

# Replica exchange with nonequilibrium switches

Andrew J. Ballard<sup>a</sup> and Christopher Jarzynski<sup>a,b,1</sup>

<sup>a</sup>Institute for Physical Science and Technology, University of Maryland, College Park, MD 20742; and <sup>b</sup>Department of Chemistry and Biochemistry, University of Maryland, College Park, MD 20742

Edited by Bruce J. Berne, Columbia University, New York, NY, and approved May 7, 2009 (received for review January 14, 2009)

**We introduce a replica exchange (parallel tempering) method in which attempted configuration swaps are generated using nonequilibrium work simulations. By effectively increasing phase space overlap, this approach mitigates the need for many replicas. We illustrate our method by using a model system and show that it is able to achieve the computational efficiency of ordinary replica exchange, using fewer replicas.**

parallel tempering | nonequilibrium work simulations | molecular dynamics | Monte Carlo | rough energy landscapes

**E**quilibrium sampling is the engine that drives computational statistical mechanics, and since the early years of numerical simulation many schemes have been developed for drawing from equilibrium ensembles. The efficient sampling of large, complex systems remains a challenge, however, not only due to the inherent cost of simulating many degrees of freedom, but also because such systems often exhibit rough energy landscapes that hinder the exploration of phase space. Replica exchange (or parallel tempering) (1, 2) has emerged as a powerful tool to address this challenge. The gist of the method is to simulate  $M$  independent copies of the system of interest, typically ordered by increasing temperature, and to perform “swaps” in which an attempted exchange of configurations between adjacent replicas is accepted or rejected according to a Metropolis-like scheme (3). By means of these swaps, the low-temperature replicas gain access to the broader expanses of phase space explored by the high-temperature replicas.

If there is an Achilles’ heel to the replica exchange method, it is the way that it scales with system size,  $N$ : The number of replicas needed to simulate the system typically grows as  $N^{1/2}$  (4). This scaling is rooted in a phase space overlap requirement—in order to achieve a reasonable frequency of accepted swaps, neighboring replicas should overlap in phase space, and with increasing system size more replicas are needed to satisfy this requirement. The difficulty is particularly acute in simulations of large molecules in an explicit solvent where the poor overlap is due mostly to the large number of solvent molecules. With the growing popularity of replica exchange, particularly for the simulation of biomolecules (5), the overlap problem has received increasing attention. Strategies proposed to address this problem include tempering or perturbing only a subset of degrees of freedom (6–8); and the use of Tsallis-like (9–11), multicanonical (12–14), or expanded-ensemble (15, 16) probability distributions.

In the present paper, we develop a replica exchange method that uses nonequilibrium simulations to increase the overlap between replicas. Our approach, illustrated in Fig. 1, is summarized as follows. When attempting a swap between replicas A and B, we first perform a finite-time “work simulation” in replica A, during which we drag the system toward the region of phase space characteristic of equilibrium ensemble B, and vice-versa in replica B. We then attempt to swap the configurations thus generated by using a work-based acceptance criterion (Eq. 2). Although the system is driven away from equilibrium during the work simulations, the acceptance/rejection step ensures that detailed balance is preserved and that the replicas’ equilibrium states are undisturbed. Although the CPU time devoted to the work simulations represents an added computational cost, the return on this investment is an increased acceptance probability. As a result, our method is able to achieve

the same sampling efficiency as ordinary replica exchange but with fewer replicas.

In *Description of Method* and *Derivation*, we describe and derive our method in detail, and then we illustrate it by using a model system adapted from ref. 3. We will use the acronym REM to refer to the usual replica exchange method, and RENS to denote our method based on nonequilibrium work simulations.

## Description of Method

Let  $\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_M$  denote the collection of replicas, and let  $H_i(x)$  and  $T_i$  denote the Hamiltonian (energy function) and the temperature, respectively, of the canonical ensemble simulated in  $\mathcal{R}_i$ . Here,  $x$  denotes a point in configuration space or phase space. Often, the  $H_i$ ’s are identical and only the temperatures differ, or vice-versa, but we need not assume this is the case. As in ref. 17, let us define a reduced Hamiltonian  $h_i(x) = H_i(x)/k_B T_i$ , so that the equilibrium distribution in  $\mathcal{R}_i$  takes the form  $p_i^{\text{eq}} \propto \exp(-h_i)$ . When implementing REM, if 2 replicas  $\mathcal{R}_A$  and  $\mathcal{R}_B$  are found in configurations  $x$  and  $y$  at the time of an attempted swap, then the swap  $x \leftrightarrow y$  is accepted with probability  $P_{\text{acc}} = \min\{1, e^{-\Delta h}\}$ , where

$$\Delta h(x, y) \equiv h_B(x) + h_A(y) - h_B(y) - h_A(x). \quad [1]$$

If there is little overlap between the distributions  $p_A^{\text{eq}}$  and  $p_B^{\text{eq}}$ , then typically  $P_{\text{acc}} \ll 1$ , and the swap is most likely rejected.

