

SMAI-JCM

SMAI JOURNAL OF COMPUTATIONAL MATHEMATICS

Some Remarks on Preconditioning Molecular Dynamics

HOUSSAM ALRACHID, LETIFMONES & CHRISTOPH ORTNER
Volume 4 (2018), p. 57-80.

http://smai-jcm.cedram.org/item?id=SMAI-JCM_2018__4__57_0

© Société de Mathématiques Appliquées et Industrielles, 2018
Certains droits réservés.

cedram

Article mis en ligne dans le cadre du
Centre de diffusion des revues académiques de mathématiques
<http://www.cedram.org/>





Some Remarks on Preconditioning Molecular Dynamics

HOUSSAM ALRACHID ¹
LETIFMONES ²
CHRISTOPH ORTNER ³

¹ IDP Laboratory, Orleans University, Orleans, France

E-mail address: housam.alrachid@univ-orleans.fr

² Mathematics Institute, University of Warwick, Coventry, United Kingdom

E-mail address: l.mones@warwick.ac.uk

³ Mathematics Institute, University of Warwick, Coventry, United Kingdom

E-mail address: c.ortner@warwick.ac.uk.

Abstract. We consider a Preconditioned Overdamped Langevin algorithm that does not alter the invariant distribution (up to controllable discretisation errors) and ask whether preconditioning improves the standard model in terms of reducing the asymptotic variance and of accelerating convergence to equilibrium. We present a detailed study of the dependence of the asymptotic variance on preconditioning in some elementary toy models related to molecular simulation. Our theoretical results are supported by numerical simulations.

Keywords. Preconditioned Overdamped Langevin algorithm, Asymptotic Variance, Central Limit Theorem, Model Hamiltonians, Lattice Model.

1. Introduction

The problem of convergence to equilibrium for diffusion processes has attracted considerable attention in recent years. Due to the significant computational cost associated with MCMC type algorithms it is important to understand, and where possible accelerate, the convergence to equilibrium of systems in statistical physics, materials science, biochemistry, machine learning and many other areas.

In such applications it is typically necessary to compute expectations of the form

$$\mu(f) = \int_{\mathbb{R}^N} f(x) \mu(dx) \quad (1.1)$$

of an observable f with respect to a target probability distribution $\mu(dx)$ on \mathbb{R}^N , where μ is of the form

$$\mu(dx) = Z_\mu^{-1} e^{-E_N} dx,$$

where $E_N : \mathbb{R}^N \rightarrow \mathbb{R}$ is e.g. a potential energy and Z_μ^{-1} the normalisation constant. Note that we have absorbed temperature into E_N . The usual difficulty is that the integration (1.1) cannot be performed directly due to the high dimensionality of the problem and MCMC methods are instead employed.

Ill-conditioning of E_N , which can be induced by a variety of mechanisms, but in particular high-dimensionality (large N), is a common challenge to overcome in order to construct an efficient sampling scheme. An attractive approach is to *precondition* the MCMC algorithm. The algorithm is transformed by a well-chosen operator (the preconditioner) in a way that does not alter the invariant measure but (hopefully) accelerates convergence.

The purpose of this paper is to explore to what extent (or, whether at all) preconditioning of a Langevin-type algorithm helps to accelerate the computation of expectations. Our study is motivated by recent advances such as the Riemannian Manifold MALA [citegiro09cite](#), Stochastic Newton Methods [22], non-reversible diffusions [7, 21], optimal scaling for Langevin algorithms [2, 30] and affine-invariant sampler [14].

All of these references present different variants of preconditioning or related modifications of MCMC methods and result in an improved performance. Reviewing all these different approaches would go beyond the scope of this introduction, however, we mention the Riemannian Manifold Monte Carlo method [8], [34] as the main motivation for the present study. In this method, preconditioning is understood as performing MCMC on a Riemannian manifold defined via a “metric” $P(x)$, the preconditioner. A range of preconditioners, ideally the hessian $\nabla^2 E_N(x)$ or a positive definite matrix $P(x)$ closely approximating the Hessian, are tested. In a broad range of examples (including, e.g., logistic regression, a stochastic volatility model, and an ODE inference example) it has been shown in numerical tests that a well-chosen preconditioner improves both the mixing time as well as the convergence of the probability density functions to the target measure, with speed-ups ranging from moderate $O(1)$ factors to orders of magnitude depending on the application.

Motivated by these promising results we applied analogous preconditioned sampling algorithms to some model molecular systems, but did not always (rarely in fact) observe the speed-ups we expected. For instance, in Figure 1.1 we show the convergence of reconstructed free energy profiles from metadynamics [19] simulations comparing the unpreconditioned and preconditioned dynamics using a Hessian-based preconditioner with varying parameter sets. This type of preconditioner usually yields at least an order of magnitude speed-up in typical geometry optimisation tasks, but fails to accelerate the assembly of the free energy surface (see Appendix A.2 for the precise setup of the test). The question naturally arises whether this failure is due to a lack of fine-tuning or due to a more fundamental limitation.

Thus, to understand better these observations we will study some elementary analytical and numerical examples, which capture some essential characteristics of typical molecular systems, but where explicit results can still be obtained. The origin of the difficulty comes from the fact that it is highly dependent not only on μ but also on the observable f whether preconditioning can achieve a significant (or, any) speed-up. We will demonstrate that for some typical observables f , even moderate preconditioning can achieve significant speed-ups while for other, equally common, observables no speed-up should be at all expected even if E_N is highly ill-conditioned. Moreover, we will show that the dimensionality (1D, 2D, 3D) of the molecular structure plays a crucial role.

While our discussion is primarily motivated by applications in molecular simulations, it should straightforwardly adapt the arguments and findings to other application areas.

1.1. Langevin Algorithm

The most commonly employed algorithms in molecular simulation are based on discretising the Langevin equation, but for the sake of simplicity we will focus on the overdamped Langevin equation,

$$dX_t = -\nabla E_N(X_t)dt + \sqrt{2}dW_t, \quad (1.2)$$

where W_t is a N -dimensional standard Brownian motion. Under mild technical conditions on E_N and X_t it is known that the dynamics $(X_t)_{t \geq 0}$ is ergodic with respect to the measure μ (see [32]).

Discretising in time,

$$X_{m+1} = X_m - \delta \nabla E_N(X_m) + \sqrt{2\delta}R_m, \quad (1.3)$$

where $R_m \sim N(0, I_{N \times N})$ and δ is a parameter quantifying the size of the discrete time increment (time-step), we obtain a Markov chain with invariant measure $\tilde{\mu}$, where typically $\tilde{\mu} - \mu \sim O(t^{-1/2}) + O(\delta)$. Here, $O(t^{-1/2})$ represents the statistical error due to the finite length of the simulation, while $O(\delta)$ represents the bias due to time discretisation.

Adding metropolisation to (1.3) leads to the Metropolis Adjusted Langevin Algorithm (MALA). It is well known that the measure μ is invariant for the MALA [31, 32]. In molecular simulation it is common to assume that the bias in $\tilde{\mu}$ is negligible compared to modelling and statistical errors and




FIGURE 1.1. RMS errors of the reconstructed free energy surface profiles from metadynamics simulations without (black dashed line) and with preconditioning (coloured solid lines). Error bars represent one standard deviation based on 10 independent simulations. See A for the detailed setup.

therefore we will not consider metropolisation in the present paper. The resulting algorithm is called the Unadjusted Langevin Algorithm (ULA).

As a matter of fact, we will argue in Section 2.4 that the issues we are addressing are largely unrelated to the time-discretisation, hence we will focus on the continuous process (1.2) instead of the Markov chain (1.3).

Thus, given an observable $f \in L^1(\mu)$, we are interested in quantifying the convergence

$$\epsilon_T(f) := \frac{1}{T} \int_0^T f(X_t) dt \rightarrow \mu(f) \quad \text{for } \mu\text{-a.e } X_0. \quad (1.4)$$

Under additional assumptions on μ and f , this convergence result is accompanied by a central limit theorem which characterises the asymptotic distribution of the fluctuations, i.e.

$$\sqrt{t}(\epsilon_t(f) - \mu(f)) \xrightarrow{D} \mathcal{N}(0, \sigma_f^2), \quad (1.5)$$

where σ_f^2 is known as the asymptotic variance for the observable f (see [6, 18]).

In Section 4 we will explicitly compute σ_f^2 for some simple energy functionals and observables, which mimic typical objects of interest in molecular simulations and demonstrate that for some typical observables, σ_f may be strongly dependent on the conditioning of E_N , in particular on system size N , while for others the dependence on N is negligible.

1.2. Preconditioned Langevin Algorithm

For a (fixed) preconditioner $P \in \mathbb{R}^{N \times N}$, symmetric positive definite, we consider the preconditioned Overdamped Langevin dynamics (P -Langevin),

$$dX_t^P = -P^{-1} \nabla E_N(X_t^P) dt + \sqrt{2} P^{-1/2} dW_t. \quad (1.6)$$

The standard overdamped Langevin dynamics is recovered by taking $P = I$. It is in principle possible to allow $P = P(X)$ but for simplicity we will not consider this in the present work. Note that, if P is fixed then the coordinate transform $Z = P^{1/2} X$ allows us to easily lift results from (1.2) to (1.6); see Section 2.1 for more details.

The preconditioner does not affect the invariance property of the diffusion process, i.e., the target measure μ is still invariant for the P -Langevin process (1.6). However, it can affect the convergence of the process to the invariant measure. That is, for $X_t^P \sim \mu_t^P = \varphi_t^P dx$, we will characterise in Section 3 the rate of convergence of φ_t^P to φ_∞ (φ_∞ denotes the density of μ) and demonstrate how preconditioning improves this rate.

Moreover, we also obtain

$$\epsilon_T^P(f) := \frac{1}{T} \int_0^T f(X_t^P) dt \rightarrow \mu(f), \quad \text{for } \mu\text{-a.e } X_0. \quad (1.7)$$

and analogously to the standard Langevin dynamics, a central limit theorem characterizes the asymptotic distribution of the fluctuations,

$$\sqrt{t} \left(\epsilon_t^P(f) - \mu(f) \right) \xrightarrow{D} \mathcal{N}(0, \sigma_{f,P}^2), \quad (1.8)$$

where $\sigma_{f,P}^2$ is the asymptotic variance of f under P -Langevin dynamics.

