Efficient transition path sampling for nonequilibrium stochastic dynamics

Gavin E. Crooks and David Chandler
Department of Chemistry, University of California at Berkeley, Berkeley, California 94720
(Received 11 December 2000; published 19 July 2001)

The transition path sampling methodology is adapted to the efficient sampling of large fluctuations in nonequilibrium systems evolving according to Langevin's equations of motion. This technique is used to simulate the behavior of the bistable Maier-Stein system at noise intensities much lower than those previously possible.

DOI: 10.1103/PhysRevE.64.026109 PACS number(s): 05.10.−a, 05.40.−a, 02.70.Rr

Recently, approximate theories have been developed that describe large, rare fluctuations in systems with Langevin dynamics that have been driven from equilibrium by a time-dependent or nongradient (i.e., not the gradient of a potential) force field [1–8]. These theories are only good approximations in the zero noise limit, and computer simulations are needed to explore the behavior of the system and the accuracy of the approximations at finite noise intensities. However, the straightforward simulation of the dynamics is inherently inefficient, since the majority of the computation time is taken watching small, uninteresting fluctuations about the stable states, rather than the interesting and rare excursions away from those states. Transition path sampling [9–13] has been developed as a Monte Carlo importance sampling of trajectories that can efficiently harvest rare transitions between stable or metastable states in equilibrium systems. Only trajectories that undergo the desired transition in a short time are sampled. In this paper, transition path sampling is adapted to nonequilibrium, dissipative, stochastic dynamics. The principle innovation is the development of a new algorithm to generate trial trajectories.

The high-friction limit of Langevin's equations of motion describe overdamped Brownian motion in a force field

$$\dot{\mathbf{x}}_i = \mathbf{f}_i(\mathbf{x},t) + \boldsymbol{\xi}_i(t). \tag{1}$$

The state of the system is specified by the vector $\mathbf{x}$. The system is subjected to a systematic force $\mathbf{F}(\mathbf{x},t)$, and a stochastic force $\boldsymbol{\xi}(t)$, resulting from δ-function-correlated white noise with variance $\boldsymbol{\epsilon}$

$$\langle \boldsymbol{\xi}_i(t) \rangle = \mathbf{0}, \quad \langle \boldsymbol{\xi}_i(t) \cdot \boldsymbol{\xi}_j(0) \rangle = \boldsymbol{\epsilon}\, \delta_{ij} \delta t. \tag{2}$$

In this paper we are interested in systems that are not in equilibrium, either because the force field $\mathbf{F}(\mathbf{x},t)$ is time dependent, or because it is nongradient. Dynamics of this class can model a large range of interesting problems, including chemical reactions [14], thermal ratchets [15], and computer networks [16]. As a particular example, we adopt the following two-dimensional system $[\mathbf{x}=(x,y)]$ proposed by Maier and Stein [4], namely,

$$\mathbf{F}(x,y) = (x - x^3 - \alpha x y^2, -\mu y(1+x^2)). \tag{3}$$

This field is not the gradient of a potential energy unless $\alpha = \mu$. The potential-energy surface for the gradient field $\alpha = \mu = 1$, is shown in Fig. 1, which should serve to orient the reader. Of primary interest are the rare transitions between the stable states. For weak noise, almost all transitions closely follow the optimal trajectory, the most probable exit path (MPEP) (see Fig. 2). There are extensive theoretical predictions [7,19] and simulation results [17–19] for this system against which the algorithms developed in this paper can be tested.

Exploring the weak noise behavior of these systems has pushed conventional simulation techniques to their limits, even for the very simple, low-dimensional dynamics so far considered. A single, very long trajectory is generated, and one is obliged to wait for interesting events to occur. Therefore, it is desirable to construct a simulation that runs as quickly as possible. The very fastest simulations utilize an analog electronic model of the system of interest, which is then driven by a zero-mean quasiwhite-noise generator [17,20]. Another approach has been to optimize the speed of the pseudorandom number generator [21], since this computational effort often dominates the total simulation time for simple stochastic dynamics. However, such simulations cannot incorporate any importance sampling of interesting events. The total simulation time necessarily increases with the rarity of the event under study, which typically increases exponentially as the noise intensity decreases.