Fig. 1 illustrates RENS, the method we propose as an alternative. The time interval from  $t_0$  to  $t_1$  corresponds to independent, equilibrium sampling in each of the  $M$  replicas (two of which are depicted), using the reduced Hamiltonians  $h_1, \dots, h_M$ . At time  $t_1$ , we decide to attempt a swap between replicas A and B. In lieu of an instantaneous exchange  $x \leftrightarrow y$ , we first perform a pair of work simulations. In  $\mathcal{R}_A$ , the system evolves from time  $t_1$  to  $t_2$ , as the reduced Hamiltonian is parametrically “switched” from  $h_A$  to  $h_B$  (Eq. 3.1). Let  $x'$  denote the configuration at the end of this switching process, and  $w_A$  the reduced work performed on the system (Eq. 4.1). In  $\mathcal{R}_B$ , we simulate the reverse process, parametrically switching from  $h_B$  to  $h_A$ , and define  $y'$  and  $w_B$  analogously. We then attempt a swap between the 2 replicas, with acceptance probability

$$P_{\text{acc}} = \min\{1, e^{-w}\}, \quad [2]$$

where  $w = w_A + w_B$ . If accepted, the configuration  $y'$  is copied into replica A, and  $x'$  into replica B, whereas, if rejected, the configurations revert to  $x$  and  $y$ . Subsequently, equilibrium sampling continues in  $\mathcal{R}_A$  and  $\mathcal{R}_B$  (using  $h_A$  and  $h_B$ , respectively) and in all other replicas until the next attempted swap. Fig. 1 depicts a successful replica swap at time  $t_2$  (open circles); if the attempt had been rejected, the replicas would have been reset to the states  $x, y$ , depicted as filled circles. Note that while work simulations

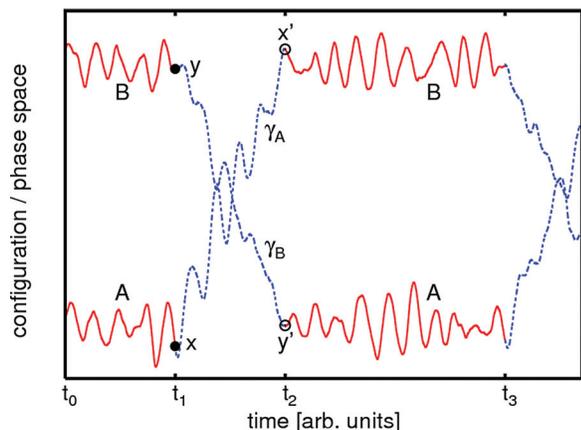
Author contributions: A.J.B. and C.J. designed research; A.J.B. and C.J. performed research; A.J.B. analyzed data; and A.J.B. and C.J. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

<sup>1</sup>To whom correspondence should be addressed. E-mail: cjarzyns@umd.edu.

This article contains supporting information online at [www.pnas.org/cgi/content/full/0900406106/DCSupplemental](http://www.pnas.org/cgi/content/full/0900406106/DCSupplemental).



**Fig. 1.** Schematic depiction of RENS. The solid red lines depict intervals of equilibrium sampling in replicas A and B, and the dashed blue lines represent work simulations.

are performed in  $\mathcal{R}_A$  and  $\mathcal{R}_B$ , the remaining replicas continue to sample at fixed  $h_i$ .

Thus, each replica alternates between sampling intervals at fixed  $h_i$  (Fig. 1, solid red lines), and work intervals (dashed blue lines). We next show that RENS satisfies detailed balance, in the following sense: In each replica  $\mathcal{R}_i$ , if we discard the data generated during the work intervals and stitch together the remaining sampling intervals, we obtain a long trajectory that samples the distribution  $p_i^{\text{eq}}$ . In effect, the acceptance criterion compensates for the fact that the system is driven out of equilibrium during the work simulations.

Our method is quite general and ultimately traces its validity to Crooks' extension of detailed balance to nonequilibrium trajectories (18). In a given implementation, however, the definition of reduced work depends on the dynamics chosen to model the evolution of the system. For discrete-time Monte Carlo dynamics, RENS is equivalent to the annealed swapping method of ref. 19 and closely related to the C-walking algorithm of ref. 20. In the following section, we derive our method for deterministic, reversible molecular dynamics (MD), and in the *SI Appendix* we extend this derivation to include stochastic evolution.

## Derivation

For a pair of replicas  $\mathcal{R}_A$  and  $\mathcal{R}_B$ , we introduce a parametrized Hamiltonian  $h(x; \lambda)$  that interpolates from  $h(x; 0) = h_A(x)$  to  $h(x; 1) = h_B(x)$ . To implement an attempted swap between these replicas, we first specify a switching protocol  $\lambda_t^A$ , with  $\lambda_0^A = 0$  and  $\lambda_\tau^A = 1$ . In  $\mathcal{R}_A$ , starting from state  $x_0 = x$ , we generate a trajectory  $\gamma_A$  during which the system evolves under the specified dynamics as the parameter  $\lambda$  is varied from 0 to 1 according to the switching protocol.

$$\gamma_A : x = x_0 \xrightarrow{\lambda \rightarrow 1} x_\tau = x'. \quad [3.1]$$

Simultaneously, in  $\mathcal{R}_B$  we generate a trajectory  $\gamma_B$ , varying  $\lambda$  from 1 to 0 using the time-reversed protocol,  $\lambda_t^B = \lambda_{\tau-t}^A$ ,

$$\gamma_B : y' = y_\tau \xleftarrow{0 \leftarrow \lambda} y_0 = y. \quad [3.2]$$

The arrows denote the direction of time. The 2 trajectories  $\gamma_A$  and  $\gamma_B$  are illustrated by the dashed blue segments between times  $t_1$  and  $t_2$  in Fig. 1. We assume these trajectories are generated by deterministic equations of motion that are symmetric under time-reversal; for any trajectory  $\gamma_A = (x_0 \rightarrow x_\tau)$  that evolves under the protocol  $\lambda_t^A$ , the time-reversed trajectory  $\bar{\gamma}_B = (\bar{x}_0 \leftarrow \bar{x}_\tau)$  evolves under  $\lambda_t^B$ , where  $\bar{x}$  denotes inversion of momenta,  $\mathbf{p} \rightarrow -\mathbf{p}$ .

