

METHODOLOGICAL AND COMPUTATIONAL ASPECTS OF PARALLEL TEMPERING METHODS IN THE INFINITE SWAPPING LIMIT

JIANFENG LU AND ERIC VANDEN-EIJNDEN

ABSTRACT. A variant of the parallel tempering method is proposed in terms of a stochastic switching process for the coupled dynamics of replica configuration and temperature permutation. This formulation is shown to facilitate the analysis of the convergence properties of parallel tempering by large deviation theory, which indicates that the method should be operated in the infinite swapping limit to maximize sampling efficiency. The effective equation for the replica alone that arises in this infinite swapping limit simply involves replacing the original potential by a mixture potential. The analysis of the geometric properties of this potential offers a new perspective on the issues of how to choose of temperature ladder, and why many temperatures should typically be introduced to boost the sampling efficiency. It is also shown how to simulate the effective equation in this many temperature regime using multiscale integrators. Finally, similar ideas are also used to discuss extensions of the infinite swapping limits to the technique of simulated tempering.

1. INTRODUCTION

Techniques such as parallel [1, 2, 17] and simulated tempering [12] have become standard tools to accelerate the sampling of systems with complicated energy landscapes, mostly because of their effectiveness and minimal requirement of detailed knowledge of the system's specifics. In a nutshell, the basic idea of these methods is to temporarily raise the temperature of the system (or that of a copy thereof) to allow it to explore its energy landscape faster. When this is done in a specific way, the properties of the system at the physical temperature can still be calculated without bias. As a result tempering techniques provide us with importance sampling schemes to analyze the statistical mechanics properties of a system at a given temperature even in situations where direct simulations at this temperature are too slow to be practical. The aim of this paper is to give some mathematical analysis of these methods via their formulation as stochastic switching processes, which will permit us to conveniently explore their properties in various parameter regimes. In so doing, we will also discuss implementation issues.

Consider first parallel tempering, in which one evolves multiple copies (or replicas) of the system at different temperatures which are occasionally swapped among these replicas (which is why the method is also known as replica exchange). Typically, these swaps are attempted at a given frequency ν with a Metropolis-Hastings acceptance/rejection criterion to guarantee that the replicas sample a known equilibrium distribution over which expectations at any of the temperatures in the stack (including the physical one) can be calculated. In this context it was found empirically [14–16] that the larger the swapping frequency, the higher the sampling efficiency is. This observation was mathematically justified in the work of Dupuis et al. [6] by showing that the large deviation rate functional for the empirical measure is monotonically increasing with ν . This observation also motivated the development of infinite swapping replica exchange dynamics [3, 6, 11, 13], in which one attempts to simulate directly the limiting dynamics at $\nu \rightarrow \infty$.

How to do so in practice is non-trivial, however. When parallel tempering is applied to complicated systems, often times a large number ($N \sim 50 - 100$) of temperatures and replicas is needed to make the sampling efficient. In such cases, the simulation of the infinite swapping limit becomes challenging because the limiting dynamics involves coefficients given in terms of sums with as many terms as there are permutations of the replicas, $N!$. An exhaustive evaluation of such sums quickly become undoable, even numerically, and alternative strategies must be used. One possibility, proposed in [21], is to resort to multiscale integrators like those in the heterogeneous multiscale methods (HMM) to effectively simulate the limiting dynamics.

One of the purposes of this paper is to present the mathematical foundation behind the algorithm proposed in [21]. Specifically, we show that the parallel tempering can be formulated as a diffusion of the replica coupled with a Markov jumping process in the space of permutations. Together this gives a stochastic switching process. Using the approach introduced in [6], the infinite swapping limit can be justified by the monotonicity of the large deviation rate functional as a function of the swapping frequency. The large swapping frequency introduces a scale separation between the dynamics of the replica and that of the permutation variable, and thus naturally calls for multiscale integrators, such as those based on HMM, which justifies the integrator introduced in [21].

Similar ideas can also be used in the context of simulated tempering, in which a single replica is used but the temperature itself is made a dynamical variable of the system. In this setup too, it is useful to consider the coupled dynamics of the system and the temperature as a stochastic switching process, and this allows one to again justify via large deviation theory that making the temperature evolve infinitely fast (the ‘infinite switch limit’) is optimal. As we will show below, the limiting equation that emerges in this infinite-switch limit is easier to simulate in principle, as it only involves a sum over the number of possible temperature states rather than its factorial. In practice, however, the question becomes how to appropriately weight the different temperatures states, since this is an input in the method which is unknown *a priori* but can greatly influence its efficiency.

The organization of the remainder of this paper is as follows. In Section 2 we present a formulation of parallel tempering as a stochastic switching process, which facilitates the discussion of infinite swapping limit and the development of integrators for the equations of motion that emerge in this limit. The infinite swapping limit is justified through large deviation theory in Section 3 and the effective equations that arise in this limit are derived in Section 4. In Section 5 we discuss the choice of the temperature ladder in parallel tempering and, in particular, explain why many temperatures are typically necessary to boost the efficiency of the method. In the infinite swapping limit, where the dynamics of the replica effectively reduce to a diffusion over a mixture potential, this question reduces to the analysis of the geometrical properties of this potential. An HMM multiscale integrator that operates in the many temperatures regime is then discussed in Section 6 and shown to perform better than the standard stochastic simulation algorithm in the regime of frequent swapping. In Section 7 we show that similar ideas can be extended to other tempering techniques such as simulate tempering. Finally, some concluding remarks are given in Section 8.

2. PARALLEL TEMPERING AS A STOCHASTIC SWITCHING PROCESS

Consider a system whose evolution is governed by the overdamped Langevin equation:

$$(1) \quad d\mathbf{x}(t) = -\nabla V(\mathbf{x}(t)) dt + \sqrt{2\beta^{-1}} d\mathbf{W}(t),$$

where $\mathbf{x}(t) \in \Omega \subset \mathbb{R}^{3n}$ denotes the instantaneous position (at time t) of the system with n particles, $-\nabla V$ is the force associated with the potential V , $\beta = 1/k_B T$ is the inverse temperature, each component of $\mathbf{W}(t) \in \mathbb{R}^{3n}$ is an independent Wiener process, and we set the friction coefficient to one for simplicity. Most of the results we present below should be straightforwardly generalizable to other dynamics, such as the inertial Langevin equation, but we will focus on (1) for the sake of simplicity.

Assuming that V is smooth and Ω is compact, it is easy to show that the overdamped dynamics (1) is ergodic with respect to the Boltzmann equilibrium distribution with density

$$(2) \quad \rho_\beta(\mathbf{x}) = Z_\beta^{-1} e^{-\beta V(\mathbf{x})},$$

where $Z_\beta = \int_\Omega e^{-\beta V(\mathbf{x})} d\mathbf{x}$. However, when the dimensionality of the system is large, $n \gg 1$, and the potential V is complicated with critical points, etc. (1) typically exhibits metastability, meaning that convergence to the equilibrium measure is very slow. In this context, the basic idea of parallel tempering is to use dynamics at higher temperature to help the sampling.

Specifically, we introduce multiple replicas of the system and make each replica evolve under a different temperature that alternatively swaps between several of them (many artificial temperatures, plus the physical one). Assuming that we use N temperatures, so that the parallel tempering correspondingly involves N replicas, this is done by constructing a stochastic process on the extended phase space $\Omega^N \times P_N$, where Ω^N is the configurational space for the N replica and P_N is the permutation group of the N temperature, such that its joint equilibrium measure is given by

$$(3) \quad \pi(d\mathbf{X}, \sigma) = \varrho(\mathbf{X}, \sigma) d\mathbf{X}, \quad \varrho(\mathbf{X}, \sigma) = \frac{1}{N!} \prod_{i=1}^N \rho_{\beta_{\sigma_i}}(\mathbf{x}_i),$$

where $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N)$ and σ_i denotes the index associated with i by the permutation σ . To this end, we assume that given any $\sigma \in P_N$ we have a diffusion with generator \mathcal{L}_σ whose equilibrium density is the product

$$(4) \quad \varrho(\mathbf{X} | \sigma) = \frac{\varrho(\mathbf{X}, \sigma)}{\int_{\Omega^N} \varrho(\mathbf{X}', \sigma) d\mathbf{X}'} = \prod_{i=1}^N \rho_{\beta_{\sigma_i}}(\mathbf{x}_i).$$

