

Direct Evaluation of Rare Events in Active Matter from Variational Path Sampling

Avishek Das^{1,*} Benjamin Kuznets-Speck^{2,*} and David T. Limmer^{1,3,4,5,†}

¹*Department of Chemistry, University of California, Berkeley, California 94720, USA*

²*Biophysics Graduate Group, University of California, Berkeley, California 94720, USA*

³*Chemical Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

⁴*Material Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

⁵*Kavli Energy NanoSciences Institute, University of California, Berkeley, California 94720, USA*



(Received 18 August 2021; accepted 9 December 2021; published 14 January 2022)

Active matter represents a broad class of systems that evolve far from equilibrium due to the local injection of energy. Like their passive analogs, transformations between distinct metastable states in active matter proceed through rare fluctuations; however, their detailed balance violating dynamics renders these events difficult to study. Here, we present a simulation method for evaluating the rate and mechanism of rare events in generic nonequilibrium systems and apply it to study the conformational changes of a passive solute in an active fluid. The method employs a variational optimization of a control force that renders the rare event a typical one, supplying an exact estimate of its rate as a ratio of path partition functions. Using this method we find that increasing activity in the active bath can enhance the rate of conformational switching of the passive solute in a manner consistent with recent bounds from stochastic thermodynamics.

DOI: 10.1103/PhysRevLett.128.028005

The constituent agents of active matter—biomolecules, colloids, or cells—autonomously consume energy to fuel their motion [1,2]. Their resultant nonequilibrium states have non-Boltzmann phase-space densities and exhibit exotic structural and dynamical collective fluctuations, including motility-induced phase separation and swarming [3–7]. Within these nonequilibrium steady states, fleeting fluctuations can free particles from external potentials [8–10], nucleate stable phases from metastable ones [11,12], and assemble passive objects [13,14]. The study of such rare dynamical events within active matter and the calculation of their associated rates is difficult. Traditional equilibrium rate theories like transition state theory and Kramers' theory require knowledge of the form of the steady-state distribution that is not in general available [15]. Further, only a few numerical methods exist that can be used to tame the exponential computational cost associated with sampling the unlikely fluctuations that lead to transitions between metastable states. Existing methods improve sampling by stratifying or branching stochastic trajectories [16–18] but do not typically employ driving forces to specifically enhance the sampling of these rare events.

Here we present a perspective and an associated numerical algorithm, termed variational path sampling (VPS), for estimating transition rates in active systems using optimized time-dependent driving forces. Our approach relies on an equality between the rate of a rare event in a reference system and a ratio of path partition functions in the reference system and with a driving force that makes the rare event occur with high probability. The VPS algorithm solves a variational problem to approximate the functional

form of an optimal time-dependent driving force for this estimate and is applicable to any stochastic dynamics. With VPS we investigate how driven fluids can direct motion into useful function. We apply this technique to study the rate of conformational changes of a passive dimer in a dense bath of active Brownian particles [19–21]. This model exemplifies how collective active fluctuations around passive solutes can drive self-assembly and speed up transitions between distinct metastable states [22,23]. We find the rate to switch between the dimer's two metastable states increases dramatically with increasing activity in the bath, which we rationalize with a recent dissipation bound from stochastic thermodynamics [24]. We study the computational efficiency of rate estimation with VPS and demonstrate its advantage over existing trajectory stratification-based methods like forward flux sampling [16].

We consider a system described by overdamped Brownian dynamics of the form

$$\gamma_i \dot{\mathbf{r}}_i(t) = \mathbf{F}_i[\mathbf{r}^N(t)] + \boldsymbol{\eta}_i(t), \quad (1)$$

where $\dot{\mathbf{r}}_i$ is the rate of change of the i th particle's position, γ_i is the corresponding friction coefficient, and $\mathbf{F}_i[\mathbf{r}^N(t)]$ is the sum of all conservative, nonconservative, and active forces exerted on the i th particle that depends on the full configuration of the N -particle system \mathbf{r}^N . The final term, $\boldsymbol{\eta}_i(t)$, is a Gaussian white noise with $\langle \eta_{ia}(t) \rangle = 0$ and

$$\langle \eta_{ia}(t) \eta_{jb}(t') \rangle = 2\gamma_i k_B T \delta_{ij} \delta_{ab} \delta(t - t') \quad (2)$$

for component (α, β) and $k_B T$ is Boltzmann's constant times the temperature. In order to study the transition rate between two long-lived metastable states, denoted A and B , we define each from a given configuration using the indicator functions,

$$h_X[\mathbf{r}^N(t)] = \begin{cases} 1 & \text{if } \mathbf{r}^N(t) \in X \\ 0 & \text{else,} \end{cases} \quad (3)$$

for either $X = A, B$. In practice this designation requires an order parameter capable of distinguishing configurations and grouping them into distinct metastable states like that illustrated in Fig. 1(a) in one dimension. Assuming there exists a separation between the time τ^\ddagger required to traverse the transition region between the two metastable states and the typical waiting time for the transition, the rate k can be evaluated from the probability to observe a transition, per unit time [25],

$$k = \frac{\langle h_B(t_f)h_A(0) \rangle}{t_f \langle h_A \rangle} = t_f^{-1} \langle h_{B|A}(t_f) \rangle, \quad (4)$$

where the angular brackets denote an average over trajectories of duration $\tau^\ddagger < t_f \ll 1/k$ started from a steady-state distribution in A and $\langle h_{B|A}(t_f) \rangle$ denotes the conditional probability for transitioning between A and B in time t_f . When t_f is chosen to satisfy the timescale separation described above, k is independent of time.

If the transition is rare, most short trajectories are nonreactive leading to difficulties in estimating the rate directly. Instead of trying to evaluate the small transition probability through stratification as other existing methods do [16,17], we instead optimize a time-dependent driving force $\lambda(\mathbf{r}^N, t)$ that constrains the transition to occur, and evaluate the probability cost associated with adding that force to the original dynamics. For a general time-dependent force λ , using the Onsager-Machlup form for the probabilities of stochastic trajectories [26], the rate expression in Eq. (4) can be rewritten as [24]

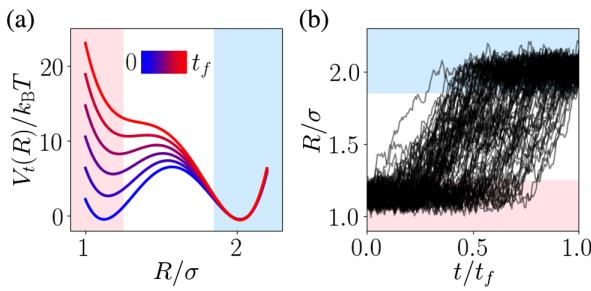


FIG. 1. Reactive trajectories with VPS. (a) Schematic representation of the total optimal time-dependent potential in an isolated passive dimer as t goes from 0 to t_f . Shaded regions are the compact (A , pink) and extended (B , light blue) states. (b) Unbiased reactive trajectories generated with $\lambda(R, t)$.

$$k = t_f^{-1} \langle e^{-\Delta U_\lambda} \rangle_{B|A,\lambda}, \quad (5)$$

where $\langle \rangle_{B|A,\lambda}$ denotes a conditioned average computed in the presence of the additional force. This relation holds for forces λ that affect the transition to occur with probability 1, such that the rate in the driven ensemble is $1/t_f$. The average is of the exponential of the change in the path action, ΔU_λ ,

$$\Delta U_\lambda[\mathbf{X}] = - \int_0^{t_f} dt \sum_i \frac{[\lambda_i^2 - 2\lambda_i \cdot (\gamma_i \dot{\mathbf{r}}_i - \mathbf{F}_i)]}{4\gamma_i k_B T}, \quad (6)$$

between trajectories generated with the added force and in its absence. The path action and all other stochastic integrals are evaluated in the Itô convention.

Equation (5) is a direct estimator for a rate employing an auxiliary control system, but it only becomes useful when the protocol $\lambda(\mathbf{r}^N, t)$ generates trajectories in a manner equivalent to the unbiased reactive trajectory distribution. This is because the expectation can be viewed as an overlap between the two reactive path distributions, and without significant overlap the exponential average is difficult to estimate. We express the optimality of λ using Jensen's inequality after taking the logarithm of Eq. (5) to obtain a variational principle:

$$\ln k \geq -\ln t_f - \langle \Delta U_\lambda \rangle_{B|A,\lambda}. \quad (7)$$

If the average change in conditioned path action $\langle \Delta U_\lambda \rangle_{B|A,\lambda}$ is minimized over all possible functional forms of λ , the rate can be obtained directly as a simple ensemble average of ΔU_{λ^*} at the minimizer $\lambda = \lambda^*$.

