partial}{\partial q} f + \frac{1}{2} \omega q f, \quad \mathcal{B} f = -\frac{\partial}{\partial p} f + \frac{1}{2} p f.$$

The operator \mathcal{L} has eigenfunctions $\rho(q, p)^{-1} \psi_{k,l}(q, p)$ and the same eigenvalues; \mathcal{F} has these same eigenfunctions and has eigenvalues $\exp(\Delta t \mu_{k,l})$. See Ref. [15] for a more general result.

For $k+l=1$, one has $\mu_{0,1} = -\gamma_-, \mu_{1,0} = -\gamma_+$ and

$$\varphi_{0,1} = (\gamma_- / \delta)^{1/2} (-p + \gamma_+ q) \rho(q, p), \quad \varphi_{1,0} = (\gamma_+ / \delta)^{1/2} (p - \gamma_- q) \rho(q, p).$$

In the case $\gamma_{\pm} = \frac{1}{2}\gamma$, \mathcal{F} has a double eigenvalue $\lambda_2 = \lambda_3 = e^{-\omega\Delta t}$ but one eigenfunction $\varphi_2(q, p) = p - \omega q$ and one generalized eigenfunction $\varphi_3(q, p) = q$, satisfying $(\mathcal{F} - \lambda_2)\varphi_3 = \varphi_2$. This results in behavior involving a linear combination of $e^{-\omega t}$ and $te^{-\omega t}$.

The eigenvalue nearest 1 depends on Δt . For small enough Δt , it is $\exp(-\gamma_-\Delta t)$, and the spectral gap is $\omega\Delta t + \mathcal{O}(\Delta t^2)$ for $\gamma \leq 2\omega$ and $\omega\Delta t(2\omega/\gamma)(1 + (1 - (2\omega/\gamma)^2)^{1/2})^{-1} + \mathcal{O}(\Delta t^2)$ otherwise. To maximize the spectral gap, choose γ to be no greater than 2ω , corresponding to an underdamped system.

5.3. Gamma for Maximum IAcT

To obtain τ_{\max} , one needs eigenelements for \mathcal{G} , given by Equation (2) with $\mathcal{F} = \exp(\Delta t\mathcal{L})$. Note that the subspace of polynomials of degree $\leq k$ is closed under application of \mathcal{L} and \mathcal{L}^\dagger . Moreover, this holds separately for subspace \mathbb{P}_k of odd polynomials of degree $\leq k$ for k odd and for subspace \mathbb{P}_k of even polynomials of degree $\leq k$ for k even. It also applies to operators \mathcal{E} , \mathcal{F} , \mathcal{F}^\dagger , and \mathcal{G} . Hence, the \mathbb{P}_k are eigenspaces for \mathcal{G} . These eigenspaces can be further decomposed as follows. Let $\mathbb{P}'_k = \mathbb{P}_k$ for $k < 2$, and let $\mathbb{P}'_k = \mathbb{P}_k \cap \mathbb{P}_{k-2}^\perp$ for $k \geq 2$. The claim is that \mathbb{P}'_k is an eigenspace, of dimension $k + 1$. To confirm this, it is enough to show that $\langle \mathcal{G}u, v \rangle = 0$ for any $u \in \mathbb{P}_k, v \in \mathbb{P}_{k-2}$, which follows almost immediately, since $\mathcal{G}^\dagger = \mathcal{G}$ and $\mathcal{G}v \in \mathbb{P}_{k-2}$. Let \mathbf{u} be a basis for \mathbb{P}'_k chosen so there are even functions of p followed by odd functions of p . In particular,

$$\begin{aligned} \mathbf{u} &= [\omega q, p]^\top, \\ \mathbf{u} &= [\omega^2 q^2 - 1, \omega q p, p^2 - 1]^\top, \\ \mathbf{u} &= [\omega^3 q^3 - 3\omega q, (\omega^2 q^2 - 1)p, \omega q(p^2 - 1), p^3 - 3p]^\top, \\ \mathbf{u} &= [\omega^4 q^4 - 6\omega^2 q^2 + 3, (\omega^3 q^3 - 3\omega q)p, (\omega^2 q^2 - 1)(p^2 - 1), \omega q(p^3 - 3p), p^4 - 6p + 3]^\top, \end{aligned}$$