This assumption is satisfied by Hamiltonian dynamics, Nosé-Hoover dynamics, and other equations of motion, provided the Hamiltonian itself is time-reversal symmetric, i.e.  $h(x; \lambda) = h(\bar{x}; \lambda)$ .

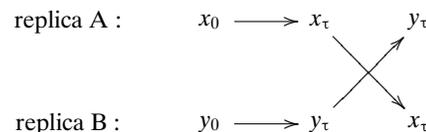
For the trajectories  $\gamma_A$  and  $\gamma_B$ , we define the reduced work as follows:

$$w_A(x_0 \rightarrow x_\tau) = h_B(x_\tau) - h_A(x_0) - \ln J_A(x_0) \quad [4.1]$$

$$w_B(y_0 \rightarrow y_\tau) = h_A(y_\tau) - h_B(y_0) - \ln J_B(y_0). \quad [4.2]$$

Here,  $J_A = |\partial x_\tau / \partial x_0|$  and  $J_B = |\partial y_\tau / \partial y_0|$  are the Jacobians associated with propagating the system from the initial to the final point. Eq. 4 is analogous to the first law of thermodynamics, with  $\ln J$  representing a heat term associated with increase of system entropy.

When both work simulations have been completed, we attempt to swap the final configurations, assigning  $x' = x_\tau$  to replica B and  $y' = y_\tau$  to replica A. Schematically,



where the parallel arrows indicate the work simulations and the crossed arrows the attempted swap. The acceptance probability for the swap is given by Eq. 2. If it is rejected, the replicas are reset to their initial states,  $x = x_0$  and  $y = y_0$ .

To analyze our method, it is useful to think of an expanded phase space containing 2 copies of the system, corresponding to  $\mathcal{R}_A$  and  $\mathcal{R}_B$ . In this space, we wish to sample the distribution  $P_{AB}^{\text{eq}}(x, y) \propto e^{-h_A(x) - h_B(y)}$ . The work simulations, followed by the attempted swap, represent an elaborate trial Monte Carlo move  $(x, y) \rightarrow (y', x')$  (21). Let  $P(y', x' | x, y)$  be the corresponding transition probability—that is, the probability that the work simulations end in configurations  $x'$  and  $y'$ , and the swap is accepted, given initial configurations  $x$  and  $y$ . To establish detailed balance, we must show that the net probability to observe a transition  $(x, y) \rightarrow (y', x')$  is equal to that of the reverse transition  $(\bar{x}, \bar{y}) \leftarrow (\bar{y}', \bar{x}')$  (3, 22), i.e. that

$$P(y', x' | x, y) P_{AB}^{\text{eq}}(x, y) = P(\bar{x}, \bar{y} | \bar{y}', \bar{x}') P_{AB}^{\text{eq}}(\bar{y}', \bar{x}'). \quad [5]$$

We will establish this result by introducing a few useful definitions and identities, which are then combined in Eq. 12.

For the work simulations in  $\mathcal{R}_A$  and  $\mathcal{R}_B$ , we can treat the final microstate as a function of the initial microstate (23),

$$x_\tau = M_A(x_0), \quad y_\tau = M_B(y_0), \quad [6]$$

obtained by integrating the deterministic equations of motion. We will use the notation

$$\pi_A(x' | x) = \delta(x' - M_A(x)) \quad [7.1]$$

to denote the probability to arrive at  $x'$  during a work simulation in  $\mathcal{R}_A$ , starting from  $x$ ; and similarly

$$\pi_B(y' | y) = \delta(y' - M_B(y)). \quad [7.2]$$

Moreover, let

$$\pi(x', y' | x, y) = \pi_A(x' | x) \pi_B(y' | y) \quad [8]$$

denote the joint probability for both events, and let

$$\alpha(x', y' | x, y) = \min\{1, e^{-w_A(x \rightarrow x') - w_B(y \rightarrow y')}\} \quad [9]$$

be the probability to accept the corresponding swap (Eq. 2). The transition probability  $P(y', x' | x, y)$  is then given by the product  $P = \pi \alpha$ .

The functions  $M_A$  and  $M_B$  introduced in Eq. 6 are related by our assumption of time-reversal symmetry. Namely, if  $x' = M_A(x)$ , then  $\bar{x} = M_B(\bar{x}')$ , which in turn implies

$$\pi_A(x'|x) = \pi_B(\bar{x}|\bar{x}')/J_A(x) \quad [10.1]$$

$$\pi_B(y'|y) = \pi_A(\bar{y}|\bar{y}')/J_B(y). \quad [10.2]$$

(Identifying  $q = \ln J$  as reduced heat, this result is equivalent to equation 9 of ref. 18.) Finally, the reduced work (Eq. 4) is odd under time-reversal,

$$w_A(x \rightarrow x') = -w_B(\bar{x}' \rightarrow \bar{x}). \quad [11]$$

Now, combining Eqs. 4 and 8–11, we get

$$\begin{aligned} P(y', x'|x, y) &= \pi(x', y'|x, y) \alpha(x', y'|x, y) \\ &= \frac{\pi(\bar{y}, \bar{x}|\bar{y}', \bar{x}')}{J_A(x) J_B(y)} \alpha(\bar{y}, \bar{x}|\bar{y}', \bar{x}') e^{-w_A(x \rightarrow x') - w_B(y \rightarrow y')} \\ &= P(\bar{x}, \bar{y}|\bar{y}', \bar{x}') e^{-h_A(y') - h_B(x') + h_A(x) + h_B(y)} \end{aligned} \quad [12]$$

Because Eq. 12 is equivalent to Eq. 5, our scheme for generating configuration swaps preserves equilibrium in each replica.