We also assume that given any $\mathbf{X} \in \Omega^N$, we have a Markov jump process with infinitesimal jump intensity $\nu h_{\sigma\sigma'}(\mathbf{X})$, where $\nu > 0$ is an overall frequency parameter and $h_{\sigma\sigma'}(\mathbf{X}) \geq 0$ satisfies the detailed balance condition

$$(5) \quad \varrho(\sigma | \mathbf{X}) h_{\sigma\sigma'}(\mathbf{X}) = \varrho(\sigma' | \mathbf{X}) h_{\sigma'\sigma}(\mathbf{X}).$$

where

$$(6) \quad \varrho(\sigma | \mathbf{X}) = \frac{\varrho(\mathbf{X}, \sigma)}{\sum_{\sigma' \in P_N} \varrho(\mathbf{X}, \sigma')} = \frac{\prod_{i=1}^N e^{-\beta_{\sigma_i} V(\mathbf{x}_i)}}{\sum_{\sigma' \in P_N} \prod_{i=1}^N e^{-\beta_{\sigma'_i} V(\mathbf{x}_i)}}.$$

The generator for the stochastic switching process $(\mathbf{X}(t), \sigma(t)) \in \Omega^N \times P_N$ is then given by

$$(7) \quad (\mathcal{L}_\nu u)(\mathbf{X}, \sigma) = \mathcal{L}_\sigma u(\mathbf{X}, \sigma) - \nu \sum_{\sigma' \neq \sigma} h_{\sigma\sigma'}(\mathbf{X}) (u(\mathbf{X}, \sigma) - u(\mathbf{X}, \sigma')),$$

where the subscript emphasizes the dependence of \mathcal{L}_ν on the overall swapping frequency ν . By construction, any stochastic switch processes constructed this way have $\varrho(\mathbf{X}, \sigma)$ as the invariant measure.

Proposition 1. *For any attempt switching frequency $\nu > 0$, $\pi(d\mathbf{X}, \sigma)$ is the invariant distribution of the process associated with \mathcal{L}_ν .*

Proof. It suffices to show that

$$(8) \quad \mathcal{L}_\nu^* \varrho = 0.$$

This follows from a direct calculation:

$$(9) \quad \begin{aligned} \left(\mathcal{L}_\nu^* \varrho \right) (\mathbf{X}, \sigma) &= \mathcal{L}_\sigma^* \varrho (\mathbf{X}, \sigma) - \nu \sum_{\sigma'} h_{\sigma\sigma'} (\mathbf{X}) \varrho (\mathbf{X}, \sigma) + \nu \sum_{\sigma'} h_{\sigma'\sigma} (\mathbf{X}) \varrho (\mathbf{X}, \sigma') \\ &= \left(\mathcal{L}_\sigma^* \varrho (\mathbf{X} | \sigma) \right) \varrho (\sigma) + \left(-\nu \sum_{\sigma'} h_{\sigma\sigma'} (\mathbf{X}) \varrho (\sigma | \mathbf{X}) + \nu \sum_{\sigma'} h_{\sigma'\sigma} (\mathbf{X}) \varrho (\sigma' | \mathbf{X}) \right) \varrho (\mathbf{X}) \\ &= 0, \end{aligned}$$

where the last equality uses the assumption that $\varrho (\mathbf{X} | \sigma)$ is stationary with respect to the dynamics corresponds to \mathcal{L}_σ and $h_{\sigma\sigma'} (\mathbf{X})$ satisfies the detailed balance condition (5). \square

It is easy to see that we can write $\varrho (\mathbf{X}, \sigma) = (N! \prod_{i=1}^N Z_{\beta_i})^{-1} \exp(-\beta \mathcal{V} (\mathbf{X}, \sigma))$ by defining the potential

$$(10) \quad \mathcal{V} (\mathbf{X}, \sigma) = \beta^{-1} \sum_{i=1}^N \beta_{\sigma_i} V (\mathbf{x}_i),$$

where β is a reference inverse temperature introduced for dimensional purpose: for example, we can take $\beta = \beta_1$. Similarly, the conditional density (4) can be written in terms of $\mathcal{V} (\mathbf{X}, \sigma)$ as

$$(11) \quad \varrho (\mathbf{X} | \sigma) = \left(\prod_{i=1}^N Z_{\beta_i}^{-1} \right) \exp(-\beta \mathcal{V} (\mathbf{X}, \sigma)),$$

which can be sampled by, for example, by the diffusion

$$(12) \quad \begin{aligned} d\mathbf{x}_i &= -\nabla_{\mathbf{x}_i} \mathcal{V} (\mathbf{X}, \sigma) + \sqrt{2\beta^{-1}} d\mathbf{W}_t^{(i)} \\ &= -\beta^{-1} \beta_{\sigma_i} \nabla V (\mathbf{x}_i) + \sqrt{2\beta^{-1}} d\mathbf{W}_t^{(i)}, \quad i = 1, \dots, N. \end{aligned}$$

We can also re-express (6) as

$$(13) \quad \varrho (\sigma | \mathbf{X}) = \frac{e^{-\beta \mathcal{V} (\mathbf{X}, \sigma)}}{\sum_{\sigma' \in P_N} e^{-\beta \mathcal{V} (\mathbf{X}, \sigma')}}.$$

Thus, to ensure the detailed balance, we may choose h as

$$(14) \quad h_{\sigma, \sigma'} (\mathbf{X}) = a_{\sigma, \sigma'} e^{-\frac{1}{2} \beta (\mathcal{V} (\mathbf{X}, \sigma') - \mathcal{V} (\mathbf{X}, \sigma))},$$

where the symmetric adjacency matrix $a_{\sigma, \sigma'} = a_{\sigma', \sigma} \in \{0, 1\}$ indicates whether the permutation σ' is allowed to be accessed from σ by a single jump. For example, we may restrict ourselves to transposition moves in which two random indices are swapped from the permutation σ , so that $\binom{N}{2}$ permutations σ' are accessible from σ . Note that the choice (14) is not unique, and $h_{\sigma, \sigma'}$ can also take other forms that satisfy the detailed balance condition with respect to $\varrho (\sigma | \mathbf{X})$.

Calculation of Expectations. To estimate the expectation of the physical observable $A(x)$ at any temperature β_k , $k = 1, \dots, N$, we can use:

$$\begin{aligned}
 \langle A \rangle_{\beta_k} &= \int_{\Omega} A(\mathbf{x}) \rho_{\beta_k}(\mathbf{x}) \, d\mathbf{x} \\
 &= \sum_{j=1}^N \sum_{\sigma \in P_N} \int_{\Omega^N} A(\mathbf{x}_j) \mathbb{1}_{\sigma_j=k} \varrho(\sigma | \mathbf{X}) \varrho(\mathbf{X}) \, d\mathbf{X} \\
 &= \sum_{j=1}^N \int_{\Omega^N} A(\mathbf{x}_j) \eta_{j,k}(\mathbf{X}) \varrho(\mathbf{X}) \, d\mathbf{X}.
 \end{aligned}
 \tag{15}$$

Here $\varrho(\mathbf{X})$ is the marginal of $\varrho(\mathbf{X}, \sigma)$ on \mathbf{X} ,

$$\varrho(\mathbf{X}) = \sum_{\sigma \in P_N} \varrho(\mathbf{X}, \sigma) = \frac{1}{N!} \sum_{\sigma \in P_N} \prod_{i=1}^N \rho_{\beta_{\sigma_i}}(\mathbf{x}_i),
 \tag{16}$$

and we defined

$$\eta_{j,k}(\mathbf{X}) = \sum_{\sigma \in P_N} \mathbb{1}_{\sigma_j=k} \varrho(\sigma | \mathbf{X}) = \frac{\sum_{\sigma \in P_N} \mathbb{1}_{\sigma_j=k} e^{-\beta^\gamma(\mathbf{X}, \sigma)}}{\sum_{\sigma \in P_N} e^{-\beta^\gamma(\mathbf{X}, \sigma)}},
 \tag{17}$$

which is the probability that the j -th replica is at the k -th temperature conditional on the configuration being fixed at \mathbf{X} . (15) means that we can estimate $\langle A \rangle_{\beta_k}$ from a sample path of the stochastic switching process using ergodicity as

$$\langle A \rangle_{\beta_k} = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \sum_{j=1}^N A(\mathbf{x}_j(t)) \eta_{j,k}(\mathbf{X}(t)) \, dt.
 \tag{18}$$

The following proposition shows that the variance of the parallel tempering estimator for the expectation of an observable at a given temperature is bounded by the variance of the observable at that temperature.