The transition path sampling methodology has been developed to efficiently sample rare events in equilibrium sys-

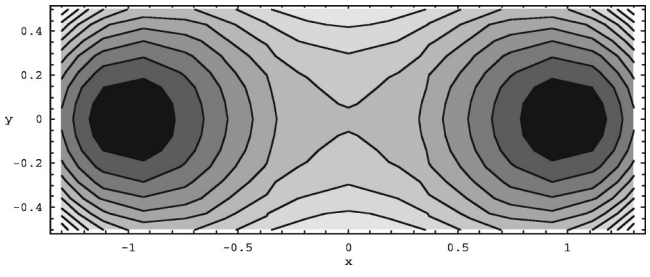

FIG. 1. The potential-energy surface of the Maier-Stein system, Eq. (3), with $\alpha=1$ and $\mu=1$. Darker shading indicates lower energies. Note the stable states at $(\pm 1,0)$, the transition state at $(0,0)$, and the surface dividing the stable states (the separatrix) at $x=0$. These general features persist for the other values of the parameters used in this paper, although the force field is no longer the gradient of a potential energy. For this equilibrium system the most probable path connecting the stable states (and therefore the path that dominates transitions in the weak noise limit) runs directly along the x axis.

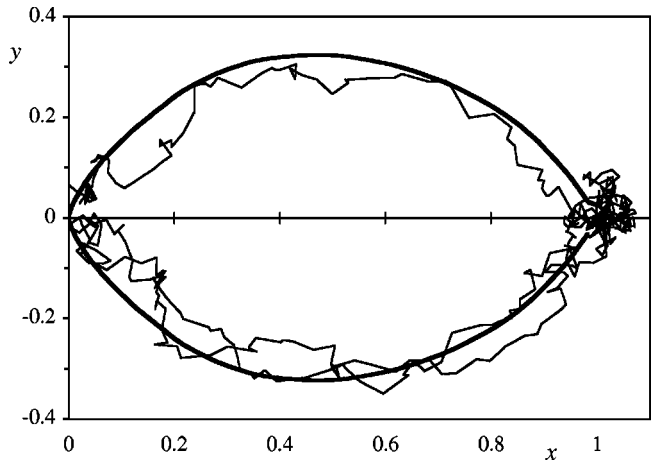

FIG. 2. A representative sample of exit paths (thin lines) for the Maier-Stein system with $\alpha=6.67$, $\mu=1.0$, and $\epsilon=0.005$, generated from a path sampling simulation. These trajectories cluster around the most probable exit paths (MPEP's) (thick lines). The MPEP's were calculated via simulated annealing of the transition paths.

tems. The main innovation is to sample path space directly using a Monte Carlo algorithm; instead of passively waiting for the dynamics to generate an interesting trajectory, a Markov chain linking different trajectories is constructed, each member of which incorporates the event of interest. This path ensemble Monte Carlo is completely analogous to conventional Monte Carlo algorithms acting on configurational ensembles. A trial trajectory is generated by a small, random change in the previous trajectory; it is immediately rejected if the desired boundary conditions are not met (typically that the path starts in region A and ends in region B), and it is accepted with a probability that generates the correct distributions of trajectories.

Unfortunately, the standard methods for efficiently sampling path space cannot be directly applied to nonequilibrium Langevin dynamics. Perhaps the most obvious method for generating new trajectories in a stochastic dynamics is the local algorithm [22]. The path is represented by a chain of states, $\mathbf{x}=[x(0),x(1),\ldots,x(L)]$, and the probability of the path, $\mathcal{P}[\mathbf{x}]$, is written as a product of single time step transition probabilities, $P[x(t)\rightarrow x(t+1)]$;

$$\mathcal{P}[\mathbf{x}]=\rho(x(0))\prod_{t=0}^{L-1} P[x(t)\rightarrow x(t+1)]. \tag{4}$$

Here, $\rho(x(0))$ is the probability of the initial state of the path. A trial trajectory $\mathbf{x}'$ is generated by changing the configuration at a single time slice, it is immediately rejected if the desired boundary conditions are not fulfilled, and it is accepted with the Metropolis probability,

$$P_{\mathrm{acc}}(\mathbf{x}\rightarrow\mathbf{x}')=\min\left[1,\frac{\mathcal{P}[\mathbf{x}']P_{\mathrm{gen}}(\mathbf{x}'\rightarrow\mathbf{x})}{\mathcal{P}[\mathbf{x}]P_{\mathrm{gen}}(\mathbf{x}\rightarrow\mathbf{x}')}\right], \tag{5}$$

which ensures a correctly weighted ensemble of paths. Here, P_{gen} is the probability of generating the trial configurations $\mathbf{x}'$.

