A GENERALIZED PARALLEL REPLICA DYNAMICS

ANDREW BINDER, TONY LELIÈVRE, AND GIDEON SIMPSON

ABSTRACT. Metastability is a common obstacle to performing long molecular dynamics simulations. Many numerical methods have been proposed to overcome it. One method is parallel replica dynamics, which relies on the rapid convergence of the underlying stochastic process to a quasi-stationary distribution. Two requirements for applying parallel replica dynamics are knowledge of the time scale on which the process converges to the quasi-stationary distribution and a mechanism for generating samples from this distribution. By combining a Fleming-Viot particle system with convergence diagnostics to simultaneously identify when the process converges while also generating samples, we can address both points. This variation on the algorithm is illustrated with various numerical examples, including those with entropic barriers and the 2D Lennard-Jones cluster of seven atoms.

1. INTRODUCTION

An outstanding obstacle for many problems modeled by *in situ* molecular dynamics (MD) is the vast separation between the characteristic time for atomic vibrations (10^{-15} s), and the characteristic time for macroscopic phenomena ($10^{-9} - 10^{-3}$ s). At the heart of this scale separation is the presence of *metastable regions* in the configuration space of the problem. Examples of metastable configurations include the defect arrangement in a crystal or the conformation of a protein. Such metastability may be due to either the energetic barriers of a potential energy driving the problem or to the entropic barriers arising from steric constraints. In the first case (energetic barriers), metastability is due to the system needing to pass through a saddle point which is higher in energy than the local minima to get from one metastable region to another. In the second case (entropic barriers), metastability is due to the system having to find a way through a narrow (but not necessarily high energy) corridor to go from one large region to another (see Section 5.3 below for an example with entropic barriers)

Motivated by the challenge of this time-scale separation, A.F. Voter proposed several methods to conquer metastability in the 1990s: Parallel Replica Dynamics (ParRep), Temperature Accelerated Dynamics (TAD) and Hyperdynamics (Hyper), [31, 33, 37–40]. These methods were derived using Transition State Theory and intuition developed from kinetic Monte Carlo models, as the latter describes the hopping dynamics between metastable regions. Indeed, the aim of all these algorithms is to efficiently generate a realization of the discrete-valued jump process amongst metastable regions. The main idea is that the details of the dynamics within each metastable region are not essential to our physical understanding. Rather, the goal should be to get the correct statistics of the so-called *state-to-state dynamics*, corresponding to jumps amongst the metastable regions. This is nontrivial in general for two reasons: (i) the original dynamics projected onto the state-to-state dynamics are not Markovian; (ii) the parameters (transition rates) of the underlying state-to-state dynamics are unknown.

In recent mathematical studies of these approaches, it has been shown that these three algorithms take advantage of *quasi-stationary distributions* (QSDs) associated with the metastable states, see [1, 24, 25, 32]. Crudely, the QSD corresponds to the distribution of the end points of trajectories conditioned on persisting in the region of interest for a very long time. This mathematical formalization clarifies the fundamental assumptions under which the algorithms will be accurate and also broadens their applicability. Indeed, one aim of this paper is to propose a modification of the original ParRep algorithm that will allow for states to be defined by generic partitions of configuration space. This bears some resemblance to *milestoning*, which also allows more general partitions of configuration space, [22, 36]. Therefore, we will not refer to “basins of attractions” or “metastable regions”, but rather simply to “states”. The only requirement is that these

Date: January 6, 2015.

states define a partition of the configuration space. The boundary at the interface between two states is called the dividing surface.

Briefly (this is detailed in Section 2.3 below), ParRep works by first allowing a single reference trajectory to explore a state. If the trajectory survives for sufficiently long, its end point will agree, in law, with the aforementioned QSD. One thus introduces a *decorrelation time*, denoted t_{corr} , as the time at which the law of the reference process will have converged to the QSD. Provided the reference process survives in the state up till t_{corr} , it is replaced by an ensemble of N independent and identically distributed replicas, each with an initial condition drawn from the QSD. The first replica to escape is then followed into the next state. As the replicas evolve independently and only a first escape is desired, they are readily simulated in parallel, providing as much as a factor of N speedup of the exit event. Thus, there are two practical challenges to implementing ParRep:

- Identifying a t_{corr} at which the law of the reference process is close to the QSD.
- Generating samples from the QSD from which to start the replicas.

In the original algorithm, t_{corr} is *a priori* chosen by the user, as the states are defined so that an approximation of the time required to get “local equilibration within the state” is available. Such a value can be estimated in the case of energetic barriers at sufficiently low temperature using harmonic transition state theory. But in the case of entropic barriers, there is, in general, no simple estimate; see, however, [23] for an example where ParRep was applied to an entropic barrier as an estimate of t_{corr} was available. In the original algorithm, the sampling of the QSD is done using a rejection algorithm, which will be inefficient if the state does not correspond to a metastable region for the original dynamics. Indeed, this will degenerate in the long time limit, as the trajectories will always exit. In this work, we propose an algorithm addressing both points, based on two ingredients:

- The use of a branching and interacting particle system called the Fleming-Viot particle process to simulate the law of the process conditioned on persisting in a state, and to sample the QSD in the longtime limit.
- The use of Gelman-Rubin statistics in order to identify the correlation time, namely the convergence time to a stationary state for the Fleming-Viot particle process.

As we state below, this modified version of ParRep, presented below in Section 3, relies on assumptions (see (A1) and (A2) below) which would require more involved analysis to fully justify. We demonstrate below in a collection of numerical experiments that this modified algorithm gives results consistent with direct simulations. We observe speedup factors up to ten in our test problems where $N = 100$ replicas were used. The aim of this paper is to present new algorithmic developments, and not to explore the mathematical foundations underpinning these ideas.

Though we focus on the ParRep algorithm, since it is the most natural setting for introducing the Fleming-Viot particle process, identifying the convergence to the QSD is also relevant to other problems, including the two other accelerated dynamics algorithms: Hyper and TAD. See [1, 25] for the relevant discussions.

Our paper is organized as follows. In Section 2, we introduce the dynamics of interest, review some properties of the QSD and recall the original ParRep algorithm. In Section 3, we then present the Fleming-Viot particle process and the Gelman-Rubin statistics, which are needed to build the modified ParRep algorithm we propose. Finally, in Section 4 we show the effectiveness and caveats of convergence diagnostics, before exploring the efficiency and accuracy of the modified ParRep algorithm on various test cases in Section 5.

1.1. Acknowledgments. A.B. was supported by a US Department of Defense NDSEG fellowship. G.S. was supported in part by the US Department of Energy Award DE-SC0002085 and the US National Science Foundation PIRE Grant OISE-0967140. T.L. acknowledges funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) grant agreement no. 614492. The authors would also like to thank C. Le Bris, M. Luskin, D. Perez, and A.F. Voter for comments and suggestions throughout the development of this work. The authors would also like to thank the referees for their helpful remarks.

2. THE ORIGINAL PARREP AND QUASI-STATIONARY DISTRIBUTIONS

2.1. Overdamped Langevin Dynamics. We consider the case of the overdamped Langevin equation

$$(2.1) \quad dX_t = -\nabla V(X_t)dt + \sqrt{2\beta^{-1}}dW_t.$$

Here, the stochastic process $(X_t)_{t \geq 0}$ takes values in $\mathbb{R}^d$, β is the inverse temperature, $V(x)$ is the driving potential and W_t a standard d -dimensional Brownian motion. In all that follows, we focus, for simplicity, on (2.1). However, the algorithm we propose equally applies to the phase-space Langevin dynamics which are also of interest. As mentioned in the introduction, for typical potentials, the stochastic process $(X_t)_{t \geq 0}$ satisfying (2.1) is metastable. Much of its trajectory is confined to particular regions of $\mathbb{R}^d$, occasionally hopping amongst them.

Let $\Omega \subset \mathbb{R}^d$ denote the region of interest (namely the state), and define

$$(2.2) \quad T = \inf \{t \geq 0 \mid X_t \notin \Omega\}$$

to be the *first exit time* from Ω , where $X_0 = x \in \Omega$. The point on the boundary, $X_T \in \partial\Omega$, is the *first hitting point*. The aim of accelerated dynamics algorithms (and ParRep in particular) is to efficiently sample (T, X_T) from the exit distribution.

2.2. Quasi-stationary Distributions. In order to present the original ParRep algorithm, it is helpful to be familiar with quasi-stationary distributions (QSD). For more details about quasi-stationary distributions, we refer the reader to, for example, [7–10, 26, 27, 34]. Reference [24] gives self-contained proofs of the results below.

Consider a smooth bounded open set $\Omega \subset \mathbb{R}^d$, that corresponds to a state. By definition, the quasi-stationary distribution ν , associated with the dynamics (2.1) and the state Ω , is the probability distribution, with support on Ω , satisfying, for all (measurable) $A \subset \Omega$ and $t \geq 0$,

$$(2.3) \quad \nu(A) = \frac{\int \mathbb{P}^x [X_t \in A, T > t] \nu(dx)}{\int \mathbb{P}^x [T > t] \nu(dx)} = \frac{\mathbb{P}^\nu [X_t \in A, T > t]}{\mathbb{P}^\nu [T > t]} = \mathbb{P}^\nu [X_t \in A \mid T > t].$$

Here and in the following, we indicate by a superscript the initial condition for the stochastic process: $\mathbb{P}^x$ indicates that $X_0 = x$ and $\mathbb{P}^\nu$ indicates that X_0 is distributed according to ν . In our setting, it can be shown that ν exists and is unique.