The main aim of our paper is to present several simplified but still realistic examples at which we can observe *whether or not* preconditioning accelerates sampling in the sense that it achieves a reduction in the asymptotic variance, i.e., $\sigma_{f,P}^2 \ll \sigma_f^2$.

Outline

The rest of the paper is organised as follows. In Section 2 we present the model Hamiltonians, E_N , to motivate some key assumptions that we make throughout our analysis. We also recall the coordinate transform that we use to reduce P -Langevin to standard Langevin, and we explain why the time-discretisation is a negligible component and can therefore be ignored for our analysis.

In Section 3 we show a long time convergence result to the invariant measure in the P -Langevin process. For some quadratic model Hamiltonians we can then precisely quantify the speed-up afforded through preconditioning.

Section 4 is devoted to describe how the central limit theorem (1.8) arises from the solution of the Poisson equation associated with the generator of the dynamics. This is then followed by a detailed analysis of σ_f^2 and $\sigma_{f,P}^2$ for some quadratic model Hamiltonians and observables.

In Section 5 we show a numerical application to illustrate the reduction in the asymptotic variance.

2. Preliminaries

2.1. Coordinate transformation

As mentioned above, a convenience afforded by our assumption that the preconditioner P is constant, is that a simple coordinate transformation can transform the P -Langevin dynamics (1.6) into standard Langevin dynamics (1.2) by taking $Z_t = P^{1/2} X_t$. We will now briefly review this transformation.

First, let us introduce the new coordinates and associated energy

$$z := P^{1/2}x \quad \text{and} \quad E_N^P(z) := E_N(x) = E_N(P^{-1/2}z).$$

The standard overdamped Langevin dynamics for z then reads

$$\begin{aligned} dZ_t &= -\nabla_z E_N(P^{-1/2}Z_t)dt + \sqrt{2}dW_t \\ &= -P^{-1/2}\nabla_x E_N(P^{-1/2}Z_t) + \sqrt{2}dW_t. \end{aligned} \quad (2.1)$$

Upon multiplying the equation with $P^{-1/2}$ we clearly recover (1.6); that is, (2.1) and (1.6) are indeed equivalent. In the new coordinate system the infinitesimal generator operator of the diffusion process Z_t is given by:

$$\mathcal{L}_P = -\nabla_z E_N^P(z) \cdot \nabla_z + \Delta_z. \quad (2.2)$$

In terms of estimating observables, we obtain that

$$e_T^P(f) = \frac{1}{T} \int_0^T f(X_t^P)dt = \frac{1}{T} \int_0^T f(P^{-1/2}Z_t)dt. \quad (2.3)$$

The key observation in analysing how preconditioning changes the properties of E_N^P and hence of the Langevin dynamics is that

$$\nabla_z^2 E_N^P(z) = P^{-1/2} \nabla_x^2 E_N(x) P^{-1/2}.$$

In the following, if H and P are symmetric positive definite, we will write $H_P := P^{-1/2}HP^{-1/2}$. We also observe, for future reference, that the spectrum of H_P satisfies

$$\inf \sigma(H_P) = \inf_{v^T P v = 1} v^T H v \quad \text{and} \quad \sup \sigma(H_P) = \sup_{v^T P v = 1} v^T H v.$$

2.2. Model Hamiltonians

Let $x \in \mathbb{R}^N$ describe a system of N/d particles at positions $(y_\alpha)_{\alpha=1}^{N/d} \subset \mathbb{R}^d$, then a simple model for potential energy is given by

$$E_N(y) = \sum_{\alpha=1}^{N/d} \sum_{\beta \neq \alpha} \phi(|y_\alpha - y_\beta|),$$

where ϕ is e.g. a Lennard-Jones type potential, $\phi(r) = r^{-12} - 2r^{-6}$. Such systems exhibit complex meta-stable behaviour, which is an issue to be entirely separated from the ill-conditioning due to high dimension.

A much simpler situation is a mass-spring model, where $u = (u_\alpha) \in \mathbb{R}^N$, with u_α e.g. denoting out-of-plane displacement of a particle, and particle connectivity described by an equivalence relation $\alpha \sim \beta$. Then the energy can be written as

$$E_N(u) = \sum_{\alpha \sim \beta} \phi_{\alpha\beta}(u_\alpha - u_\beta), \quad (2.4)$$

where $\phi_{\alpha\beta}$ could be taken as strictly convex to avoid meta-stability. We will use systems of this kind in our numerical experiments.

In order to admit explicit analytical calculations we simplify further E_N by expanding it about an equilibrium (e.g., at $u = 0$) which then yields a quadratic energy

$$E_N(u) = \sum_{\alpha \sim \beta} \frac{k_{\alpha\beta}}{2} |u_\alpha - u_\beta|^2 =: \frac{1}{2} u^T H u. \quad (2.5)$$

The spring constants $k_{\alpha\beta}$ could model how the interaction between different atomic species / environments differs. We assume throughout that there exist bounds \underline{k}, \bar{k} on the spring constants and a

bound \bar{n} on the number of neighbours, which are both independent of N . More precisely, we assume that

$$0 < \underline{k} \leq k_{\alpha\beta} \leq \bar{k} \quad \text{and} \quad \#\{\beta : \alpha \sim \beta\} \leq \bar{n}. \quad (2.6)$$

Note that we have chosen a scaling of the energies where increasing system size $N \rightarrow \infty$ does not yield a continuum limit but rather an infinite lattice system. The consequence is that if we have a bound in the spring constants $0 < k_{\alpha\beta} \leq \bar{k}$, and each atom α is connected with at most \bar{n} neighbours, then it follows readily that

$$|E_N(u)| \lesssim \bar{n}\bar{k}|u|^2 \quad \text{or, equivalently} \quad \|H\| \lesssim \bar{n}\bar{k}, \quad (2.7)$$

where $\|H\|$ denotes the $\ell^2 \rightarrow \ell^2$ operator norm.

To be even more specific let us assume that E_N is a d -dimensional lattice model, i.e., each atom index α corresponds to a coordinate $\ell_\alpha \in \{0, \dots, M+1\}^d$ and $\alpha \sim \beta$ if and only if $|\ell_\alpha - \ell_\beta| = 1$, that is

$$E_N(u) := \sum_{\alpha_1=0}^M \cdots \sum_{\alpha_d=0}^M \sum_{j=1}^d k_{\alpha, \alpha+e_j} |u(\alpha + e_j) - u(\alpha)|^2. \quad (2.8)$$

By clamping the boundary sites at $u = 0$, we obtain $N = M^d$ free lattice sites. In this case, H is a (possibly inhomogeneous) discrete elliptic operator and employing (2.6) and using the min-max characterisation of eigenvalues (see, e.g., Sect XIII.1 in ReedSimon) to compare H to the discrete Laplacian for which the spectrum can be computed explicitly [15], we can readily show that there exist constants c_0, c_1 such that

$$c_0(j/N)^{2/d} \leq \lambda_j \leq c_1(j/N)^{2/d}, \quad (2.9)$$

where $\sigma(H) = \{\lambda_j | j = 1, \dots, N\}$ is the ordered spectrum of H . This dimension-dependence of the eigenvalue distribution will be important later on.

2.3. A model preconditioner

Although for the simple model problems described in Section 2.2 it is straightforward to compute Hessians and use those as preconditioners, this would remove us from the practice of molecular simulations where Hessians are not normally computable. Instead, we will consider preconditioners that only roughly capture the structure of the energy functionals and their Hessians.

Following [29], we will use a preconditioner of the form

$$v^T P v = c \sum_{\alpha \sim \beta} |v_\alpha - v_\beta|^2, \quad (2.10)$$

where c is a free parameter to be fitted to the model. The idea is that it captures the connectivity information but not the fine details of the Hessian. For optimisation and saddle search, preconditioners of this kind have been shown to yield considerable speed-ups even for much more complex electronic structure type models [29].

For example, considering E_N given by (2.5) with hessian $H := \nabla^2 E_N$ and choosing $c := \frac{1}{2}(\underline{k} + \bar{k})$, we observe that

$$\frac{2\underline{k}}{\underline{k} + \bar{k}} v^T P v \leq v^T H v \leq \frac{2\bar{k}}{\underline{k} + \bar{k}} v^T P v, \quad (2.11)$$

which in particular implies that

$$\kappa(H_P) \leq \frac{\bar{k}}{\underline{k}}.$$

Here, $\kappa(A) := \|A\| \|A^{-1}\|$ denotes the condition number of the matrix A .

That is, provided the inhomogeneity \bar{k}/\underline{k} is not too severe, then the preconditioned energy landscape has only very moderate conditioning, independent of N , while typically $\kappa(\nabla^2 E) \rightarrow \infty$ as $N \rightarrow \infty$, with

a rate depending on the connectivity. For instance, for a lattice model (2.9) implies that $\kappa(\nabla^2 E_N) \sim N^{2/d}$.

2.4. Time Discretisation

To measure the cost/accuracy ratio of a practical sampling algorithm based on the Langevin process we need to account also for the time-discretisation,

$$X_{n+1} = X_n - \delta P^{-1} \nabla E_N(X_n) + \sqrt{2\delta} P^{-1/2} R_n, \quad (2.12)$$

where $R_n \sim N(0, I_{N \times N})$.

It is tempting to assume that one advantage afforded by preconditioning is to take larger time-steps. As we show in the following, this is a matter of scale, thus justifying our choice to focus purely on the time-continuous P-Langevin dynamics.

In the scaling that we have chosen in the model problems of Section 2.2, both the Hessians H and preconditioned Hessians H_P are bounded operators (on ℓ^2), independent of system size N . Thus, the time-steps have similar restrictions for the preconditioned and unpreconditioned Langevin processes [10]. More precisely, in the quadratic model problem (2.5) we may assume $\bar{k} = 1$ without loss of generality. Then, using the model preconditioner (2.10) with $c = \bar{k} = 1$, say, we have that $\sup \sigma(H) = \sup \sigma(H_P)$.