Although effective [23], the local algorithm suffers from several deficiencies, the most serious of which is that the relaxation time of the path scales as L^3, where L is the total number of time steps [24]. The Maier-Stein system requires on the order of thousands of time steps to make the large, rare excursions away from the stable states that are of interest, which renders the local algorithm impractical.

Several simple and efficient methods of generating trial trajectories (shooting and shifting [10]) have been developed for equilibrium dynamics. Unfortunately, they are not directly applicable to nonequilibrium dynamics, since they assume a knowledge of the initial state probability. Statistical mechanics provides simple expressions for equilibrium probabilities, but no such simple expression exists for nonequilibrium steady states.

Fortunately, there is an alternative representation of a stochastic path that admits a simple and efficient path sampling algorithm. A stochastic trajectory can be defined by the chain of states that the system visits, but it can also be represented by the initial state and the set of random numbers, the noise history, that was used to generate the trajectory. The probability of the path can then be written as

$$\mathcal{P}[\mathbf{x}]=\rho(x(0))\prod_{t=0}^{L-1}\frac{1}{\sqrt{2\pi\epsilon}}\exp\{-\xi(t)^2/2\epsilon\}, \tag{6}$$

where each ξ is a Gaussian random number of zero mean and ϵ variance. This is a convenient representation, since we normally generate a stochastic trajectory from a set of random numbers, and not random numbers from a trajectory.

Suppose that we have the initial state and the noise history of a relatively short path that undergoes the rare event in which we are interested. (We will return to the problem of creating this initial path shortly.) A trial path can be created by replacing the noise at a randomly chosen time step with a new set of Gaussian random numbers. This trial trajectory is accepted as a new member of the Markov chain of paths if it still undergoes the event of interest. The noise histories of the accepted paths become correlated in time due to this constraint. Since high-friction Langevin dynamics is highly dissipative, nearby trajectories converge rapidly, and a small change in the noise generally produces a small change in the trajectory. Therefore, most trial trajectories are accepted. Only rarely does the change in the noise produce a path that no longer executes the event under study.

This noise sampling algorithm does not suffer from the poor scaling of relaxation time with path length that renders the local algorithm impractical, since a local move in noise space induces a small but nonlocal move in path space. Consider, for a moment, an unconstrained path. Then every change in the noise history is accepted. After $\mathcal{O}(L)$ moves, almost all of the random numbers used to generate the path will have been replaced, and an entirely new path will have been generated, one that is uncorrelated with the previous path. Generating a trial trajectory from the noise history is very fast since the random numbers needed to generate the path have already been created and stored. The amount of information that must be stored scales with the number of

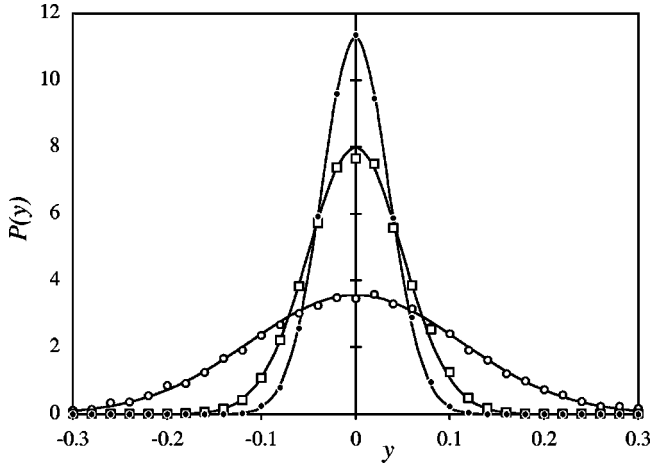

FIG. 3. Exit location distributions for the Maier-Stein system with $\alpha=6.67$, $\mu=2.0$, and $\epsilon=0.05$ (○), 0.01 (□), or 0.005 (●). Symbols are averages from a path sampling simulation (8192 samples) and lines are the theoretical predictions, $P(y) \propto \exp(-2y^2/\epsilon)$ (Refs. [7,18]).