for $k = 1, 2, 3, 4$, respectively. The polynomials in ωq and in p are modified versions He_k of Hermite polynomials. We have $\mathcal{L}\mathbf{u} = A\mathbf{u}$ for some matrix of constants A , given by

$$A = \begin{bmatrix} 0 & -k\omega & & & \\ \omega & -\gamma & \ddots & & \\ & \ddots & \ddots & -\omega & \\ & & k\omega & -k\gamma & \end{bmatrix}. \tag{16}$$

(cf. Ref. [14], Equations (10.96) and (10.97) for \mathcal{L}_K .) We have $\mathcal{E}\mathbf{u} = \mathbf{0}$ and $\mathcal{F}_0\mathbf{u} = \mathcal{F}\mathbf{u} = \exp(\Delta t A)\mathbf{u}$. Therefore, from Equations (13) and (14), one has

$$C_k = \langle \mathcal{F}_0^k \mathbf{u}, \mathbf{u}^\top \rangle = \exp(k\Delta t A) \langle \mathbf{u}, \mathbf{u}^\top \rangle,$$

and

$$K = \coth\left(-\frac{1}{2}\Delta t A\right) \langle \mathbf{u}, \mathbf{u}^\top \rangle = \coth\left(-\frac{1}{2}\Delta t A\right) C_0.$$

For each value of k , τ_{\max} is the maximum eigenvalue of

$$\frac{1}{2} \left(C_0^{-1} \coth\left(-\frac{1}{2}\Delta t A\right) C_0 + \coth\left(-\frac{1}{2}\Delta t A\right)^\top \right). \tag{17}$$

As explained in Section 3.1, if the basis functions are re-ordered so that those that are even functions of p precede those that are odd functions of p , the eigenvalue problem splits into two nearly equal parts.

For $k = 1$, $A = X\Lambda X^{-1}$ where $\Lambda = \text{diag}(-\gamma_-, -\gamma_+)$ and

$$X = \begin{bmatrix} \omega & \omega \\ \gamma_- & \gamma_+ \end{bmatrix}.$$

The covariance matrix $C_0 = \langle \mathbf{u}, \mathbf{u}^T \rangle = \text{diag}(1, 1)$, and matrix (17) is $\text{diag}(\tau_+, \tau_-)$ where

$$\tau_{\pm} = \frac{1}{\delta} \left(\pm \gamma_{\pm} \coth\left(\frac{1}{2}\Delta t \gamma_{\mp}\right) \mp \gamma_{\mp} \coth\left(\frac{1}{2}\Delta t \gamma_{\pm}\right) \right).$$

Using the identity $\coth(x \pm y) = (\sinh 2x \mp \sinh 2y) / (\cosh 2x - \cosh 2y)$, one gets that the eigenvalues of matrix (17) are

$$\tau_{\pm} = \frac{\sinh(\frac{1}{2}\Delta t \gamma) \pm \gamma \sinh(\frac{1}{2}\Delta t \delta) / \delta}{\cosh(\frac{1}{2}\Delta t \gamma) - \cosh(\frac{1}{2}\Delta t \delta)}.$$

The greater of τ_{\pm} is τ_+ , and it is a straightforward exercise to show that τ_+ is an decreasing function of ω as long as

$$(\omega \Delta t)^2 \leq \pi^2 + (\gamma \Delta t / 2)^2. \tag{18}$$

In practice, a numerical integrator is used, whose step size is chosen so that $\Delta t L = \theta$ where θ is some fractional value, e.g., $\frac{1}{2}$, and L is the magnitude of largest eigenvalue of the Jacobian matrix of the right-hand side of the system of (first-order) stochastic differential equations. In particular,