Proposition 2.1. *Suppose that $E_N(x) = \frac{1}{2}x^T Hx$ and $P \in \mathbb{R}^{N \times N}$ with H and P both positive definite. Suppose further that $\frac{\delta}{2}|P^{-1}H| < 1$. Then the invariant measure for (2.12) is Gaussian with covariance matrix*

$$C_\delta^P = (I - \frac{\delta}{2}P^{-1}H)^{-1}H^{-1}.$$

Proof. See Section 6.1. ■

It follows that the error in the covariance operator is $O(\delta)$ with constant independence of N , and independent of the choice of preconditioner. Indeed, even with an “optimal” preconditioner $P = H$, one would not obtain an improvement in the bias: if there exist eigenmodes $Hv = \lambda v$ with $\lambda \ll 1$, then

$$\begin{aligned} |C_\delta^H v - Cv| &= |(1 - \frac{\delta}{2})^{-1} - 1| |Cv|, \\ \text{whereas } |C_\delta^I v - Cv| &= |(1 - \frac{\lambda\delta}{2})^{-1} - 1| |Cv|, \end{aligned}$$

that is,

$$\frac{|C_\delta^H v - Cv|}{|C_\delta^I v - Cv|} \sim \lambda^{-1} \quad \text{for } \delta, \lambda \text{ sufficiently small,}$$

where $C = H^{-1}$ is the covariance operator of the unbiased measure. Therefore, for the remainder of the paper, we will not consider the effect of preconditioning on time-discretisation but only focus on the speed of convergence to equilibrium in the undiscretised (P-)Langevin dynamics. We only stress that this point of view is only valid as long as $\|H\|$ and $\|H_P\|$ are comparable.

3. Exponential convergence to the invariant measure

In this section, we prove exponential convergence to the equilibrium. For the sake of simplicity we represent the probability density functions μ and μ_t^P by their respective densities φ_∞ and φ_t^P . Under the transformation $z = P^{1/2}x$, we obtain transformed probability densities

$$\psi_\infty(z) := (\det P)^{-1/2} \varphi_\infty(P^{-1/2}z) \quad \text{and} \quad \psi_t^P(z) := (\det P)^{-1/2} \varphi_t^P(P^{-1/2}z).$$

Their evolution is described by the Fokker–Planck equation

$$\partial_t \psi_t^P = \mathcal{L}_P^* \psi_t^P := \nabla_z \cdot \left(\nabla_z E_N(P^{-1/2}z) \psi_t^P + \nabla_z \psi_t^P \right), \quad (3.1)$$

where \mathcal{L}_P^* is the dual operator of \mathcal{L}_P (for the $L^2(dz)$ –scalar product) defined in (2.2).

The proof of the following theorem is inspired from [21] or A.19 in Villani.

Theorem 3.1. *Suppose that $E_N \in C^2(\mathbb{R}^N)$, such that $\frac{1}{2}|\nabla E_N(x)|^2 - \Delta E_N(x) \rightarrow +\infty$ as $|x| \rightarrow +\infty$. Then there exists $\lambda_P > 0$ such that for all initial conditions $\psi_0 \in L^2(\varphi_\infty^{-1})$, and for all times $t \geq 0$*

$$\begin{aligned} \|\psi_t^P - \psi_\infty\|_{L^2(\psi_\infty^{-1})}^2 &\leq e^{-\lambda_P t} \|\psi_0 - \psi_\infty\|_{L^2(\psi_\infty^{-1})}^2, & \text{or, equivalently,} \\ \|\varphi_t^P - \varphi_\infty\|_{L^2(\varphi_\infty^{-1})}^2 &\leq e^{-\lambda_P t} \|\varphi_0 - \varphi_\infty\|_{L^2(\varphi_\infty^{-1})}^2, \end{aligned}$$

$\|\cdot\|_{L^2(\psi_\infty^{-1})}$ denotes the norm in $L^2(\mathbb{R}^N, \psi_\infty^{-1} dz)$, namely $\|f\|_{L^2(\psi_\infty^{-1})}^2 = \int_{\mathbb{R}^N} f^2(z) \psi_\infty^{-1} dz$. The exponent λ_P is the spectral gap of the Fokker–Planck operator \mathcal{L}_P^* defined in (3.1) (i.e., the smallest non-zero eigenvalue of $-\mathcal{L}_P^*$).

Proof. See Section 6.2. ■

We now consider the quadratic case, $E_N(x) = \frac{1}{2}x^T H x$ with H symmetric and positive definite. In this case, we have the following characterisation of the spectrum of \mathcal{L}_P^* and hence of λ_P . The proof is based on [12, 21, 28].

Theorem 3.2. *The spectrum of the operator \mathcal{L}_P^* is*

$$\sigma(\mathcal{L}_P^*) = \sigma(\mathcal{L}_P) = \left\{ -\sum_{\lambda \in \sigma(H_P)} k_\lambda \lambda : k_\lambda \in \mathbb{N} \right\}. \quad (3.2)$$

Proof. See Section 6.4. ■

An immediate consequence of Theorem 3.2 is that

$$\lambda_P = \inf \sigma(H_P) \setminus \{0\}, \quad (3.3)$$

in other words, the smallest non-zero generalised eigenvalue of

$$Hv = \lambda P v.$$

Returning to the quadratic model Hamiltonians and model preconditioners introduced in Sections 2.2 and 2.3, assume that $\min \sigma(H) \sim N^{-s}$ (which is consistent with (2.9)), while $\min \sigma(H_P) \sim 1$, then we obtain that $\lambda_I \sim N^{-s}$ with $\lambda_P \sim 1$. This result should give us confidence in the value of preconditioning.

Remark 3.3. The preconditioning ideas and results presented in Sections 2 and 3 are similar to the Brascamp–Lieb inequality [4]. In some sense, this inequality claims that a good preconditioner is the Hessian.

Precisely [20], if E_N is strictly convex, for any function $f \in L^2(e^{-E_N})$,

$$\int \left[f - \int f e^{-E_N} \right]^2 e^{-E_N} \leq \int \nabla f (\nabla^2 E_N)^{-1} \cdot \nabla f e^{-E_N},$$

where we assume here that the normalization $\int e^{-E_N} = 1$. This means that if one considers the Fokker–Planck equation

$$\partial_t \psi = \operatorname{div}[(\nabla^2 E_N)^{-1} e^{-E_N} \nabla(\psi e^{E_N})],$$

which is associated to the following overdamped Langevin dynamics [34]:

$$dX_t = -(\nabla^2 E_N)^{-1} \nabla E_N(X_t) dt + \operatorname{div}[(\nabla^2 E_N)^{-1}](X_t) dt + \sqrt{2(\nabla^2 E_N)^{-1}}(X_t) dW_t,$$

then, if E_N is strictly convex, we have

$$\begin{aligned} \frac{1}{2} \frac{d}{dt} \int \left(\frac{\psi}{e^{-E_N}} - 1 \right)^2 e^{-E_N} &= - \int (\nabla^2 E_N)^{-1} \nabla \left(\frac{\psi}{e^{-E_N}} \right) \cdot \nabla \left(\frac{\psi}{e^{-E_N}} \right) e^{-E_N} \\ &\leq \int \left(\frac{\psi}{e^{-E_N}} - 1 \right)^2 e^{-E_N}. \end{aligned}$$

And thus

$$\int \left(\frac{\psi}{e^{-E_N}} - 1 \right)^2 e^{-E_N} \leq \left[\int \left(\frac{\psi_0}{e^{-E_N}} - 1 \right)^2 e^{-E_N} \right] e^{-2t},$$

whatever the potential E_N is. It is in some sense “universal”. For example, it does not depend on the temperature: if we multiply E_N by a constant β (inverse of the temperature) it remains the same, whereas understanding the dependency of the spectral gap on the temperature is tricky in general.

4. Analysis of the asymptotic variance

In this section we present sufficient conditions under which the estimator

$$\epsilon_T^P(f) = \frac{1}{T} \int_0^T f(X_t^P) dt = \frac{1}{T} \int_0^T f(P^{-1/2} Z_t) dt$$

satisfies a central limit theorem of the form (1.8) and we characterise the associated asymptotic variance.

4.1. Generalities

The fundamental requirements to prove the central limit theorem is establishing the well-posedness of the Poisson equation

$$- \mathcal{L}_P \phi(z) = f(P^{-1/2} z) - \mu(f), \quad \mu(\phi) = 0, \quad (4.1)$$

for all bounded and continuous functions $f : \mathbb{R}^N \rightarrow \mathbb{R}$, where \mathcal{L}_P is defined by (2.2), and obtaining estimates on the growth of the unique solution ϕ . Recall that we shall assume that μ admits a smooth, strictly positive density denoted by $\psi_\infty(z)$, such that $\int_{\mathbb{R}^N} \psi_\infty(z) dz = 1$ and the SDE (2.1) has a unique strong solution.

Referring to results in [9, 24] we suppose that the process Z_t admits a Lyapunov function (see the Definition 6.1 in Section 6.5), which is sufficient to ensure the geometric ergodicity of Z_t (see [16, 23]). In terms of the potential energy E_N and the preconditioner P , we require that there exists $\beta \in (0, 1)$ such that

$$\liminf_{|z| \rightarrow +\infty} \left[(1 - \beta) |\nabla E_N^P(z)|^2 + \Delta E_N^P(z) \right] > 0. \quad (4.2)$$

It is straightforward to check that this condition holds whenever E_N is strongly convex and in particular if it is of the form (2.5).

If condition (4.2) holds, then the process Z_t will be geometrically ergodic. More specifically, the law of the process Z_t started from a point $z \in \mathbb{R}^N$ will converge exponentially fast in the total-variation norm to the equilibrium distribution μ (cf. (1.7)).

Assuming (4.2), we also obtain the following well-posedness result for the Poisson equation (4.1).

Theorem 4.1. *Suppose that (4.2) holds, then there exists $c > 0$, such that for any measurable observable f satisfying $|f|^2 \leq e^{-\beta E_N^P(z)}$, the Poisson equation (4.1) admits a unique strong solution satisfying the bound $|\phi(z)|^2 \leq ce^{-\beta E_N^P(z)}$. In particular, $\phi(P^{1/2} x) \in L^2(\mu)$.*

Proof. See Section 6.5. ■

The technique of using a Poisson equation to obtain a central limit theorem for an additive functional of a Markov process is widely known (see e.g. [3]). For linear and quadratic observables, we can in fact produce an analytic solution to this Poisson problem; see Section 4.2 below.