The QSD enjoys three properties. First, it is related to an elliptic eigenvalue problem. Let L be the infinitesimal generator of (2.1), defined by, for any smooth function $v : \mathbb{R}^d \rightarrow \mathbb{R}$,

$$(2.4) \quad Lv = -\nabla V \cdot \nabla v + \beta^{-1} \Delta v.$$

The operator L is related to the stochastic process through the following well-known result: if the function $u : \mathbb{R}^+ \times \Omega \rightarrow \mathbb{R}$ satisfies the Kolmogorov equation:

$$(2.5) \quad \begin{cases} \partial_t u = Lu = -\nabla V \cdot \nabla u + \beta^{-1} \Delta u, & \text{for } t > 0, x \in \Omega, \\ u(t, x) = f(x) & \text{for } t > 0, x \in \partial\Omega, \\ u(0, x) = u_0(x) & \text{for } x \in \Omega, \end{cases}$$

then u admits the probabilistic representation formula (Feynman-Kac relation):

$$(2.6) \quad u(t, x) = \mathbb{E}^x [u_0(X_t) 1_{T > t}] + \mathbb{E}^x [f(X_T) 1_{T \leq t}].$$

Recall that T defined by (2.2), is the first exit time of X_t from Ω . Provided Ω is bounded with sufficiently smooth boundary and V is smooth, L has an infinite set of Dirichlet eigenvalues and orthonormal eigenfunctions

$$(2.7) \quad L\varphi_j = -\lambda_j \varphi_j, \quad \varphi_j|_{\partial\Omega} = 0, \quad j = 1, 2, \dots$$

Here, the eigenfunctions are orthonormal with respect to the invariant measure restricted to Ω :

$$\frac{\int_{\Omega} \varphi_j(x) \varphi_k(x) \exp(-\beta V(x)) dx}{\int_{\Omega} \exp(-\beta V(x)) dx} = \delta_{jk}, \quad \forall j, k = 1, 2, \dots$$

The eigenfunction associated with the lowest eigenvalue is signed, and, taking it to be positive, the QSD is

$$(2.8) \quad \nu(dx) = \frac{\varphi_1(x) e^{-\beta V(x)} dx}{\int_{\Omega} \varphi_1(x) e^{-\beta V(x)} dx}.$$

While this expression is explicit in terms of φ_1 and V , since the problem is posed in $\mathbb{R}^d$ with d large, it is not practical to sample the QSD directly by first computing φ_1 .

The second property associated to the QSD is that for all $t \geq 0$ and $A \subset \partial\Omega$:

$$(2.9) \quad \begin{aligned} \mathbb{P}^\nu [X_T \in A, T > t] &= \mathbb{P}^\nu [T > t] \mathbb{P}^\nu [X_T \in A] \\ &= (e^{-\lambda_1 t}) \left(\int_A -\frac{1}{\beta \lambda_1} \frac{\nabla \varphi_1 e^{-\beta V} \cdot \mathbf{n}}{\int_\Omega \varphi_1(x) e^{-\beta V(x)} dx} dS \right), \end{aligned}$$

where dS is the surface Lebesgue measure on $\partial\Omega$ and $\mathbf{n}$ the unit outward normal vector to Ω . Thus, the first hitting point and first exit time are independent, and exit times are exponentially distributed. These two properties will be one of the main arguments justifying ParRep. As explained in [24], they are consequences of (2.6) and (2.8).

The third property of the QSD also plays an important role in ParRep. Let us again consider X_t satisfying (2.1) with $X_0 \sim \mu_0$ (μ_0 with support in Ω). Define the law of X_t , conditioned on non-extinction as:

$$(2.10) \quad \mu_t(\bullet) = \frac{\mathbb{P}^{\mu_0} [X_t \in \bullet, T > t]}{\mathbb{P}^{\mu_0} [T > t]} = \mathbb{P}^{\mu_0} [X_t \in \bullet \mid T > t].$$

One can check that

$$(2.11) \quad \mu_t(v_0) = \mathbb{E}^{\mu_0} [v_0(X_t) \mid T > t] = \frac{\int_\Omega v(x, t) \mu_0(dx)}{\int_\Omega \bar{v}(x, t) \mu_0(dx)},$$

where v solves (2.5) with initial condition $v(0, x) = v_0(x)$ and boundary conditions $v|_{\partial\Omega} = 0$ while $\bar{v}$ solves (2.5) with initial condition $\bar{v}(0, x) = 1$ and boundary conditions $\bar{v}|_{\partial\Omega} = 0$. Through eigenfunction expansions of the form

$$(2.12) \quad v(t, x) = \sum_{k=1}^{\infty} e^{-\lambda_k t} \varphi_k(x) \int v_0(y) \varphi_k(y) \frac{e^{-\beta V(y)}}{\int e^{-\beta V}} dy$$

we obtain: for t sufficiently large, the *total variation norm* can be bounded as

$$(2.13) \quad \|\mu_t - \nu\|_{\text{TV}} \equiv \sup_{\|f\| \leq 1} \left| \int f(x) \mu_t(dx) - \int f(x) \nu(dx) \right| \leq C(\mu_0) e^{-(\lambda_2 - \lambda_1)t}.$$

In the above expression, $\|f\| = \|f\|_{L^\infty(\Omega)}$. This shows that if the process remains in Ω for a sufficiently large amount of time (typically of the order of $1/(\lambda_2 - \lambda_1)$), then its law at time t is close to the QSD ν .

Since we are interested in ensuring that the state to state dynamics are accurate, we observe that (2.13) implies agreement of the exit distribution of (T, X_T) in total variation norm between processes initially distributed according to μ_t and ν . Indeed, starting from the probability measure μ_t , given any measurable $g : \mathbb{R}^+ \times \partial\Omega \rightarrow \mathbb{R}$, we see that exit distribution observables can be reformulated as observables on Ω against μ_t :

$$(2.14) \quad \mathbb{E}^{\mu_t} [g(T, X_T)] = \int \underbrace{\mathbb{E}^x [g(T, X_T)]}_{\equiv G(x)} \mu_t(dx) = \int G(x) \mu_t(dx).$$

Therefore,

$$(2.15) \quad \sup_{\|g\| \leq 1} |\mathbb{E}^{\mu_t} [g(T, X_T)] - \mathbb{E}^\nu [g(T, X_T)]| \leq \|\mu_t - \nu\|_{\text{TV}} \leq C(\mu_0) e^{-(\lambda_2 - \lambda_1)t},$$

where $\|g\| = \|g\|_{L^\infty(\mathbb{R}^+ \times \partial\Omega)}$. Thus, convergence of μ_t to ν implies agreement of the exit distributions, starting from μ_t and ν .

We are now in position to introduce the original ParRep algorithm.

2.3. Parallel Replica Dynamics. The goal of the ParRep algorithm is to rapidly generate a physically consistent first hitting point and first exit time for each visited state. Information about where, precisely, the trajectory is within each state will be sacrificed to more rapidly obtain this information.

In the following, we assume that we are given a partition of the configuration space $\mathbb{R}^d$ into states, and we denote by Ω one generic element of this partition. We also assume that we have N CPUs available for parallel computation.

The original ParRep algorithm [39] is implemented in three steps, repeated as the process moves from one state to another. It requires the specification, *a priori*, of two times to equilibrate to each state, t_{corr} and t_{phase} . Let us consider a single reference process, X_t^{ref} , with $X_0^{\text{ref}} \sim \mu_0$ evolving under (2.1), and set the simulation clock, corresponding to the physical time, to zero, $t_{\text{sim}} = 0$. The simulation clock t_{sim} will be updated during the algorithm.

Decorrelation Step: Let Ω denote the state in which $X_{t_{\text{sim}}}^{\text{ref}}$ currently resides. If the trajectory has not left Ω after running for t_{corr} amount of time, the algorithm proceeds to the dephasing step, the simulation clock being advanced as

$$t_{\text{sim}} \mapsto t_{\text{sim}} + t_{\text{corr}}.$$

Otherwise a new decorrelation starts from the new state, the simulation clock being advanced as

$$t_{\text{sim}} \mapsto t_{\text{sim}} + T^{\text{ref}}, \quad T^{\text{ref}} = \inf \{t \geq 0 \mid X_{t_{\text{sim}}+t}^{\text{ref}} \notin \Omega\}.$$

Dephasing Step: In this step, N independent and identically distributed samples of the QSD of Ω are generated. These samples will be distributed over the N CPUs and will be used as initial conditions in the subsequent parallel step. During the dephasing step, the counter t_{sim} is not advanced.

In the original ParRep, the sampling of the QSD is accomplished by a rejection algorithm. For $k = 1, \dots, N$, generate a starting point $\tilde{X}_0^k \sim \eta_0$ and integrate it under (2.1) until either time t_{phase} or $\tilde{X}_t^k$ leaves Ω . Here, η_0 denotes any distribution with support in Ω (for example, a Dirac mass at the end point of the reference trajectory, after the decorrelation step). If $\tilde{X}_t^k$ has not exited before time t_{phase} , set the k -th replica's starting point $X_0^k = \tilde{X}_{t_{\text{phase}}}^k$ and advance $k \mapsto k + 1$. Otherwise, reject the sample, and start a new trajectory with $\tilde{X}_0^k \sim \eta_0$. Since these samples are independent, they can be generated in parallel.