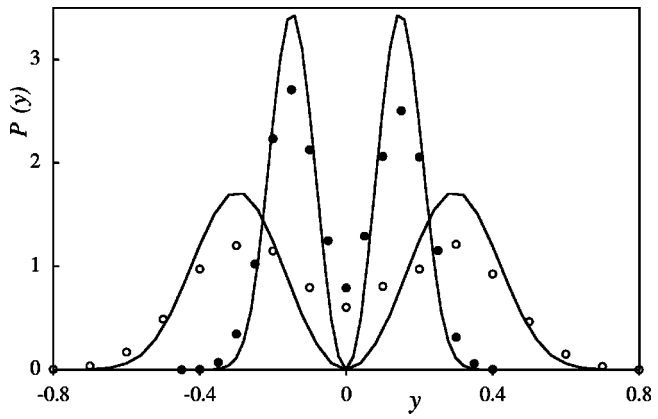

FIG. 4. Exit location distributions for the Maier-Stein system with $\alpha=10$, $\mu=0.67$, and $\epsilon=0.04$ (○) or 0.005 (●). Symbols are averages from path sampling simulations (8192 samples) and lines are the symmetrized Weibull distribution, $P(y)=|y|^{(2/\mu)-1} \times \exp(-|y/A|^{2/\mu}/\epsilon)$ (Refs. [7,19]). The parameter $A\approx 1$ is determined from the behavior of the most probable exit path near the saddle point $y=\pm Ax^{\mu}$.

time steps, but this is a trivial amount of memory for the low-dimensional systems considered here.

Unlike the local path sampling algorithm, sampling the noise history allows a choice of methods for integrating the dynamics. For compatibility with previous digital simulations [25] we used the second-order Runge-Kutta method. Compared to a simple finite difference equation, this integrator is more stable and allows longer time steps. The maximum total time of the trajectories was $\tau=16$, with a time step of $\Delta t=1/512$, for a total of 8192 time slices. This time step is small enough to ensure better than 90% trial move acceptance rate at any one time slice for the noise intensities studied. It requires approximately 1 s of CPU time to generate a statistically independent path. Unlike simulations (digital or analog) without importance sampling, these simulation times are largely independent of the noise intensity. (There is a logarithmic increase of the transition time with decreasing noise [21], which would eventually require longer paths.) The smallest noise intensities used to generate trajectories in this paper are typically an order of magnitude smaller than the smallest values that can be practically studied with an analog simulation.

An initial path can be generated using the following procedure. The initial point of the path is fixed, an entirely random initial noise history is generated, and the end point of the corresponding trajectory is computed. A small change is then made in the noise history, and this move is accepted only if the new end point of the trajectory is closer to the desired final region than the previous path. In this manner the final point of the trajectory can be dragged into the desired region, and a valid initial path obtained. It is then necessary to relax this initial path so that the correct transition path ensemble is generated.

A separate Monte Carlo move is used to sample the initial configuration. A trial configuration is selected from an entirely separate, nonpath sampling simulation that has been relaxed to the steady state. A long trajectory ensures that the final state is largely insensitive to the initial state, and therefore that this trial move is often accepted, even if the change in the initial state is large. Alternatively, the initial state can simply be fixed at some representative point of the steady-state ensemble. The simulation results will not be altered if the trajectory is significantly longer than the relaxation time of the system.

Figure 2 shows several representative trajectories that carry the system from the stable region around (1,0) to the separatrix at $x=0$. For $\mu=1.0$ and $\alpha>4$ the set of exit paths bifurcates [3,5]. Instead of following the x axis to the transition state, trajectories instead make large excursions away from the axis, and approach the transition state from the top- or bottom-right quadrants. For weak noise, a single path sampling simulation of this system would lock into either the top or bottom set of trajectories and equilibration in path space would be very slow. This is analogous to the behavior of glasses and procedures developed to study such systems could be used to aid path sampling. For the current system this is not an issue, since this bifurcation is known to exist and the paths are symmetric about the x axis.

The finite noise trajectories cluster around the most probable exit paths, which are the transition paths in the zero noise limit. These can be calculated directly from theory, but here they were generated via gradually annealing the system to very weak noise intensities ($\epsilon=10^{-5}$), which is analogous to the path quench used in Ref. [10]. The acceptance rate for parts of the path approached 0% at $\epsilon\approx 0.0005$, effectively freezing the trajectory in place. This represents the lower noise limit for the current implementation. To study weaker noise it would be necessary to use smaller time steps (which would increase the total number of time slices) or smaller changes in the noise.