Proposition 2. *For each $k = 1, \dots, N$ we have*

$$\int_{\Omega^N} \left(\sum_{j=1}^N A(\mathbf{x}_j) \eta_{j,k}(\mathbf{X}) \right)^2 \varrho(\mathbf{X}) \, d\mathbf{X} \leq \int_{\Omega} |A(\mathbf{x})|^2 \rho_{\beta_k}(\mathbf{x}) \, d\mathbf{x}.
 \tag{19}$$

Proof. By definition

$$\sum_{j=1}^N \eta_{j,k}(\mathbf{X}) = \sum_{j=1}^N \sum_{\sigma \in P_N} \mathbb{1}_{\sigma_j=k} \varrho(\sigma | \mathbf{X}) = \sum_{\sigma \in P_N} \varrho(\sigma | \mathbf{X}) = 1.
 \tag{20}$$

Therefore, by Jensen's inequality, we have

$$\left(\sum_{j=1}^N A(\mathbf{x}_j) \eta_{j,k}(\mathbf{X}) \right)^2 \leq \sum_{j=1}^N |A(\mathbf{x}_j)|^2 \eta_{j,k}(\mathbf{X}),
 \tag{21}$$

and hence

$$\int_{\Omega^N} \left(\sum_{j=1}^N A(\mathbf{x}_j) \eta_{j,k}(\mathbf{X}) \right)^2 \varrho(\mathbf{X}) \, d\mathbf{X} \leq \int \sum_{j=1}^N |A(\mathbf{x}_j)|^2 \eta_{j,k}(\mathbf{X}) \varrho(\mathbf{X}) \, d\mathbf{X} = \int_{\Omega} |A(\mathbf{x})|^2 \rho_{\beta_k}(\mathbf{x}) \, d\mathbf{x},
 \tag{22}$$

where the last equality follows from (15). This establishes (19). \square

Note that since the bound (19) follows from Jensen's equality, we expect it to be sharp for generic observables. Therefore, the efficiency of the sampling scheme is determined by the convergence to equilibrium of the parallel tempering scheme, as we will discuss in the following sections.

3. LARGE DEVIATION PRINCIPLE FOR THE EMPIRICAL MEASURE OF THE STOCHASTIC SWITCHING PROCESS

In this section we derive a large deviation principle for the empirical measure of the stochastic switching process $(\mathbf{X}(t), \sigma(t))$ marginalized on $\mathbf{X}(t)$, that is:

$$(23) \quad \pi_T(d\mathbf{X}) = \frac{1}{T} \int_0^T \delta_{\mathbf{X}(t)}(d\mathbf{X}) dt.$$

By construction of the stochastic switching process its infinitesimal generator splits into the two parts (see (7)):

$$(24) \quad \mathcal{L}_v = \mathcal{L}_\sigma + v \mathcal{L}_{\text{jump}},$$

As a consequence, the large deviation rate functional for the empirical measure also has an additive structure:

Proposition 3. *The empirical measure (23) satisfies a large deviation principle with rate functional given by*

$$(25) \quad I^v(\mu) = J_\sigma(\mu) + v J_{\text{jump}}(\mu),$$

where μ is a probability measure on $\Omega^N \times P_N$. In particular, both J_σ and J_{jump} are independent of v . While the specific form of J_σ depends on the choice of \mathcal{L}_σ , J_{jump} is fully determined by the jump intensity $h_{\sigma\sigma'}(\mathbf{X})$ as:

$$(26) \quad J_{\text{jump}}(\mu) = \frac{1}{2} \sum_{\sigma, \sigma' \in P_N} \int_{\Omega^N} h_{\sigma\sigma'}(\mathbf{X}) \left[1 - \sqrt{\frac{\theta(\mathbf{X}, \sigma')}{\theta(\mathbf{X}, \sigma)}} \right]^2 \mu(d\mathbf{X}, \sigma),$$

where we have assumed that μ is absolutely continuous with respect to π and denote $\theta(\mathbf{X}, \sigma) := [d\mu/d\pi](\mathbf{X}, \sigma)$.

The additive structure was first observed in the work of Dupuis et al. [6, 13] in the context of parallel tempering.

Proof. The additive structure comes from the following representation of the rate functional [4, 5]:

$$(27) \quad \begin{aligned} I^v(\mu) &= - \sum_{\sigma \in P_N} \int_{\Omega^N} \theta(\mathbf{X}, \sigma)^{1/2} \mathcal{L}_v(\theta(\mathbf{X}, \sigma)^{1/2}) \pi(d\mathbf{X}, \sigma) \\ &= - \sum_{\sigma \in P_N} \int_{\Omega^N} \theta(\mathbf{X}, \sigma)^{1/2} \mathcal{L}_\sigma(\theta(\mathbf{X}, \sigma)^{1/2}) \pi(d\mathbf{X}, \sigma) \\ &\quad - v \sum_{\sigma \in P_N} \int_{\Omega^N} \theta(\mathbf{X}, \sigma)^{1/2} \mathcal{L}_{\text{jump}}(\theta(\mathbf{X}, \sigma)^{1/2}) \pi(d\mathbf{X}, \sigma). \end{aligned}$$

Using the definition of $\mathcal{L}_{\text{jump}}$, we calculate

$$(28) \quad \begin{aligned} &- \sum_{\sigma \in P_N} \int_{\Omega^N} \theta(\mathbf{X}, \sigma)^{1/2} \mathcal{L}_{\text{jump}}(\theta(\mathbf{X}, \sigma)^{1/2}) \pi(d\mathbf{X}, \sigma) \\ &= \sum_{\sigma, \sigma' \in P_N} \int_{\Omega^N} \theta(\mathbf{X}, \sigma)^{1/2} h_{\sigma\sigma'}(\mathbf{X}) (\theta(\mathbf{X}, \sigma)^{1/2} - \theta(\mathbf{X}, \sigma')^{1/2}) \varrho(\sigma | \mathbf{X}) \varrho(\mathbf{X}) d\mathbf{X}. \end{aligned}$$

Since $h_{\sigma\sigma'}(\mathbf{X})$ satisfies the detailed balance condition, we have

$$\begin{aligned}
 & \sum_{\sigma, \sigma' \in P_N} \int_{\Omega^N} \theta(\mathbf{X}, \sigma)^{1/2} h_{\sigma\sigma'}(\mathbf{X}) (\theta(\mathbf{X}, \sigma)^{1/2} - \theta(\mathbf{X}, \sigma')^{1/2}) \rho(\sigma | \mathbf{X}) \rho(\mathbf{X}) d\mathbf{X} = \\
 (29) \quad & = \sum_{\sigma, \sigma' \in P_N} \int_{\Omega^N} \theta(\mathbf{X}, \sigma)^{1/2} h_{\sigma'\sigma}(\mathbf{X}) (\theta(\mathbf{X}, \sigma)^{1/2} - \theta(\mathbf{X}, \sigma')^{1/2}) \rho(\sigma' | \mathbf{X}) \rho(\mathbf{X}) d\mathbf{X} \\
 & = \sum_{\sigma, \sigma' \in P_N} \int_{\Omega^N} \theta(\mathbf{X}, \sigma')^{1/2} h_{\sigma\sigma'}(\mathbf{X}) (\theta(\mathbf{X}, \sigma')^{1/2} - \theta(\mathbf{X}, \sigma)^{1/2}) \rho(\sigma | \mathbf{X}) \rho(\mathbf{X}) d\mathbf{X},
 \end{aligned}$$

where we swapped the indices σ and σ' to get the second equality. Combined with the previous identity, we get

$$\begin{aligned}
 & - \sum_{\sigma \in P_N} \int_{\Omega^N} \theta(\mathbf{X}, \sigma)^{1/2} \mathcal{L}_{\text{jump}}(\theta(\mathbf{X}, \sigma)^{1/2}) \pi(d\mathbf{X}, \sigma) \\
 (30) \quad & = \frac{1}{2} \sum_{\sigma, \sigma' \in P_N} \int_{\Omega^N} h_{\sigma\sigma'}(\mathbf{X}) \left[\theta(\mathbf{X}, \sigma)^{1/2} - \theta(\mathbf{X}, \sigma')^{1/2} \right]^2 \pi(d\mathbf{X}, \sigma) \\
 & = \frac{1}{2} \sum_{\sigma, \sigma' \in P_N} \int_{\Omega^N} h_{\sigma\sigma'}(\mathbf{X}) \left[1 - \theta(\mathbf{X}, \sigma')^{1/2} \theta(\mathbf{X}, \sigma)^{-1/2} \right]^2 \mu(d\mathbf{X}, \sigma).
 \end{aligned}$$