Parallel Step: Let the N samples obtained after the dephasing step evolve under (2.1) in parallel (one on each CPU), driven by independent Brownian motions, until one escapes from Ω . Let us denote

$$k_\star = \text{argmin}_k T^k$$

the index of the first replica which exits Ω . During the time interval $[t_{\text{sim}}, t_{\text{sim}} + NT^{k_\star}]$, the reference process is defined as trapped in Ω . Accordingly, the simulation clock is advanced as

$$(2.16) \quad t_{\text{sim}} \mapsto t_{\text{sim}} + NT^{k_\star},$$

The first replica to escape becomes the new reference process. A new decorrelation step now starts, applied to the new reference process with starting point $X_{t_{\text{sim}}}^{\text{ref}} = X_{T^{k_\star}}^{k_\star}$.

The justifications underlying the ParRep algorithm are the following. Using the third property (2.13) of the QSD, it is clear that if t_{corr} is chosen sufficiently large, then, at the end of the decorrelation step, the reference process is such that $X_{t_{\text{sim}}}^{\text{ref}}$ is approximately distributed according to the QSD. This same property explains why the rejection algorithm used in the dephasing step yields (approximately) N i.i.d. samples distributed according to the QSD, at least if t_{phase} is chosen sufficiently large. Finally the second property (2.9) justifies the parallel step; since the N replicas are i.i.d. and drawn from the QSD, they have exponentially distributed exit times and thus $\mathbb{P}^\nu[\min\{T^1, \dots, T^N\} > t] = \mathbb{P}^\nu[T^1 > Nt]$. Moreover, by the independence property in (2.9), the exit points $X_{T^{k_\star}}^{k_\star}$ and $X_{T^1}^1$ have the same distribution.

Notice that it is the magnification of the first exit time by a factor of N in the parallel step that yields the speedup in terms of wall clock time. If the partition of the configuration space is chosen in such a way that, most of the time, the stochastic process exits from the state before having reached the QSD (namely before t_{corr}), there is no speedup. In this case, ParRep essentially consists in following the reference process. There is no error, but no gain in performance, and computational resources are wasted. To observe a significant speedup, the partition of the configuration space should be such that most of the defined states

are metastable, in the sense that the typical exit time from the state is much larger than the time required to approximate the QSD.

Of course, the QSD is only sampled approximately, and this introduces error in ParRep. The time t_{corr} (resp. t_{phase}) must be sufficiently large such that $\mathbb{P}^{\mu_0} [X_{t_{\text{corr}}}^{\text{ref}} \in \bullet \mid T^{\text{ref}} > t_{\text{corr}}] \approx \nu$ (resp. such that $\mathbb{P}^{\eta_0} [X_{t_{\text{phase}}} \in \bullet \mid T > t_{\text{phase}}] \approx \nu$). The mismatch between the distributions at times t_{corr} and t_{phase} , directly, and independently, contribute to the overall error of ParRep; see [32]. Also note that these parameters are state dependent. In view of (2.13), one may think that a good way to choose t_{corr} and t_{phase} is to consider a multiple of $1/(\lambda_2 - \lambda_1)$. This is unsatisfactory for two reasons. First, it is difficult to numerically compute the spectral gap $\lambda_2 - \lambda_1$ because of the high-dimensionality of the associated elliptic problem. Second, the pre-factors $C(\mu_0)$ and $C(\eta_0)$ in (2.13) are also difficult to evaluate, and could be large.

In view of the preceding discussion, ParRep can be applied to a wide variety of problems and for any predefined partition of the configuration space into states provided one has:

- A way to construct an adequate t_{corr} (or more precisely to assess the convergence of μ_t to ν) for each state;
- A way to sample the QSD of each state.

The aim of the next section is to provide a modified ParRep algorithm to deal with these two difficulties.

3. THE MODIFIED PARREP ALGORITHM

We propose to use a branching and interacting particle system (the Fleming-Viot particle process) together with convergence diagnostics to simultaneously and dynamically determine adequate values t_{corr} and t_{phase} , while also generating an ensemble of N samples from a distribution close to that of the QSD.

3.1. The Fleming-Viot Particle Process. In this section, we introduce a branching and interacting particle system which will be one of the ingredients of the modified ParRep algorithm. This process is sometimes called the Fleming-Viot particle process [15].

Let us specify the Fleming-Viot particle process; see also the illustration in Figure 1. Let us consider i.i.d. initial conditions X_0^k ($k \in \{1, \dots, N\}$) distributed according to μ_0 , a probability distribution with support in Ω . The process is as follows:

- (1) Integrate N realizations of (2.1) with independent Brownian motions until one of them, say X_t^1 , exits;
- (2) Kill the process that exits;
- (3) With uniform probability $1/(N-1)$, randomly choose one of the survivors, $X_t^2, \dots, X_t^N$, say X_t^2 ;
- (4) Branch X_t^2 , with one copy persisting as X_t^2 , and the other becoming the new X_t^1 (and thus evolving in the future independently from X_t^2).

We denote this branching and interacting particle process by $\mathbf{X}_{t,N} = (X_t^1, \dots, X_t^N)$, and define the associated empirical distribution

$$(3.1) \quad \mu_{t,N} \equiv \frac{1}{N} \sum_{k=1}^N \delta_{X_t^k}.$$

The Fleming-Viot particle process can be implemented in parallel, with each replica X_t^k evolving on distinct CPUs. The communication cost (due to the branching step) will be small, provided the state under consideration is such that the exit events are relatively rare; *i.e.*, it is metastable.

The Fleming-Viot particle process has been studied for a variety of underlying stochastic processes; see, for example, [13, 15, 29] and the references therein. In [29], the authors prove that for a problem in dimension one, the following relation holds: for any $A \subset \Omega$,

$$(3.2) \quad \lim_{N \rightarrow \infty} \mu_{t,N}(A) = \mu_t(A).$$

From (2.13) and (3.2), we infer that $\lim_{t \rightarrow \infty} \lim_{N \rightarrow \infty} \mu_{t,N}(A) = \nu(A)$. This result is anticipated to hold for general dynamics, including (2.1).

The property (3.2) of the Fleming-Viot particle process is instrumental in our modified ParRep algorithm. It will be used in two ways:

FIGURE 1. The branching & interacting particle system used to sample the QSD in the case $N = 3$ at three times: $t_1 < t_2 < t_3$. The trajectories run independently until one exits, as in (b). The process that has reached the boundary is killed, as in (c). Then a survivor is instantaneously branched to maintain a constant number of trajectories, as in (d).

- Since an ensemble of realizations distributed according to μ_t can be generated using the Fleming-Viot particle process, we will assess convergence of μ_t to the stationary distribution ν by applying convergence diagnostics to the ensemble $(X_t^1, \dots, X_t^N)$. This will give a practical way to estimate the time required for convergence to the QSD in the decorrelation step, by *simultaneously* running the decorrelation step (on the reference process) and the dephasing step using a Fleming-Viot particle process on N other samples, starting from the same initial condition as the reference process: the decorrelation time is estimated as the time required for convergence to a stationary state for the Fleming-Viot particle process.
- In addition, the Fleming-Viot particle process introduced in the procedure described above gives a simple way to sample the QSD. We use the replicas generated by the modified dephasing step at the time of stationarity.

The modified ParRep algorithm will be based on the two following assumptions on the Fleming-Viot particle process. While we do not make this rigorous, we believe it could be treated in specific cases, and our numerical experiments show consistency with direct numerical simulation.

Assumption (A1): For sufficiently large N , $\mu_{t,N}$ is a good approximation of μ_t ;

Assumption (A2): The realizations generated by the Fleming-Viot particle process are sufficiently weakly correlated so as to allow the use of both the convergence diagnostics presented below and the temporal acceleration expression (2.16), which both assume independence.

As already mentioned above, the first assumption is likely satisfied in our setting, though we were not able to find precisely this result in the literature. See [29] for such a result in a related problem.

The second assumption is more questionable. We make two comments on this. First, our numerical experiments show that the modified ParRep algorithm (which is partly based on (A2)) indeed yields correct results compared to direct numerical simulation; thus, the assumption is not grossly wrong, at least in these settings. Second, the correlations introduced by the Fleming-Viot particle process are most likely a concern for problems where the state is only weakly metastable. In truly metastable states, exits will be infrequent,

so the correlations amongst the replicas will be weak. For states which are not metastable, the reference process will likely exit before stationarity can be reached, rendering the concern moot. It is therefore in problems between the two cases that practitioners may have some cause for concern.

There are several ways to ameliorate reservations about the second assumption. First, it is known that for such branching and interacting particle systems, a propagation of chaos result holds, [35]. This means that if we run Fleming-Viot with $M \gg N$ processes, then, as M tends to infinity, with N fixed, the first N trajectories in the process become i.i.d. Second, one could run a separate Fleming-Viot particle process for each of the N replicas, retaining only the first trajectory of each of the N i.i.d. Fleming-Viot particle processes. Finally, the Fleming-Viot particle process, together with convergence diagnostics, could be used to identify appropriate values of t_{phase} and t_{corr} , as explained below in Section 3.2. This value t_{phase} could then be used in the original rejection sampling algorithm, run in tandem with the Fleming-Viot particle process. This would provide independent samples from the QSD.

3.2. Convergence Diagnostics. While the Fleming-Viot particle process gives us a process that will converge to the QSD, there is still the question of how long it must be run in order for $\mu_{t,N}$ (and thus μ_t according to (A1)) to be close to equilibrium. This is a ubiquitous problem in applied probability and stochastic simulation: when sampling a distribution via Markov Chain Monte Carlo, how many iterations are sufficient to be close to the stationary distribution? For a discussion on this general issue, see for example, [4, 6, 11]. We propose to use convergence diagnostics to test for the stationarity of $\mu_{t,N}$. When a user specified convergence criterion is satisfied, $\mu_{t,N}$ is declared to be at its stationary value and the time at which this occurs is taken to be t_{corr} (and t_{phase}).