### Illustrative Dynamics

A simple dynamical scheme nicely illustrates our method. We suppose  $\mathcal{R}_A$  and  $\mathcal{R}_B$  are described by different temperatures,  $T_A < T_B$ , but the same  $H$ , and we construct Hamilton's equations augmented by a term proportional to  $\dot{\lambda}$ :

$$\dot{\mathbf{q}}_i = \frac{\partial H}{\partial \mathbf{p}_i}, \quad \dot{\mathbf{p}}_i = -\frac{\partial H}{\partial \mathbf{q}_i} + \dot{\lambda} s_{\lambda} \mathbf{p}_i, \quad [13]$$

with  $s_{\lambda} = (1/2T_{\lambda})(dT_{\lambda}/d\lambda)$ . Here,  $T_{\lambda}$  interpolates from  $T_0 = T_A$  to  $T_1 = T_B$ . The extra term in Eq. 13 provides rudimentary temperature control during the work simulations. As  $\lambda$  is varied from 0 to 1 in replica  $A$ , the momenta are scaled up, effectively heating up the system; in replica  $B$ , the system is cooled. An equivalent rescaling of momenta is standard practice in REM, where it is performed instantaneously rather than over the course of a trajectory; see equation 12 of ref. 24.

These dynamics do not preserve phase space volume  $\nabla \cdot \dot{x} = N\dot{\lambda}s_{\lambda} \neq 0$ , where  $N$  is the number of degrees of freedom. The Jacobian for a work simulation in  $\mathcal{R}_A$  is then

$$J_A = \exp\left(\int_0^{\tau} dt \nabla \cdot \dot{x}\right) = \left(\frac{T_B}{T_A}\right)^{N/2}, \quad [14]$$

and in  $\mathcal{R}_B$  we have  $J_B = J_A^{-1}$ . (The fact that  $J_A$  and  $J_B$  do not depend on initial conditions is specific to these dynamics.)

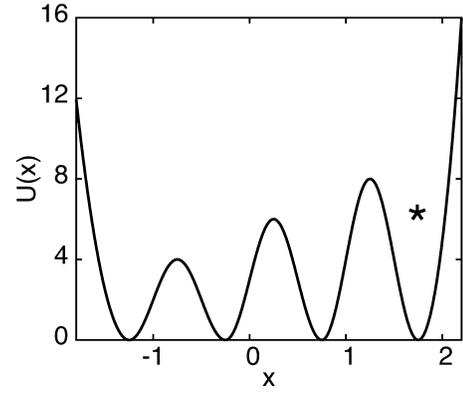
By heating the system during the work simulation in  $\mathcal{R}_A$  and cooling it in  $\mathcal{R}_B$ , the scaling term  $\dot{\lambda}s_{\lambda}\mathbf{p}_i$  increases the probability for accepting the configuration swap. (Indeed, for a system of ideal gas particles, evolution under Eq. 13 exactly transforms a Maxwell–Boltzmann distribution from temperature  $T_A$  to  $T_B$ , or vice versa. In this special case,  $w_A + w_B = 0$ , and  $P_{\text{acc}} = 1$ .)

Even when RENS is used with stochastic equations of motion, such as Langevin dynamics or the Andersen thermostat (see *SI Appendix*), it is useful to include the scaling term in Eq. 13, as this term dynamically adjusts the momenta in response to the changing temperature  $T_{\lambda}$ .

### Efficiency Considerations and Numerical Results

While our method is valid for arbitrary choice of switching time,  $\tau$ , it is instructive to consider two limiting cases. In the sudden limit,  $\tau = 0$ , we have  $w = \Delta h$  (Eqs. 1, 3, 4), as noted by Wyczalkowski and Pappu (25).<sup>†</sup> In this case, RENS reduces to REM, and if there

<sup>†</sup> If the scaling term of Eq. 13 is included, then only the potential energy contributes to  $w$ , exactly as with REM when momenta are rescaled; see equations 12 and 15 of ref. 24.



**Fig. 2.** Mock-up of a rough potential energy landscape adapted from ref. 3, Chapter 14. An asterisk marks the fourth well ( $x \geq 1.25$ ). Ordinary replica exchange works well for 1 particle, but encounters difficulties when  $n_p = 10$ .

is little overlap between  $p_A^{\text{eq}}$  and  $p_B^{\text{eq}}$ , then  $P_{\text{acc}} \ll 1$ . For a properly thermalized system in the opposite quasi-static limit,  $\tau \rightarrow \infty$ , the system evolves reversibly as  $\lambda$  is varied infinitely slowly; the reduced work is the corresponding reduced free energy difference,

$$w_A = \Delta f = f_B - f_A = -w_B, \quad [15]$$

where  $f_i = -\ln \int dx e^{-h_i}$ ; hence  $w = 0$  and  $P_{\text{acc}} = 1$  (Eq. 2). Thus we can manipulate the acceptance probability  $P_{\text{acc}}$  by adjusting the switching time  $\tau$ . Generically, we expect that the more slowly we perform the work simulation, the greater the probability to accept the configuration swap (see Fig. 4). This expectation implies a computational tradeoff—what is the optimal value of  $\tau$ ?—which we now address with a simple analysis.