This is exactly J_{jump} defined in (26). □

4. INFINITE SWAPPING LIMIT

As J_{jump} is non-negative, it is natural to consider taking the limit $\nu \rightarrow \infty$ to maximize the rate functional, which we refer to as the infinite swapping limit. To derive the effective dynamics that emerges in this limit, consider the forward Kolmogorov equation for the stochastic switching process $(\mathbf{X}(t), \sigma(t))$:

$$\begin{aligned}
 (31) \quad & \partial_t \rho(t, \mathbf{X}, \sigma) = \mathcal{L}_\nu^* \rho(t, \mathbf{X}, \sigma) \\
 & = \mathcal{L}_\sigma^* \rho(t, \mathbf{X}, \sigma) - \nu \sum_{\sigma'} h_{\sigma\sigma'}(\mathbf{X}) \rho(t, \mathbf{X}, \sigma) + \nu \sum_{\sigma'} h_{\sigma'\sigma}(\mathbf{X}) \rho(t, \mathbf{X}, \sigma').
 \end{aligned}$$

To take the limit $\nu \rightarrow \infty$, let us introduce the ansatz

$$(32) \quad \rho(t, \mathbf{X}, \sigma) = \rho_0(t, \mathbf{X}, \sigma) + \nu^{-1} \rho_1(t, \mathbf{X}, \sigma) + \nu^{-2} \rho_2(t, \mathbf{X}, \sigma) + \dots$$

Matching orders of ν , we get

$$(33) \quad \mathcal{L}_{\text{jump}}^* \rho_0 = 0;$$

$$(34) \quad \mathcal{L}_{\text{jump}}^* \rho_1 = \partial_t \rho_0 - \mathcal{L}_\sigma^* \rho_0,$$

and similarly for the higher order expansions. The leading order (33) reads

$$(35) \quad (\mathcal{L}_{\text{jump}}^* \rho_0)(t, \mathbf{X}, \sigma) = - \sum_{\sigma'} h_{\sigma\sigma'}(\mathbf{X}) \rho_0(t, \mathbf{X}, \sigma) + \sum_{\sigma'} h_{\sigma'\sigma}(\mathbf{X}) \rho_0(t, \mathbf{X}, \sigma') = 0,$$

which implies that, since by construction $\rho(\sigma | \mathbf{X})$ is the invariant measure of the jumping process associated with $\{h_{\sigma\sigma'}(\mathbf{X})\}$

$$(36) \quad \rho_0(t, \mathbf{X}, \sigma) = f_0(t, \mathbf{X}) \rho(\sigma | \mathbf{X}),$$

for some f_0 to be determined. Substitution this into the next order gives

$$(37) \quad \rho(\sigma | \mathbf{X}) \partial_t f_0(t, \mathbf{X}) = \mathcal{L}_\sigma^* \rho_0 - \sum_{\sigma'} h_{\sigma\sigma'}(\mathbf{X}) \rho_1(t, \mathbf{X}, \sigma) + \sum_{\sigma'} h_{\sigma'\sigma}(\mathbf{X}) \rho_1(t, \mathbf{X}, \sigma').$$

Summing over σ , we obtain the following solvability condition for this equation (recall that $\sum_{\sigma} \rho(\sigma | \mathbf{X}) = 1$)

$$(38) \quad \partial_t f_0(t, \mathbf{X}) = \sum_{\sigma} \mathcal{L}_{\sigma}^* (\rho(\sigma | \mathbf{X}) f_0(t, \mathbf{X})) =: \overline{\mathcal{L}}^* f_0(t, \mathbf{X}),$$

where the last identity defines $\overline{\mathcal{L}}$, the infinitesimal process of the ‘‘averaged process’’ in the infinite swapping limit.

The above asymptotic derivation can be made rigorous as in usual averaging, and we conclude:

Proposition 4. *As $\nu \rightarrow \infty$, $\mathbf{X}(t)$ converges to the averaged process defined by the infinitesimal generator $\overline{\mathcal{L}}$.*

Let us now make these formula explicit. Since

$$(39) \quad \mathcal{L}_{\sigma} = \sum_{k=1}^N \left(-\beta^{-1} \beta_{\sigma_k} \nabla V(\mathbf{x}_k) \cdot \nabla_{\mathbf{x}_k} + \beta^{-1} \Delta_{\mathbf{x}_k} \right),$$

we have

$$(40) \quad \begin{aligned} \overline{\mathcal{L}}^* f(\mathbf{X}) &= \sum_{\sigma \in P_N} \mathcal{L}_{\sigma}^* (\rho(\sigma | \mathbf{X}) f(\mathbf{X})) \\ &= \sum_{\sigma \in P_N} \sum_{k=1}^N \left[\nabla_{\mathbf{x}_k} \cdot \left((\nabla_{\mathbf{x}_k} \mathcal{V}(\mathbf{X}, \sigma)) \rho(\sigma | \mathbf{X}) f(\mathbf{X}) \right) + \beta^{-1} \Delta_{\mathbf{x}_k} (\rho(\sigma | \mathbf{X}) f(\mathbf{X})) \right] \\ &= \sum_{\sigma \in P_N} \sum_{k=1}^N \nabla_{\mathbf{x}_k} \cdot \left((\nabla_{\mathbf{x}_k} \mathcal{V}(\mathbf{X}, \sigma)) \rho(\sigma | \mathbf{X}) f(\mathbf{X}) \right) + \beta^{-1} \sum_{k=1}^N \Delta_{\mathbf{x}_k} f(\mathbf{X}), \end{aligned}$$

where the last equality uses $\sum_{\sigma} \rho(\sigma | \mathbf{X}) = 1$ for any \mathbf{X} . Therefore, the corresponding stochastic differential equations are given by

$$(41) \quad \begin{aligned} d\mathbf{x}_j &= - \sum_{\sigma \in P_N} \rho(\sigma | \mathbf{X}) \nabla_{\mathbf{x}_j} \mathcal{V}(\mathbf{X}, \sigma) dt + \sqrt{2\beta^{-1}} d\mathbf{W}_t^{(j)} \\ &= -R_j(\mathbf{X}) \nabla V(\mathbf{x}_j) dt + \sqrt{2\beta^{-1}} d\mathbf{W}_t^{(j)} \end{aligned}$$

with the scaling factor $R_j(\mathbf{X})$ given by

$$(42) \quad R_j(\mathbf{X}) = \beta^{-1} \sum_{\sigma \in P_N} \beta_{\sigma_j} \rho(\sigma | \mathbf{X}) = \frac{\sum_{\sigma \in P_N} \beta_{\sigma_j} e^{-\beta \mathcal{V}(\mathbf{X}, \sigma)}}{\sum_{\sigma \in P_N} e^{-\beta \mathcal{V}(\mathbf{X}, \sigma)}}.$$

The dynamics we get after averaging is similar to the original overdamped dynamics under inverse temperature β , except that the drift is scaled by the factor $R_j(\mathbf{X})$. Notice that in terms of the factor $\eta_{j,k}(\mathbf{X})$ defined in (17), we can write (42) as

$$(43) \quad R_j(\mathbf{X}) = \beta^{-1} \sum_{k=1}^N \beta_k \eta_{j,k}(\mathbf{X}).$$

Finally, note that (41) can be written as

$$(44) \quad d\mathbf{x}_j = -\nabla_{\mathbf{x}_j} \mathcal{V}(\mathbf{X}) dt + \sqrt{2\beta^{-1}} d\mathbf{W}_t^{(j)},$$

where we defined the mixture potential

$$(45) \quad \mathcal{V}(\mathbf{X}) = -\beta^{-1} \log \sum_{\sigma \in P_N} e^{-\beta \mathcal{V}(\mathbf{X}, \sigma)} = -\beta^{-1} \log \sum_{\sigma \in P_N} e^{-\sum_{j=1}^N \beta_{\sigma(j)} V(\mathbf{x}_j)}$$

As a result, the convergence properties of parallel tempering in the infinite swapping limit (and the boost in efficiency this method provides on vanilla sampling via (1)) can be analyzed via examination of the geometrical properties of $\mathcal{V}(\mathbf{X})$. This is what we will do in the next section on an example.