We have found Gelman-Rubin statistics to be effective for this purpose, [5, 6, 17]. In the simplest form, such statistics compute the ratio of two estimates of the asymptotic variance of a given observable. Since numerator and denominator estimate the same quantity, the ratio converges to one.

The statistic can be defined as follows. Let $\mathcal{O} : \Omega \rightarrow \mathbb{R}$ be some observable, and let

$$(3.3) \quad \bar{\mathcal{O}}_t^k \equiv t^{-1} \int_0^t \mathcal{O}(X_s^k) ds, \quad \bar{\mathcal{O}}_t \equiv \frac{1}{N} \sum_{k=1}^N \bar{\mathcal{O}}_t^k = \frac{1}{N} \sum_{k=1}^N t^{-1} \int_0^t \mathcal{O}(X_s^k) ds,$$

be the average of an observable along each trajectory and the average of the observable along all trajectories. Then the statistic of interest for observable $\mathcal{O}$ is

$$(3.4) \quad \hat{R}_t(\mathcal{O}) = \frac{\frac{1}{N} \sum_{k=1}^N t^{-1} \int_0^t (\mathcal{O}(X_s^k) - \bar{\mathcal{O}}_t)^2 ds}{\frac{1}{N} \sum_{k=1}^N t^{-1} \int_0^t (\mathcal{O}(X_s^k) - \bar{\mathcal{O}}_t^k)^2 ds}.$$

Notice that $\hat{R}_t(\mathcal{O}) \geq 1$, and as all the trajectories explore Ω , $\hat{R}_t(\mathcal{O})$ converges to one as t goes to infinity. We also observe that since (3.4) is a trajectory average, both its bias and variance will be $O(t^{-1})$, [3].

These statistics were not developed with the intention of handling branching interacting particle systems. The authors had in mind that the N trajectories would be completely independent, which is not the case for the Fleming-Viot particle process. This is one reason why we introduced Assumption (A2) above. However, we will demonstrate in the numerical experiments below that this convergence diagnostic indeed provides meaningful results for the Fleming-Viot particle process (see in particular Section 4.1).

Here, we caution the reader that all convergence diagnostics are susceptible to the phenomena of *pseudo-convergence*, which occurs when a particular observable or statistic appears to have reached a limiting value, and yet the empirical distribution of interest remains far from stationarity; see [4]. This can occur, for instance, if the state has an internal barrier that obstructs the process from migrating from one mode to the other. A computational example of this is given below, in Section 4.2.

There is still the question of what observables to use in computing the statistics. Candidates include:

- Moments of the coordinates;
- Energy $V(x)$;
- Distances to reference points in configuration space.

Assuming they are not costly to evaluate, as many such observables should be used; see the example in Section 4.2.

Our test for stationarity is as follows. Given some collection of observables $\{\mathcal{O}_j : \mathbb{R}^d \rightarrow \mathbb{R}\}_{j \in \{1, \dots, J\}}$, their associated statistics $\{\hat{R}_t(\mathcal{O}_j)\}_{j \in \{1, \dots, J\}}$, and a tolerance $\text{TOL} > 0$, we take as a stationarity criterion:

$$(3.5) \quad \forall j \in \{1, \dots, J\}, \hat{R}_t(\mathcal{O}_j) < 1 + \text{TOL}.$$

In other words, the dephasing and decorrelation times are set as

$$(3.6) \quad t_{\text{phase}} = t_{\text{corr}} = \inf \left\{ t \geq 0 \mid \hat{R}_t(\mathcal{O}_j) < 1 + \text{TOL}, \forall j \right\}.$$

3.3. The Modified ParRep Algorithm. We now have the ingredients needed to present the modified ParRep algorithm (which should be compared to the original ParRep given in Section 2.3). Let us consider a single reference process, X_t^{ref} , with $X_0^{\text{ref}} \sim \mu_0$ evolving under (2.1), and let us set $t_{\text{sim}} = 0$.

Decorrelation and Dephasing Step: Denote by Ω the state in which $X_{t_{\text{sim}}}^{\text{ref}}$ lives. The decorrelation and dephasing steps are carried out at the same time, in parallel: the reference process X_t^{ref} and the Fleming-Viot particle process $\mathbf{X}_{t,N}$ begin at the same time from the same point in Ω . Convergence diagnostics are assessed on $\mu_{t,N}$, and when the stationarity criterion (3.5) is satisfied, in the case that the reference process has never left Ω , both decorrelation and dephasing steps are terminated, and one proceeds to the Parallel Step, after advancing the simulation clock as

$$t_{\text{sim}} \mapsto t_{\text{sim}} + t_{\text{corr}},$$

t_{corr} being defined by (3.6). In this case, the decorrelation/dephasing step is said to be successful.

If at any time before reaching stationarity the reference process leaves Ω , the Fleming-Viot particle process terminates, $\mathbf{X}_{t,N}$ is discarded, the simulation clock is advanced as

$$t_{\text{sim}} \mapsto t_{\text{sim}} + T^{\text{ref}}, \quad T^{\text{ref}} = \inf \left\{ t \geq 0 \mid X_{t_{\text{sim}}+t}^{\text{ref}} \notin \Omega \right\}.$$

The process X_t^{ref} then proceeds into the new state, where a new decorrelation/dephasing step starts. In this case, the decorrelation/dephasing step is said to be unsuccessful.

Parallel Step: The parallel step is similar to the original parallel step. Consider the N positions of $\mathbf{X}_{t_{\text{corr}},N}$ obtained at the end of the dephasing step as initial conditions. These are then evolved in parallel following (2.1), driven by independent Brownian motions, until one replica, with index k_* , escapes from Ω . The simulation clock is then advanced according to (2.16),

$$t_{\text{sim}} \mapsto t_{\text{sim}} + NT^{k_*}$$

The replica which first exits becomes the new reference process, and a new decorrelation/dephasing step starts.

Our modified ParRep algorithm differs from the original in two essential ways, which merit comment. First, as already mentioned, the main attraction of the modified ParRep algorithm is that the convergence times t_{corr} and t_{phase} do not need to be chosen *a priori*, but are instead computed on the fly using physically informed observables. This is why the Fleming-Viot particle process (combined with a convergence diagnostic) is essential to our algorithm. This is the major improvement of modified ParRep, which broadens the applicability of the algorithm to a general partition of the configuration space, as *a priori* estimates for t_{corr} and t_{phase} are, in general, unlikely to be available.

Second, in the modified ParRep, the Fleming-Viot particle process is also used to sample the QSD (the dephasing step). Part of the appeal of the Fleming-Viot particle process is that it is robust enough to sample states which are not strongly metastable (the typical exit time is not dramatically larger than the time to converge to the QSD). For highly metastable regions, rejection sampling and the Fleming-Viot particle process will yield similar results, but in more general scenarios, this will not be the case. Indeed, while the Fleming-Viot particle process is well defined in the limit of time tending to infinity, rejection sampling will degenerate since all replicas will eventually exit. This is critical when applying convergence diagnostics, as the time to stop is found dynamically by condition (3.5). Therefore, if one requires N replicas distributed according to the QSD before proceeding to the parallel step, the Fleming-Viot particle process may, therefore, be more efficient than rejection sampling.

Our discussion of the Fleming-Viot particle process, as a QSD sampling strategy, would not be complete without two comments. First, some implementations of the rejection algorithm in ParRep do not require all N replicas to be dephased before proceeding to the parallel step. Replicas are run asynchronously, and

FIGURE 2. The Gelman-Rubin statistics as a function of time for the observables (4.1) for the Fleming-Viot particle process. The potential is $V(x) = -2\cos(\pi x)$, the state is $(-1, 1)$ and the number of replicas is $N = 10^4$.

as soon as one has reached the *a priori* value of t_{phase} , it is immediately promoted to the parallel step. This is how the algorithm is implemented in, for instance, [23], and it should be taken into account when comparing QSD sampling algorithms. Second, as already mentioned above, the robustness of the Fleming-Viot process comes at the cost of generating correlated samples, which are not easy to control. Two methods for overcoming correlations amongst the replicas are proposed at the end of Section 3.1.

We will illustrate this modified ParRep algorithm on various numerical examples in Section 5, after a discussion of the performance and limitations of the convergence diagnostics in Section 4.

4. ILLUSTRATION OF CONVERGENCE DIAGNOSTICS

In this section we present two numerical examples showing the subtleties of the Gelman-Rubin statistics and the broader problems raised by stationarity testing. These are “offline” in the sense that they are not used as part of the ParRep algorithm here. They show that the Gelman-Rubin statistics are consistent with our expectations, but also susceptible to pseudo-convergence.

In both of these examples, $N = 10^4$ replicas are used, and the stochastic differential equation (2.1) is discretized using Euler-Maruyama with a time step $\Delta t = 10^{-4}$. The Mersenne Twister algorithm is used as a pseudo-random number generator in these two examples, as implemented in [16].

4.1. Periodic Potential in 1D. For the first example, consider the 1D periodic potential $V(x) = -2\cos(\pi x)$ at $\beta = 1$ and the state $\Omega = (-1, 1)$. The initial condition is $X_0 = 0.99$. Running the Fleming-Viot particle process algorithm, we examine the Gelman-Rubin statistics for the observables:

$$(4.1) \quad x, \quad V(x), \quad |x - x_{\text{ref}}|$$

where x_{ref} is the local minima of the current basin; $x_{\text{ref}} = 0$ in this case.