Theorem 4.2. *Under the conditions of Theorem 4.1, there exists a constant $0 < \sigma_{f,P}^2 < \infty$ such that the asymptotic distribution of the fluctuations of $\epsilon_t^P(f)$ about $\mu(f)$ are given by the central limit theorem*

$$\sqrt{t} \left(\epsilon_t^P(f) - \mu(f) \right) \xrightarrow{D} \mathcal{N}(0, \sigma_{f,P}^2), \text{ as } t \rightarrow +\infty, \quad (4.3)$$

where $\sigma_{f,P}^2$ (the asymptotic variance for the observable f) is given by

$$\begin{aligned} \sigma_{f,P}^2 &= 2 \int |\nabla_z \phi(z)|^2 \mu^P(dz) \\ &= 2 \int |\nabla_z \phi(P^{1/2}x)|^2 \mu(dx), \end{aligned} \quad (4.4)$$

where $\mu^P(dz) = \psi_\infty(z)dz$.

Proof. See Section 6.6. ■

4.2. Explicit solution for linear observables

In this section we exhibit explicit solutions ϕ of the Poisson equation (4.1) when E_N is quadratic and f is linear, and compute the associated asymptotic variance $\sigma_{f,P}^2$. This simplest possible case is of course well-known but we summarise it nevertheless to prepare for more interesting cases.

Suppose, therefore, that

$$P = I, \quad E_N(x) = \frac{1}{2}x^T Hx \quad \text{and} \quad f(x) = v \cdot x,$$

where $v \in \mathbb{R}^N$. From symmetry it follows that $\mu(f) = 0$, hence the Poisson equation (4.1) becomes

$$Hx \cdot \nabla \phi(x) - \Delta \phi(x) = v \cdot x. \quad (4.5)$$

Seeking a solution of the form $\phi(x) = d \cdot x$ with $d \in \mathbb{R}^N$, we obtain $d = H^{-1}v$, i.e.,

$$\phi(x) = x \cdot H^{-1}v,$$

and in particular,

$$\sigma_f^2 = \sigma_{f,I}^2 = 2 \int |H^{-1}v|^2 \mu(dx) = 2|H^{-1}v|^2.$$

In particular, choosing v to be a normalised eigenmode of H with associated eigenvalue λ , we obtain $\sigma_f^2 = 2\lambda^{-2}$.

Focusing specifically on a d -dimensional lattice model, we know from (2.9) that $\min \sigma(H) \sim N^{-2/d}$ while $\max \sigma(H) \sim 1$. Hence, it follows that σ_f^2 is moderate for the high-frequency eigenmodes, but large for the observables corresponding to low-frequency eigenmodes.

Next, we turn to the preconditioned dynamics. In this case we effectively replace H with $H_P = P^{-1/2}HP^{-1/2}$ and v with $P^{-1/2}v$ and thus obtain

$$\sigma_{f,P}^2 = 2 \int |H_P^{-1}P^{-1/2}v|^2 \mu(dx) = 2|P^{1/2}H^{-1}v|^2 =: 2|H^{-1}v|_P^2. \quad (4.6)$$

If we assume that $c_0 = \min \sigma(H_P)$, $c_1 = \max \sigma(H_P)$, then a simple rewrite yields

$$2c_0 v^T H^{-1} v \leq \sigma_{f,P}^2 \leq 2c_1 v^T H^{-1} v. \quad (4.7)$$

Comparing with $\sigma_f^2 = 2v^T H^{-2} v$ and recalling our standing assumption (2.11) that $c_0 \sim 1$, $c_1 \sim 1$ (the spectrum is bounded above and below independently of N), we conclude that preconditioning does not entirely remove ill-conditioning but it is (potentially) diminished.

More concretely, for a d -dimensional lattice system, we obtain that

$$N^{-2/d} \lesssim \frac{\sigma_{f,P}^2}{\sigma_f^2} \lesssim 1,$$

and both bounds are attained for specific observables. We conclude that preconditioned Langevin *can* be significantly more efficient than standard Langevin (low-frequency observables) but that it will be *comparable in efficiency* for high-frequency observables.

Intuitively, low-frequency observables are “macroscopic” in nature and include e.g. energy, average bond-length etc., while high-frequency observables include in particular single bonds, bond angles and dihedral angles (in a large molecule) or a bond-length near a crack-tip. In the next sections, we consider three toy models mimicking “realistic” observables of these kinds, occurring in real-world simulations, to further substantiate our remarks.

4.3. Example 1: Energy per particle

We now consider $E_N(x) = \frac{1}{2} x^T H x$ and $f(x) = N^{-1} E_N(x)$. A straightforward computation yields

$$\langle x \cdot B x \rangle_\mu = \text{Tr}(H^{-1} B) \quad \text{for } B \in \mathbb{R}^{N \times N}, \quad (4.8)$$

which in particular implies that

$$\mu(f) = \frac{1}{N} \frac{\int E_N(x) e^{-E_N(x)}}{\int e^{-E_N(x)}} = \frac{\text{Tr} I}{2N} = \frac{1}{2}.$$

Thus, the Poisson equation becomes

$$H x \cdot \nabla \phi(x) - \Delta \phi(x) = \frac{1}{2N} x^T H x - \frac{1}{2}$$

We seek a solution of the form $\phi(x) = \frac{1}{2} x^T B x + l \cdot x - \text{Tr} B$ (to ensure that $\mu(\phi) = 0$), then this yields the equation

$$H x \cdot (B x + l) - \text{Tr} B = \frac{1}{2N} x^T H x - \frac{1}{2}.$$

This is satisfied for $B = \frac{1}{N} I$, $l = 0$, hence $\phi(x) = \frac{1}{2N} |x|^2 - \frac{1}{2}$.

We can now compute the asymptotic variance as

$$\sigma_f^2 = \sigma_{f,I}^2 = 2 \int \left| \frac{1}{N} x \right|^2 \mu(dx) = \frac{2}{N^2} \text{Tr} H^{-1}.$$

Repeating the same argument in transformed coordinates $z = P^{1/2} x$, we also obtain the asymptotic variance of the energy for the preconditioned Langevin dynamics:

$$\sigma_{f,P}^2 = \frac{2}{N^2} \text{Tr} H_P^{-1} = \frac{2}{N^2} \text{Tr}(H^{-1} P).$$

Let us now focus on a lattice model, where we have (2.9). Then we obtain that

$$\frac{1}{2} \sigma_f^2 \approx N^{-2} \sum_{j=1}^N (j/N)^{-2/d} \approx N^{-1} \int_{1/N}^1 s^{-2/d} ds \approx \begin{cases} 1, & d = 1, \\ N^{-1} \log N, & d = 2, \\ N^{-1}, & d = 3, \end{cases}$$

while, clearly, $\frac{1}{2}\sigma_{f,P}^2 \approx N^{-1}$. In summary,

$$\frac{\sigma_f^2}{\sigma_{f,P}^2} \approx \begin{cases} N, & d = 1, \\ \log N, & d = 2, \\ 1, & d = 3; \end{cases}$$

that is, preconditioning only gives a significant speed-up in one-dimensional lattices but not in two- or three-dimensional lattices.

4.4. Example 2: Bond-length

In our second example we observe a single bond in the crystal or molecule. That is, we still use $E_N(x) = \frac{1}{2}x^T Hx$ but the observable is now given by

$$f(x) = x_i - x_j \quad \text{for some fixed bond } i \sim j.$$

This is a linear observable, hence a special case of the discussion in Section 4.2. Hence, we obtain

$$\sigma_f^2 = 2|H^{-1}l|^2 \quad \text{where } l_n = \begin{cases} 1, & n = i \\ -1, & n = j, \\ 0, & \text{otherwise.} \end{cases}$$

In order to estimate σ_f^2 further we consider again the d -dimensional lattice model (2.8) and P given by (2.10). For $d \geq 2$, since P is a homogeneous discrete elliptic operator, we know from [27] that

$$|[P^{-1}]_{ni} - [P^{-1}]_{nj}| \leq C(1 + |n - i|)^{-d},$$

where we note that i, j are now neighbouring lattice sites; i.e., $[P^{-1}]_{ni} - [P^{-1}]_{nj}$ denotes a discrete gradient of the lattice Green's function.

Therefore, we obtain that

$$\sigma_f^2 = 2|H^{-1}l|^2 \lesssim 2|P^{-1}l|^2 \lesssim 2 \sum_{n \in \mathbb{Z}^d} (1 + |n - i|)^{-2d} < \infty; \quad (4.9)$$

that is, σ_f^2 has an upper bound that is independent of N . A lower bound follows simply from the fact that $\|H\| \leq 1$ and hence $v^T H^{-1}v \geq |v|^2$, which implies

$$\sigma_f^2 \geq |l|^2 = 1.$$

To obtain bounds on $\sigma_{f,P}^2$, we use (4.7) to estimate

$$2c_0 = 2c_0|l|^2 \leq \sigma_{f,P}^2 \leq 2c_1 l^T H^{-1} l \leq 2c_1 |l| \|H^{-1}l\|,$$

and we have already shown in (4.9) that this is bounded above, independently of N .

In summary, we obtain that

$$\frac{\sigma_f^2}{\sigma_{f,P}^2} \lesssim 1 \quad \text{for } d \geq 2,$$

that is, we expect no substantial (if any) speed-up for the bond-length observable from preconditioning for $d \geq 2$.

By contrast, for $d = 1$, the system $P^{-1}l$ can be solved explicitly, and in this case one obtains $\sigma_f^2 \sim N$ as $N \rightarrow \infty$ (specifically, $|P^{-1}l|^2 = N/12$), that is,

$$\frac{\sigma_f^2}{\sigma_{f,P}^2} \sim N \quad \text{for } d = 1.$$

Thus, we conclude that preconditioning helps to accelerate the computation of the bond-length observable only for one-dimensional structures. The intuitive explanation of this effect is that far-away regions of space have little influence on a single bond and hence only local equilibration matters. The difference in 1D is that elastic interaction is naturally more long-ranged than in dimension $d > 1$.