5. EFFICIENCY ANALYSIS IN THE HARMONIC CASE

From the discussion above, we know that, given any set of temperatures, the optimal way to operate parallel tempering is in fact to simulate the effective equation in (41) that emerges in the $\nu \rightarrow \infty$ limit. This is clearly doable to the extent that we can calculate the factors $R_j(\mathbf{X})$ defined in (42). Since these factors contained sums over $\sigma \in P_N$, i.e., over $N!$ terms, calculating $R_j(\mathbf{X})$ by exhaustive evaluation of these sums is only an option if the number N of temperatures remains moderately low. In practice, however, one often needs to use $N = 50 - 100$, for which these exhaustive calculations are not doable. How to proceed in these situations will be explained in Section 6. The question we address here is why there is a practical need to use a rather larger number of temperatures rather than, say, $N = 2$. This will also allow us to investigate how to choose the temperature ladder in the method.

Insight about this issue can be gained by looking at the special case when the potential is harmonic

$$(46) \quad V(\mathbf{x}) = \frac{1}{2}|\mathbf{x}|^2, \quad \mathbf{x} \in \mathbb{R}^D,$$

in which case the canonical density reduces to

$$(47) \quad \rho_\beta(\mathbf{x}) = (2\pi\beta^{-1})^{-\frac{D}{2}} e^{-\frac{1}{2}\beta|\mathbf{x}|^2}.$$

Even though this example is very simple, it has the advantage to be amenable to analysis and it illustrates difficulties that are shared with more realistic (and complex) systems. We can always set $\beta = 1$ by rescaling $\mathbf{x} \rightarrow \beta^{-1}\mathbf{x}$, so let us focus on this case here. As we have seen in Proposition 2, the variance of the estimator for parallel tempering can be controlled by the variance of the observable, and thus we shall just focus on the convergence of the dynamics to the invariant measure.

5.1. The case with two temperatures. Let us first consider the situation where we only add one artificial temperature, so that we have two replicas with $\beta_1 = \beta = 1$ and $\beta_2 = \bar{\beta} < 1$. In this case, the effective equation we obtain in the $\nu \rightarrow \infty$ limit reads

$$(48) \quad \begin{aligned} d\mathbf{x}_1 &= -r_1(V(\mathbf{x}_1), V(\mathbf{x}_2))\nabla V(\mathbf{x}_1) dt + \sqrt{2}dW_1 = -r_1\left(\frac{1}{2}|\mathbf{x}_1|^2, \frac{1}{2}|\mathbf{x}_2|^2\right)\mathbf{x}_1 dt + \sqrt{2}dW_1 \\ d\mathbf{x}_2 &= -r_2(V(\mathbf{x}_1), V(\mathbf{x}_2))\nabla V(\mathbf{x}_2) dt + \sqrt{2}dW_2 = -r_2\left(\frac{1}{2}|\mathbf{x}_1|^2, \frac{1}{2}|\mathbf{x}_2|^2\right)\mathbf{x}_2 dt + \sqrt{2}dW_2 \end{aligned}$$

where

$$(49) \quad r_1(E_1, E_2) = \frac{e^{-E_1 - \bar{\beta}E_2} + \bar{\beta}e^{-E_2 - \bar{\beta}E_1}}{e^{-E_1 - \bar{\beta}E_2} + e^{-E_2 - \bar{\beta}E_1}}, \quad r_2(E_1, E_2) = r_1(E_2, E_1)$$

Therefore, we can write down a closed evolution equation for the law of $\mathcal{E}_1 = \frac{1}{2D}|\mathbf{x}_1|^2$ and $\mathcal{E}_2 = \frac{1}{2D}|\mathbf{x}_2|^2$. A few simple manipulations show that this equation can be written as

$$(50) \quad d \begin{pmatrix} \mathcal{E}_1 \\ \mathcal{E}_2 \end{pmatrix} = \begin{pmatrix} 2\mathcal{E}_1 & 0 \\ 0 & 2\mathcal{E}_2 \end{pmatrix} \begin{pmatrix} \partial_{\mathcal{E}_1} \log g(\mathcal{E}_1, \mathcal{E}_2) \\ \partial_{\mathcal{E}_2} \log g(\mathcal{E}_1, \mathcal{E}_2) \end{pmatrix} dt + D^{-1} \begin{pmatrix} 2 \\ 2 \end{pmatrix} dt + \sqrt{2D^{-1}} \begin{pmatrix} \sqrt{2\mathcal{E}_1} & 0 \\ 0 & \sqrt{2\mathcal{E}_2} \end{pmatrix} \begin{pmatrix} dW_1 \\ dW_2 \end{pmatrix},$$

where

$$(51) \quad g(\mathcal{E}_1, \mathcal{E}_2) = (\mathcal{E}_1\mathcal{E}_2)^{\frac{1}{2} - \frac{1}{D}} \left(e^{-D(\mathcal{E}_1 + \bar{\beta}\mathcal{E}_2)} + e^{-D(\mathcal{E}_2 + \bar{\beta}\mathcal{E}_1)} \right)^{1/D}.$$

Writing compactly $\mathcal{E} = (\mathcal{E}_1, \mathcal{E}_2)^\top$, this equation is of the form

$$(52) \quad d\mathcal{E} = M(\mathcal{E}) \nabla_{\mathcal{E}} \log g(\mathcal{E}) dt + D^{-1} \operatorname{div} M(\mathcal{E}) dt + \sqrt{2D^{-1}} M^{1/2}(\mathcal{E}) dW$$

with $M(\mathcal{E}) = \operatorname{diag}(2\mathcal{E}_1, 2\mathcal{E}_2)$, which indicates that its invariant density is proportional to g^D , i.e. it is given by

$$(53) \quad \varrho(\mathcal{E}_1, \mathcal{E}_2) = \mathcal{C}_D^{-1} (\mathcal{E}_1 \mathcal{E}_2)^{\frac{D}{2}-1} \left(e^{-D(\mathcal{E}_1 + \bar{\beta} \mathcal{E}_2)} + e^{-D(\mathcal{E}_2 + \bar{\beta} \mathcal{E}_1)} \right),$$

where the normalization constant is given by

$$\begin{aligned} \mathcal{C}_D &= \int (\mathcal{E}_1 \mathcal{E}_2)^{\frac{D}{2}-1} \left(e^{-D(\mathcal{E}_1 + \bar{\beta} \mathcal{E}_2)} + e^{-D(\mathcal{E}_2 + \bar{\beta} \mathcal{E}_1)} \right) d\mathcal{E}_1 d\mathcal{E}_2 \\ &= 2D^{-D/2} (\bar{\beta} D)^{-D/2} \Gamma\left(\frac{D}{2}\right)^2 \\ &\sim \frac{2}{D} (2e)^{-D} \bar{\beta}^{-D/2} \end{aligned}$$

as $D \rightarrow \infty$. Thus, (50) and equivalently (52) describe diffusion on the energy landscape

$$(54) \quad -\log g(\mathcal{E}_1, \mathcal{E}_2).$$

As long as $\bar{\beta} < 1$, this landscape possesses two minima with a saddle point in between (see Fig. 1). For large N , the minima are approximately (that is, to leading order in N^{-1}) located at




FIGURE 1. Energy landscape of \mathcal{E}_1 and \mathcal{E}_2 for $D = 100$ and $\bar{\beta} = 0.2$. The white circles show the location of the two minima and the saddle points, located approximately at (55) and (56), respectively. The potential is also capped at 200 for clarity.

$$(55) \quad (\mathcal{E}_1, \mathcal{E}_2) = \frac{1}{2}(1, \bar{\beta}^{-1}), \quad \text{and} \quad (\mathcal{E}_1, \mathcal{E}_2) = \frac{1}{2}(\bar{\beta}^{-1}, 1)$$

with energy approximately given by $E_m = 1 + \log 2 + \frac{1}{2} \log \bar{\beta}$ and the saddle point is approximately located at

$$(56) \quad (\mathcal{E}_1, \mathcal{E}_2) = ((1 + \bar{\beta})^{-1}, (1 + \bar{\beta})^{-1})$$

with energy approximately given by $E_s = 1 + \log(1 + \bar{\beta})$. Therefore, for large D , we can use results from large deviation theory to deduce that the smallest nonzero eigenvalue of the generator of (50) (that is, roughly, the inverse of mean first passage time between the minima) is log-asymptotically given by

$$(57) \quad \bar{\lambda}_1 \asymp \exp(-D\Delta E), \quad \Delta E = E_s - E_m = \log(1 + \bar{\beta}) - \log 2 + \frac{1}{2} \log \bar{\beta}^{-1} \quad \text{as } D \rightarrow \infty.$$

Therefore, for any $\bar{\beta} < 1$ fixed, the system becomes increasingly metastable as D gets larger. In particular, if we start with a configuration with $\mathcal{E}_1 < \mathcal{E}_2$, it will take exponentially (in D) long time for the dynamics to reach the configurations with $\mathcal{E}_1 > \mathcal{E}_2$. This means exponentially slow convergence to equilibrium, and shows that using two temperature is not sufficient in general.