We consider  $M$  replicas and assume for specificity that we are mainly interested in sampling from one of them, which we denote as the primary replica; the remaining replicas serve only to enhance sampling in the primary replica. (A similar analysis can be carried out if we are equally interested in sampling all  $M$  ensembles.) The term output trajectory will denote the trajectory obtained by concatenating the sampling intervals generated in the primary replica after discarding the work intervals. We let  $\bar{\tau}_{\text{eq}}$  denote the average duration of a sampling interval. The output trajectory samples the equilibrium distribution of interest and is in effect the end product of our method. Let  $t_c$  denote a characteristic correlation time associated with this output trajectory, and define

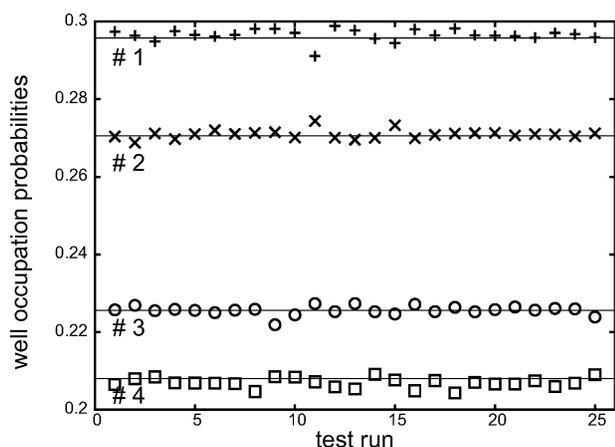
$$X \equiv \tau/\bar{\tau}_{\text{eq}}. \quad [16]$$

In terms of these quantities, the sample cost

$$t^* = (1 + X)Mt_c \quad [17]$$

is a measure of the total computational cost, summed over all  $M$  replicas (26), of producing a single, statistically independent sample in the primary replica. The factor  $(1 + X)$  accounts for the overhead cost of the work intervals: For every unit of sampling time,  $X$  units of time were devoted to the discarded work simulations. The sample cost provides a figure of merit; the smaller the value of  $t^*$ , the more efficiently we are using the computational resources. While the correlation time  $t_c$  generally decreases with increasing  $M$  or  $\tau$ —through the randomizing effect of successful replica exchanges—in Eq. 17 this trend competes with the overhead factors  $M$  and  $1 + X$ .

To investigate these issues, we simulated a model system of  $n_p = 10$  particles, moving independently in the potential shown in Fig. 2. We took  $M = 2$  replicas, at  $T_A = 0.30$  and  $T_B = 2.0$  (arbitrary units). In the primary replica at  $T_A = 0.30$ , sampling is hindered by the barriers separating the local minima of  $U(x)$ ; whereas at  $T_B = 2.0$  the particles are able to jump from well to well.



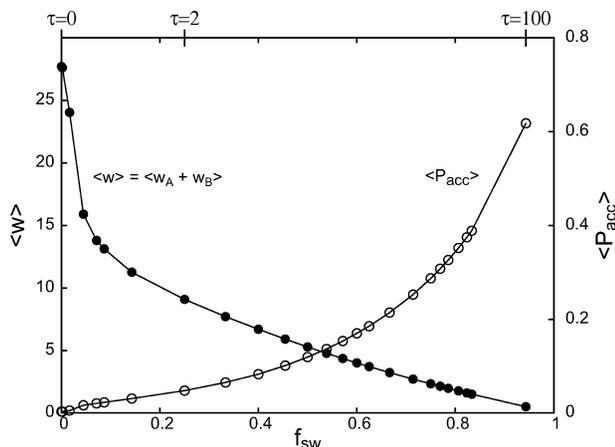
**Fig. 3.** Observed occupations of wells 1–4, obtained by following a single tagged particle in the output trajectory, for each of the 25 test runs. The solid horizontal lines are exact values determined by integration of the single-particle Boltzmann distribution.

We performed MD simulations by using Eq. 13 in combination with an Andersen thermostat (27) (see *SI Appendix*).

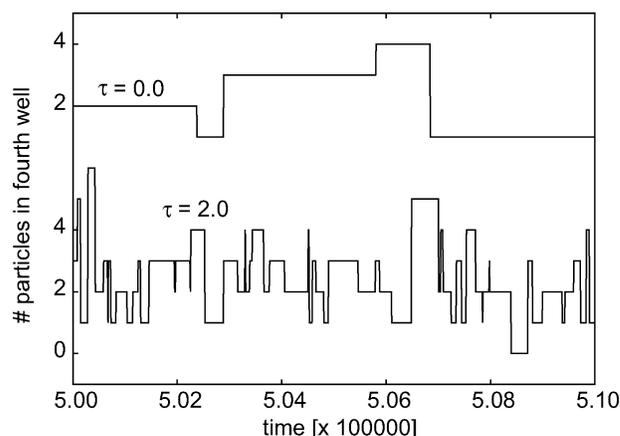
While REM performs well for a single particle in this potential (3), when  $n_p = 10$  it encounters difficulties due to poor phase space overlap, as can be understood using an argument of Kofke's (28). Typically, in  $\mathcal{R}_A$  each particle is found near a local minimum of  $U(x)$ , whereas in  $\mathcal{R}_B$  they are distributed more uniformly. Thus a configuration swap is likely to be accepted only if all particles in  $\mathcal{R}_B$  are found very near to the minima of  $U(x)$ , which is unlikely when  $n_p \gg 1$ . With RENS, the work simulation in  $\mathcal{R}_B$  increases the swap acceptability by shepherding the particles closer to the minima of  $U(x)$ .