4.5. Example 3: Umbrella sampling

Our final example is inspired by a technique called umbrella sampling [17]. Given a potential energy $E_N(x)$ and a reaction coordinate $\xi(x)$, we wish to compute

$$A(\xi_0) := -\log \int e^{-E_N(x)} \delta(\xi(x) - \xi_0) dx.$$

Umbrella sampling achieves this by placing a restraint on the potential energy,

$$E_{N,K}(x) := E_N(x) + \frac{K}{2}(\xi(x) - \xi_0)^2,$$

for some $K > 0$, $\xi_0 \in \mathbb{R}$. Let μ_K denote the corresponding equilibrium measure, then it can be shown [17] that, for K large,

$$\left. \frac{\partial A}{\partial \xi} \right|_{\xi=\bar{\xi}_0} \approx -K(\bar{\xi}_0 - \xi_0) \quad \text{where } \bar{\xi}_0 = \langle \xi \rangle_{\mu_K}.$$

Thus, $\partial_\xi A$ and hence A can be reconstructed in this way. More sophisticated variations of the idea exist of course, but for the sake of simplicity of presentation we will focus on this particularly simple variant.

To construct an analytically accessible toy problem mimicking umbrella sampling we consider again a quadratic energy $E_N(x) = \frac{1}{2}x^T H x$ and a linear reaction coordinate $\xi(x) := l \cdot x$ where, for simplicity, we assume that $|l| = 1$ (for $|l| = O(1)$, the argument is analogous). The restrained potential with penalty parameter $K > 0$ is then given by

$$E_{N,K}(x) = \frac{1}{2}x^T H x + \frac{K}{2}(\xi(x) - \xi_0)^2,$$

for some $\xi_0 \in \mathbb{R}$, while the observable from which we can reconstruct the mean force is simply

$$f(x) = K(l \cdot x - \xi_0).$$

We are again in the context of Section 4.2 and therefore obtain

$$\sigma_f^2 = 2K^2 |H_K^{-1} l|^2,$$

where

$$H_K = \nabla^2 E_{N,K} = H + K l l^T.$$

The Sherman–Morrison formula yields

$$H_K^{-1} l = \left(H^{-1} - \frac{K H^{-1} l l^T H^{-1}}{1 + K l^T H^{-1} l} \right) l = \frac{H^{-1} l}{1 + K l^T H^{-1} l}, \quad (4.10)$$

and hence,

$$\sigma_f^2 = 2K^2 |H_K^{-1} l|^2 = \frac{2K^2 |H^{-1} l|^2}{(1 + K l^T H^{-1} l)^2}.$$

Since our focus in the present example is the ill-conditioning induced by large K rather than ill-conditioning induced by H (e.g. through system size N), let us assume that $\max \sigma(H) = 1$ (as always) while $\min \sigma(H) \geq c_0$, for some moderate constant c_0 . This would, e.g., be the typical situation for a small molecule, or if we preconditioned H but without accounting for the umbrella. We then obtain

$$c_0^2 \leq \frac{2K^2 c_0^2}{(1 + K)^2} \leq \sigma_f^2 \leq \frac{2K^2}{(1 + K c_0)^2} \leq \frac{2}{c_0^2};$$

that is, $\frac{1}{2}\sigma_f^2 \sim 1$ as $K \rightarrow \infty$.

By contrast, suppose now that we choose a preconditioner

$$P_K := P + Kl^T,$$

where P is a preconditioner for H satisfying $c_0^P v^T P v \leq v^T H v \leq v^T P v$. Then a straightforward calculation yields

$$c_0^P \leq \frac{v^T H v + K(l \cdot v)^2}{v^T P v + K(l \cdot v)^2} = \frac{v^T H_K v}{v^T P_K v} \leq 1.$$

It follows from (4.7) that

$$\sigma_{f,P}^2 = 2K^2 |H_K^{-1} l|_{P_K}^2 \approx 2K^2 |H_K^{-1/2} l|^2,$$

where \approx now indicates upper and lower bounds with constants that are independent of K .

Using (4.10) we obtain

$$\sigma_{f,P}^2 \approx 2K^2 \frac{l^T H^{-1} l}{1 + Kl^T H^{-1} l} \sim K \quad \text{as } K \rightarrow \infty.$$

We can therefore conclude that

$$\frac{\sigma_{f,P}^2}{\sigma_f^2} \approx \frac{(l^T H^{-1} l)(1 + Kl^T H^{-1} l)}{|H^{-1} l|^2} \sim 1 + K \quad \text{as } K \rightarrow \infty;$$

that is, preconditioning the umbrella actually achieves a significant *deterioration* of the asymptotic variance and thus the P -Langevin dynamics actually becomes *less efficient* than the standard Langevin dynamics. However note the following crucial remark:

Remark 4.3 (Step-sizes revisited). The surprising result of the present section does not in fact fully fall within our starting assumptions. While $\|H\| = 1$, $\|H_K\|$ is in fact of order $O(1 + K)$ which means that the time-step for the discretisation of the Langevin equation should be of order $O(K^{-1})$, which exactly balances the lower mixing of the preconditioned dynamics and make the two schemes again comparable. Indeed, in our computational examples we will need to choose $\Delta t = O(K^{-1})$ to prevent instability.

In practice, the restraint parameter K is chosen of the same order of magnitude of the stiffest bond in a molecule, while the reaction coordinate will normally be a function of the softest bonds and hence it would create no additional time-step restriction. In such a situation, it is indeed preferable to *not* precondition the restraint.

However, we emphasise again that the interaction between preconditioning and time-stepping is an issue that we do not properly address in the present work and which will require further attention in the future.

5. Numerical Tests

We conclude our discussion by demonstrating the extension of our explicit computations to a mildly non-linear lattice model. As potential energy $E_{M^d} : \mathbb{R}^{\{1, \dots, M\}^d} \rightarrow \mathbb{R}$, we choose

$$E_N(u) := \sum_{\alpha_1=0}^M \cdots \sum_{\alpha_d=0}^M \sum_{j=1}^d \phi\left(\frac{1}{\sqrt{d}}(u(\alpha + e_j) - u(\alpha))\right),$$

where $u(\alpha) := 0$ if any $\alpha_j \in \{0, M + 1\}$, and with convex nearest-neighbour pair potential

$$\phi(r) = \frac{1}{8}(r^2 + \sin(r)^2).$$

Upon choosing an arbitrary linear labelling of indices $\alpha \in \{1, \dots, M\}^d$, this is a special case of (2.4). As preconditioner we choose (2.10), which we can also write as

$$\langle Pv, v \rangle = \frac{1}{4d} \sum_{\alpha_1=0}^M \cdots \sum_{\alpha_d=0}^M \sum_{j=1}^d |u(\alpha + e_j) - u(\alpha)|^2.$$

The occurrence of d in the definitions of E_N and P ensures that $\|\nabla E_N\| \approx \|P\| \approx 1$; moreover, since ϕ is strictly convex, we have that $\text{cond}((\nabla E_N(x))_P)$ is bounded above independently of N ; cf. (2.11).

For all simulations (with small modifications for the umbrella sampling example) we choose

$$\delta = 0.1, \quad N_{\text{steps}} = 10^5, \quad N_{\text{runs}} = 400.$$

We then use a Cholesky factorisation to compute P , i.e. $P = LL^T$, followed by

$$X_{n+1}^P = X_n^P - \delta P^{-1} \nabla E_N(X_n^P) + \sqrt{2\delta} L^{-T} R_n, \quad \text{for } n = 1, \dots, N_{\text{steps}},$$

where $R_n \sim N(0, I)^N$. The estimate of the observable value is then given by $\bar{f} = N_{\text{steps}}^{-1} \sum_{n=1}^{N_{\text{steps}}} f(X_n^P)$. We compute N_{runs} trajectories in order to estimate the asymptotic variance from N_{runs} independent samples of \bar{f} .

For the umbrella sampling test, we choose $\xi(x) := l \cdot x$ to be the bond-length observable again with $\xi_0 = 0.33$. The modification to the energy and observable is then as described in Section 4.5. With $\delta = 0.1$, the discretised Langevin dynamics turns out to be unstable, hence we had to choose $\delta_K = 0.1/K$ instead. To account for this (see also Remark 4.3), we study $K\sigma_f^2$ instead of σ_f^2 in our tests. We perform the umbrella sampling test only for $d = 2, N = 8^2$, since we focus here on the magnitude on the restraint parameter K rather than the system size.

The results of the simulations are shown in table 5.1. The numbers closely match the analytical predictions of Sections 4.3, 4.4 and 4.5. To conclude this discussion we only remark that we did *not* fine-tune step-sizes, which means one could likely make small improvements to both the preconditioned and unpreconditioned processes. However, we believe that the *trends* across dimension, system size and restraint parameter are reliable.

In particular, we stress that even though the preconditioned variants often have a smaller asymptotic variance, often (in particular for $d = 3$) this improvement is only by a moderate constant factor. Because only the trends are reliable indicators in these tests only a successive improvement with increasing N or K (e.g. as in the $d = 1$ tests) can be considered a success for the preconditioned algorithm.

Conclusion

In this paper we strived to develop an intuition what the effect of preconditioning has on molecular simulations. The results are very mixed: it is clear that preconditioning accelerates convergence of the probability density functions to equilibrium (see Theorem 3.1 as well as the discussion in Section 4.2), and this necessarily implies accelerated convergence for *some* observables. However, for many concrete observables of practical importance little (if any) benefit can be gained. This was a surprising outcome for us and indicates that alternative avenues need to be explored on how a priori information about the analytical structure of configuration space should be exploited in molecular simulation.

We emphasize again that our (partially negative) conclusion, contrary to much of the existing literature, is due to the fact that we test the convergence of specific observables. Moreover, we stress that we have only performed a limited set of tests on highly simplified toy models and a limited set of observables, while more realistic models may exhibit many features that we neglected.