The Gelman-Rubin statistics, as a function of time, appear for several observables in Figure 2. As expected, the statistics tend to one as time goes to infinity. They also display the aforementioned $O(t^{-1})$ bias. We also examine the empirical distributions in Figure 3, compared to the density of the QSD (which can be precisely computed by solving an eigenvalue problem using the formula (2.8) in this simple 1D situation). By $t = 1.0$, the qualitative features of the distribution are good, and the Gelman-Rubin statistics are less than 1.1 for the three observables. We will see that that $\text{TOL} = 0.1$ gives reasonable results in many cases.

From this simple experiment, we first observe that the Gelman-Rubin statistics seems to yield sensible results for assessing the convergence of the Fleming-Viot particle process. Further inspection of Figures 2 and 3 suggests (3.5) may be conservative: if the tolerance is too small, the convergence time may be overestimated compared to what can be observed on the empirical distribution.

FIGURE 3. Empirical distributions for the Fleming-Viot particle process on the 1D periodic potential $V(x) = -2\cos(\pi x)$ over the domain $(-1, 1)$ with $N = 10^4$ replicas, see also Figure 2. The red curve is the reference density of the QSD.

4.2. Double Well Potential in 2D. As a second example, we consider the potential

$$(4.2) \quad V(x, y) = \frac{1}{6} [4(1 - x^2 - y^2)^2 + 2(x^2 - 2)^2 + ((x + y)^2 - 1)^2 + ((x - y)^2 - 1)^2],$$

plotted in Figure 4. Notice that there are two minima, near $(\pm 1, 0)$, along with an internal barrier, centered at the origin, separating them. Thus, there are two channels joining the two minima, with a saddle point in each of these channels. We study this problem at inverse temperature $\beta = 5$. The aim of this example is to illustrate the possibility of pseudo-convergence when using convergence diagnostics, even when it is applied to independent replicas. We thus concentrate on the sampling of the canonical distribution with density $Z^{-1}e^{-\beta V(x, y)}$, using independent realizations, instead of the sampling of the QSD using the Fleming-Viot particle process.

FIGURE 4. Contours of potential (4.2). Note the two minima, near $(\pm 1, 0)$.

FIGURE 5. The Gelman-Rubin statistics as a function of time for the observables (4.3) for $N = 10^4$ independent replicas following the dynamics (2.1). The potential is given by (4.2). Since the replicas begin in the left mode, the internal barrier makes it challenging to adequately sample the x observable.

The observables used in this problem are:

$$(4.3) \quad x, y, V(x, y), \|\mathbf{x}\|_{\ell^2}.$$

Starting our trajectories at $X_0 = (-1, 0)$, the Gelman-Rubin statistics as a function of time appear in Figure 5, and the empirical distributions at two specific times are shown in Figure 6. We make the following remarks.

First, were we to have neglected the x observable, the Gelman-Rubin statistics of the other observables would have fallen below 1.1 by $t = 20$. But as we can see in Figure 6, this is completely inadequate for sampling $Z^{-1}e^{-\beta V(x, y)}$. Thus, if the tolerance is set to 0.1, the convergence criterion may be fulfilled before actual convergence if the observables are poorly chosen. This is a characteristic problem of convergence diagnostics; they are necessary, but not sufficient to assess convergence.

Let us now consider the x observable, which is sensitive to the internal barrier. From Figure 5, we see that it is not monotonic and that, even after running till $t = 2000$, the associated statistic still exceeds 1.1. On the other hand, if we consider the ensemble at $t = 500$, the empirical distribution appears to equally sample both modes. Indeed 48% of the replicas are in the right basin. Despite this, the Gelman-Rubin statistic for x is still relatively large. This is due to the conservative nature of (3.4), already mentioned in Section 4.1. Once the ensemble has an appreciable number of samples in each mode, it will only reach one after all the trajectories have adequately sampled both modes.

FIGURE 6. Empirical distributions from $N = 10^4$ independent replicas following the dynamics (2.1) for the potential (4.2), see also Figure 5.

FIGURE 7. Box plots of t_{phase} for the 1D periodic problem at different values of N with $\text{TOL} = 0.1$. Each distribution includes 10^4 samples. The box encompasses the middle 50% of the data, the median is the line across the box, and the diamond is the mean. The stems extend out 1.5 times the Interquartile Range (the length of the box). Outliers, outside the stems, are shown with circles.

4.3. Distribution of t_{corr} and t_{phase} . In the modified ParRep algorithm, $t_{\text{corr}} = t_{\text{phase}}$, determined by (3.5), is a random variable due to N being finite. This introduces some uncertainty in the decorrelation time; it may be that, for a particular realization, (3.5) is satisfied at a time which is too small for the replicas, or the reference process, to have sampled the QSD. This can be mitigated by using more replicas. As N tends to infinity, the empirical distribution of the Fleming-Viot particle process will converge to the law of the process, conditioned on non-extinction, μ_t , as in (3.2). Hence, $\hat{R}_t(\mathcal{O}_j)$ tends to a deterministic evolution, and t_{phase} will approach a deterministic value, for the given observables and tolerance.

To illustrate this, we explored the 1D periodic problem of Section 4.1 with $N = 10^1, 10^2, 10^3$, and 10^4 replicas at $\text{TOL} = 0.1$. We ran the process until it satisfied (3.5), recording the value of t_{phase} , in each of 10^4 independent realizations, for each of the values of N . The results, shown in Figure 7, show the reduction of variability as N becomes large.

In the rest of our examples, appearing in Section 5, we use $N = 100$ replicas. The typical variations of $t_{\text{corr}} = t_{\text{phase}}$, measured as the ratio of the standard deviation to the mean of the samples, are 10%–20% of the mean. Both the mean and variance are reported in our data tables, below.

4.4. Remarks on Numerical Experiments. The problems presented above, though simple, demonstrate both the effectiveness and the caveats of the convergence diagnostics.

First, the convergence diagnostics can be conservative, which is computationally wasteful. Second, there is the possibility of pseudo-convergence: as is the case with all convergence diagnostics, they cannot guarantee stationarity. Thus, some amount of heuristic familiarity with the underlying problem is essential to obtain reasonable results. One should be careful when choosing the observables for (3.4), selecting degrees of freedom which are associated with the metastable features of the dynamics, as revealed by the x observable in the preceding example. Again, the burden is on the practitioner to be familiar with the system and have some sense of the relevant observables to test for stationarity. However, to set, *a priori*, suitable values of t_{corr} and t_{phase} requires more precise knowledge of the system than is needed to apply convergence diagnostics. In some sense, this is comparable to the relationship between *a priori* and *a posteriori* estimates in other fields of numerical analysis. *a priori* estimates are generally insufficient for assessing convergence, since they require some knowledge of the (unknown) solution, such as its regularity. This is why *a posteriori* estimates have been sought out, to provide error bounds based on computable quantities.

Hence, the reader may wonder why we did not make use of the *a priori* convergence estimate (2.13) to determine the convergence time to the QSD. The reasons are that we lack estimates of both the exponential rate of convergence, $(\lambda_2 - \lambda_1)$, and the prefactor, $C(\mu_0)$. This is a general challenge to the application of many Markov Chain Monte Carlo algorithms: under very weak assumptions, it can be shown that the law of a Markov Chain converges to equilibrium exponentially fast in time (see, for example, [30]). However, such results cannot be used to estimate the convergence, or *burn-in*, time since the error bound involves difficult to estimate quantities. The need to identify such a time has motivated the broad investigation of convergence diagnostics for stochastic algorithms.

5. MODIFIED PARREP EXAMPLES

In this section we present a number of numerical results obtained with the modified ParRep algorithm. Before presenting the examples, we review the numerical methods and the parameters used in the experiments.

5.1. Common Parameters and Methods. For each problem we compare the direct, serial simulation to our proposed algorithm combining ParRep with the Fleming-Viot particle process and convergence diagnostics, with several different values of TOL for the stationarity criterion (3.5). For each case of each problem, we perform 10^5 independent realizations of the experiment to ensure we have adequate data from which to make statistical comparisons. For these ensembles of experiments, which are performed in parallel, SPRNG 2.0, with the linear congruential generator, is used as the pseudo-random number generator, [28]. In all examples, $N = 100$ replicas are used.

The stochastic differential equation (2.1) is discretized using Euler-Maruyama with a time step of $\Delta t = 10^{-4}$ in all cases except for the experiments in Section 5.4, where, due to the computational complexity, $\Delta t = 10^{-3}$. To minimize numerical distortion of the exit times due to the time step discretization, we make use of a correction to time discretized ParRep presented in [2].

Exit distributions obtained from our modified ParRep algorithm and from the direct serial runs are compared using the two sample nonparametric *Kolmogorov-Smirnov test* (K.-S. Test), [18]. Working at the $\alpha = 0.05$ level of significance, this test allows us to determine whether or not we can accept the null hypothesis, which, in this case, is that the samples of exit time distributions arose from the same underlying distribution. Though the test was formulated for continuous distributions, it can also be applied to discrete ones, where it will be conservative with respect to Type 1 statistical errors, [12, 20]. In our tables, we report PASS for instances where the null hypothesis is not rejected (the test cannot distinguish the ParRep results from the direct serial results), and FAIL for cases where the null hypothesis is rejected. We also report the p -values, measuring the probability of observing a test statistic at least as extreme as that which is computed.