When simulating this system using RENS, the replicas “toggle” between sampling and work intervals (Fig. 1). We implemented this as follows. During an interval of sampling, a work simulation was initiated at random, with an attempt rate  $r = 0.166$ . Once initiated, the work interval lasted for the prescribed switching time  $\tau$ , after which the replicas reverted to sampling, and so on. Thus the average duration of a sampling interval was  $\bar{\tau}_{\text{eq}} = 1/r \approx 6.0$ , which is roughly 3 times the relaxation rate within one of the local wells of  $U$ .

With these parameters, we performed 25 test runs, with  $\tau$  ranging from 0 to 100. To establish proof of principle, we tabulated



**Fig. 4.** Average reduced work (filled circles) and observed acceptance frequency (open circles) as functions of the fraction of simulation time devoted to work intervals.  $\langle P_{\text{acc}} \rangle$  is defined as the fraction of attempted replica exchanges that were accepted in a given test run. At  $f_{\text{sw}} = 0$ , corresponding to ordinary (instantaneous) replica exchange,  $\langle P_{\text{acc}} \rangle \approx 0.003$ .



**Fig. 5.**  $n_4(t)$  is plotted over an interval of time for output trajectories obtained with REM (upper trace, shifted for clarity) and RENS (lower trace). The 2 traces represent roughly the same number of attempted replica exchanges ( $\approx 1700$ ) but reveal substantially different acceptance rates.

empirical occupation probabilities for the 4 wells by following a tagged particle in the output trajectory. For each test run, we found that the relative amount of time the particle spent in each well was in agreement with the equilibrium distribution, within statistical error (Fig. 3).

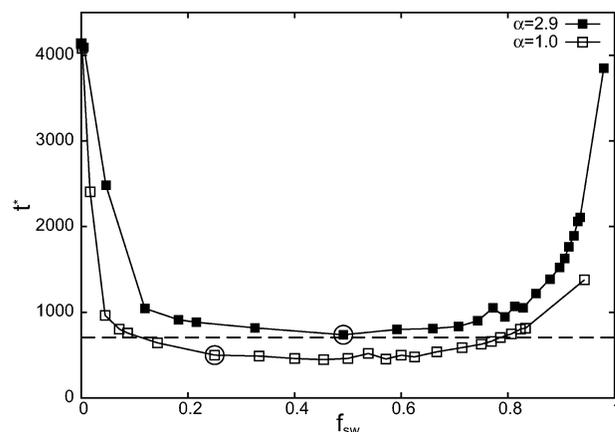
For the same set of test runs, in Fig. 4 we plot the observed swap acceptance frequency and average reduced work as functions of the fraction of simulation time devoted to the work intervals,

$$f_{\text{sw}} = \frac{X}{1+X} = \frac{\tau}{\bar{\tau}_{\text{eq}} + \tau}. \quad [18]$$

As anticipated, with increasing  $f_{\text{sw}}$  (or  $\tau$ ) we approach the reversible limit of  $w = 0$  and  $P_{\text{acc}} = 1$  (Eq. 15).

To illustrate the accelerated sampling achieved with our method, we considered  $n_4(t)$ , the number of particles found in the fourth well of  $U(x)$  at time  $t$  of the output trajectory. Fig. 5 shows  $n_4(t)$  for segments of the  $\tau = 0.0$  and  $\tau = 2.0$  test runs. For the relatively modest cost of setting aside 25% of the simulation time to the work intervals, transitions into and out of the fourth well are greatly facilitated.

Next, for each test run we used block-averaging (29) to evaluate a correlation time  $t_c = (1/\sigma^2) \int_{-\infty}^{+\infty} dt c(t)$ , where  $\sigma^2$  and  $c(t)$  are the variance and autocorrelation of  $n_4(t)$ . In Fig. 6, we plot the sample



**Fig. 6.** Sample cost  $t^*$  plotted against  $f_{\text{sw}}$  both with and without adjustment for increased relative cost of work simulations (see *Efficiency Considerations and Numerical Results*). The circles identify the run at  $\tau = 2.0$ . The dashed line is the sample cost of REM with  $M = 4$  replicas.

cost  $t^*$  for each test run (empty squares). At  $f_{sw} = 0$  (that is, when using REM), this cost is high,  $t^* > 4,000$ ; few swaps are accepted, and particles are trapped in the fourth well for long times. As we increase  $f_{sw}$ , the sample cost drops significantly, reaching a broad minimum  $t^* \sim 450 - 500$  for  $f_{sw} \sim 0.2 - 0.6$ ; here the allocation of CPU time to work intervals delivers a clear benefit. For  $f_{sw} > 0.6$ , we enter a regime of diminishing returns:  $\langle P_{acc} \rangle$  continues to increase with  $\tau$  (Fig. 4), but not enough to justify the expense of increasingly long work simulations.