The optimal control force λ^* that saturates Eq. (7) is unique and given by the solution of the backward Kolmogorov equation [27–29], as detailed in the Supplemental Material (SM) [30]. Specifically, the optimal force is $2k_B T$ times the gradient of the logarithm of the committor probability [41] of ending in state B at t_f . A schematic illustration of the optimal effective time-dependent potential $V_t(R)$ added to a double-well potential is illustrated in Fig. 1(a). The resultant force gradually destabilizes the reactant well to ensure the transition almost surely within the short duration t_f . Viewed in the backward direction of time, the potential follows the negative logarithm of the relaxation of an initially localized distribution in B to its steady state. The force is thus optimal in the sense that reactive trajectories, like those in Fig. 1(b), generated with it are drawn from the reference path ensemble with the correct statistical weights. Generically, $\lambda^*(\mathbf{r}^N, t)$ is a function of all particle coordinates, so it is not typically tractable to compute. We demonstrate here that one- and two-body representations of λ can be sufficiently close to optimal as to estimate the rate accurately even in cases where the rare

event is collective, similar to related observations in large-deviation sampling [34,42–44].

We study the accuracy and utility of this formalism in a system composed of an active bath and a passive dimer that can undergo conformational changes between two metastable states. All particles interact pairwise via a Weeks-Chandler-Andersen (WCA) repulsive potential [45],

$$V_{\text{WCA}}(\mathbf{r}) = \left\{ 4\epsilon \left[\left(\frac{\sigma}{r} \right)^{12} - \left(\frac{\sigma}{r} \right)^6 \right] + \epsilon \right\} \Theta(r_{\text{WCA}} - r), \quad (8)$$

with energy scale ϵ , and particle diameter σ , truncated at $r_{\text{WCA}} \equiv 2^{1/6}\sigma$ with the Heaviside function Θ . Active particles experience an additional self-propulsion force of magnitude v_0 , $\mathbf{F}_i^a(t) = v_0 \mathbf{e}[\theta_i(t)]$ where the director is $\mathbf{e}(\theta_i) = (\cos \theta_i, \sin \theta_i)$ and θ_i obeys $\dot{\theta}_i(t) = \xi_i(t)$, with

$$\langle \xi_i(t) \rangle = 0, \quad \langle \xi_i(t) \xi_j(t') \rangle = 2D_\theta \delta_{ij} \delta(t - t'), \quad (9)$$

for angular diffusion constant D_θ . Passive solutes separated by distance R are bound by a double-well potential,

$$V_{\text{dw}}(R) = \Delta V [1 - (R - r_{\text{WCA}} - w)^2/w^2]^2, \quad (10)$$

with an energy barrier of height ΔV between the compact and extended states at $R = r_{\text{WCA}}$ and $R = r_{\text{WCA}} + 2w$, respectively [46]. We study the transition rates between these states, employing indicator functions $h_A(t) = \Theta(R_A - R)$ and $h_B(t) = \Theta(R - R_B)$ for $R_A = 1.25\sigma$ and $R_B = 1.85\sigma$. Conformation transitions like these in dense fluids are collective in origin [46] and serve as a sensitive probe of the bath.

The VPS algorithm estimates an optimal force using a low-rank ansatz by iteratively solving the variational problem in Eq. (7), and uses this force to directly obtain a rate estimate. For computing the rate of isomerization of the passive dimer, we approximate λ^* with a time-dependent interaction along the dimer bond vector \mathbf{R} , expressed as a sum of Gaussians,

$$\lambda(\mathbf{R}, t) = \hat{\mathbf{R}} \sum_{p,q=1}^{M_R, M_t} c_{pq}^{(i)} e^{-(R - \mu_{R,p})^2/2\nu_R^2 - (t - \mu_{t,q})^2/2\nu_t^2}, \quad (11)$$

where $c_{pq}^{(1)} = -c_{pq}^{(2)}$ are variational parameters to be tuned, and the locations and widths $\mu_{R,p}$, $\mu_{t,q}$, ν_R , and ν_t are held fixed. To impose the conditioning while minimizing $\langle \Delta U_\lambda \rangle_{B|A,\lambda}$, we use a Lagrange multiplier s to construct a loss function $\Omega_\lambda = \langle \Delta U_\lambda \rangle_\lambda + s(\langle h_{B|A} \rangle_\lambda - 1)$. For a general force that does not ensure the transition with unit probability, there is a multiplicative contribution to the estimate of the rate in Eq. (5) from $\langle h_{B|A} \rangle_\lambda$, which for most optimized forces is negligible.

The optimization problem maps onto the computation of a cumulant generating function for the statistics of the

indicator $h_B(t_f)$ studied previously [29,33], with the short trajectories starting from a steady-state distribution in the initial state. As such we can employ generalizations of recent reinforcement learning procedures to efficiently estimate the gradients of the loss function with respect to the variational parameters [47]. Specifically, we modify the Monte Carlo value baseline (MCVB) algorithm [33], which performs a stochastic gradient descent to optimize $c_{pq}^{(i)}$. We add two preconditioning steps over the MCVB algorithm. First, we generate an initial reactive trajectory using a routine reminiscent of well-tempered metadynamics [36]. Then we symmetrize the learned force to ensure time translational invariance of the transition paths. We denote this preconditioning algorithm MCVB-T. Further information is available in the SM [30].

We first illustrate the systematic convergence of VPS by estimating the isomerization rate of an isolated passive dimer. Such a simplified system allows us to compare to numerically exact results, and study convergence of the force ansatz in the complete basis limit, where $M_R, M_t \rightarrow \infty$ and the Gaussians cover the thermally sampled region in R and t . For this simple system, we take $k_B T = \gamma = \sigma = \epsilon = 1$, $w = 0.25\sigma$, with diffusive timescale $\tau = \sigma^2 \gamma / k_B T$. We simulate the one-dimensional version of Eq. (1) along R , with $V_{\text{dw}}(R)$ only. For simplicity we define state A by the initial condition $R(0) = r_{\text{WCA}}$ and state B via $R_B = 1.45\sigma$. To provide a steady-state value in Eq. (4) [33,48], we use an Euler method and take in this example $t_f = \gamma w \sigma / \sqrt{8k_B T \Delta V}$. We choose $\mu_{R,p}$ and $\mu_{t,q}$ evenly distributed in $R/\sigma \in [0.9, 1.77]$ and $t \in [0, t_f]$, respectively, and ν_R, ν_t to be half the distance between Gaussian centers. We consider basis sizes $M_R = M_t = 2-40$, each optimized independently and used to sample $\sim 10^5$ transition paths.

Figure 2(a) illustrates a typical learning curve for the control force, showing convergence of the variational rate bound toward the numerically exact rate. The variational estimate requires a basis of $M_R, M_t > 40$ to approach the rate to within the statistical uncertainty of the estimate; however, alternative estimates with small basis sets can be refined using a cumulant expansion approximation to Eq. (5). Specifically, truncating the exact exponential relation at the ℓ th cumulant as

$$\ln k \approx -\ln t_f + \sum_{n=1}^{\ell} \frac{1}{n!} \frac{d^n \ln \langle e^{-\Delta U_\lambda} \rangle_{B|A,\lambda}}{d\Delta U_\lambda^n} \quad (12)$$

provides an approximation to the rate that converges in the limit that ℓ is large. Figure 2(b) illustrates this convergence, where we find that even coarse representations of the control force can yield close estimates of the rate with only the first few cumulants, illustrating a trade-off between basis set completeness and statistical efficiency. Sweeping across a wide range of barrier heights in Fig. 2(c), we find excellent agreement between the log rate from brute-force

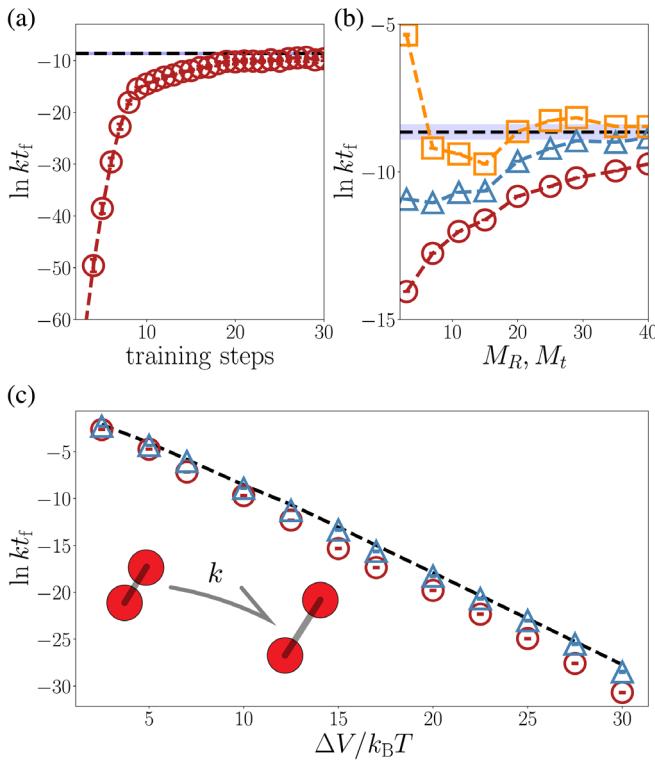


FIG. 2. Convergence of isomerization rates for an isolated passive dimer. (a) Learning curve for $\Delta V = 10k_BT$ and $M_R, M_t = 20$. (b) Convergence of the variational rate estimate (circles) and cumulant corrections for $\ell = 2$ (triangles) and $\ell = 4$ (squares) with basis size as compared to the numerically exact answer (dashed line). (c) Variational (circles) and $\ell = 2$ (triangles) estimate of the rate compared to the exact value (dashed line) with increasing barrier height.

simulations and a truncation of the cumulant expansion to $\ell = 2$ using $M_R = 80$ and $M_t = 30$.