Asymptotic Variance: Energy

$d = 1$			$d = 2$			$d = 3$		
$N^{\frac{1}{d}}$	σ_f^2	$\sigma_{f,P}^2$	$N^{\frac{1}{d}}$	σ_f^2	$\sigma_{f,P}^2$	$N^{\frac{1}{d}}$	σ_f^2	$\sigma_{f,P}^2$
8	3.7e-01	6.4e-02	4	1.1e-01	4.4e-02	4	2.8e-02	9.8e-03
16	3.6e-01	3.7e-02	8	3.2e-02	9.3e-03	6	8.0e-03	3.5e-03
32	3.0e-01	1.8e-02	16	1.0e-02	2.3e-03	9	2.3e-03	9.4e-04
128	1.7e-01	4.5e-03	32	2.8e-03	5.4e-04	13	7.9e-04	3.0e-04

Asymptotic Variance: Bond-length

$d = 1$			$d = 2$			$d = 3$		
$N^{\frac{1}{d}}$	σ_f^2	$\sigma_{f,P}^2$	$N^{\frac{1}{d}}$	σ_f^2	$\sigma_{f,P}^2$	$N^{\frac{1}{d}}$	σ_f^2	$\sigma_{f,P}^2$
8	2.4e+01	7.4e+00	4	3.2e+01	8.5e+00	4	1.9e+01	7.5e+00
16	4.2e+01	7.2e+00	8	5.4e+01	9.5e+00	6	1.8e+01	7.6e+00
32	7.6e+01	7.4e+00	16	7.2e+01	1.1e+01	9	1.9e+01	7.5e+00
128	2.8e+02	7.0e+00	32	1.1e+02	1.1e+01	13	2.2e+01	7.7e+00

Asymptotic Variance: Umbrella Sampling

$d = 2, N^{1/2} = 8$		
K	$K\sigma_f^2$	$\sigma_{f,P}^2$
10	1.6e+01	1.8e+01
20	3.5e+01	4.0e+01
40	6.7e+01	7.7e+01
80	1.3e+02	1.8e+02

TABLE 5.1. Numerically estimated asymptotic variances of the energy observable (Section 4.3), the bond-length observable (Section 4.4) and the restraint observable occurring in umbrella sampling (Section 4.5). The nonlinear potential energy used in these tests is described in Section 5. All results match the analytical predictions of Sections 4.3, 4.4 and 4.5.

6. Proofs

6.1. Proof of Proposition 2.1

Proof. The covariance of the invariant measure associated to the dynamics (2.12) is given by the following identity:

$$\begin{aligned}
C_\delta^P &= (I - \delta P^{-1}H)C_\delta^P(I - \delta HP^{-1}) + 2\delta P^{-1} \\
C_\delta^P &= C_\delta^P - \delta P^{-1}HC_\delta^P - \delta C_\delta^P HP^{-1} + \delta^2 P^{-1}HC_\delta^P HP^{-1} + 2\delta P^{-1} \\
P^{-1}HC_\delta^P + C_\delta^P HP^{-1} &= 2P^{-1} + \delta P^{-1}HC_\delta^P HP^{-1}.
\end{aligned}$$

Expanding C_δ^P , we have

$$C_\delta^P \sim C_0 + \delta C_1 + \delta^2 C_2 + O(\delta^3),$$

then one gets

$$\begin{aligned}
P^{-1}HC_0 + C_0HP^{-1} &= 2P^{-1} &\Rightarrow C_0 &= H^{-1} \\
P^{-1}HC_1 + C_1HP^{-1} &= P^{-1}HP^{-1} &\Rightarrow C_1 &= \frac{1}{2}P^{-1} \\
P^{-1}HC_2 + C_2HP^{-1} &= \frac{1}{2}P^{-1}HP^{-1}HP^{-1} &\Rightarrow C_2 &= \frac{1}{4}P^{-1}HP^{-1}.
\end{aligned}$$

Proceeding by induction, one can therefore obtain:

$$\forall k \in \mathbb{N}, \quad C_k = 2^{-k} (P^{-1}H)^k H^{-1}.$$

Therefore C_δ^P can be rewritten as:

$$C_\delta^P = \sum_{k=0}^{+\infty} \left(\frac{\delta}{2} P^{-1}H \right)^k H^{-1}.$$

Indeed, the sum $\sum_{k \geq 0} C_k$ converges since $\frac{\delta}{2} |P^{-1}H| < 1$. The covariance operator can be rewritten as

$$\begin{aligned} C_\delta^P &= \left(\sum_{k=0}^{\infty} \left(\frac{\delta}{2} P^{-1}H \right)^k \right) H^{-1} \\ &= \left(I - \frac{\delta}{2} P^{-1}H \right)^{-1} H^{-1} \\ &= \left(H - \frac{\delta}{2} HP^{-1}H \right)^{-1}, \end{aligned}$$

which concludes the proof. ■

6.2. Proof of Theorem 3.1

Proof. Under the assumptions on the potential E_N , see A.19 in [33], the density ψ_∞ satisfies a Poincaré inequality: there exists $\lambda_P > 0$ such that for all probability density functions ϕ , we have:

$$\int_{\mathbb{R}^N} \left| \frac{\phi}{\psi_\infty} - 1 \right|^2 \psi_\infty dz \leq \frac{1}{\lambda_P} \int_{\mathbb{R}^N} \left| \nabla \left(\frac{\phi}{\psi_\infty} \right) \right|^2 \psi_\infty dz. \quad (6.1)$$

The optimal parameter λ_P in (6.1) is the opposite of the smallest (in absolute value) non-zero eigenvalue of the Fokker-Planck operator \mathcal{L}_P^* , which is self-adjoint in $L^2(\mathbb{R}^N, \psi_\infty^{-1} dx)$. Thus the exponent λ_P is the spectral gap of \mathcal{L}_P^* .

If ψ_t^P is a solution of (3.1), therefore for all initial condition $\psi_0 \in L^2(\psi_\infty^{-1})$, $\forall t \geq 0$:

$$\frac{d}{dt} \|\psi_t^P - \psi_\infty\|_{L^2(\psi_\infty^{-1})}^2 = -2 \int_{\mathbb{R}^N} \left| \nabla \left(\frac{\psi_t^P}{\psi_\infty} \right) \right|^2 \psi_\infty dz. \quad (6.2)$$

Indeed,

$$\begin{aligned} \frac{d}{dt} \|\psi_t^P - \psi_\infty\|_{L^2(\psi_\infty^{-1})}^2 &= \frac{d}{dt} \int |\psi_t^P - \psi_\infty|^2 \psi_\infty^{-1} dz \\ &= 2 \int \partial_t \psi_t^P (\psi_t^P - \psi_\infty) \psi_\infty^{-1} dz \\ &= 2 \int \nabla \cdot (\nabla E_N^P(z) \psi_t^P + \nabla \psi_t^P) \left(\frac{\psi_t^P}{\psi_\infty} - 1 \right) dz \\ &= -2 \int (\nabla E_N^P(z) \psi_t^P + \nabla \psi_t^P) \cdot \nabla \left(\frac{\psi_t^P}{\psi_\infty} \right) dz. \end{aligned}$$

But we have

$$\begin{aligned} \nabla E_N^P(z) \psi_t^P + \nabla \psi_t^P &= -\nabla (\ln(\psi_\infty)) \psi_t^P + \nabla \psi_t^P \\ &= -\frac{\nabla \psi_\infty \psi_t^P}{\psi_\infty} + \nabla \psi_t^P \\ &= \nabla \left(\frac{\psi_t^P}{\psi_\infty} \right) \psi_\infty, \end{aligned}$$

which yields (6.2). Therefore, using (6.1),

$$\begin{aligned} \frac{d}{dt} \|\psi_t^P - \psi_\infty\|_{L^2(\psi_\infty^{-1})}^2 &\leq -2\lambda_P \int \left| \frac{\psi_t^P}{\psi_\infty} - 1 \right|^2 \psi_\infty dz \\ &= -2\lambda_P \int \left| \psi_t^P - \psi_\infty \right|^2 \psi_\infty^{-1} dz, \end{aligned}$$

then

$$\|\psi_t^P - \psi_\infty\|_{L^2(\psi_\infty^{-1})}^2 \leq e^{-\lambda_P t} \|\psi_0 - \psi_\infty\|_{L^2(\psi_\infty^{-1})}^2.$$

■

6.3. Preliminaries for spectral analysis

In the linear case, i.e $E_N^P(z) = \frac{1}{2}z^T H_P z$, where H and P are symmetric positive definite and $H_P := P^{-1/2} H P^{-1/2}$, the analysis will be carried out in a suitable system of coordinates which simplifies the calculations and the proofs of the main theorems. For this reason, we will perform one conjugation and one additional change of variables.

From the partial differential equation point of view and in order to use standard techniques from the spectral analysis of partial differential equations, then it appears to be useful to work in $L^2(\mathbb{R}^N, dz; \mathbb{C})$ instead of $L^2(\mathbb{R}^N, \psi_\infty dz; \mathbb{C})$. The mapping $\phi \mapsto \psi_\infty^{-1/2} \phi$ maps unitarily $L^2(\mathbb{R}^N, dz; \mathbb{C})$ into $L^2(\mathbb{R}^N, \psi_\infty dz; \mathbb{C})$ with the associated transformation rules for the differential operators:

$$e^{-\frac{1}{2}H_P z} \nabla_z e^{\frac{1}{2}H_P z} = \nabla_z + \frac{1}{2} \nabla_z H_P.$$

Thus, the operator $\mathcal{L}_P = -\nabla_z E_N^P(z) \cdot \nabla_z + \Delta_z$ is transformed into

$$\begin{aligned} \bar{\mathcal{L}}_P &= e^{-\frac{1}{2}H_P z} \mathcal{L}_P e^{\frac{1}{2}H_P z} \\ &= \Delta_z - \frac{1}{4} |\nabla_z E_N^P(z)|^2 + \frac{1}{2} \Delta_z E_N^P(z) \\ &= \Delta_z - \frac{1}{4} z^T H_P^2 z + \frac{1}{2} \text{Tr}(H_P). \end{aligned} \tag{6.3}$$

The kernel of $\bar{\mathcal{L}}_P$ is $\mathbb{C} e^{-\frac{z^T H_P z}{4}}$ and the operator $\bar{\mathcal{L}}_P$ is unitarily equivalent to the operator \mathcal{L}_P .