To this point, we have neglected the computational cost of the “acceptance/rejection” step itself, as well as that of the possible subsequent exchange of configurations (which typically involves communication between different processors). Moreover, we have assumed identical costs, per unit simulation time, for the work and the sampling intervals. It is easy enough to drop the latter assumption: We replace  $X$  by  $\alpha X$  in Eqs. 17 and 18, where  $\alpha$  is the observed CPU cost of generating a work simulation relative to that of a sampling trajectory of equal duration. In our test runs, we found  $\alpha = 2.9$ , and the points shown as filled squares in Fig. 6 have been adjusted for this value. (If our model had included particle–particle interactions,  $\alpha$  would have been closer to unity.)

Whether or not we make the adjustment to account for  $\alpha \neq 1$ , Fig. 6 clearly shows that for a fixed set of replicas, it can be highly advantageous to use nonequilibrium work simulations to generate attempted configuration swaps. The benefits of increased acceptance substantially outweigh the overhead cost of generating the trial configurations.

With RENS, we improve efficiency by tuning the switching time,  $\tau$ , as in Fig. 6. With REM, one can instead vary the number of replicas. To compare these 2 options, we performed test runs of REM ( $\tau = 0$ ) at  $M = 2, 3, \dots, 11$ . (In each run, we set  $T_1 = 0.30$  and  $T_M = 2.0$ , with intermediate replicas spaced evenly in  $T^{-1}$ .) Among these runs, the smallest sample cost,  $t^* = 706$ , was achieved with  $M = 4$  replicas, and is shown as a straight line in Fig. 6. This value is comparable to the optimal sample cost achieved with RENS using  $M = 2$ . Thus for this simple system, RENS is able to match the efficiency of REM with fewer replicas.

## Discussion

When applying REM to a problem of interest, the phase space overlap requirement dictates a minimum number of replicas,  $M^*$ , needed to achieve a reasonable swap acceptance frequency. With RENS, the work simulations have the effect of increasing phase space overlap, thus allowing for fewer replicas,  $M < M^*$ . There are several reasons why one might wish to exploit this flexibility.

- (i) Most obviously, if we perform simulations using a cluster of  $P$  processors, then RENS allows us to assign 1 replica per processor—the easiest and most natural (and traditional) allocation—even if  $P < M^*$ .
- (ii) It is often useful to picture replica exchange as a diffusion process in which trajectories hop randomly along the chain  $\mathcal{R}_1, \dots, \mathcal{R}_M$ . In this picture,  $\sim M^2$  successful swaps are needed for a given trajectory to accomplish an entire transit between  $\mathcal{R}_1$  to  $\mathcal{R}_M$ . Thus using fewer replicas (with

RENS) can significantly reduce the cost of interprocessor communication associated with attempted configuration swaps.

- (iii) REM is often implemented synchronously: Swaps are attempted only after every replica completes a predetermined duration of equilibrium sampling. With 1 replica per processor, this can be highly inefficient, limited by the speed of the slowest processor. RENS lends itself naturally to asynchronous implementation. A master process initiates work simulations in a randomly chosen replica pair, whereas the remaining replicas, unaffected, continue sampling.
- (iv) With any replica exchange strategy, there are parameters we adjust to optimize efficiency, such as the number of replicas,  $M$ , and the choice of intermediate temperatures or Hamiltonians. It is potentially very useful to improve efficiency adaptively, during the actual production run (30). RENS offers a relatively painless way to accomplish this, namely by adjusting the durations of the work simulations. For example, if it is observed that a low  $P_{acc}$  between  $\mathcal{R}_n$  and  $\mathcal{R}_{n+1}$  poses a bottleneck for efficient sampling, then the switching time for that replica pair,  $\tau_{n,n+1}$ , can be increased.
- (v) To this point, we have treated the data generated during the work simulations as “junk” to be discarded after the attempted configuration swap. However, by a trick of statistical reweighting one can scavenge equilibrium information from such nonequilibrium trajectories (see equation 4 of ref. 31). The work simulations themselves can then contribute to the equilibrium sampling in each replica, thus increasing the efficiency of RENS. For Monte Carlo sampling, Frenkel (32) has developed an analogous, thrifty algorithm that relies on the “waste-recycling” of otherwise rejected trial moves.

These considerations suggest that RENS offers a flexible, efficient and useful sampling strategy. We expect that it can further be enhanced through combination with other approaches such as solute tempering (7) and generalized effective potentials (11), or by the use of large time steps (33) or artificial flow fields (34) during the work simulations.

The simple model we have borrowed (3) is well suited as an initial test case of our method: It exhibits the difficulties faced by REM for large ( $n_p = 10$ ) systems, its equilibrium properties can be evaluated exactly (Fig. 3), and its efficiency can be computed with high statistical accuracy for many values of  $\tau$  (Fig. 6). Further assessments of our method will come both from applications to problems of genuine physical interest and from analytical and semianalytical treatments that have provided useful insight into the performance of replica exchange strategies (17, 28, 30, 35–37).

**ACKNOWLEDGMENTS.** We gratefully acknowledge useful discussions with Jordan Horowitz and Suri Vaikuntanathan. This material is based upon work supported by the National Science Foundation Grant CHE-0841557 and by the University of Maryland.