In the tables below, we also report the mean dephasing time ($\langle t_{\text{phase}} \rangle$), its variance ($\text{Var}(t_{\text{phase}})$), the average exit time ($\langle T \rangle$), and the percentage of realizations for which the decorrelation/dephasing step were successful (% Dephased). The average and variance of t_{phase} are computed only over realizations for which the reference process decorrelates.

Finally, we report the speedup of the algorithm, which we define as

$$(5.1) \quad \text{Speedup} \equiv \frac{\text{Physical Exit Time}}{\text{Computational Time}}.$$

The computational time is either the exit time of the reference process (if the exit occurs during the decorrelation step), or the sum of the time spent in the decorrelation/dephasing step and in the parallel step (if the exit occurs during the parallel step). Formula (5.1) measures the speedup of the algorithm over direct numerical simulation of the exit event. Let us heuristically consider the mean behavior of (5.1). Let p denote the probability that the system decorrelates, so that $1 - p$ corresponds to the probability that the reference process exits before (3.5) is satisfied. For realizations where the reference process exited before decorrelating, the speedup is unity. For realizations which exit in the parallel step, let $\langle t_{\text{phase}} \rangle$ and $\langle t_{\text{par}} \rangle$ denote the average dephasing and parallel exit times. Roughly, these values relate to the eigenvalues, introduced in (2.7), and the choice of N as

$$(5.2) \quad \langle t_{\text{phase}} \rangle \approx \frac{C_{\text{phase}}}{\lambda_2 - \lambda_1}, \quad \langle t_{\text{par}} \rangle \approx \frac{1}{N\lambda_1}$$

To obtain $\langle t_{\text{par}} \rangle$, we have assumed that the dephased distribution is a good approximation of the QSD. The constant C_{phase} is determined by the initial distribution, the tolerance, the observables, and the temperature of the system.

Thus, assuming these random variables are not too strongly correlated, we have that, roughly,

$$(5.3) \quad \langle \text{Speedup} \rangle \approx (1 - p) + p \frac{\langle t_{\text{phase}} \rangle + N \langle t_{\text{par}} \rangle}{\langle t_{\text{phase}} \rangle + \langle t_{\text{par}} \rangle} \approx (1 - p) + p \frac{\frac{C_{\text{phase}}}{\lambda_2 - \lambda_1} + \frac{1}{\lambda_1}}{\frac{C_{\text{phase}}}{\lambda_2 - \lambda_1} + \frac{1}{N\lambda_1}}.$$

As N tends to infinity, the maximum speedup monotonically saturates:

$$(5.4) \quad \lim_{N \rightarrow \infty} \langle \text{Speedup} \rangle \approx 1 + p \frac{1}{C_{\text{phase}}} \frac{\lambda_2 - \lambda_1}{\lambda_1}.$$

Alternatively, if

$$(5.5) \quad N \ll \frac{1}{C_{\text{phase}}} \frac{\lambda_2 - \lambda_1}{\lambda_1},$$

we can expand the expression to recover the linear speedup regime,

$$(5.6) \quad \langle \text{Speedup} \rangle \approx (1 - p) + pN \left(1 + C_{\text{phase}} \frac{\lambda_1}{\lambda_2 - \lambda_1} \right).$$

For strongly metastable systems, we expect that p will be close to one (the reference process almost always decorrelates) and that the relative spectral gap will be large, $\lambda_2 - \lambda_1 \gg \lambda_1$ (the typical exit time is much larger than the time to converge to the QSD). The constant C_{phase} will also depend on the degree of metastability, though the dependence is less clear. Thus, we expect there will be a large range of N satisfying (5.5), and a linear speedup will be observed. However, there is also an intrinsic limitation on the performance of ParRep that depends on the strength of metastability of each state.

The ratio in (5.4) has a natural interpretation. As N becomes large in ParRep, the cost of finding the first exit during the parallel step is driven to zero, and what remains is the cost of decorrelation and dephasing. These pieces of the algorithm have a time scale of $(\lambda_2 - \lambda_1)^{-1}$. For a direct serial simulation, in a sufficiently metastable system, the cost is to find the first exit, which has a time scale of λ_1^{-1} .

Since dephasing and decorrelation terminate at a finite time, there will always be some bias in the exit distributions, which is in general difficult to distinguish from the statistical noise. When the number of independent realizations of the experiment becomes sufficiently large, the statistical tests may detect this bias and identify two distinct distributions, even if the difference between the two distributions is tiny. Two things are thus to be expected when comparing the exit time distributions obtained by ParRep and direct numerical simulations:

- For a fixed number of independent realizations, as the tolerance is reduced to zero, the p -value of the K.-S. test goes to 1, and the null hypothesis is accepted; the distributions will be the same. This is because, as the tolerance becomes more stringent, it takes longer to reach stationarity. On the one hand, there is then a higher chance that the reference process exits during the joint

FIGURE 8. The partition of the configuration space into states for the potential (5.7). The circle $\circ$ indicates the initial condition, $(.5, .5)$.

decorrelation/dephasing step. On the other, for the realizations which manage to satisfy convergence criterion (3.5), the distribution will be closer to that of the QSD. Both these effects induce ParRep to better replicate the true, unaccelerated distributions.

- For fixed tolerance, as the number of realizations tends to $+\infty$, the p -value of the K.-S. tends to zero, and the null hypothesis is rejected. This is because as the statistical noise is reduced, the bias due to a finite stationarity time for a given tolerance becomes apparent. This is born out in numerical experiments.

In addition, due to sampling variability, the p -value may not vary monotonically as a function of TOL or of the number of realizations

5.2. Periodic Potential in 2D. As a first example, we consider the two dimensional periodic potential

$$(5.7) \quad V(x, y) = -\cos(\pi x) - \cos(\pi y),$$

simulated at $\beta = 3$. The states are defined as the translates of $(-1, 1)^2$, seen in Figure 8. The observables used with the Gelman-Rubin statistics in (3.4) are

$$(5.8) \quad x, \quad y, \quad V(x, y), \quad \|\mathbf{x} - \mathbf{x}_{\text{ref}}\|_{\ell^2},$$

where $\mathbf{x}_{\text{ref}}$ will be the local minimum of the state under consideration.

5.2.1. Escape from a Single State. We first consider the problem of escaping from the state $\Omega = (-1, 1)^2$ with initial condition $X_0 = (0.5, 0.5)$. We run 10^5 independent realizations of Serial and ParRep simulation, with quantitative data appearing in Table 1, and the distributions in Figures 9.

At the tolerance value $\text{TOL} = 0.2$, there is already good agreement in terms of first hitting point distributions, as seen in Table 5.7. To compare the first hitting point distributions, the boundary of $(-1, 1)^2$ is treated as a one dimensional manifold composed of the four edges (top, right, bottom, and left); this is the third plot of Figure 9. We indeed have good agreement.

Unlike the exit point distribution, the exit time distribution seems to require a more stringent tolerance. There are statistical errors at $\text{TOL} = 0.2$, which are mitigated as the tolerance is reduced. At $\text{TOL} = 0.05$, the K.-S. test shows agreement. However, based on the first two of Figure 9, the exit time distribution is in good qualitative agreement even for $\text{TOL} = 0.1$, in the sense that the L^1 distance between the two distributions is small and the decay rates are in agreement.

Notice that for $\text{TOL} = 0.05$, there is good performance in the sense that 84% of the realizations exited during the parallel step, and the statistical agreement is high. The corresponding speedup is 6.25.

With regard to our earlier comment that, at fixed tolerance, as the sample size increases, the p -values tend to decrease, we consider the case of $\text{TOL} = 0.1$ here. If we compare the first 10^2 , 10^3 , 10^4 and 10^5

TABLE 1. **Periodic Potential in 2D–Single Escape:** Comparison of ParRep and an unaccelerated serial process escaping from $(-1, 1)^2$ with $X_0 = (0.5, 0.5)$ for potential (5.7).

Method	TOL	$\langle t_{\text{phase}} \rangle$	$\text{Var}(t_{\text{phase}})$	$\langle T \rangle$	$\langle \text{Speedup} \rangle$	% Dephased	X_T K.-S. Test (p)	T K.-S. Test (p)
Serial	–	–	–	34.8	–	–	–	–
ParRep	0.2	1.12	0.0267	35.3	20.8	93.5%	PASS (0.59)	FAIL (2.0×10^{-5})
ParRep	0.1	2.43	0.103	35.0	11.6	90.2%	PASS (0.48)	FAIL (0.0037)
ParRep	0.05	5.10	0.393	34.8	6.25	83.6%	PASS (0.52)	PASS (0.33)
ParRep	0.01	26.2	9.08	34.8	1.63	46.5%	PASS (0.27)	PASS (0.42)

FIGURE 9. **Periodic Potential in 2D–Single Escape:** Exit distributions from $\Omega = (-1, 1)^2$ for the 2D Periodic potential (5.7). On the right is a magnification of the figure from the left; notice that the confidence intervals do not overlap at $\text{TOL} = .1$. For the hitting point distribution, the boundary of Ω has been treated as a 1D manifold parametrized clockwise, with the labels “Top”, “Right”, “Bottom”, and “Left” corresponding respectively to the edges $(-1, 1) \times \{1\}$, $\{1\} \times (-1, 1)$, etc.

experimental realizations against the corresponding serial results, we obtain p -values of 0.68, 0.18, 0.0099, and 0.0037.

TABLE 2. **Periodic Potential in 2D–Multiple Escapes:** Comparison of ParRep and an unaccelerated serial process escaping from $(-3, 3)^2$ with $X_0 = (0.5, 0.5)$ for potential (5.7).