We next compute the isomerization rate with VPS when the dimer is immersed in an explicit solvent of active Brownian particles with $N = 80$ and a total density of $0.6/\sigma^2$. The dimer particles have a friction $\gamma_d = 2\gamma$ and the solvent particles have $\gamma_s = 4\gamma$. We take $\gamma = \sigma = \epsilon = 1$, $k_BT = 0.5$, $\Delta V = 7k_BT$, $\tau = \sigma^2\gamma/2k_BT = 1$, $D_\theta = 1/\tau$, and time step $10^{-5}\tau$. We also change $w = 0.45\sigma$ such that the collisional cross section of the dimer is large. Collisions with active particles transduce energy along the dimer bond and we study the change in the isomerization rate as the bath activity $v_0\sigma/k_BT$ is varied from 0 to 18. We use a basis size of $M_R = M_t = 50$ distributed between $R/\sigma \in [0.9, 2.3]$ and $t \in [0, t_f]$, where $t_f = 0.2\tau$. The optimization starts by learning forces $\lambda(\mathbf{R}, t)$ for the isolated dimer with WCA interactions between monomers, followed by the MCVB-T algorithm. Then, $\lambda(\mathbf{R}, t)$ is optimized in the presence of the bath for $v_0 = 0$ and higher values of v_0 are initialized from converged forces at the previous v_0 .

The rate is a strong function of activity, increasing 20-fold over the range of v_0 's considered. While the

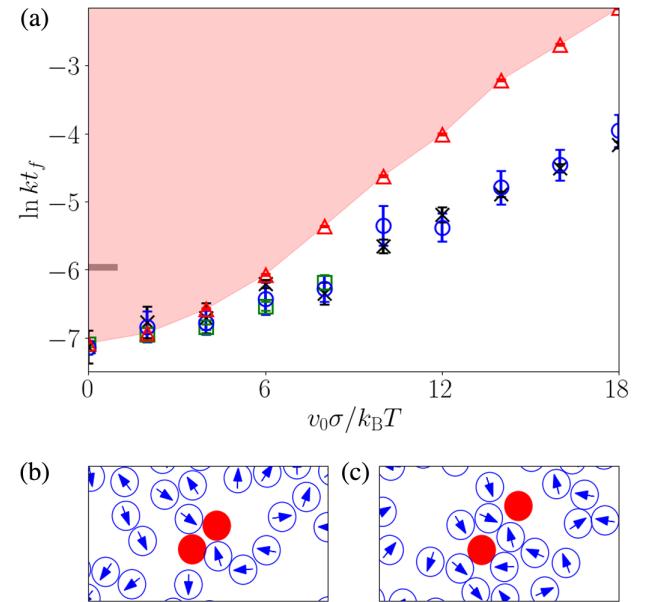


FIG. 3. Rate enhancement of isomerization in an active fluid. (a) Change in the rate as estimated from direct unbiased simulations (crosses), from exponential estimate (squares), and from histogram reweighting (circles). The excess dissipated heat (triangles) bounds the rate enhancement achievable marked by the red shaded region. The thick tick mark on the left denotes the rate for the isolated dimer. (b), (c) Typical snapshots of reactive trajectories of the active bath (blue) and passive dimer (red) at $t = 0$ and $t = t_f$.

variational rate estimate from Eq. (7) is closest for the passive bath, it weakens with increasing v_0 , indicating a growing importance of solvent degrees of freedom in the optimal control force. With converged forces at each v_0 , we run 10^6 trajectories of length t_f to compute k from Eq. (5). This estimate correctly predicts the suppression of k due to passive solvation and can be converged statistically for $v_0\sigma/k_BT < 9$, which is supported by direct rate estimates from unbiased simulations in Fig. 3(a). Above $v_0\sigma/k_BT = 9$, the optimized force is not close enough to λ^* to estimate k directly through the exponential average or a low order cumulant expansion.

Provided we have access to the transition path ensemble from direct unbiased simulations or methods like transition path sampling [49–51], we can supplement the estimate of k using histogram reweighting [52]. k satisfies a reweighting relation of the form

$$k = \frac{e^{-\Delta U_\lambda} P_{B|A,\lambda}(\Delta U_\lambda)}{t_f P_{B|A,0}(\Delta U_\lambda)}, \quad (13)$$

where we have defined $P_{B|A,\lambda}(\Delta U_\lambda) = \langle \delta(\Delta U_\lambda | \mathbf{X}) - \Delta U_\lambda \rangle_{B|A,\lambda}$ and similarly for its undriven counterpart $\lambda = 0$. We evaluate k with this estimator by sampling 10^4 driven and only 6–100 unbiased reactive paths, using

the Bennett acceptance ratio [37] to evaluate the ratio of probabilities. Compared with the brute-force estimate in Fig. 3(a), we find this reweighting predicts k accurately across all values of v_0 with significantly higher statistical efficiency than a brute-force calculation, which validates the accuracy and utility of the control forces. We have compared the VPS rate estimates in the SM [30], using either Eq. (5) or Eq. (13), to the Rosenbluth variant of forward flux sampling [16], and find that VPS is statistically more efficient and converges more quickly with the number of reactive trajectories.

Access to an ensemble of transition paths in this active system provides us with mechanistic insight into the process. The rate enhancement observed for the compact to extended state transition of the passive dimer with bath activity can be understood using recent results from stochastic thermodynamics. Specifically, the rate enhancement achievable by coupling a reactive mode to a nonequilibrium driving force is bounded from above by the heat dissipated over the course of the transition [24]. In this case the nonequilibrium driving is afforded by the interactions between the dimer and the active bath, so the bound takes the form

$$\ln k \leq \ln k_0 + \frac{1}{2k_B T} \langle Q - Q_0 \rangle_{B|A}, \quad (14)$$

where k_0 is the rate at $v_0 = 0$ and $\langle Q - Q_0 \rangle_{B|A}$ is the dissipative heat less its average at $v_0 = 0$ given by

$$Q = \int_0^{t_f} dt \sum_{i \in d} \sum_{j \in s} (\dot{\mathbf{r}}_i - \dot{\mathbf{r}}_j) \cdot \mathbf{F}_{\text{WCA}}(\mathbf{r}_{ij}), \quad (15)$$

which is a sum of the total force from the WCA potential of the solvent particles (s) on the dimer (d) times the difference in their velocities in an ensemble at fixed v_0 (Supplemental Material [30]). This bound is verified in Fig. 3(a) for all v_0 , and is saturated at small v_0 . The specific mechanism of energy transfer from bath to dimer that promotes transitions is clarified by examining reactive trajectories driven by the biasing force and are typical, after removal of the bias from the incomplete basis set. Figures 3(b) and 3(c) show typical snapshots of the solvated dimer at the start and end of the reaction. Energy transfer results from active particles accumulating around the dimer, and preferentially in its cross section, pushing it apart into an extended state. This mechanism of action is reminiscent of how nonequilibrium agents collect in the corners of mesoscopic gears to power their directed rotation [22,23]. At low v_0 , we find the driven isomerization process is efficient, while deviation from the bound at large v_0 demonstrates that energy is additionally funneled into nonreactive modes. Further studies showing the unbiased nature of the VPS-sampled transition path ensemble in terms of duration and distribution of transition paths,

and quantification of the changing solvation environment with v_0 , are provided in the SM [30].

In conclusion, we developed a novel formalism and corresponding algorithm termed variational path sampling to compute rate constants in nonequilibrium systems by optimally driving the systems to transition between metastable states. VPS can be used to compute rates in arbitrary stochastic systems and extends the use of optimal control forces in large deviation sampling to transient rare events [29,34,42,43,53]. VPS complements trajectory-level importance sampling methods by generating the rare reactive event through a time series of driving forces instead of a sequence of rare noise histories. We expect this approach to find broad use in rate computations for rare events in dissipative systems throughout the physical sciences and across scales.

The source code and data that reproduce the findings of this study are openly available online [54].

A. D., B. K.-S., and D. T. L. were supported by NSF Grant No. CHE1954580. The authors thank Dominic Rose, Juan Garrahan, and Phillip Geissler for useful discussions.