5.2. The case with many temperatures. Consider now what happens with multiple temperatures, in which case a similar analysis as in the previous section can be done. In particular, the mixture potential (for the energies) has local minima at

$$(58) \quad \mathcal{E}_i = \frac{1}{2\beta_{\sigma_i}} \quad \text{for any permutation } \sigma,$$

whose energies scale as

$$(59) \quad \frac{1}{2} \sum_i \log \beta_i + \frac{N}{2} (\log 2 + 1).$$

The saddle points are located at

$$(60) \quad \begin{cases} \mathcal{E}_i = \frac{1}{2\beta_{\sigma_i}} & \text{for } i \neq j, k \\ \mathcal{E}_j = \mathcal{E}_k = (\beta_{\sigma_j} + \beta_{\sigma_k})^{-1} \end{cases}$$

for any permutation σ and any pair $j \neq k$, whose energies scale as

$$(61) \quad \frac{1}{2} \sum_{i \neq j, k} \log \beta_j + \log(\beta_j + \beta_k) + \frac{N-2}{2} \log 2 + \frac{N}{2}.$$

Therefore the barrier scale as

$$(62) \quad \Delta E_{j,k} = \log\left(\frac{\beta_j + \beta_k}{(\beta_j \beta_k)^{1/2}}\right) - \log 2.$$

Assuming that $\beta_1 > \beta_2 > \dots > \beta_N$, this indicates that we are more likely to see transition between minima lying on adjacent β_j s, and in order to keep all the rates equal (if we use Arrhenius formula such that rate is given by exponential of the energy barrier), we must have

$$(63) \quad \frac{\beta_j + \beta_{j+1}}{(\beta_j \beta_{j+1})^{1/2}} = \text{cst}$$

for all $j = 1, \dots, N$. This agrees with the analysis based on transition state theory performed in our previous work [11], and shows that the optimal choice of temperature on the ladder should satisfy

$$(64) \quad \frac{\beta_{j+1}}{\beta_j} = \left(\frac{\beta_N}{\beta_1}\right)^{1/(N-1)}$$

and that this ratio should be of order 1. This indicates that using multiple temperature removes the energy barrier and hence makes convergence much faster.

6. SIMULATION OF THE STOCHASTIC SWITCHING DYNAMICS

6.1. Direct simulation based on stochastic simulation algorithm. Here we present an exact implementation of the stochastic switching dynamics, assuming that we can integrate the Markov process corresponds to \mathcal{L}_σ exactly. In practice, this integration often needs to be discretized, which will introduce some additional errors. We will discuss these later.

Given that at time t the system is in the state $(\mathbf{X}(t), \sigma(t))$, the algorithm updates the state via the following steps:

1. Draw a random number $r \sim U([0, 1])$;
2. Integrate \mathbf{X} with the fixed $\sigma := \sigma(t)$ to time $t + \tau$, such that

$$r = \exp\left(-\nu \sum_{\sigma' \neq \sigma} \int_t^{t+\tau} h_{\sigma\sigma'}(\mathbf{X}(s)) ds\right),$$

where r is the random number drawn in Step 1;

3. Choose $\sigma(t + \tau)$ with probability

$$\mathbb{P}(\sigma(t + \tau) = \sigma' \mid \sigma, \tau) = \frac{h_{\sigma\sigma'}(\mathbf{X}(t + \tau))}{\sum_{\tilde{\sigma} \neq \sigma} h_{\sigma\tilde{\sigma}}(\mathbf{X}(t + \tau))};$$

4. Update the time to $t + \tau$ and repeat till the target time.

If a large swapping rate ν is employed, the time lags τ will likely be rather small as their expectation is proportional to ν^{-1} . In these situations, the stochastic switching dynamics exhibits two scales: the time scale of the diffusion process and that of the jumping process which is much faster. As a result, direct simulation by the algorithm discussed above is no longer efficient. This difficulty is in fact very similar to the one we encounter when numerically integrating SDEs with multiple time scales, for which a by-now standard idea is to explore the asymptotic limit explicitly to design a more efficient multiscale integrator. In what follows, we explore such an integrator, which was proposed in [21] and fits the framework of the heterogeneous multiscale methods (HMM) [7, 9, 19].

6.2. A multiscale integrator. The basic idea of the HMM integrator is to compute on-the-fly the scaling factor $R_j(\mathbf{X})$ in the limiting equation (41) for the slow variables \mathbf{X} using short bursts of simulation of the fast process σ – these simulations are performed in what is called the microsolver in HMM, whereas the routine used to evolve the slow variables is referred to as the macrosolver. The data generated by the microsolver is passed to the unknown coefficients for the macrosolver via an estimator. In the present context, the HMM scheme we will use works as follows:

Choose a time step Δt appropriate for the simulation of the limiting SDEs (41) and pick a frequency ν such that $\nu \gg 1/\Delta t$ (for example $\nu = 10^2/\Delta t - 10^3/\Delta t$). Start the simulation with $\mathbf{X}^{(0)} = \mathbf{X}(0)$ and $\sigma^{(0)} = \sigma(0)$, and set $t_0 = 0$. Then for $k \geq 0$, we use the following 3-step iteration algorithm to update:

1. *Microsolver*: Evolve $\sigma^{(k)}$ via SSA from t_k to $t_{k+1} := t_k + \Delta t$ using the rate in (14) with $\mathbf{X} = \mathbf{X}^{(k)}$ fixed, that is: Set $\sigma^{(k,0)} = \sigma^{(k)}$, $t_{k,0} = t_k$, and for $l \geq 1$, do:
 - (a) Compute the lag to the next event via

$$\tau_l = -\frac{\ln r}{\nu h_{\sigma^{(k,l-1)}}(\mathbf{X}^{(k)})}$$

where r is a random number uniformly distributed in the interval $(0, 1)$ and $h_\sigma(\mathbf{X}^{(k)}) = \sum_{\sigma' \neq \sigma} h_{\sigma, \sigma'}(\mathbf{X}^{(k)})$;

(b) Pick $\sigma_{k,l}$ with probability

$$\mathbb{P}_{\sigma^{(k,l)}} = \frac{h_{\sigma^{(k,l-1)}, \sigma^{(k,l)}}(\mathbf{X}^{(k)})}{h_{\sigma^{(k,l-1)}}(\mathbf{X}^{(k)})};$$

(c) Set $t_{k,l} = t_{k,l-1} + \tau_l$ and repeat till the first L such that $t_{k,L} > t_k + \Delta t$; then set $\sigma^{(k+1)} = \sigma^{(k,L)}$ and reset $\tau_L = t_k + \Delta t - t_{k,L-1}$.

2. *Estimator*: Given the trajectory of σ , estimate $\eta_{j,i}(\mathbf{X}^{(k)})$ and $R_j(\mathbf{X}^{(k)})$ via

$$\begin{aligned} \hat{\eta}_{j,i}(\mathbf{X}^{(k)}) &= \frac{1}{\Delta t} \int_{t_k}^{t_k + \Delta t} \mathbb{1}_{i=\sigma_j(s)} \, ds \\ &= \frac{1}{\Delta t} \sum_{l=1}^L \tau_l \mathbb{1}_{i=\sigma_j^{(k,l)}}, \end{aligned}$$

and

$$\begin{aligned} \hat{R}_j(\mathbf{X}^{(k)}) &= \beta^{-1} \sum_i \frac{\beta_i}{\Delta t} \int_{t_k}^{t_k + \Delta t} \mathbb{1}_{i=\sigma_j(s)} \, ds \\ &= \beta^{-1} \sum_i \beta_i \hat{\eta}_{j,i}(\mathbf{X}^{(k)}). \end{aligned}$$

3. *Macrosolver*: Evolve $\mathbf{X}^{(k)}$ to $\mathbf{X}^{(k+1)}$ using one time-step of size Δt in the numerical integrator for (41) with $R_j(\mathbf{X}^{(k)})$ replaced by the factor $\hat{R}_j(\mathbf{X}^{(k)})$ calculated in the estimator. Then repeat the three steps above for each time-step k .