$$\omega \Delta t = \theta \quad \text{if } \gamma \leq 2\omega; \tag{19}$$

otherwise, it holds that $\omega \Delta t < \gamma \Delta t / 2$. Hence, Inequality (18) is more than satisfied. It is more complicated to analyze the behavior of values of τ for $k > 1$, so we exploit the smallness of $\omega \Delta t$ and $\gamma \Delta t$ to do an asymptotic analysis. From Expression (17), one sees that for \mathbb{P}'_k , its τ_{\max} equals $\tau_{\max}^{(k)} + \mathcal{O}(\Delta t)$ where $\tau_{\max}^{(k)}$ is the greatest eigenvalue of

$$-\frac{1}{\Delta t} (C_0^{-1} A^{-1} C_0 + A^{-T}). \tag{20}$$

For $k = 1$, one has $C_0 = \text{diag}(1, 1)$ and $\tau_{\max}^{(1)} = 2\gamma / (\omega^2 \Delta t)$, corresponding to the eigenfunction ωq . For $k = 2$, one has

$$A^{-1} = \frac{1}{2\gamma\omega^2} \begin{bmatrix} -\gamma^2 - \omega^2 & 2\gamma\omega & -\omega^2 \\ -\gamma\omega & 0 & 0 \\ -\omega^2 & 0 & -\omega^2 \end{bmatrix},$$

$C_0 = \text{diag}(2, 1, 2)$, and the largest eigenvalue for Expression (20) is

$$\tau_{\max}^{(2)} = \frac{1}{2\Delta t} \left(\frac{\gamma}{\omega^2} + \frac{2}{\gamma} + \sqrt{\frac{\gamma^2}{\omega^4} + \frac{4}{\gamma^2}} \right),$$

corresponding to an eigenfunction that is some linear combination of $\omega^2 q^2 - 1$ and $p^2 - 1$. Again $\tau_{\max}^{(k)}$ is a decreasing function of ω for small enough Δt . For general k , the size of the elements of A in Equation (16) increase with ω , which suggests that for practical values of Δt , the largest eigenvalue for Expression (20) decreases with ω . This can be confirmed numerically. To reduce the number of parameters, write $A = -\omega B(\gamma/\omega)$, and rewrite Expression (20) as

$$\frac{1}{\omega \Delta t} (C_0^{-1} B(\gamma/\omega)^{-1} C_0 + B(\gamma/\omega)^{-T}). \tag{21}$$

The value $\tau_{\max}^{(k)}$ is the largest eigenvalue of Matrix (21), so $\gamma\Delta t \tau_{\max}^{(k)}$ is a function only of the ratio ω/γ . Figure 1 plots this quantity for $k = 1, 2, 3, 4$, confirming that $\tau_{\max}^{(k)}$ is decreasing function of ω for fixed values of γ and Δt .

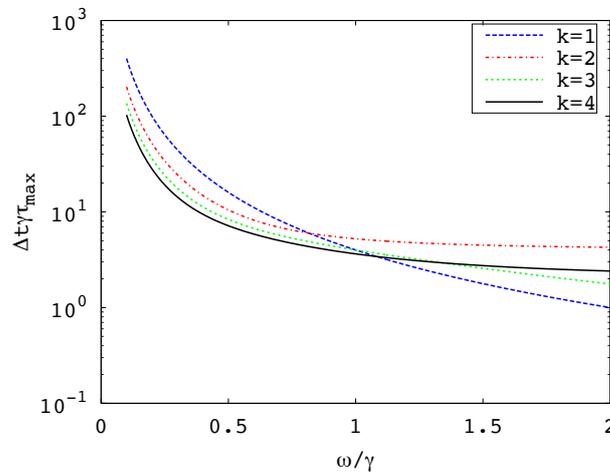


Figure 1. $\gamma\Delta t \tau_{\max}^{(k)}$ vs. ω/γ for $k = 1, 2, 3, 4$.

Therefore, for a multidimensional quadratic potential, τ_{\max} is very likely to be determined by the lowest frequency mode. (This is readily established for the spectral gap for reasonable values of Δt .)