1. Swendsen RH, Wang JS (1986) Replica Monte-Carlo simulation of spin-glasses. *Phys Rev Lett* 57:2607–2609.
2. Earl DJ, Deem MW (2005) Parallel tempering: Theory, applications, and new perspectives. *Phys Chem Chem Phys* 7:3910–3916.
3. Frenkel D, Smit B (2002) *Understanding Mol Simul* (Academic, San Diego).
4. Hukushima K, Nemoto K (1996) Exchange Monte Carlo method and application to spin glass simulations. *J Phys Soc Jpn* 65:1604–1608.
5. Gnanakaran S, Nymeyer H, Portman J, Sanbonmatsu KY, Garcia AE (2003) Peptide folding simulations. *Curr Opin Struct Biol* 13:168–174.
6. Fukunishi H, Watanabe O, Takada S (2002) On the Hamiltonian replica exchange method for efficient sampling of biomolecular systems: Application to protein structure prediction. *J Chem Phys* 116:9058–9067.
7. Liu P, Kim B, Friesner RA, Berne BJ (2005) Replica exchange with solute tempering: A method for sampling biological systems in explicit water. *Proc Natl Acad Sci USA* 102:13749–13754.
8. Liu P, Huang X, Zhou R, Berne BJ (2006) Hydrophobic aided replica exchange: An efficient algorithm for protein folding in explicit solvent. *J Phys Chem B* 38:19018–19022.
9. Hansmann UHE (1997) Parallel tempering algorithm for conformational studies of biological molecules. *Chem Phys Lett* 281:140–150.
10. Whitfield TW, Bu L, Straub JE (2002) Generalized parallel sampling. *Physica A* 305:157–171.
11. Jang S, Shin S, Pak Y (2003) Replica-exchange method using the generalized effective potential. *Phys Rev Lett* 91:058305.
12. Calvo F, Doye JPK (2000) Entropic tempering: A method for overcoming quasiergodicity in simulation. *Phys Rev E* 63:010902.
13. Faller R, Yan Q, de Pablo JJ (2002) Multicanonical parallel tempering. *J Chem Phys* 116:5419–5423.
14. Mitsutake A, Sugita Y, Okamoto Y (2001) Generalized-ensemble algorithms for molecular simulations of biopolymers. *Biopolymers* 60:96–123.
15. Fenwick MK, Escobedo FA (2003) Expanded ensemble and replica exchange methods for simulation of protein-like systems. *J Chem Phys* 119:11998–12010.
16. Rick SW (2007) Replica exchange with dynamical scaling. *J Chem Phys* 16:054102.
17. Park S (2008) Comparison of the serial and parallel algorithms of generalized ensemble simulations: An analytical approach. *Phys Rev E* 77:016709.

18. Crooks GE (1998) Nonequilibrium measurements of free energy differences for microscopically reversible Markovian systems. *J Stat Phys* 90:1481–1487.
19. Opps S, Schofield J (2001) Extended state-space Monte Carlo methods. *Phys Rev E* 63:056701.
20. Brown S, Head-Gordon T (2003) Cool-walking: A new Markov chain Monte Carlo sampling method. *J Comput Chem* 24:68–76.
21. Stern H (2007) A molecular simulation with variable protonation states at constant pH. *J Chem Phys* 126:164112.
22. Chandler D (1987) *Introduction to Modern Statistical Mechanics* (Wiley, New York).
23. Oberhofer H, Dellago C, Geissler PL (2005) Biased sampling of nonequilibrium trajectories: Can fast switching simulations outperform conventional free energy calculation methods? *J Phys Chem B* 109:6902–6915.
24. Sugita Y, Okamoto Y (1999) Replica-exchange molecular dynamics method for protein folding. *Chem Phys Lett* 314:141–151.
25. Wyczalkowski MA, Pappu RV (2008) Satisfying the fluctuation theorem in free-energy calculations with Hamiltonian replica exchange. *Phys Rev E* 77:026104.
26. Zuckerman DM, Lyman E (2006) A second look at canonical sampling of biomolecules using replica exchange simulation. *J Chem Theory Comput* 2:1200–1202.
27. Andersen HC (1980) Molecular dynamics simulations at constant pressure and/or temperature. *J Chem Phys* 72:2384–2393.
28. Kofke DA (2002) On the acceptance probability of replica-exchange Monte Carlo trials. *J Chem Phys* 117:6911–6914.
29. Flyvbjerg H, Petersen HG (1989) Error estimates on averages of correlated data. *J Chem Phys* 91:461–466.
30. Trebst S, Troyer M, Hansmann UHE (2006) Optimized parallel tempering simulations of proteins. *J Chem Phys* 124:174903.
31. Hummer G, Szabo A (2001) Free energy reconstruction from nonequilibrium single-molecule pulling experiments. *Proc Natl Acad Sci USA* 98:3658–3661.
32. Frenkel D (2004) Speed-up of Monte Carlo simulations by sampling of rejected states. *Proc Natl Acad Sci USA* 101:17571–17575.
33. Lechner W, Oberhofer H, Dellago C, Geissler PL (2006) Equilibrium free energies from fast-switching trajectories with large time steps. *J Chem Phys* 124:044113.
34. Vaikuntanathan S, Jarzynski C (2008) Escorted free energy simulations: Improving convergence by reducing dissipation. *Phys Rev Lett* 100:190601.
35. Zheng W, Andrec M, Gallicchio E, Levy RM (2007) Simulating replica exchange simulation of protein folding with a kinetic network model. *Proc Natl Acad Sci USA* 104:15340–15345.
36. Sindhikara D, Meng Y, Roitberg AE (2008) Exchange frequency in replica exchange molecular dynamics. *J Chem Phys* 128:024103.
37. Nymeyer H (2008) How efficient is replica exchange molecular dynamics? An analytic approach. *J Chem Theory Comput* 4:626–636.