TOL	X_T K.-S. Test (p)	T K.-S. Test (p)
0.2	PASS (0.56)	FAIL (5.2×10^{-3})
0.1	PASS (0.51)	PASS (0.24)
0.05	PASS (0.53)	PASS (0.59)
0.01	PASS (0.11)	PASS (0.87)

FIGURE 10. **Periodic Potential in 2D–Multiple Escapes:** Exit time distribution for the 2D periodic potential (5.7) when ParRep is used to escape from $(-3, 3)^2$. On the right is a magnification of the figure from the left; notice that the confidence intervals do not overlap at $\text{TOL} = .2$. For the hitting point distribution, the boundary of $(-3, 3)^2$ has been treated as a 1D manifold parametrized clockwise, with the labels “Top”, “Right”, “Bottom”, and “Left” corresponding respectively to the edges $(-3, 3) \times \{3\}$, $\{3\} \times (-3, 3)$, etc.

5.2.2. *Escape from a Region Containing Multiple States.* Next, we consider the problem of escaping from the region $(-3, 3)^2$, running our modified ParRep algorithm over the states indicated in Figure 8. The results are reported in Table 2 and Figure 10. Beneath the value $\text{TOL} = 0.1$, the algorithm is statistically consistent, both in terms of first hitting point and first exit time distributions.

FIGURE 11. A domain with entropic barriers. The trajectory begins at $(5.5, 0.1)$ for the experiment in Section 5.3.1 and at $(0.1, 0.1)$ for the experiment in Section 5.3.2 (see the $\circ$ symbols). The width of the necks between states is 0.1, and the dividing surfaces between states are indicated by the short segments. A separate reference point is used in each state for the observable $\|\mathbf{x} - \mathbf{x}_{\text{ref}}\|$.

TABLE 3. **Entropic Barrier in 2D—Single Escape:** Comparison of ParRep and an un-accelerated serial process escaping from state 2 of Figure 11 with initial condition $(5.5, .1)$.

Method	TOL	$\langle t_{\text{phase}} \rangle$	$\text{Var}(t_{\text{phase}})$	$\langle T \rangle$	$\langle \text{Speedup} \rangle$	% Dephased	$\mathbb{P}[\text{state 3}]$	T K.-S. Test (p)
Serial	—	—	—	45.0	—	—	$(0.493, 0.500)$	—
ParRep	0.2	12.7	2.31	44.8	3.46	77.9%	$(0.495, 0.502)$	PASS (0.90)
ParRep	0.1	25.4	8.75	45.1	1.95	58.0%	$(0.493, 0.499)$	PASS (0.51)
ParRep	0.05	50.3	33.3	45.1	1.27	32.3%	$(0.497, 0.503)$	PASS (0.41)
ParRep	0.01	236.0	658.0	45.1	1.00	0.342%	$(0.496, 0.502)$	PASS (0.31)

5.3. Entropic Barriers in 2D. Let us now consider a test case with entropic barriers. Consider pure Brownian motion in the domain represented in Figure 11, with reflecting boundary conditions. Here, the observables are

$$(5.9) \quad x, y, \|\mathbf{x} - \mathbf{x}_{\text{ref}}\|_{\ell^2},$$

with the reference points indicated in Figure 11.

5.3.1. Escape from State 2. As a first experiment, we look for first escapes from state 2, with initial condition $(5.5, 0.1)$. Quantitative results appear in Table 3 and the exit time distributions are plotted in Figure 12. For this problem, there is good statistical agreement even at $\text{TOL} = 0.2$. As the hitting point distributions across each of the two channels are nearly uniform, we only report the probability of passing into state 3 versus state 1. Note that in the $\text{TOL} = 0.01$ case, the exit is almost always due to the reference process escaping before stationarity is achieved. This is an example of setting the parameter so stringently as to render ParRep inefficient.

5.3.2. Getting to State 3 from State 1. As a second experiment, we start the trajectories in state 1 at $(0.1, 0.1)$, and examine how long it takes to get to state 3, running the modified ParRep algorithm over the states represented on Figure 11. The results of this experiment are given in Table 4 and Figure 13. Again, there is a very good statistical agreement in all cases.

FIGURE 12. **Entropic Barrier in 2D–Single Escape:** Exit time distribution from state 2 of Figure 11 with initial conditions $(5.5, .1)$, along with confidence intervals.

TABLE 4. **Entropic Barrier in 2D–Multiple Escapes:** Comparison of ParRep and an unaccelerated serial process getting from state 1 to state 3 with initial condition $(.1, .1)$; see Figure 11.

TOL	T K.-S. Test (p)
0.2	PASS (0.070)
0.1	PASS (0.14)
0.05	PASS (0.090)
0.01	PASS (0.74)

FIGURE 13. **Entropic Barrier in 2D–Multiple Escapes:** Exit time distributions to get to state 3 from state 1 in the entropic barrier problem pictured in Figure 11, along with confidence intervals. The full algorithm is applied in each state the trajectory visits.

FIGURE 14. Low energy conformations of the Lennard-Jones cluster of seven atoms in 2D. Also indicated are the associated energy levels.

5.4. Lennard-Jones Clusters. For a more realistic problem, we consider a Lennard-Jones cluster of seven atoms in 2D, denoted $\text{LJ}_7^{2\text{D}}$. The potential used in this problem is then:

$$(5.10) \quad V_7^{2\text{D}}(\mathbf{x}) = \frac{1}{2} \sum_{i \neq j} \phi(\|\mathbf{x}_i - \mathbf{x}_j\|), \quad \mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_7), \quad \phi(r) = r^{-12} - 2r^{-6}.$$

Due to the computational cost associated with this problem, the time step size is increased to $\Delta t = 10^{-3}$. A smaller time step could have been used, but we would have been precluded from obtaining large samples of the problem.

The initial condition in this problem is near the closest packed configuration, with six atoms located $\pi/3$ radians apart on a unit circle and one in the center, as shown in Figure 14 (a). We then look at first exits from this configuration at inverse temperature $\beta = 6$. Exits correspond to transitions into basins of attraction for a conformation other than the lowest energy, closest packed, conformation; see Figure 14. Here, basins of attraction are those associated to the simple gradient dynamics $\dot{x} = -\nabla V(x)$: these basins define the states on which ParRep is applied. For the numbering of the conformations, we adopt the notation found in [14], where the authors also explore transitions in $\text{LJ}_7^{2\text{D}}$. We also refer to [19] for an algorithm for computing transitions within Lennard-Jones clusters.

Before proceeding to our results, we make note of several things that are specific to this problem. States are identified and compared as follows:

- (1) Given the current configuration X_t , a gradient descent is run (following the dynamics $\dot{x} = -\nabla V(x)$ with $x_0 = X_t$) until the norm of the gradient becomes sufficiently small. This is accomplished using the RK45 time discretization scheme, up to the time when the ℓ^2 norm of the gradient, relative to the value of $|V|$, is smaller than 10^{-5} ;
- (2) Once the local minima is found, in order to identify the conformation, the lengths of the twenty one “bonds” between pairs of atoms are computed. Let b_{ij} denote the distance between the atoms i and j .
- (3) The current and previous states are compared using the ℓ^1 norm between the 21 bond lengths:

$$\text{Distance to } C_0 = \sum_{1 \leq i < j \leq 7} |b_{ij} - b_{ij}^0|$$

where b_{ij}^0 are the bond lengths of the C_0 conformation.

- (4) If this distance exceeds .5, the states are determined to be different.

Next, when examining the output, we group conformations according to energy levels. In Figure 14, we present for each conformation C_0 , C_1 , C_2 and C_3 a particular numbering. Other permutations of the atoms also correspond to the same energy level, and thus to the same conformation.

TABLE 5. **LJ₇^{2D}-Single Escape:** Comparison of the exit times for ParRep and an unaccelerated serial process escaping from the closest packed LJ₇^{2D} configuration.

Method	TOL	$\langle t_{\text{phase}} \rangle$	$\text{Var}(t_{\text{phase}})$	$\langle T \rangle$	$\langle \text{Speedup} \rangle$	% Dephased	T K.-S. Test (p)
Serial	–	–	–	17.0	–	–	–
ParRep	0.2	0.411	0.0336	19.1	29.3	98.5%	FAIL (3.6×10^{-235})
ParRep	0.1	.976	0.125	18.0	14.9	95.3%	FAIL (5.0×10^{-55})
ParRep	0.05	2.08	0.433	17.6	7.83	90.0%	FAIL (9.4×10^{-28})
ParRep	0.01	10.8	5.67	17.0	1.82	52.1%	PASS (0.92)

TABLE 6. **LJ₇^{2D}-Single Escape:** Comparison of the hitting points for ParRep and an unaccelerated serial process escaping from the closest packed LJ₇^{2D} configuration.

Method	TOL	$\mathbb{P}[C_1]$	$\mathbb{P}[C_2]$
Serial	–	(0.502, 0.508)	(0.491, 0.498)
ParRep	0.2	(0.508, 0.514)	(0.485, 0.492)
ParRep	0.1	(0.506, 0.512)	(0.488, 0.494)
ParRep	0.05	(0.505, 0.512)	(0.488, 0.495)
ParRep	0.01	(0.504, 0.510)	(0.490, 0.496)

The observables that are used in this problem are:

$$(5.11a) \quad \text{Energy: } V(x)$$

$$(5.11b) \quad \text{Square distance to the center of mass: } \sum_{i=1}^7 \|x_i - x_{\text{cm}}\|_{\ell^2}^2$$

$$(5.11c) \quad \text{Distance to } C_0: \sum_{1 \leq i < j \leq 7} |b_{ij} - b_{ij}^0|$$

It is essential to use observables that are rotation and translation invariant, as there is nothing to prevent the cluster from drifting or rotating (*i.e.* no atom is pinned).