Having established that the lowest frequency mode is most likely responsible for τ_{\max} , let ω denote the lowest frequency ω_1 , and consider the optimal choice of γ . From Matrix (21), observe that $\omega\Delta t \tau_{\max}^{(k)}$ is a function only of γ/ω . Accordingly, we choose γ/ω to minimize $\max_k \omega\Delta t \tau_{\max}^{(k)}$. However, Δt may depend on γ , but only if $\gamma > 2\omega_v$, according to Equation (19). Let us initially ignore this possibility and revisit it after determining the optimal γ .

Assuming that the maximum of $\omega\Delta t \tau_{\max}^{(k)}$ is attained for $k = 1$ or $k = 2$, there are two possible locations for the best γ . One is the choice $\gamma = \sqrt{2}\omega$, for which $\omega\Delta t \tau_{\max}^{(2)} = 1 + \sqrt{2}$ and $\omega\Delta t \tau_{\max}^{(1)} = 2\sqrt{2}$. The other is the choice $\gamma = (\sqrt{6}/2)\omega$ where $\omega\Delta t \tau_{\max}^{(k)} = \sqrt{6}$ for both $k = 1$ and $k = 2$. The latter is the better choice. This and the assumption that the maximum occurs for $k < 3$ is supported by Figure 2, which plots τ_{\max} as a function of γ for fixed values of ω and Δt . Since $(\sqrt{6}/2)\omega_1 < 2\omega_v$, we see from Equation (19) that there is no concern about Δt limiting the value of γ .

Note that for the optimal γ , the dynamics is underdamped in every mode.

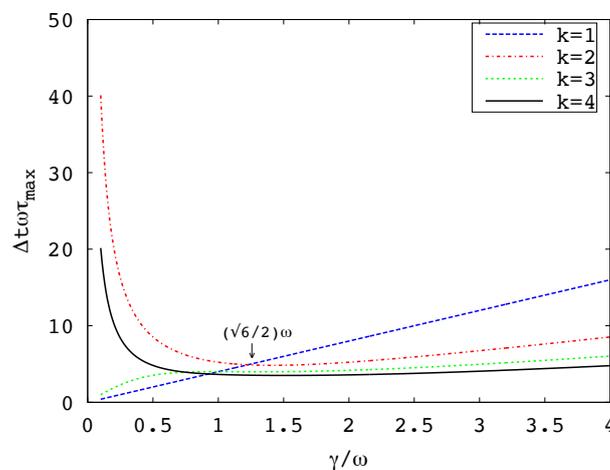


Figure 2. $\omega\Delta t \tau_{\max}^{(k)}$ vs. γ/ω for $k = 1, 2, 3, 4$.

6. Discussion and Conclusions

It is asserted in the literature [1,4] that, for comparing MCMC propagators and for ensuring reasonably reliable simulation results, it is of great value to estimate the maximum autocorrelation time over all possible observables. Certainly, this would seem desirable for comparing propagators, since it is best to assume as little as possible about the actual observables that are to be estimated from the Markov chain. In addition, the longest IAcT would be more reliably estimated than any other IAcT, because its accurate estimation would call for a greatest number of samples; moreover, the accuracy of estimates of other IAcTs might be influenced by the longest IAcT. Unfortunately, estimating IAcTs and assessing such estimates is a difficult task, which can benefit greatly from further research.

In this article, it is suggested that the maximum IAcT be estimated by considering an arbitrary linear combination of basis functions chosen to capture the lowest frequency motions of the dynamics of the MCMC chain. It is also worth considering the approach based on a set of well chosen indicator functions [11].

Another quantity for measuring performance of an MCMC propagator is the spectral gap. It is related to equilibration time, but as a predictor of IAcTs, it can be arbitrarily too pessimistic or it can be too optimistic. Such is the case for irreversible propagators, especially for energy landscapes where entropy barriers dominate energy barriers.

Shorter integrated autocorrelation times are thus a primary consideration in the design of MCMC propagators. Simple examples are constructed to explain how irreversibility might help, though their relevance to practical propagators is unclear. Nonetheless, it is advantageous not to insist on reversibility. In particular, the ballistic component of propagators like hybrid Monte Carlo and Langevin dynamics may offer dramatic speedups for overcoming entropy barriers, though no advantage for energy barriers.