Remark. It is useful to compare the above scheme with a time-splitting integrator for the dynamics of $(\mathbf{X}(t), \sigma(t))$, that is for time step size Δt , we alternate between evolving $\mathbf{X}(t)$ and $\sigma(t)$ while freezing the other component. More precisely, from t_k to $t_{k+1} = t_k + \Delta t$ the algorithm updates the state via

1. Evolve $\sigma^{(k)}$ via SSA from t_k to $t_{k+1} := t_k + \Delta t$ using the rate in (14) with $\mathbf{X} = \mathbf{X}^{(k)}$ to get $\sigma^{(k+1)}$, using the similar procedure as the *Microsolver* in the multiscale integrator;
2. Evolve $\mathbf{X}^{(k)}$ to $\mathbf{X}^{(k+1)}$ using $\mathcal{L}_{\sigma^{(k+1)}}$, namely, if Euler-Maruyama scheme is used

$$(65) \quad \mathbf{x}_j^{(k+1)} = \mathbf{x}_j^{(k)} - \Delta t \beta^{-1} \beta_{\sigma_j^{(k+1)}} \nabla V(\mathbf{x}_j^{(k)}) + \sqrt{2\beta^{-1} \Delta t} \eta_j^{(k)},$$

where $\eta_j^{(k)}$ are i.i.d. standard Gaussian random variables.

The difference with the multiscale integrator lies in the evolution of \mathbf{X} , where the splitting scheme uses the current instance $\sigma^{(k+1)}$ while the multiscale integrator incorporates into the average of the coefficients. Since we are interested in the regime that the dynamics of σ is much faster of \mathbf{X} , using the average is preferred as it is consistent with the $\nu \rightarrow \infty$ limit.

We will not go into details of the numerical analysis of the multiscale integrator here. Numerical analysis of HMM integrators in similar spirit can be found in [7, 8, 10, 18–20], the basic idea is to show that the scheme is consistent with the averaging limit as $\nu \rightarrow \infty$ by showing that the estimator gives a good approximation to $R_j(\mathbf{X})$.

7. EXTENSION TO SIMULATED TEMPERING

The analysis we have performed in the context of parallel tempering can be generalized to simulated tempering, as we briefly explain in this section. For a performance comparison between parallel and simulated tempering, we refer the readers to [22].

We begin by noting that simulated tempering, in which we use one single replica but make the temperatures dynamically switch between N values $\beta_1, \beta_2, \dots, \beta_N$, can be cast in a framework similar to that of the stochastic switching process discussed in Section 2. This amounts to assuming that the state space of the process is $\Omega^N \times \{1, \dots, N\}$, and requiring that this process samples the joint equilibrium measure

$$(66) \quad \pi(d\mathbf{x}, i) = \varrho(\mathbf{x}, i) d\mathbf{x}, \quad \text{with} \quad \varrho(\mathbf{x}, i) = \frac{n_i e^{-\beta_i V(\mathbf{x})}}{\sum_j n_j Z_{\beta_j}}.$$

where $\{n_i\}_{i=1, \dots, N}$ is a set of positive weights to be chosen *a priori*. To this end, the dynamics in x given i can be chosen as a time rescaled overdamped equation:

$$(67) \quad d\mathbf{x}(t) = -\beta^{-1} \beta_i \nabla V(\mathbf{x}(t)) + \sqrt{2\beta^{-1}} d\mathbf{W}_t,$$

so that it samples

$$(68) \quad \varrho(\mathbf{x} | i) = \frac{\varrho(\mathbf{x}, i)}{\int_{\Omega} \varrho(\mathbf{y}, i) d\mathbf{y}} = Z_{\beta_i}^{-1} e^{-\beta_i V(\mathbf{x})}.$$

The jumping rate matrix between the temperature indices i given x can be chosen e.g.

$$(69) \quad \nu k_{ij}(\mathbf{x}) = \nu \min\left(\frac{n_j}{n_i} e^{-(\beta_j - \beta_i)V(\mathbf{x})}, 1\right), \quad i \neq j$$

with $k_{ii}(\mathbf{x}) = -\sum_{j \neq i} k_{ij}$ so that the jumping process samples

$$(70) \quad \varrho(i | \mathbf{x}) = \frac{\varrho(\mathbf{x}, i)}{\sum_{j=1}^N \varrho(\mathbf{x}, j)} = \frac{n_i e^{-\beta_i V(\mathbf{x})}}{\sum_j n_j e^{-\beta_j V(\mathbf{x})}}.$$

Here ν is again an overall scaling parameter for the switching frequency. Other choices such as

$$(71) \quad k_{ij}(\mathbf{x}) = \frac{n_j}{n_i e^{-(\beta_i - \beta_j)V(\mathbf{x})} + n_j}, \quad i \neq j,$$

or

$$(72) \quad k_{ij}(\mathbf{x}) = \sqrt{\frac{n_j}{n_i}} e^{-\frac{1}{2}(\beta_j - \beta_i)V(\mathbf{x})}, \quad i \neq j,$$

are possible as well. The infinitesimal generator of the stochastic switching process $(\mathbf{x}(t), i(t))$ defined this way is given by

$$(73) \quad (\mathcal{L}_{\nu} u)(\mathbf{x}, i) = -\beta^{-1} \beta_i \nabla V(\mathbf{x}) \cdot \nabla u(\mathbf{x}, \beta) + \beta^{-1} \Delta u(\mathbf{x}, \beta) + \nu \sum_{j \neq i} k_{ij}(\mathbf{x}) (u(\mathbf{x}, j) - u(\mathbf{x}, i)),$$

and it is an easy exercise to show that its invariant measure is indeed (66).

In terms of calculating expectations, in the context of simulated tempering we can use

$$(74) \quad \begin{aligned} \langle A \rangle_{\beta_k} &\equiv \int_{\Omega} A(\mathbf{x}) \rho_{\beta_k}(\mathbf{x}) d\mathbf{x} \\ &= \int_{\Omega} A(\mathbf{x}) g_k(\mathbf{x}) \varrho(\mathbf{x}) d\mathbf{x} \end{aligned}$$

where $\varrho(\mathbf{x})$ is the marginal of $\varrho(\mathbf{x}, i)$ on \mathbf{x} :

$$(75) \quad \varrho(\mathbf{x}) = \sum_{i=1}^N \varrho(\mathbf{x}, i) = \frac{\sum_{i=1}^N n_i e^{-\beta_i V(\mathbf{x})}}{\sum_{i=1}^N n_i Z_{\beta_i}},$$

and we defined

$$(76) \quad g_k(\mathbf{x}) = \frac{\rho_{\beta_k}(\mathbf{x})}{\varrho(\mathbf{x})} = \frac{Z_{\beta_k}^{-1} e^{-\beta_k V(\mathbf{x})} \sum_{j=1}^N n_j Z_{\beta_j}}{\sum_{j=1}^N n_j e^{-\beta_j V(\mathbf{x})}}.$$

These expressions show the main (and well-known) difficulty one is faced when using simulated tempering: even though it looks simpler than simulated tempering because it involves a single replica, it requires one to learn the partition functions Z_{β_i} to calculate expectations (in contrast, parallel tempering does not require this).