A related remark is that, to not introduce additional auxiliary parameters, no confining potential is used. Because of this, the cluster does not always change into one of the other conformations. In some cases, an atom can drift away from the cluster, and the quenched configuration corresponds to the isolated atom and the five remaining atoms. In others, the quenched configuration corresponds to C_0 , but with two exterior atoms exchanged. These scenarios occurred in less than .1% of the 10^5 realizations of all experiments, both ParRep and unaccelerated.

Over 99.9% of the transitions from C_0 are to C_1 or C_2 . A transition from C_0 to C_3 has not been observed. Since there are thus only two ways to leave C_0 , we use the 95% Clopper-Pearson confidence intervals [21] (well adapted to binomial random variables) to assess the quality of the exit point distribution in our statistical tests.

Our statistical assessment of ParRep is given in Tables 5 and 6. It is only at the most stringent tolerance of $\text{TOL} = 0.01$ that an excellent agreement is obtained in the exit time distribution, though some amount of speedup is still gained. As in the other problems, the Kolmogorov-Smirnov test may be an overly conservative measure of the quality of ParRep. As shown in Figure 15, the exit time distributions are already in good qualitative agreement, even at less stringent tolerances. In contrast, the transition probabilities to the other conformations are in very good agreement at tolerances beneath 0.1.

FIGURE 15. $\text{LJ}_7^{2\text{D}}$ -Single Escape: Exit time distributions for $\text{LJ}_7^{2\text{D}}$ at $\beta = 6$.

5.5. Conclusions from the Numerical Experiments. From these numerical experiments, we observe that the modified ParRep is indeed an efficient algorithm on various test cases, for which no a priori knowledge on the decorrelation/dephasing time has been used. We thus trade setting t_{corr} and t_{phase} , *a priori*, for selecting physically informed observables together with a tolerance. In conclusion, a tolerance on the order of 0.1 in (3.4) seems to yield sensible results.

We also observe that the tolerance criteria used to assess stationarity does not need to be very stringent to get the correct distribution for the hitting points (thus to get the correct state-to-state Markov chains). In contrast, the K-S. test on the exit time distribution requires smaller tolerances to predict statistical agreement. However, even when the K-S. test fails to reject the null hypothesis, the agreement is very good. Overall, it is easier to obtain statistical agreement in the sequence of visited states than it is to get statistical agreement of the time spent within each state.

REFERENCES

- [1] D. Aristoff and T. Lelièvre. Mathematical Analysis of Temperature Accelerated Dynamics. *MMS*, 12(1):290–317, 2014.
- [2] D. Aristoff, T. Lelièvre, and G. Simpson. The parallel replica method for simulating long trajectories of Markov chains. *Appl. Math. Res. Express*, 2014:332–352, 2014.
- [3] S. Asmussen and P. W. Glynn. *Stochastic Simulation: Algorithms and Analysis*. Springer, November 2010.
- [4] S. Brooks, A. Gelman, G. Jones, and X.-L. Meng. *Handbook of Markov Chain Monte Carlo*. Chapman and Hall/CRC, 2011.
- [5] S.P. Brooks and A. Gelman. General methods for monitoring convergence of iterative simulations. *J. Comput. Graph. Stat.*, pages 434–455, 1998.
- [6] SP Brooks and GO Roberts. Convergence assessment techniques for Markov chain Monte Carlo. *Stat. Comp.*, 8(4):319–335, 1998.
- [7] P. Cattiaux, P. Collet, A. Lambert, S. Martínez, S. Méléard, and J. San Martín. Quasi-stationary distributions and diffusion models in population dynamics. *Ann. Probab.*, 37(5):1926–1969, 2009.
- [8] P. Cattiaux and S. Méléard. Competitive or weak cooperative stochastic Lotka-Volterra systems conditioned on non-extinction. *J. Math. Biol.*, 60(6):797–829, 2010.
- [9] P. Collet, S. Martínez, and J. San Martín. Asymptotic laws for one-dimensional diffusions conditioned to nonabsorption. *Ann. Probab.*, 23(3):1300–1314, 1995.
- [10] P. Collet, S. Martínez, and J. San Martín. *Quasi-Stationary Distributions*. Springer, 2013.
- [11] Mary Kathryn Cowles and Bradley P. Carlin. Markov chain Monte Carlo convergence diagnostics: a comparative review. *J. Amer. Statist. Assoc.*, 91(434):883–904, 1996.
- [12] H. L. Crutcher. A note on the possible misuse of the Kolmogorov-Smirnov test. *J. Appl. Meteorol.*, 14(8):1600–1603, 1975.
- [13] P. Del Moral. *Feynman-Kac Formulae*. Springer, 2004.
- [14] C Dellago, P G Bolhuis, and D Chandler. Efficient transition path sampling: Application to Lennard-Jones cluster rearrangements. *J. Chem. Phys.*, 108(22):9236–9245, 1998.
- [15] P.A. Ferrari and N. Maric. Quasi stationary distributions and Fleming-Viot processes in countable spaces. *Electron. J. Probab.*, 12(24):684–702, 2007.
- [16] M. Galassi, J. Davies, J. Theiler, B. Gough, G. Jungman, P. Alken, M. Booth, and F. Rossi. GNU Scientific Library, March 2013.

- [17] A. Gelman and D.B. Rubin. Inference from iterative simulation using multiple sequences. *Stat. Sci.*, 7(4):457–472, 1992.
- [18] J. D. Gibbons. *Nonparametric Methods for Quantitative Analysis*. American Sciences Press, Third edition, 1997.
- [19] M. Hairer and J. Weare. Improved diffusion Monte Carlo and the Brownian fan. arXiv:1207.2866, 2012.
- [20] S. D. Horn. Goodness-of-fit tests for discrete data: a review and an application to a health impairment scale. *Biometrics*, 33(1):237, 1977.
- [21] N. L. Johnson, A. W. Kemp, and S. Kotz. *Univariate discrete distributions*. Wiley, Third edition, 2005.
- [22] S. Kirmizialtin and R. Elber. Revisiting and computing reaction coordinates with directional milestoning. *J. Phys. Chem. A*, 115(23):6137–6148, 2011.
- [23] O. Kum, B. M. Dickson, S. J. Stuart, B. P. Uberuaga, and A. F. Voter. Parallel replica dynamics with a heterogeneous distribution of barriers: Application to n-hexadecane pyrolysis. *The Journal of Chemical Physics*, 121(20):9808, 2004.
- [24] C. Le Bris, T. Lelièvre, M. Luskin, and D. Perez. A mathematical formalization of the parallel replica dynamics. *Monte Carlo Meth. Appl.*, 18(2):119–146, 2012.
- [25] T. Lelièvre and F. Nier. Low temperature asymptotics for quasi-stationary distribution in a bounded domain. arXiv:1309.3898, 2013.
- [26] S. Martínez and J. San Martín. Quasi-stationary distributions for a Brownian motion with drift and associated limit laws. *J. Appl. Probab.*, 31(4):911–920, 1994.
- [27] S. Martínez and J. San Martín. Classification of killed one-dimensional diffusions. *Ann. Probab.*, 32(1A):530–552, 2004.
- [28] M. Mascagni and A. Srinivasan. Algorithm 806: SPRNG: a scalable library for pseudorandom number generation. *TOMS*, 26(3), September 2000.
- [29] S. Méléard and D. Villemonais. Quasi-stationary distributions and population processes. *Probability Surveys*, 2012.
- [30] S. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Cambridge University Press, second edition, 2009.
- [31] D. Perez, B.P. Uberuaga, Y. Shim, J.G. Amar, and A.F. Voter. Accelerated molecular dynamics methods: introduction and recent developments. *Ann. Rep. Comp. Chem.*, 5:79–98, 2009.
- [32] G. Simpson and M. Luskin. Numerical analysis of parallel replica dynamics. *M2AN*, 47(5):1287–1314, 2013.
- [33] M. R. Sørensen and A. F. Voter. Temperature-accelerated dynamics for simulation of infrequent events. *J. Chem. Phys.*, 112(21):9599–9606, 2000.
- [34] D. Steinsaltz and S. N. Evans. Quasistationary distributions for one-dimensional diffusions with killing. *T. Am. Math. Soc.*, 359(3):1285–1324 (electronic), 2007.
- [35] A.-S. Sznitman. Topics in propagation of chaos. In *École d’Été de Probabilités de Saint-Flour XIX—1989*, volume 1464 of *Lecture Notes in Math.*, pages 165–251. Springer, 1991.
- [36] E. Vanden-Eijnden, M. Venturoli, G. Ciccotti, and R. Elber. On the assumptions underlying milestoning. *J. Chem. Phys.*, 129(17):174102, 2008.
- [37] A.F Voter. Hyperdynamics: Accelerated molecular dynamics of infrequent events. *Phys. Rev. Lett.*, 78(20):3908–3911, 1997.
- [38] A.F Voter. A method for accelerating the molecular dynamics simulation of infrequent events. *J. Chem. Phys.*, 106(11):4665–4677, 1997.
- [39] A.F. Voter. Parallel replica method for dynamics of infrequent events. *Phys. Rev. B*, 57(22):13985–13988, 1998.
- [40] A.F. Voter, F. Montalenti, and T.C. Germann. Extending the time scale in atomistic simulation of materials. *Ann. Rev. Mater. Sci.*, 32:321–346, 2002.