In the goal of modifying the kernel of the operator $\bar{\mathcal{L}}_P$ into a centered Gaussian with identity covariance matrix, we perform a second change of variables. In the following, we introduce the new coordinates $y = H_P^{1/2} z$, so that $\nabla_z = H_P^{1/2} \nabla_y$. Then the operator $\bar{\mathcal{L}}_P$ becomes:

$$\tilde{\mathcal{L}}_P = \nabla_y^T H_P \nabla_y - \frac{1}{4} y^T H_P y + \frac{1}{2} \text{Tr}(H_P). \tag{6.4}$$

The operator $\tilde{\mathcal{L}}_P$ is still acting in $L^2(\mathbb{R}^N, dz; \mathbb{C})$. In the new coordinate system ($Y_t = H_P^{1/2} Z_t$), the corresponding stochastic process is:

$$dY_t = -H_P Y_t dt + \sqrt{2} H_P^{1/2} dW_t,$$

so that $\text{Ker}(\tilde{\mathcal{L}}_P) = \frac{1}{(2\pi)^{N/4}} e^{-\frac{|y|^2}{4}}$. The last conjugation and change of variables are used to compute the spectrum of \mathcal{L}_P^* needed to proof Theorem 3.2 (see Section 6.4).

Let us now introduce some additional notations. Recall the space of rapidly decaying complex valued \mathcal{C}^∞ functions

$$\mathcal{S}(\mathbb{R}^N) = \left\{ f \in \mathcal{C}^\infty(\mathbb{R}^N), \forall \alpha, \beta \in \mathbb{N}^N, \exists C_{\alpha, \beta} \in \mathbb{R}_+, \sup_{x \in \mathbb{R}^N} |x^\alpha \partial_x^\beta f(x)| \leq C_{\alpha, \beta} \right\},$$

and its dual is denoted $\mathcal{S}'(\mathbb{R}^N)$.

The Weyl-quantization $q^W(x, D_x)$ of a symbol $q(x, \xi) \in \mathcal{S}'(\mathbb{R}^N)$ is an operator defined by its Schwartz-kernel

$$\left[q^W(x, D_x) \right] (x, y) = \int_{\mathbb{R}^N} e^{i(x-y) \cdot \xi} q \left(\frac{x+y}{2}, \xi \right) \frac{d\xi}{(2\pi)^N}.$$

For instance, the Weyl symbol of the operator

$$-\tilde{\mathcal{L}}_P + \frac{1}{2} \text{Tr}(H_P) = -\nabla_y^T H_P \nabla_y + \frac{1}{4} y^T H_P y$$

is

$$q(y, \xi) = \xi^T H_P \xi + \frac{y^T H_P y}{4}. \quad (6.5)$$

Those tools are essential to proof Theorem 3.2 (see Section 6.4). For more details on Weyl-quantization, one can refer to [28].

6.4. Proof of Theorem 3.2

Proof. Referring to Theorem 1.2.2 in [12], the spectrum of the operator $q^W(y, D_y) = -\tilde{\mathcal{L}}_P + \frac{1}{2} \text{Tr}(H_P)$ associated with the elliptic quadratic Weyl symbol $q(y, \xi)$ defined by (6.5) is given by

$$\sigma(q^W(y, D_y)) = \left\{ \sum_{\substack{\lambda \in \sigma(G) \\ \text{Im} \lambda \geq 0}} -i\lambda(r_\lambda + 2k_\lambda), k_\lambda \in \mathbb{N} \right\}.$$

where G is the so-called Hamilton map associated with q , and r_λ is the algebraic multiplicity of $\lambda \in \sigma(G)$ (the dimension of the characteristic space). The Hamilton map is the \mathbb{C} -linear map $G : \mathbb{C}^{2N} \rightarrow \mathbb{C}^{2N}$ associated with the matrix

$$G = \begin{bmatrix} 0 & H_P \\ -\frac{1}{4} H_P & 0 \end{bmatrix} \in \mathbb{C}^{2N \times 2N}.$$

The matrix G is similar to another matrix denoted \bar{G} and defined by

$$\bar{G} = \begin{bmatrix} \frac{1}{\sqrt{2}} & 0 \\ 0 & \sqrt{2} \end{bmatrix} G \begin{bmatrix} \sqrt{2} & 0 \\ 0 & \frac{1}{\sqrt{2}} \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 0 & H_P \\ -H_P & 0 \end{bmatrix}.$$

Now, the characteristic polynomial of G can be computed by

$$\begin{aligned} \det(G - \lambda I) &= \det(\bar{G} - \lambda I) = 2^{-2N} \begin{vmatrix} -2\lambda I & H_P \\ -H_P & -2\lambda I \end{vmatrix} \\ &= 2^{-2N} \begin{vmatrix} -2\lambda I & H_P \\ -H_P - i2\lambda I & i(H_P + i2\lambda I) \end{vmatrix} \\ &= 2^{-2N} \begin{vmatrix} i(-H_P + i2\lambda I) & H_P \\ 0 & i(H_P + i2\lambda I) \end{vmatrix} \\ &= 2^{-2N} \det(H_P - i2\lambda I) \det(H_P + i2\lambda I). \end{aligned}$$

Since $\text{Re}(\sigma(H_P)) \geq 0$, one thus obtains that

$$\sigma(G) \cap \{\lambda, \text{Im} \lambda \geq 0\} = \frac{i}{2} \sigma(H_P).$$

In particular,

$$\sum_{\substack{\lambda \in \sigma(G) \\ \text{Im}\lambda \geq 0}} -i\lambda 2k_\lambda = \sum_{\mu \in \sigma(H_P)} k_{\frac{i}{2}\mu} \mu$$

and

$$\sum_{\substack{\lambda \in \sigma(G) \\ \text{Im}\lambda \geq 0}} -i\lambda r_\lambda = \frac{\text{Tr}(H_P)}{2},$$

which concludes the proof of the theorem. ■

6.5. Proof of Theorem 4.1

Definition 6.1. (Foster-Lyapunov criterion)

We say that the Foster-Lyapunov criterion holds for (2.1) if there exists a function $U : \mathbb{R}^N \rightarrow \mathbb{R}$ and constants $C > 0$ and $b \in \mathbb{R}$ such that $\mu(U) < \infty$,

$$\mathcal{L}_P U(z) \leq -cU(z) + b\mathbf{1}_C \quad (6.6)$$

and $U(z) \geq 1$, $z \in \mathbb{R}^N$, where $\mathbf{1}_C$ is the indication function over a petite Borel subset C of \mathbb{R}^N (refer to tve for more details).

For the generator \mathcal{L}_P corresponding to (2.2), compact sets are always petite. In the following we prove that the Foster-Lyapunov criterion holds for (2.1). But first we need an assumption on the potential E_N^P .

Assumption A There exists $k > 0$ such that $\mu^P(z)$ is bounded from above for all $|z| \geq k$ and, for some $0 < \beta < 1$,

$$\liminf_{|z| \rightarrow +\infty} \left[(1 - \beta) |\nabla E_N^P(z)|^2 + \Delta E_N^P(z) \right] > 0.$$

Lemma 6.2. *Under Assumption A, the Foster-Lyapunov criterion holds for (2.1) with:*

$$U(z) = e^{\beta E_N^P(z)}, \quad 0 < \beta < 1. \quad (6.7)$$

Proof. Recall the generator of (2.1)

$$\mathcal{L}_P = -\nabla_z E_N^P(z) \nabla + \Delta.$$

For $U(z) = e^{\beta E_N^P(z)}$, one obtains:

$$\begin{aligned} \mathcal{L}_P U(z) &= -\nabla U(z) \cdot \nabla E_N^P(z) + \Delta U(z) \\ &= -\nabla e^{\beta E_N^P(z)} \cdot \nabla E_N^P(z) + \nabla \cdot (\nabla e^{\beta E_N^P(z)}) \\ &= -\beta |\nabla E_N^P(z)|^2 e^{\beta E_N^P(z)} + \beta \Delta E_N^P(z) e^{\beta E_N^P(z)} \\ &\quad + \beta^2 |\nabla E_N^P(z)|^2 e^{\beta E_N^P(z)} \\ &= -\beta \left[(1 - \beta) |\nabla E_N^P(z)|^2 + \Delta E_N^P(z) \right] U(z). \end{aligned} \quad (6.8)$$

Therefore, by Assumption A, for $\varepsilon > 0$, $\exists k > 0$ such that $\forall |z| > k$:

$$(1 - \beta) |\nabla E_N^P(z)|^2 + \Delta E_N^P(z) > \varepsilon,$$

and so also,

$$\mathcal{L}_P U(z) \leq -\beta \varepsilon U(z) + b\mathbf{1}_{C_k},$$

where $C_k = \{z \in \mathbb{R}^N; |z| \leq k\}$ and $b > 0$.

Finally, since $\psi_\infty(z)$ is bounded, then $U(z)$ is bounded away from zero uniformly. Then $U(z)$ can be rescaled to satisfy the condition $U(z) \geq 1$. Thus, we valid the Foster-Lyapunov criterion for (2.1). ■

Using this lemma, the proof of the well-posedness result for the Poisson equation (4.1) come straightforward using Theorem 3.2 in [3].

6.6. Proof of Theorem 4.2

Proof. We start the proof by decomposing $\mu_t^P(f) - \mu(f)$ into a martingale and a remainder terms: Using (4.1), (2.2) and (2.1)

$$\begin{aligned} \epsilon_t^P(f) - \mu(f) &= \frac{1}{t} \int_0^t f(P^{-1/2}Z_s) ds - \mu(f) \\ &= \frac{1}{t} \int_0^t \left(f(P^{-1/2}Z_s) - \mu(f) \right) ds \\ &= \frac{1}{t} \int_0^t -\mathcal{L}_P \phi(Z_s) ds \\ &= \frac{\phi(Z_0) - \phi(Z_t)}{t} + \frac{\sqrt{2}}{t} \int_0^t \nabla \phi(Z_s) dW_s \\ &:= R_t + M_t. \end{aligned}$$

Consider now the rescaling $\sqrt{t}(\epsilon_t^P(f) - \mu(f))$. Using the central limit theorem for the martingale term $\sqrt{t}M_t$ (see [11], Theorem 5.3), one obtains the following convergence in distribution

$$\sqrt{t}M_t \xrightarrow{D} \mathcal{N}(0, \sigma_{f,P}^2),$$

with

$$\sigma_{f,P}^2 = 2 \int |\nabla_z \phi(z)|^2 \mu^P(dz) = 2 \int |\nabla_z \phi(z)|^2 \mu(dx).$$

It remains to study the remainder term $\sqrt{t}R_t$. We consider the two cases: If $Z_0 \sim \mu$, then since $\phi \in L^2(\mu)$, we have that

$$\sqrt{t}R_t \xrightarrow{D} 0 \text{ in } L^2(\mu).$$

In the more general case, we must refer to a "propagation of chaos" argument (see for example [6], Section 8 and [7]), to obtain the same result. ■

Acknowledgments

We thank Tony Lelièvre and Jonathan Weare for helpful discussions. The work was supported by ERC Starting Grant 335120.