Setting this difficulty inside, an argument similar to the one presented in Section 3 indicates that the large deviation rate functional for the empirical measure of the process with generator (73) is the sum of the terms, the first of which is independent of ν and the second proportional to ν . This indicates that, in the context of simulating tempering too, it is optimal to take the limit as $\nu \rightarrow \infty$ (which we will refer to as the ‘infinite switching limit’). In this limit, it is easy to show using averaging arguments similar to those presented in Section 4, that the process $\mathbf{x}(t)$ converges to the solution of the following effective equation

$$(77) \quad d\mathbf{x}(t) = -\beta^{-1} \bar{\beta}(\mathbf{x}(t)) \nabla V(\mathbf{x}(t)) + \sqrt{2\beta^{-1}} d\mathbf{W}_t,$$

where we defined

$$(78) \quad \bar{\beta}(\mathbf{x}) = \sum_{i=1}^N \beta_i \varrho(i | \mathbf{x}) = \frac{\sum_{i=1}^N \beta_i n_i e^{-\beta_i V(\mathbf{x})}}{\sum_{i=1}^N n_i e^{-\beta_i V(\mathbf{x})}}.$$

This equation is in principle much easier to simulate than the corresponding effective equation (41) obtained in the infinite switching limit of parallel tempering, since the sum in (78) involves N terms instead of $N!$ (and hence can be evaluated straightforwardly even for large N). However, the efficiency boost of the method depends crucially on the choice of the weights n_i . It has been argued that the optimal choice is to take $n_i = Z_{\beta_i}^{-1}$, but, unfortunately, these partition functions are unknown *a priori* (which brings us back to the issue with the estimator (74) for the expectations). How to go around these difficulties (and justify the optimality of the choice $n_i = Z_{\beta_i}^{-1}$ starting from the effective equation in (77)) will be discussed elsewhere.

8. CONCLUDING REMARKS

We have shown how a formulation of parallel tempering as a stochastic switching process for the coupled dynamics of replica configuration and temperature permutation naturally leads to considering these equations in the infinite swapping limit. Indeed this can be justified from the additive structure of the rate functional for the empirical measure of the process, using tools from large deviation theory as originally proposed in [13] and further justified in [5]. This observation suggests to simulate directly the effective equation that emerges in this limit. Unfortunately, this is no trivial task because the direct calculation of the coefficients in this equation involves sums over all possible permutations of the temperatures. Clearly this is only doable if the number of temperatures/replicas is small, which is a set-up in which parallel tempering leads to no significant boost in efficiency. This observation is common knowledge for practitioners of parallel tempering, and is usually justified by showing that many temperatures are required to get a significant acceptance rate of the temperature swaps. Within the infinite swapping limit, we

arrive at the same conclusion, but from a different viewpoint: as we showed here, it is necessary to use many temperatures in order to eliminate the free energy barriers on the mixture potential.

Once one realizes that the effective equation emerging in the infinite swapping limit must be operated with many temperatures, the question arises as how to do so in practice. Following the suggestion made in [21], we showed here that it can be done using multiscale integrators such as HMM. These schemes are straightforward generalization of those traditionally used in parallel tempering, and we therefore they should be useful to improve the efficiency of this method with rather minimal modifications of the existing codes. Alternatively, we showed that one can use similar ideas in the context of simulated tempering, in which case the limiting equation that arise in the infinite switching limit is simpler.

ACKNOWLEDGMENTS

We thank C. Abrams, B. Leimkuhler, and A. Martinsson for interesting discussions. The work of JL is supported in part by the National Science Foundation (NSF) under grant DMS-1454939. The work of EVE is supported in part by the Materials Research Science and Engineering Center (MRSEC) program of the NSF under award number DMR-1420073 and by NSF under award number DMS-1522767.

REFERENCES

- [1] Giovanni Bussi, Francesco Luigi Gervasio, Alessandro Laio, and Michele Parrinello, *Free-energy landscape for β hair-pin folding from combined parallel tempering and metadynamics*, J. Am. Chem. Soc. **128** (2006), no. 41, 13435–13441.
- [2] M. Cecchini, F. Rao, M. Seeber, and A. Caflisch, *Replica exchange molecular dynamics simulations of amyloid peptide aggregation*, J. Chem. Phys. **121** (2004), no. 21, 10748–10756.
- [3] J. D. Doll and Paul Dupuis, *On performance measures for infinite swapping Monte Carlo methods*, J. Chem. Phys. **142** (2015), 024111.
- [4] M.D. Donsker and S.R.S. Varadhan, *Asymptotic evaluation of certain Markov process expectations for large time. I.*, Comm. Pure Appl. Math. **28** (1975), 1–47.
- [5] P. Dupuis and Y. Liu, *On the large deviation rate function for the empirical measures of reversible jump Markov processes*, Ann. Probab. **43** (2015), 1121–1156.
- [6] Paul Dupuis, Yufei Liu, Nuria Plattner, and J. D. Doll, *On the Infinite Swapping Limit for Parallel Tempering*, Multiscale Model. Simul. **10** (2012), no. 3, 986–1022.
- [7] W. E and B. Engquist, *The heterogeneous multi-scale methods*, Commun. Math. Sci. **1** (2003), 87–133.
- [8] W. E, D. Liu, and E. Vanden-Eijnden, *Analysis of multiscale methods for stochastic differential equations*, Comm. Pure Appl. Math. **58** (2005), 1544–1585.
- [9] Weinan E, Bjorn Engquist, Xiantao Li, Weiqing Ren, and Eric Vanden-Eijnden, *The Heterogeneous Multiscale Method: A Review*, Commun. Comput. Phys. **2** (2007), no. 3, 367–450.
- [10] J. Lu and K. Spiliopoulos, *Multiscale integrators for irreversible langevin samplers*, 2016. preprint, arXiv:1606.09539.
- [11] Jianfeng Lu and Eric Vanden-Eijnden, *Infinite swapping replica exchange molecular dynamics leads to a simple simulation patch using mixture potentials.*, J. Chem. Phys. **138** (February 2013), no. 8, 084105.
- [12] E. Marinari and G. Parisi, *Simulated Tempering: A New Monte Carlo Scheme*, Eur. Lett. **19** (1992), 451–458.
- [13] Nuria Plattner, J D Doll, Paul Dupuis, Hui Wang, Yufei Liu, and J E Gubernatis, *An infinite swapping approach to the rare-event sampling problem.*, J. Chem. Phys. **135** (October 2011), no. 13, 134111.
- [14] Edina Rosta and Gerhard Hummer, *Error and efficiency of replica exchange molecular dynamics simulations*, J. Chem. Phys. **131** (2009), no. 16, 165102.
- [15] Daniel Sindhikara, Yilin Meng, and Adrian E. Roitberg, *Exchange frequency in replica exchange molecular dynamics*, J. Chem. Phys. **128** (2008), no. 2, 0–10.
- [16] Daniel J. Sindhikara, Daniel J. Emerson, and Adrian E. Roitberg, *Exchange often and properly in replica exchange molecular dynamics*, J. Chem. Theory Comput. **6** (2010), no. 9, 2804–2808.
- [17] Yuji Sugita and Yuko Okamoto, *Replica exchange molecular dynamics method for protein folding*, Chem. Phys. Lett. **314** (1999), no. November, 141–151.
- [18] M. Tao, H. Owhadi, and J. E. Marsden, *Nonintrusive and structure preserving multiscale integration of stiff ODEs, SDEs, and Hamiltonian systems with hidden slow dynamics via flow averaging*, Multiscale Modeling and Simulation **8** (2010), 1269–1324.

- [19] E. Vanden-Eijnden, *Numerical techniques for multiscale dynamical systems with stochastic effects*, Commun. Math. Sci. **1** (2003), 385–391.
- [20] ———, *On HMM-like integrators and projective integration methods for systems with multiple time scales*, Commun. Math. Sci. **5** (2007), 495–505.
- [21] Tang-Qing Yu, Jianfeng Lu, Cameron F. Abrams, and Eric Vanden-Eijnden, *Multiscale implementation of infinite-swap replica exchange molecular dynamics*, Proc. Natl. Acad. Sci. USA **113** (2016), 11744–11749.
- [22] Cheng Zhang and Jianpeng Ma, *Comparison of sampling efficiency between simulated tempering and replica exchange*, J. Chem. Phys. **129** (2008), no. 2008.

DEPARTMENT OF MATHEMATICS, DEPARTMENT OF PHYSICS, AND DEPARTMENT OF CHEMISTRY, DUKE UNIVERSITY,
BOX 90320, DURHAM NC 27708, USA

E-mail address: jianfeng@math.duke.edu

COURANT INSTITUTE OF MATHEMATICAL SCIENCES, NEW YORK UNIVERSITY, 251 MERCER STREET, NEW YORK, NY
10012, USA

E-mail address: eve2@cims.nyu.edu