^{*}These authors contributed equally to this work.

[†]dlimmer@berkeley.edu

- [1] M. C. Marchetti, J.-F. Joanny, S. Ramaswamy, T. B. Liverpool, J. Prost, M. Rao, and R. A. Simha, *Rev. Mod. Phys.* **85**, 1143 (2013).
- [2] M. R. Shaebani, A. Wysocki, R. G. Winkler, G. Gompper, and H. Rieger, *Nat. Rev. Phys.* **2**, 181 (2020).
- [3] M. E. Cates and J. Tailleur, *Annu. Rev. Condens. Matter Phys.* **6**, 219 (2015).
- [4] T. Speck, *Eur. Phys. J. Special Topics* **225**, 2287 (2016).
- [5] T. Nemoto, É. Fodor, M. E. Cates, R. L. Jack, and J. Tailleur, *Phys. Rev. E* **99**, 022605 (2019).
- [6] T. GrandPre, K. Klymko, K. K. Mandadapu, and D. T. Limmer, *Phys. Rev. E* **103**, 012613 (2021).
- [7] G. Gompper, R. G. Winkler, T. Speck, A. Solon, C. Nardini, F. Peruani, H. Löwen, R. Golestanian, U. B. Kaupp, L. Alvarez *et al.*, *J. Phys. Condens. Matter* **32**, 193001 (2020).
- [8] E. Woillez, Y. Zhao, Y. Kafri, V. Lecomte, and J. Tailleur, *Phys. Rev. Lett.* **122**, 258001 (2019).
- [9] A. Militaru, M. Innerbichler, M. Frimmer, F. Tebbenjohanns, L. Novotny, and C. Dellago, *Nat. Commun.* **12**, 2446 (2021).
- [10] E. Woillez, Y. Kafri, and V. Lecomte, *J. Stat. Mech.* (2020) 063204.
- [11] A. K. Omar, K. Klymko, T. GrandPre, and P. L. Geissler, *Phys. Rev. Lett.* **126**, 188002 (2021).
- [12] D. Richard, H. Löwen, and T. Speck, *Soft Matter* **12**, 5257 (2016).
- [13] J. Stenhammar, R. Wittkowski, D. Marenduzzo, and M. E. Cates, *Phys. Rev. Lett.* **114**, 018301 (2015).
- [14] S. A. Mallory, C. Valeriani, and A. Cacciuto, *Annu. Rev. Phys. Chem.* **69**, 59 (2018).
- [15] A. Nitzan, *Chemical Dynamics in Condensed Phases: Relaxation, Transfer and Reactions in Condensed Molecular Systems* (Oxford University Press, New York, 2006).

[16] R. J. Allen, C. Valeriani, and P. R. Ten Wolde, *J. Phys. Condens. Matter* **21**, 463102 (2009).

[17] A. Warmflash, P. Bhimalapuram, and A. R. Dinner, *J. Chem. Phys.* **127**, 154112 (2007).

[18] F. Cérou and A. Guyader, *Stoch. Anal. Appl.* **25**, 417 (2007).

[19] Y. Fily and M. C. Marchetti, *Phys. Rev. Lett.* **108**, 235702 (2012).

[20] G. S. Redner, M. F. Hagan, and A. Baskaran, *Phys. Rev. Lett.* **110**, 055701 (2013).

[21] J. Bialké, H. Löwen, and T. Speck, *Europhys. Lett.* **103**, 30008 (2013).

[22] S. Mallory, M. Bowers, and A. Cacciuto, *J. Chem. Phys.* **153**, 084901 (2020).

[23] A. Sokolov, M. M. Apodaca, B. A. Grzybowski, and I. S. Aranson, *Proc. Natl. Acad. Sci. U.S.A.* **107**, 969 (2010).

[24] B. Kuznets-Speck and D. T. Limmer, *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2020863118 (2021).

[25] D. Chandler, *J. Chem. Phys.* **68**, 2959 (1978).

[26] L. Onsager and S. Machlup, *Phys. Rev.* **91**, 1505 (1953).

[27] J. Szavits-Nossan and M. R. Evans, *J. Stat. Mech.* (2015) P12008.

[28] R. Chetrite and H. Touchette, *Ann. Inst. Henri Poincaré* **16**, 2005 (2015).

[29] R. Chetrite and H. Touchette, *J. Stat. Mech.* (2015) P12001.

[30] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.128.028005> for simulation details, comparison to forward flux sampling, and analysis of the reactive transition paths, which includes Refs. [6,16,24,31–40].

[31] S. N. Majumdar and H. Orland, *J. Stat. Mech.* (2015) P06039.

[32] F. Carollo, J. P. Garrahan, I. Lesanovsky, and C. Pérez-Espigares, *Phys. Rev. A* **98**, 010103(R) (2018).

[33] A. Das, D. C. Rose, J. P. Garrahan, and D. T. Limmer, *J. Chem. Phys.* **155**, 134105 (2021).

[34] A. Das and D. T. Limmer, *J. Chem. Phys.* **151**, 244123 (2019).

[35] P. Warren and R. Allen, *Entropy* **16**, 221 (2013).

[36] A. Barducci, G. Bussi, and M. Parrinello, *Phys. Rev. Lett.* **100**, 020603 (2008).

[37] M. R. Shirts and J. D. Chodera, *J. Chem. Phys.* **129**, 124105 (2008).

[38] R. Zwanzig, *Nonequilibrium Statistical Mechanics* (Oxford University Press, New York, 2001).

[39] R. J. Allen, P. B. Warren, and P. R. Ten Wolde, *Phys. Rev. Lett.* **94**, 018104 (2005).

[40] R. J. Allen, D. Frenkel, and P. R. ten Wolde, *J. Chem. Phys.* **124**, 024102 (2006).

[41] W. E, and E. Vanden-Eijnden, *Annu. Rev. Phys. Chem.* **61**, 391 (2010).

[42] U. Ray, Garnet Kin-Lic Chan, and D. T. Limmer, *Phys. Rev. Lett.* **120**, 210602 (2018).

[43] T. Nemoto, F. Bouchet, R. L. Jack, and V. Lecomte, *Phys. Rev. E* **93**, 062123 (2016).

[44] D. Jacobson and S. Whitelam, *Phys. Rev. E* **100**, 052139 (2019).

[45] J. D. Weeks, D. Chandler, and H. C. Andersen, *J. Chem. Phys.* **54**, 5237 (1971).

[46] C. Dellago, P. G. Bolhuis, and D. Chandler, *J. Chem. Phys.* **110**, 6617 (1999).

[47] D. C. Rose, J. F. Mair, and J. P. Garrahan, *New J. Phys.* **23**, 013013 (2021).

[48] M. Delarue, P. Koehl, and H. Orland, *J. Chem. Phys.* **147**, 152703 (2017).

[49] P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler, *Annu. Rev. Phys. Chem.* **53**, 291 (2002).

[50] U. Ray, G. K.-L. Chan, and D. T. Limmer, *J. Chem. Phys.* **148**, 124120 (2018).

[51] P. Buijsman and P. Bolhuis, *J. Chem. Phys.* **152**, 044108 (2020).

[52] D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications* (Elsevier, New York, 2001), Vol. 1.

[53] J. Dolezal and R. L. Jack, *J. Stat. Mech.* (2019) 123208.

[54] A. Das, B. Kuznets-Speck, and D. T. Limmer, 10.5281/zenodo.5763101.

Supplemental Material:

Direct evaluation of rare events in active matter from variational path sampling

Avishek Das,^{1,*} Benjamin Kuznets-Speck,^{2,*} and David T. Limmer^{1, 3, 4, 5, †}

¹*Department of Chemistry, University of California, Berkeley, CA, 94720, USA*

²*Biophysics Graduate Group, University of California, Berkeley, CA, 94720, USA*

³*Chemical Sciences Division, LBNL, Berkeley, CA, 94720, USA*

⁴*Material Sciences Division, LBNL, Berkeley, CA, 94720, USA*

⁵*Kavli Energy NanoSciences Institute, University of California, Berkeley, CA, 94720, USA*

(Dated: December 7, 2021)

DEMONSTRATION OF OPTIMAL CONTROL FORCE

For ease of notation, we consider states A and B being specific phase space points, \mathbf{r}_A^N and \mathbf{r}_B^N , respectively. Let $\lambda^*(\mathbf{r}^N, t) = 2k_B T \nabla \Phi$, where $\Phi(\mathbf{r}_B^N, t_f | \mathbf{r}^N, t)$ is the log of the probability to end at a single target configuration, \mathbf{r}_B^N , conditioned on being at \mathbf{r}^N at time t . This conditioned probability satisfies the logarithmic transform of the backward Kolmogorov equation,^{S1-S5}

$$\partial_t \Phi + \sum_i \left[D_i (\nabla_i \Phi)^2 + \nabla_i \Phi \cdot \mathbf{F}_i / \gamma_i + D_i \nabla_i^2 \Phi \right] = 0 \quad (\text{S1})$$

where $D_i = k_B T / \gamma_i$ is the diffusion constant and the gradients act on \mathbf{r}_i . The optimal force achieves the reactive transition by construction, rendering $\langle h_{B|A} \rangle_{\lambda^*} = 1$. The change in path action with the optimal force is

$$\Delta U_{\lambda^*}[\mathbf{X}] = - \int_0^{t_f} dt \sum_i D_i (\nabla_i \Phi)^2 + \nabla_i \Phi \cdot \mathbf{F}_i / \gamma_i - \nabla_i \Phi \cdot \dot{\mathbf{r}}_i \quad (\text{S2})$$