There are a variety of predictions regarding the distribution of exit locations [7,19], the point on the y axis where the

transition path first crosses from one stable state to the other. Figure 3 shows path sampling simulation results and theoretical predictions for parameters where it is known that the theoretical predictions are accurate. Excellent agreement is observed, validating the path sampling algorithm. Finally, Fig. 4 displays exit location distributions and theoretical predictions for parameters where the agreement between the two is known to be poor [19]. Path sampling was used to study the exit location distribution at a noise intensity approximately 10 times smaller than previously possible. Even at this very low noise intensity, the agreement between theory and simulation remains unsatisfactory.

It is a pleasure to thank D. G. Luchinsky, R. Mannella, and P. V. E. McClintock for their helpful correspondence, and for providing example code for a conventional simulation of the Maier-Stein system. This work was initiated with support from the National Science Foundation under Grant No. CHE-0078458, and completed with support from the U.S. Department of Energy, Basic Energy Sciences Grant No. FDDE-FG03-99ER14987.

[1] M.I. Dykman *et al.*, Phys. Rev. Lett. **68**, 2718 (1992).
[2] R.S. Maier and D.L. Stein, Phys. Rev. Lett. **69**, 3691 (1992).
[3] R.S. Maier and D.L. Stein, Phys. Rev. Lett. **71**, 1783 (1993).
[4] R.S. Maier and D.L. Stein, Phys. Rev. E **48**, 931 (1993).
[5] R.S. Maier and D.L. Stein, J. Stat. Phys. **83**, 291 (1996).
[6] M.I. Dykman, D.G. Luchinsky, P.V.E. McClintock, and V.N. Smelyanskiy, Phys. Rev. Lett. **77**, 5229 (1996).
[7] R.S. Maier and D.L. Stein, SIAM (Soc. Ind. Appl. Math.) J. Appl. Math. **57**, 752 (1997).
[8] V.N. Smelyanskiy, M.I. Dykman, and R.S. Maier, Phys. Rev. E **55**, 2369 (1997).
[9] C. Dellago, P.G. Bolhuis, F.S. Csajka, and D. Chandler, J. Chem. Phys. **108**, 1964 (1998).
[10] C. Dellago, P.G. Bolhuis, and D. Chandler, J. Chem. Phys. **108**, 9236 (1998).
[11] P.G. Bolhuis, C. Dellago, and D. Chandler, Faraday Discuss. Chem. Soc. **110**, 421 (1998).
[12] D. Chandler, in *Computer Simulation of Rare Events and Dynamics of Classical and Quantum Condensed-Phase Systems–Classical and Quantum Dynamics in Condensed Phase Simulations*, edited by J. B. Berne, G. Ciccotti, and D. F. Coker (World Scientific, Singapore, 1998), pp. 51–66.
[13] C. Dellago, P.G. Bolhuis, and D. Chandler, J. Chem. Phys. **110**, 6617 (1999).
[14] M.I. Dykman, E. Mori, J. Ross, and P.M. Hunt, J. Chem. Phys. **100**, 5735 (1994).
[15] M.O. Magnasco, Phys. Rev. Lett. **71**, 1477 (1993).
[16] J.-T. Lim and S.M. Meerkov, IEEE Trans. Commun. **35**, 1278 (1987).
[17] D.G. Luchinsky and P.V.E. McClintock, Nature (London) **389**, 463 (1997).
[18] D.G. Luchinsky *et al.*, Phys. Rev. Lett. **79**, 3109 (1997).
[19] D.G. Luchinsky *et al.*, Phys. Rev. Lett. **82**, 1806 (1999).
[20] D.G. Luchinsky, P.V.E. McClintock, and M.I. Dykman, Rep. Prog. Phys. **61**, 889 (1998).
[21] R. Mannella, in *Supercomputation in Nonlinear and Disordered Systems*, edited by L. Vàzquez, F. Tirando, and I. Martin (World Scientific, Singapore, 1997), pp. 100–130.
[22] L.R. Pratt, J. Chem. Phys. **85**, 5045 (1986).
[23] M. Paniconi and M.F. Zimmer, Phys. Rev. E **59**, 1563 (1999).
[24] D.M. Ceperley, Rev. Mod. Phys. **67**, 279 (1995).
[25] D. G. Luchinsky (private communication).