Appendix A. Setup of the Metadynamics example

In the example shown in Figure 1.1, we investigated how preconditioning the *second order Langevin dynamics* affects the sampling of different observables when applied to adaptive potential of mean force (PMF) techniques such metadynamics (MTD) [19]. Unlike the previously discussed umbrella sampling method, adaptive PMF techniques build their bias and mean force or potential of mean force estimates on-the-fly during the dynamics. In a general form, the corresponding biased and preconditioned equations of motion with preconditioner $P = P(q)$ can be written as

$$dq_t = M^{-1}p_t dt \quad (\text{A.1a})$$

$$dp_t = -\nabla E_N(q_t)dt + \nabla \xi(q_t)F_b(\xi(q_t), t) - PM^{-1}p_t dt + \sqrt{\frac{2}{\beta}}P^{1/2}dW_t \quad (\text{A.1b})$$

where $M \in \mathbb{R}^{N \times N}$ is the diagonal mass matrix, β is the inverse temperature, ξ is the collective variable (reaction coordinate) and $F_b(\xi, t)$ is the biasing force. In the case of metadynamics, $F_b^{\text{MTD}}(\xi, t) = -\nabla_{\xi} E_b^{\text{MTD}}(\xi, t)$, where $E_b^{\text{MTD}}(\xi, t)$ is some history dependent biasing potential composed by Gaussians regularly deposited in the collective variable space:

$$E_b^{\text{MTD}}(\xi, t) = \sum_{t' < t} \delta_{t'} \exp\left(-\frac{1}{2}(\xi - \xi(q_{t'}))^T w^{-2}(\xi - \xi(q_{t'}))\right) \quad (\text{A.2})$$

where $\delta_{t'}$ is the height of the corresponding Gaussian and w is a diagonal matrix including the widths of the collective variable components.

The molecular system we chose for this test was the 2-(formylamino)propionaldehyde in gas phase using the Amber99SB [13] force field.

The molecule has a single slow degree of freedom, a dihedral angle, that was selected as the one-dimensional collective variable (we note that this variable is highly associated to one of the two dihedrals of alanine dipeptide, a test system widely used in the computational chemistry field).

We performed 200 ps long molecular dynamics simulations at 300 K using the BBK integrator scheme [5] with 0.5 fs timestep. In the case of unpreconditioned dynamics $P = \gamma M$ was used with $\gamma = 5.0 \text{ ps}^{-1}$, while for the preconditioned dynamics we applied $P = \gamma M + \tau \tilde{H}$, where \tilde{H} is a Hessian-based preconditioner, whose positiveness is guaranteed by rebuilding the matrix using the spectral decomposition of the Hessian with the absolute values of the eigenvalues [26]. We used γM as diagonal a stabiliser and varied γ and τ parameters.

A deposition frequency of $1/50 \text{ fs}^{-1}$ and a starting height of $\delta = 0.004 \text{ eV}$ were used in a well-tempered variant of MTD [1] with $T_w = 10000 \text{ K}$.

Free energy profiles were reconstructed simply as the negative of the actual history dependent biasing potential. The reference for computing the RMS error of the profiles was obtained from a 200 ns long unpreconditioned constrained dynamics simulation [25].

The result of this test is shown in Figure 1.1, where we plotted the RMS error of the reconstructed profiles based on 10 independent MTD simulations for each parameter set. We observe that in general, preconditioning does not improve the convergence.

As a matter of fact, we performed a variety of similar tests, e.g. with different parameters, or different observables, all of which led to similar conclusions. Indeed, depending on the choice of observable, preconditioning often has an even larger *negative* impact.

References

- [1] A. Barducci, G. Bussi, and M. Parrinello. Well-tempered metadynamics: A smoothly converging and tunable free-energy method. *Phys. Rev. Lett.*, 100(2):020603, 2008.

SOME REMARKS ON PRECONDITIONING MOLECULAR DYNAMICS

- [2] A. Beskos, G. Roberts, and A. Stuart. Optimal scalings for local metropolis–hastings chains on nonproduct targets in high dimensions. *The Annals of Applied Probability*, 19(3):863–98, 2009.
- [3] R. Bhattacharya. A central limit theorem for diffusions with periodic coefficients. *The Annals of Probability*, 13(2):385–396, 1985.
- [4] H.J. Brascamp and E.H. Lieb. On extensions of the Brunn-minkowski and Prékopa-Leindler theorems, including inequalities for log concave functions, and with an application to the diffusion equation. *Journal of functional analysis*, 22(4):366–389, 1976.
- [5] A. Brünger, C. L. Brooks, and M. Karplus. Stochastic boundary conditions for molecular dynamics simulations of ST2 water. *Chemical Physics Letters*, 105(5):495–500, 1984.
- [6] P. Cattiaux, D. Chafai, and A. Guillin. Central limit theorems for additive functionals of ergodic Markov diffusions processes. *ALEA*, 9(2):337–382, 2012.
- [7] A. B. Duncan, T. Lelièvre, and G. A. Pavliotis. Variance reduction using nonreversible Langevin samplers. *Journal of Statistical Physics*, 163(3):457–491, 2016.
- [8] M. Girolami and B. Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.
- [9] P. W. Glynn and S.P. Meyn. A Liapounov bound for solutions of the Poisson equation. *The Annals of Probability*, pages 916–931, 1996.
- [10] E. Hairer, C. Lubich, and G. Wanner. Geometric numerical integration illustrated by the Stormer-Verlet method. *Acta numerica*, 12(12):399–450, 2003.
- [11] I. S. Helland. Central limit theorems for martingales with discrete or continuous time. *Scandinavian Journal of Statistics*, 9(2):79–94, 1982.
- [12] M. Hitrik and K. Pravda-Starov. Spectra and semigroup smoothing for non-elliptic quadratic operators. *Math. Ann.*, 344(4):801–846, 2009.
- [13] W. Hornak, R. Abel, A. Okur, B. Strockbine, A. Roitberg, and C. Simmerling. Comparison of multiple amber force fields and development of improved protein backbone parameters. *Prot. Struct. Funct. Bioinf.*, 65(3):712–725, 2006.
- [14] F. Hou, J. Goodman, D.W. Hogg, J. Weare, and C. Schwab. An affine-invariant sampler for exoplanet fitting and discovery in radial velocity data. *The Astrophysical Journal*, 745(2), 2012.
- [15] A. Iserles. A first course in the numerical analysis of differential equations. *Cambridge Texts in Applied Mathematics*, Cambridge University Press, 2008.
- [16] A. M. Stuart J. C. Mattingly and D. J. Higham. Ergodicity for SDEs and approximations: locally Lipschitz vector fields and degenerate noise. *Stochastic processes and their applications*, 101(2):185–232, 2002.
- [17] J. Kästner and W. Thiel. Bridging the gap between thermodynamic integration and umbrella sampling provides a novel analysis method: Umbrella integration. *J. Chem. Phys.*, 123, 2005.
- [18] C. Kipnis and S.R.S. Varadhan. Ocentral limit theorem for additive functionals of reversible Markov processes and applications to simple exclusions. *Communications in Mathematical Physics*, 104(1):1–19, 1986.
- [19] A. Laio and M. Parrinello. Escaping free-energy minima. *Proceedings of the National Academy of Sciences*, 99:12562–12566, 2002.
- [20] T. Lelièvre. Personal communication. 2017.
- [21] T. Lelièvre, F. Nier, and G. Pavliotis. Optimal non-reversible linear drift for the convergence to equilibrium of a diffusion. *Journal of Statistical Physics*, 152(2):237–274, 2013.
- [22] J. Martin, L. C. Wilcox, C. Burstedde, and O. Ghattas. A stochastic newton MCMC method for large-scale statistical inverse problems with application to seismic inversion. *SIAM Journal on Scientific Computing*, 34(3):A1460–A1487, 2012.

- [23] S. P. Meyn and R. L. Tweedie. Stability of Markovian processes III: Foster-Lyapunov criteria for continuous-time processes. *Advances in Applied Probability*, 25(3):518–548, 1993.
- [24] S. P. Meyn and R. L. Tweedie. A survey of Foster-Lyapunov techniques for general state space Markov processes. *In Proceedings of the Workshop on Stochastic Stability and Stochastic Stabilization, Metz, France. Citeseer*, 1993.
- [25] L. Mones, N. Bernstein, and G. Csanyi. Exploration, sampling, and reconstruction of free energy surfaces with gaussian process regression. *Journal of Chemical Theory and Computation*, 12(10):5100–5110, 2016.
- [26] L. Mones, G. Csanyi, and C. Ortner. in preparation. 2017.
- [27] D. Olson and C. Ortner. Regularity and locality of point defects in multilattices. *to appear in AMRX, arXiv:1608.08930*, 2016.
- [28] M. Ottobre, G. A. Pavliotis, and K. Pravda-Starov. Exponential return to equilibrium for hypoelliptic quadratic systems. *Journal of Functional Analysis*, 262(9):4000–4039, 2012.
- [29] D. Packwood, J. Kermode, L. Mones, N. Bernstein, N. I. M. Gould J. Woolley, C. Ortner, and G. Csanyi. universal preconditioner for simulating condensed phase materials. *J. Chem. Phys.*, 144(16), 2016.
- [30] N. S. Pillai, A. M. Stuart, and A. H. Thiéry. Optimal scaling and diffusion limits for the langevin algorithm in high dimensions. *The Annals of Applied Probability*, 22(6):2320–2356, 2017.
- [31] C. P. Robert and G. Casella. Monte Carlo statistical methods. *2nd ed. Springer, New York*, 2004.
- [32] G. O. Roberts and R. L. Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341–363, 1996.
- [33] V. Villani. Hypocoercivity. *Bull. Am. Math. Soc.*, 46(2), 2009.
- [34] T. Xifara, C. Sherlock, S. Livingstone, S. Byrne, and M. Girolami. Langevin diffusions and the metropolis-adjusted langevin algorithm. *Statistics & Probability Letters*, 91:14–19, 2014.