Using Ito's Lemma for the total time derivative of Φ to eliminate the final term,

$$\Delta U_{\lambda^*}[\mathbf{X}] = - \int_0^{t_f} dt \left\{ \sum_i [D_i (\nabla_i \Phi)^2 + \nabla_i \Phi \cdot \mathbf{F}_i / \gamma_i + D_i \nabla_i^2 \Phi] - \dot{\Phi} + \partial_t \Phi \right\} \quad (\text{S3})$$

and substituting the backward Kolmogorov equation

$$\Delta U_{\lambda^*}[\mathbf{X}] = \int_0^{t_f} dt \dot{\Phi} = \Phi(\mathbf{r}_B^N, t_f | \mathbf{r}_B^N, t_f) - \Phi(\mathbf{r}_B^N, t_f | \mathbf{r}_A^N, 0) \quad (\text{S4})$$

the change of action can be evaluated exactly. The boundary terms from the exact integration are 0 and the log of the transition probability between A and B in the reference system. The latter can be identified with $\ln \langle h_{B|A} \rangle$, resulting in the equality in Eq. 7 using the definition of the rate k in Eq. 4. This reasoning extends linearly to cases where A and B are collections of configurations.

PROTOCOL FOR LEARNING OPTIMAL FORCE

We optimize the time-dependent control force $\lambda(\mathbf{R}, t)$ by minimizing $\langle \Delta U_{\lambda} \rangle_{B|A, \lambda}$ over variational parameters $c_{pq}^{(i)}$ using Lagrange multiplier s to impose the $B|A$ conditioning. If s is chosen to be more negative than an approximately estimated threshold value $s^* = \ln[\langle h_{B|A} \rangle / (1 - \langle h_{B|A} \rangle)]$, the optimized forces drive the reaction with unit probability and s need not be individually optimized. For a rare transition, any choice of s with a magnitude an order or more larger than the energy barrier height will robustly provide forces that always satisfy the $B|A$ conditioning.^{S6}

For optimization we use an extension of a reinforcement learning algorithm called Monte Carlo Value Baseline (MCVB).^{S6} This algorithm computes the correlation of the gradient of the log of the trajectory probability, called Malliavin weights,^{S7-S8} with the instantaneous change in ΔU_{λ} and h_B over the course of the trajectory. These yield the exact gradients of the loss-function Ω_{λ} with respect to the tunable parameters, with which a stochastic

Algorithm 1 Monte-Carlo Value Baseline with Time-translation invariance (MCVB-T)

```

1: inputs Gaussian coefficients for a general force  $\lambda_\chi(\mathbf{r}^N, t)$  and value function  $\mathcal{V}_\psi(\mathbf{r}^N, t)$ 
2: parameters learning rates  $\alpha^\chi, \alpha^\psi$ ; total optimization steps  $I$ ; trajectory length  $t_f$  consisting of  $J$  timesteps of duration  $\Delta t$  each; number of trajectories  $N$ 
3: initialize choose initial weights  $\chi$  and  $\psi$ , define iteration variables  $i$  and  $j$ , force and value function gradients  $\delta_P, \delta_V$ , define functional form for stepwise increments (rewards)  $\xi$  to the loss-function  $\Delta U_\lambda + sh_{B|A}$ 
4:  $i \leftarrow 0$ 
5: repeat
6:   Generate trajectories  $[\mathbf{X}(t)]$  with first-order Euler propagation starting from uncorrelated steady-state configurations in state  $A$ . Every trajectory starts experiencing the force  $\lambda$  from a random time  $t_{\text{mid}}$  which is sampled uniformly from  $[0, t_f]$ . Configurations, times, noises (with variance  $2\gamma k_B T \Delta t$ ), Malliavin weights, integral of value function gradients, and rewards are denoted by  $\mathbf{r}_j^N, t_j, \eta_j, y_\chi(t_j), z_\psi(t_j)$  and  $\xi(t_j) = \xi_j$  respectively.
7:    $j \leftarrow 0$ 
8:    $\delta_P \leftarrow 0$ 
9:    $\delta_V \leftarrow 0$ 
10:   $y_\chi(t_0) \leftarrow 0$ 
11:   $z_\psi(t_0) \leftarrow 0$ 
12:  repeat
13:     $\dot{y}_\chi(t_j) \leftarrow \eta_j \cdot \nabla_\chi \lambda_\chi(\mathbf{r}_j^N, t_j) / 2k_B T \Delta t$ 
14:     $y_\chi(t_{j+1}) \leftarrow y_\chi(t_j) + \Delta t \dot{y}_\chi(t_j)$ 
15:     $\dot{z}_\psi(t_j) \leftarrow \nabla_\psi \mathcal{V}_\psi(\mathbf{r}_j^N, t_j)$ 
16:     $z_\psi(t_{j+1}) \leftarrow z_\psi(t_j) + \Delta t \dot{z}_\psi(t_j)$ 
17:     $\delta_P \leftarrow \delta_P + \xi_j y_\chi(t_{j+1}) - \mathcal{V}_\psi(\mathbf{r}_j^N, t_j) \dot{y}_\chi(t_j)$ 
18:     $\delta_V \leftarrow \delta_V + \xi_j z_\psi(t_{j+1}) - \mathcal{V}_\psi(\mathbf{r}_j^N, t_j) \dot{z}_\psi(t_j)$ 
19:     $j \leftarrow j + 1$ 
20:  until  $j = J$ 
21:  average  $\delta_P, \delta_V$  over  $N$  trajectories to get  $\bar{\delta}_P, \bar{\delta}_V$ 
22:   $\chi \leftarrow \chi - \alpha^\chi \bar{\delta}_P$ 
23:   $\psi \leftarrow \psi + \alpha^\psi \bar{\delta}_V$ 
24:   $i \leftarrow i + 1$ 
25: until  $i = I$ 

```

gradient descent is performed. The MCVB algorithm simultaneously learns the driving force and a corresponding value function, $\mathcal{V}(\mathbf{R}, t) = \langle \Delta U_{\lambda, t} + sh_{B|A} \rangle_{\lambda|R(t)=R}$, where $\Delta U_{\lambda, t}$ contains the integrated action difference only within $[t, t_f]$ and the expectation is conditioned on starting from R at time t . The value function greatly reduces the the variance of the gradients at zero cost, allowing better convergence. Our modification to this algorithm, refered to as MCVB-T, is a preconditioning step that enforces time translational symmetry for the log of the bridge probability, $\Phi(\mathbf{r}_f^N, t_f | \mathbf{r}^N, t) = \Phi(\mathbf{r}_f^N, t_f - t | \mathbf{r}^N, 0)$, as the reference forces are not explicit functions of time. This is achieved by randomly choosing a $t_{\text{mid}} \in [0, t_f]$ for every trajectory used for averaging the force gradient, and applying the force $\lambda(\mathbf{R}, t \in [t_{\text{mid}}, t_f])$ on it only for a duration $[t_{\text{mid}}, t_f]$. This ensures that trajectories undergoing the transition at late times are accounted for while training the force.

Details of the MCVB-T algorithm are available in the pseudocode in Algorithm 1. The set of Gaussian coefficients parametrizing the force and the value function are denoted in short by χ and ψ respectively. The MCVB algorithm is a special case of MCVB-T with fixed $t_{\text{mid}} = 0$.

Figure S1 illustrates all learning curves that led to the results in Fig. 3, plotting $\ln h_{B|A} - \Delta U_\lambda|_{B|A}$ as a function of training steps. The averages of this estimator and the gradients are computed over 40 trajectories simulated at each training step. The trajectories are initialized with coordinates randomly chosen from a collection of 10000 steady-state dimer and bath configurations in state A , collected once in every 0.1τ time units from a long trajectory without any driving forces. We first learn optimal forces in the absence of the explicit bath in Fig. S1(a), and then optimize these forces in the presence of the bath in Fig. S1(b). We start our optimization by first finding an arbitrary force that ensures the transition with a nonzero probability. We learn initial parameters $c_{pq}^{(i)}$ with a routine similar to well-tempered metadynamics.^{S9} Starting with $c_{pq}^{(i)} = 0$, at fixed frequency we add $c_{pq}^{(i)} \mapsto c_{pq}^{(i)} + \tau_m \omega T_m / [T_m + \omega \mathcal{N}(R, t)]$, where $\mathcal{N}(R, t)$ is a running histogram of order parameter R up to the current time t , and hyperparameters T_m, ω and τ_m determine how quickly the force landscape is filled. The blue learning curve in Fig. S1(a) refers to 100 steps of metadynamics run with $\tau_m = 10t/M_t$, $\omega = 4000$ and $T_m = 9000$. We find that the force solely from metadynamics is highly suboptimal compared to the rate bound, indicated by the black dashed line. Starting with a λ averaged over

all metadynamics steps and with $\mathcal{V} = 0$, we next optimize both sets of parameters with MCVB-T and then MCVB each over 1000 steps with learning rates $\alpha^\chi = 40$, $\alpha^\psi = 200$ and $s = -100$. We find that the variational estimate converges tightly to the exact rate bound.