Irreversible propagators typically have a parameter that determines the extent of diffusive behavior. For discrete Langevin dynamics applied to a multidimensional quadratic potential, the optimal value corresponds to slightly underdamped dynamics for the lowest frequency mode, with other modes experiencing even lighter damping. However, this conclusion is not necessarily indicative of the optimal damping for the usual situation, in which there is a multiplicity of energy minima, and further studies are needed.

In a nutshell, general-purpose sampling benefits greatly from the use of irreversible propagators in extended state space, with “diffusion parameter(s)” chosen to minimize the maximum integrated autocorrelation time, which for a quadratic potential energy corresponds to moderately underdamped dynamics in the lowest frequency mode.

Author Contributions: Youhan Fang performed the analysis for Section 5.2 and the computation for Section 5.3. Robert D. Skeel did the analysis for Sections 3.1, 4, 5.1 and 5.3 and wrote the paper. Both authors have read and approved the final manuscript.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Fang, Y.; Cao, Y.; Skeel, R. Quasi-Reliable Estimates of Effective Sample Size. *arXiv* **2017**, arXiv:1705.03831v2.
2. Schütte, C.; Huisinga, W. Biomolecular Conformations as Metastable Sets of Markov Chains. In Proceedings of the 38th Annual Allerton Conference on Communication, Control, and Computing, Monticello, IL, USA, 4–6 October 2000; pp. 1106–1115.
3. Schütte, C.; Sarich, M. *Metastability and Markov State Models in Molecular Dynamics*; American Mathematical Society: Providence, RI, USA, 2013.
4. Lyman, E.; Zuckerman, D.M. On the Structural Convergence of Biomolecular Simulations by Determination of the Effective Sample Size. *J. Phys. Chem. B* **2007**, *111*, 12876–12882.
5. Rey-Bellet, L.; Spiliopoulos, K. Irreversible Langevin Samplers and Variance Reduction: A Large Deviations Approach. *Nonlinearity* **2015**, *28*, 2081–2104.

6. Leimkuhler, B.; Matthews, C. Rational Construction of Stochastic Numerical Methods for Molecular Sampling. *Appl. Math. Res. Express* **2013**, *2013*, 34–56.
7. Leimkuhler, B.; Matthews, C.; Stoltz, G. The Computation of Averages from Equilibrium and Nonequilibrium Langevin Molecular Dynamics. *IMA J. Numer. Anal.* **2016**, *36*, 13–79.
8. Geyer, C.J. Practical Markov Chain Monte Carlo. *Stat. Sci.* **1992**, *7*, 473–511.
9. Priestly, M.B. *Spectral Analysis and Time Series*; Academic Press: Cambridge, MA, USA, 1981.
10. Goodman, J. Acor, Statistical Analysis of a Time Series, 2009. Available online: <http://www.math.nyu.edu/faculty/goodman/software/acor/> (accessed on 4 August 2014).
11. Zhang, X.; Bhatt, D.; Zuckerman, D.M. Automated Sampling Assessment for Molecular Simulations Using the Effective Sample Size. *J. Chem. Theory Comput.* **2010**, *6*, 3048–3057.
12. Cancés, E.; Legoll, F.; Stoltz, G. Special Issue on Molecular Modelling: Theoretical and Numerical Comparison of Some Sampling Methods for Molecular Dynamics. *ESAIM Math. Model. Numer. Anal.* **2007**, *41*, 351–389.
13. Diaconis, P.; Holmes, S.; Neal, R.M. Analysis of a Nonreversible Markov Chain Sampler. *Ann. Appl. Probab.* **2000**, *10*, 726–752.
14. Risken, H. *The Fokker-Planck Equation. Methods of Solution and Applications*; Springer Series in Synergetics; Springer: New York, NY, USA, 1984; Volume 18.
15. Kozlov, S. Effective diffusion for the Fokker-Planck equation. *Math. Notes* **1989**, *45*, 360–368.



© 2017 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).