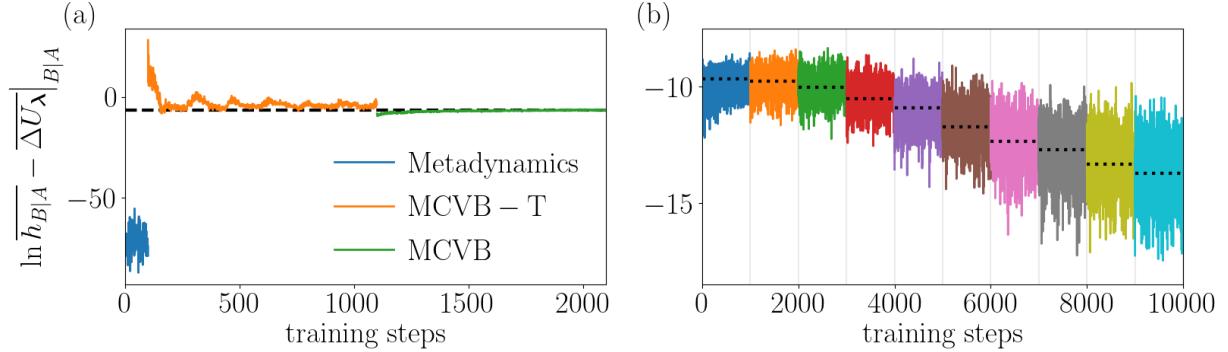


FIG. S1. Learning curves for variational bound. (a) Optimization for the isolated dimer with 100 steps of well-tempered metadynamics(blue), 1000 steps MCVB-T(orange) and 1000 steps MCVB(green). Black dashed line is $\ln k_{tf}$ for the isolated dimer. (b) Learning curves for 1000 steps each, in presence of the explicit bath with $v_0 \sigma / k_B T = 0, 2, 4, \dots, 18$. Black dotted lines denote the corresponding converged values.

Next we use the converged λ and \mathcal{V} to start the optimization in presence of the bath, as illustrated in Figure S1(b). We successively optimize for each $v_0 \sigma / k_B T \in \{0, 2, 4, \dots, 18\}$ starting from the previous converged result, each over 1000 steps. Each time we choose $(\alpha^\chi, \alpha^\psi) = (0, 200)$ for the first 200 steps and $(40, 200)$ for the remaining 800 steps. Learning the value function before starting to change the force in this way avoids a brief period of divergence at the beginning of each optimization run.^{S6} The results are robust towards changing the learning rates as long as α^ψ is kept about 5-10 times of α^χ , such that the value function is always approximately accurate whenever the force is being changed.

Results in Fig. 2 were also obtained similar to this protocol, but with no value function. For Fig. 2(c), the initialization parameters τ_m , ω and T_m are chosen at each barrier height so that at least half of the transitions are reactive.

UNBIASED REACTIVE EVENTS FROM VPS

We use Eq. 5 and 13 in the main text to obtain rate estimates from direct simulations using the low-rank optimized force λ . For the passive dimer in an active bath with $v_0 \sigma / k_B T = 8$, we have illustrated in Figure S2(a) overlap of the driven distribution $P_{B|A,\lambda}(\Delta U_\lambda)$ with the unbiased distribution $P_{B|A}(\Delta U_\lambda)$ after tilting to correct the systematic error. The scaling constant k_{rwt} , which is our estimate for the rate k , has been evaluated from Bennett Acceptance Ratio method^{S10} by using the tilting exponent as the reduced potential. This overlap is observable only when the driving force λ is near-optimal. If the tilted distribution does not contain enough statistics to represent the unbiased distribution, the estimate k_{exp} from Eq. 5 given by the area under the tilted distribution will underestimate the rate. If the basis set is complete and the exact optimal force λ^* can be obtained, ΔU_{λ^*} will follow a Dirac delta distribution $P_{B|A,\lambda^*}(\Delta U_{\lambda^*}) = \delta(\Delta U_{\lambda^*} + \ln k_{tf})$, and the first cumulant will be sufficient to describe the log of the average of the exponential. This is also evident in Eq. S4 where $\lambda = \lambda^*$ makes $\Delta U_\lambda[\mathbf{X}]$ trajectory independent. In that case, all three estimates of k from Eqs. 5, 7 and 13 will be equal and the unbiased reactive events will be entirely force-assisted rather than being driven by thermal fluctuations.

Figure S2(b) shows the systematic and statistical errors in $\ln k_{tf}$ calculated as VPS estimates k_{exp} and k_{rwt} from Equations 5 and 13 respectively in the main text. We have computed the errors by comparing to direct unbiased simulation, as the number N_w of uncorrelated trajectories of duration t_f is varied, expressed through the total number of simulation timesteps $N_F = N_w t_f / \delta t$ where δt is a single timestep. Given the optimized driving force, k_{exp} is computed by averaging over N_w trajectories and k_{rwt} is obtained by reweighting $P_{B|A,\lambda}(\Delta U_\lambda)$ and $P_{B|A}(\Delta U_\lambda)$ each computed with $N_w/2$ total trajectories, of which only a fraction are reactive without the driving force. At small N_F , k_{exp} systematically underestimates the rate due to the full area under $e^{-\Delta U_\lambda} P_{B|A}(\Delta U_\lambda)$ not being accessible because of incomplete overlap, making k_{exp} formally unbiased but statistically biased. This error disappears with large N_F . However, the full rate can still be successfully obtained by comparing segments of incomplete distributions. Thus

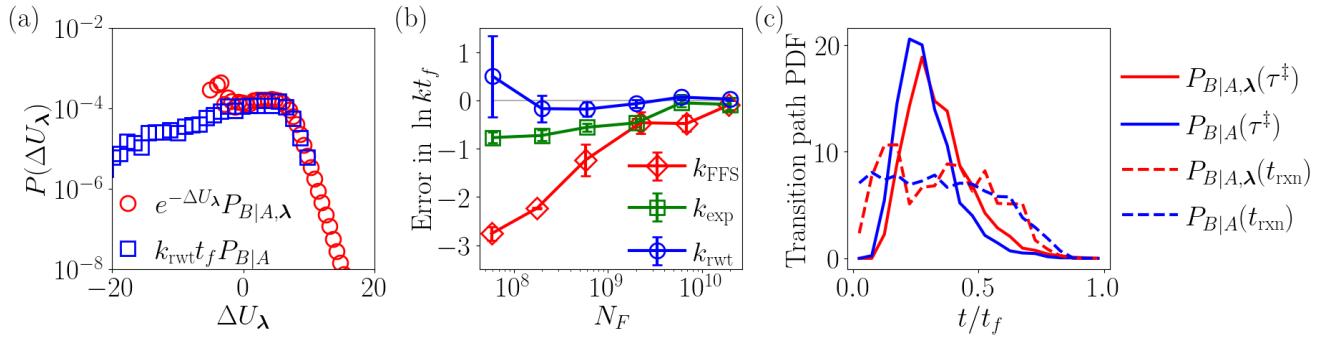


FIG. S2. Unbiased rates, statistical convergence and transition path ensemble with VPS for $v_0\sigma/k_B T = 8$

. (a) Overlap of the incomplete tilted biased and the unbiased distributions, with the scaling coefficient computed from Bennett Acceptance Ratio. (b) Errors in the rate estimates from Equations 5 (k_{exp}) and 13 (k_{rwt}) and from Forward Flux Sampling (k_{FFS}) as the amount of total simulation timesteps N_F is varied. (c) Probability Density Functions (PDF) of transition path times and reactive escape times in the transition path ensemble, computed from the driven trajectories and unbiased reactive trajectories.

even when the undriven trajectory ensemble has ≤ 10 reactive trajectories at smaller values of N_F , k_{rwt} incurs much less error and provides a rate estimate that is both formally and statistically unbiased.

Figure S2(c) demonstrates convergence of the transition path ensemble obtained from direct simulations with the optimized forces even before the tilting correction. $P_{B|A,\lambda}(\tau^\ddagger)$ and $P_{B|A}(\tau^\ddagger)$ are distributions of the transition path time τ^\ddagger measured as the time after leaving state A and before reaching state B without returning to A . $P_{B|A}(\tau^\ddagger)$ is from 2000 reactive trajectories obtained from 10^6 total unbiased simulated trajectories, while $P_{B|A,\lambda}(\tau^\ddagger)$ is from a total of 2000 driven trajectories all of which were reactive. We find convergence in the distribution of transition path times signifying direct access to the nearly unbiased transition path ensemble by using the optimal force. Similarly we compare the distribution of the start time of the reaction $t_{\text{rxn}} \in [0, t_f]$ measured as the time the trajectory last leaves A before arriving in B . We again find convergence in the driven ensemble compared to the unbiased reactive ensemble indicating the forces $\lambda(\mathbf{R}, t)$ are near-optimal at all values of t .

The directly evaluated rates without additional forces used to compare VPS estimates have in most cases been computed from 5 trajectories, each of duration $10^4 \tau$ with τ being the diffusive timescale. We compute k using Eq. 4 by evaluating the expectation with a rolling window over the trajectories after relaxing to a steady-state. We deviate from this protocol only for Figure 2(c), where the barrier heights are too large to estimate the rate from direct simulations. Here we use a numerically exact escape rate obtained from Kramer's theory.^{S11}

COMPARISON WITH FORWARD FLUX SAMPLING

In Figure S2(b) we have compared the numerical cost of VPS at an active self-propulsion $v_0\sigma/k_B T = 8$ with that of a Rosenbluth-like variant of Forward Flux Sampling (RB-FFS).^{S12} Starting from an ensemble of steady-state configurations in A , RB-FFS uses multiple interfaces between A and B to sequentially generate the transition path ensemble and compute the nonequilibrium reaction rate without an additional driving force.^{S13 S14} The transition paths generated from RB-FFS are unbranched and each has an associated weight as part of the transition path ensemble, analogous to VPS, from which the rate is estimated.

In order to apply RB-FFS, we define the interfaces along R as $R/\sigma \in \{1.25, 1.29, 1.33, 1.38, 1.43, 1.50, 1.57, 1.65, 1.72, 1.77, 1.81, 1.85\}$ with the first and the last interfaces corresponding to R_A and R_B respectively. We start RB-FFS trajectories from the same ensemble of steady-state configurations in A that we have used for VPS, and record the configurations whenever the trajectories cross R_A from the A side. Every time a trajectory reaches B , we replace it in A at a random steady-state configuration. From each of \mathcal{M}_0 recorded configurations located at R_A , we generate reactive paths by shooting $\mathcal{M} = 100$ trajectories from each interface sequentially and randomly choosing one out of those that reach the next interface instead of coming back to A . The rate estimate is given by product of the forward flux of crossing R_A and the conditional probabilities of reaching subsequent interfaces, computed from an average over weighted reactive trajectories from the RB-FFS simulation.^{S12} We have varied \mathcal{M}_0 between 20 and 7000 to study the convergence of the RB-FFS rate estimate k_{FFS} as a function of the total number of simulation timesteps N_F , as shown in Figure S2(b). Statistical errors have been

estimated over 3 independent parallel runs of the entire RB-FFS procedure.

We find that the VPS rate estimates k_{exp} and k_{rwt} incur much smaller errors than k_{FFS} at small N_F , though at large N_F all estimates converge to the same rate. Specifically, k_{rwt} converges to the true rate fastest among the three estimates, and k_{exp} incurs much smaller systematic errors than k_{FFS} even before convergence. This demonstrates that the use of the optimized force in a simple low-dimensional basis in VPS reduces the computational cost of estimating the exact rate by an order of magnitude or more compared to a trajectory stratification based method like RB-FFS. Further, we find that in RB-FFS, obtaining sufficient statistics given by a large \mathcal{M}_0 required the use of a long serial simulation to converge the flux of crossing the first interface at R_A . Parallelization of the \mathcal{M} trajectory segments starting from each interface scaled poorly due to a broad distribution of durations over the trajectory segments, each of which must continue till they reach either the next interface, or A . Since shooting from the next interface can only start after the slowest of the previous trajectory segment has concluded, parallel implementations of RB-FFS scaled poorly and required overall a very long serial simulation. Our attempts to parallelize RB-FFS in an alternate fashion by reducing the serial configurations \mathcal{M}_0 worsened the systematic error in k_{FFS} even when averaged over fully independent RB-FFS implementations over many parallel threads. In contrast, every step of VPS is trivially parallelizable because of all trajectories being of the same duration t_f , which corresponds to $N_F = 2 \times 10^4$. As a result, a parallel implementation of VPS reduced its cost linearly and the overall computation required much shorter serial simulations.

DISSIPATIVE RATE BOUNDS

Stochastic thermodynamics provides a fundamental speed-limit on the enhancement achievable of the rate k of a rare nonequilibrium transition over a reference equilibrium rate k_0 in terms of the excess heat dissipation in the reactive path ensemble,^{S15}

$$\ln k \leq \ln k_0 + \frac{1}{2k_B T} \langle Q \rangle_{B|A} \quad (\text{S5})$$

where $\langle Q \rangle_{B|A}$ is defined as the time-reversal asymmetric contribution from the change in path action between the equilibrium reference and the nonequilibrium system in which it is measured. This bound holds under mild assumptions of instantonic or diffusive transitions and follows from a similar change of measure as leads to Eq. 7, with the additional observation of the time-reversal symmetric contribution of the change in path action being negligible near equilibrium. Here we show how to arrive at Eq. 14 employing this bound. As in the main text, we assume a separation of timescales between local relaxation and typical transitions so that the rate problem is well-posed.

We consider the rate enhancement afforded by coupling the dimer to an active solvent over the equilibrium passive bath isomerization rate. However, if we simply compute the excess heat dissipated as resulting from the time reversal asymmetric change in path action in turning v_0 from 0 to some finite value, the heat will be extensive in the number of solvent degrees of freedom and thus not have a well-defined thermodynamic limit. To mitigate this, we note that the isomerization rate of the dimer would be independent of v_0 if the dimer and solvent did not interact. Denoting k^{ni} and k_0^{ni} the rates of isomerization when the dimer is uncoupled to the solvent at finite or zero v_0 , respectively, then $k^{\text{ni}} = k_0^{\text{ni}}$ and

$$\ln \frac{k}{k_0} = \ln \frac{k}{k^{\text{ni}}} \frac{k_0^{\text{ni}}}{k_0} \leq \frac{1}{2k_B T} (\langle Q \rangle_{B|A} - \langle Q \rangle_{B|A,0}) \quad (\text{S6})$$

where $\langle Q \rangle_{B|A}$ is the excess dissipation resulting from turning on interactions between the dimer and solvent at finite v_0 , and $\langle Q \rangle_{B|A,0} = Q_0$ results from turning on interactions between the dimer and solvent at $v_0 = 0$. The inequality is preserved even though a difference of heats is taken since the second ratio of rates k_0^{ni}/k_0 are both evaluated at equilibrium and thus the symmetric part of the action is zero. This second heat subtracts out the dissipation that is uncorrelated with the isomerization, and the remaining excess dissipation is left finite even when the number of solvent particles is large, so long as the dimer is correlated with the solvent over a finite distance.

The full path action for a system at finite v_0 in the presence of dimer-solvent interactions is

$$U_{v_0} = -\frac{1}{4k_B T} \int_0^{t_f} dt \gamma_d^{-1} \sum_{i \in d} \left[\gamma_d \dot{r}_i + \nabla_i V_{dw} - \sum_{j \in d} \mathbf{F}_{WCA}(\mathbf{r}_{ij}) - \sum_{j \in s} \mathbf{F}_{WCA}(\mathbf{r}_{ij}) \right]^2 + \gamma_s^{-1} \sum_{i \in s} \left[\gamma_s \dot{r}_i - v_0 \mathbf{e}[\theta_i] - \sum_{j \in d} \mathbf{F}_{WCA}(\mathbf{r}_{ij}) - \sum_{j \in s} \mathbf{F}_{WCA}(\mathbf{r}_{ij}) \right]^2 - \frac{1}{4D_\theta} \int_0^{t_f} dt \sum_{i \in s} \dot{\theta}_i^2 \quad (S7)$$

and using the convention that v_0 is invariant under time-reversal,^{S16} the dissipated heat associated with turning on interactions between the solvent and dimer is

$$Q(t_f) = \int_0^{t_f} dt \left[\sum_{i \in d} \sum_{j \in d} \dot{r}_i \mathbf{F}_{WCA}(\mathbf{r}_{ij}) + \sum_{j \in s} \sum_{i \in d} \dot{r}_j \mathbf{F}_{WCA}(\mathbf{r}_{ji}) \right] \quad (S8)$$

and is the same if $v_0 = 0$ or is nonzero. Since $\mathbf{F}_{WCA}(\mathbf{r}_{ji}) = -\mathbf{F}_{WCA}(\mathbf{r}_{ij})$ we find the definition of the Q in Eq. 15.

We measure $\langle Q \rangle_{B|A}(t_f)$ by averaging Eq. S8 over reactive trajectories of length $t_f = .2\tau$ sampled from long, 2×10^8 time-step, simulations in the nonequilibrium steady state at fixed v_0 , with all other parameters as in the main text. Assuming transitions are uncorrelated, we compile Q samples from 24–96 independent simulations, and use this data to calculate a mean and standard error, as depicted in Fig. 3 (red triangles).

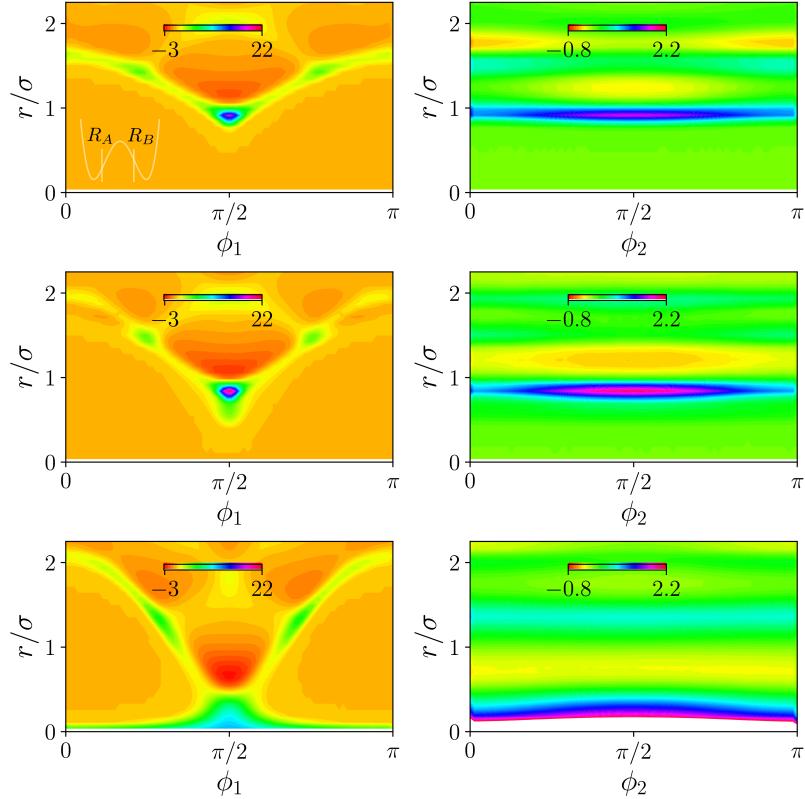


FIG. S3. Solvation structure of the dimer by the active bath. Difference in pair distributions $\Delta g_X(r, \phi_1)$ (Left) and $\Delta g_X(r, \phi_2)$ (Right). Configurations are conditioned such that the bath is sampled with the dimer in the collapsed state $\Delta g_A(r, \phi_{1,2})$ (Top), the transition region $\Delta g_{AB}(r, \phi_{1,2})$ (Middle) or the extended state $\Delta g_B(r, \phi_{1,2})$ (Bottom).

NONEQUILIBRIUM SOLVATION STRUCTURE

We have studied how the solvation structure around the dimer evolves with activity. We compute the two-dimensional pair distribution functions for the position of the solvent around the dimer bond within a conditioned

steady state ensemble average

$$g_X(r, \phi_1) = \frac{1}{\rho_s \rho_d} \frac{\langle \sum_{i \in s} \delta(\mathbf{r}_{cm}^d) \delta(r - |\mathbf{r}_i - \mathbf{r}_{cm}^d|) \delta(\phi_1 - \arccos(\mathbf{r}_i \cdot \mathbf{R})) h_X(R) \rangle}{\langle h_X(R) \rangle}, \quad (S9)$$

with the center of the dimer bond $\mathbf{r}_{cm}^d = (\mathbf{r}_1 + \mathbf{r}_2)/2$ as a reference. Here, r is the radial distance between \mathbf{r}_{cm}^d and surrounding bath particles, which make an angle ϕ_1 with the bond vector $\mathbf{R} = \mathbf{r}_1 - \mathbf{r}_2$. The indicator function $h_X(R)$ restricts configurations where the bond length R falls into state X . Similarly, we probe the orientation of active solute particles around the dimer bond vector with the pair distribution

$$g_X(r, \phi_2) = \frac{1}{\rho_s \rho_d} \frac{\langle \sum_{i \in s} \delta(\mathbf{r}_{cm}^d) \delta(r - |\mathbf{r}_i - \mathbf{r}_{cm}^d|) \delta(\phi_2 - \arccos(\mathbf{e}[\theta_i] \cdot \mathbf{R})) h_X(R) \rangle}{\langle h_X(R) \rangle}. \quad (S10)$$

where ϕ_2 is the angle between a bath director and the dimer bond, and $\rho_s \sigma^2 = 0.6$ is the density of the solvent, and $\rho_d \sigma^2 = 0.008$ is the density of the dimers. To compute $g_X(r, \phi_i)$ we average over configurations sampled from 24 simulations each with a length of 2×10^8 time-steps.

In Fig. S3, we consider the change in the pair distributions $\Delta g_X(r, \phi_{i=1,2}) = g_X(r, \phi_{i=1,2}, v_0 = 9) - g_X(r, \phi_{i=1,2}, v_0 = 0)$ in an active bath with $v_0 = 9$ and its equilibrium counterpart at $v_0 = 0$. The different rows impose different conditions for the dimer bond distance $R = |\mathbf{R}|$ to be either primarily in state $X = A$ (top row), with $R < 1.55\sigma$, in the transition region $X = AB$ between states $1.55\sigma < R < 1.65\sigma$ (center row), or mostly in state $X = B$ (bottom row), with $R > 1.65\sigma$.

These results demonstrate that the rate enhancement is correlated with active particles dynamically wedging within the cross section of the dimer, pushing it apart into an extended state. The left column of Fig. S3 demonstrates that activity greatly enhances the packing of bath particles between the two bonded dimer particles, while the right column illustrates that bath particles preferentially orient perpendicular to the bond vector once far enough within the cross section. In state A , the active nature of the bath causes particles to push the dimer apart along R , as evidenced by the depletion for $\phi_2 = \pi/2$ and $r/\sigma > 1$ in the top-right panel of Fig. S3. The transition region, center-left, shows a significantly higher peak in the radial distribution function around $\phi_2 = \pi/2$, marking a decrease in the height of the effective free energy barrier along R . This analysis also illustrates the mechanism of increased stability in the active dimer extended state. Namely, Fig. S3 bottom-left shows that the driven bath particles act to inhibit the extended state from closing.

* These authors contributed equally

† dlimmer@berkeley.edu

[S1] S. N. Majumdar and H. Orland, *Journal of Statistical Mechanics: Theory and Experiment* **2015**, P06039 (2015).
 [S2] J. Szavits-Nossan and M. R. Evans, *Journal of Statistical Mechanics* **2015**, P12008 (2015).
 [S3] R. Chetrite and H. Touchette, *Annales Henri Poincaré* **16**, 2005 (2015).
 [S4] R. Chetrite and H. Touchette, *Journal of Statistical Mechanics* **2015**, P12001 (2015).
 [S5] F. Carollo, J. P. Garrahan, I. Lesanovsky, and C. Pérez-Espigares, *Physical Review A* **98**, 010103 (2018).
 [S6] A. Das, D. C. Rose, J. P. Garrahan, and D. T. Limmer, arXiv:2105.04321 (2021).
 [S7] A. Das and D. T. Limmer, *Journal of Chemical Physics* **151**, 244123 (2019).
 [S8] P. Warren and R. Allen, *Entropy* **16**, 221 (2013).
 [S9] A. Barducci, G. Bussi, and M. Parrinello, *Physical Review Letters* **100**, 020603 (2008).
 [S10] M. R. Shirts and J. D. Chodera, *Journal of chemical physics* **129**, 124105 (2008).
 [S11] R. Zwanzig, *Nonequilibrium statistical mechanics* (Oxford university press, 2001).
 [S12] R. J. Allen, C. Valeriani, and P. R. Ten Wolde, *Journal of physics: Condensed matter* **21**, 463102 (2009).
 [S13] R. J. Allen, P. B. Warren, and P. R. Ten Wolde, *Physical review letters* **94**, 018104 (2005).
 [S14] R. J. Allen, D. Frenkel, and P. R. ten Wolde, *The Journal of chemical physics* **124**, 024102 (2006).
 [S15] B. Kuznets-Speck and D. T. Limmer, *Proceedings of the National Academy of Sciences* **118** (2021).
 [S16] T. GrandPre, K. Klymko, K. K. Mandadapu, and D. T. Limmer, *Physical Review E* **103**, 012613 (2021).