Automatic error control during forward flux sampling of rare events in master equation models

Cite as: J. Chem. Phys. 152, 035102 (2020); doi: 10.1063/1.5129461

Submitted: 29 September 2019 • Accepted: 29 December 2019 •

Published Online: 16 January 2020

Max C. Klein and Elijah Roberts^{a)}

AFFILIATIONS

Department of Biophysics, Johns Hopkins University, Baltimore, Maryland 21218, USA

^{a)} Author to whom correspondence should be addressed: eroberts@jhu.edu. Telephone: 410-516-2384.

Current address: Johns Hopkins University, Jenkins Hall 110, 3400 N Charles St., Baltimore, MD 21218, USA.

ABSTRACT

Enhanced sampling methods, such as forward flux sampling (FFS), have great capacity for accelerating stochastic simulations of nonequilibrium biochemical systems involving rare events. However, the description of the tradeoffs between simulation efficiency and error in FFS remains incomplete. We present a novel and mathematically rigorous analysis of the errors in FFS that, for the first time, covers the contribution of every phase of the simulation. We derive a closed form expression for the optimally efficient count of samples to take in each FFS phase in terms of a fixed constraint on sampling error. We introduce a new method, forward flux pilot sampling (FFPilot), that is designed to take full advantage of our optimizing equation without prior information or assumptions about the phase weights and costs along the transition path. In simulations of both single and multidimensional gene regulatory networks, FFPilot is able to completely control sampling error. We then discuss how memory effects can introduce additional error when relaxation along the transition path is slow. This extra error can be traced to correlations between the FFS phases and can be controlled by monitoring the covariance between them. Finally, we show that, in sets of simulations with matched error, FFPilot is on the order of tens-to-hundreds of times faster than direct sampling and noticeably more efficient than previous FFS methods.

Published under license by AIP Publishing. <https://doi.org/10.1063/1.5129461>

I. INTRODUCTION

Understanding how inanimate biochemical molecules come together and interact to form a living cell is one of the fundamental goals of biology.¹ The cell's state can be described as a point in a high dimensional space, where each dimension corresponds to the concentration of a different molecule. Complex, nonequilibrium regulatory and signaling networks connect these molecules through positive and negative interactions, naturally resulting in a large number of metastable states within the space,² representing different cellular phenotypes. Projected into two dimensions, the phase space appears as a rough landscape, sometimes referred to as a quasipotential,² epigenetic,³ or phenotypic⁴ landscape. Stochastic fluctuations cause the cell's state to move along the landscape and occasionally jump between metastable states.

Metastable systems, because they depend upon random fluctuations, are typically modeled using a formulation of the chemical master equation (CME)^{5,6} or using stochastic differential equations.⁷ In the former case, the models are often numerically studied using the stochastic simulation algorithm^{8,9} (SSA) or one of its many varieties.^{10–12} Such simulations have provided insight into diverse biological processes, including the lysis/lysogeny decision in bacteriophage λ ,¹³ the *lac* operon in *Escherichia coli*,¹⁴ check-pointing during the cell cycle,¹⁵ differentiation of stem cells,¹⁶ the binding of an intrinsically disordered peptide to a protein,¹⁷ macrophage regulation,¹⁸ and gradient detection during yeast mating.¹⁹

Transitions between metastable states in stochastic biochemical systems are infrequent in that one must wait a long time to observe a large fluctuation that causes the system to switch states.²⁰ The time spent in the transition region is also very short relative to the waiting

time.²¹ Such dynamics are known as rare events and are expensive to simulate as most of the computational effort is spent simply simulating the waiting state. Numerical methods for improving the efficiency of simulating rare events, generally known as enhanced sampling (ES) techniques, have a long history. The earliest work in the field is typically credited to Kahn and Harris²² but has since been applied to the study of many systems in molecular mechanics. The key assumption is that metastable systems exhibit a large barrier separating the states on some free energy landscape. By biasing the simulation toward the transition path between the states, one can more efficiently recover its free energy profile and, thus, the transition rates. For example, umbrella sampling^{23,24} uses overlapping biasing potentials to confine multiple simulations to narrow windows along the path; metadynamics^{25,26} gradually adds a repulsive potential to low points on the free energy surface causing the system to explore low probability regions of phase space and to eventually cross the barrier; transition path sampling^{27–30} generates a statistically correct set of transition paths starting from an initial trial path using acceptance and rejection criteria; weighted ensemble^{31,32} runs multiple independent trajectories while dynamically splitting and merging them, with careful accounting of the trajectory weights, to balance simulations along the transition path. In general, a final unbiasing and/or recombination step is always needed to calculate unbiased statistics.

Although stochastic biochemical systems are described by quasipotential rather than free energy landscapes, the underlying physics is compatible with ES. Consequently, many varieties of ES can be applied to stochastic biochemical systems with rare event dynamics. Three of the most well-known candidates are forward flux sampling (FFS),^{33–36} nonequilibrium umbrella sampling,^{37–39} and weighted ensemble.^{40–44} Although much discussion has taken place regarding the strengths and weaknesses of each of these methods,^{45,46} there is no consensus as to an optimal approach. In the case of a transition between only two metastable states along a single order parameter, FFS appears to be a reasonable choice and is the focus of this work.

The FFS method can be used to calculate both the transition rate, i.e., the inverse of the mean first passage time (*MFPT*), between two metastable states and the probability distribution along the transition path.³⁵ Although it can be applied to nonstationary processes,⁴⁷ here, we investigate only stationary processes. The fundamental operation of FFS is to partition the phase space along an order parameter using a series of nonintersecting interfaces that track progress over the transition barrier. Many trial trajectories are started successively from each interface in order to compute the probability of advancing to the next interface vs returning to the initial state. The product of the advancement probabilities and the probability flux out of the initial state gives the mean transition rate.

The generation of many trial trajectories at each of the interfaces is a large part of the computational work in FFS. It is not surprising, then, that a number of authors have tried to optimize the performance of FFS by studying the interplay between the number of trajectories sampled and the statistical error. Allen *et al.*⁴⁸ first introduced a framework for studying the relationship between the error in the estimated transition rate constant and the error in the estimated interface trial probabilities. They also studied the computational efficiency of FFS, taken to be the inverse of cost

time error, and showed that it was relatively insensitive to the choice of interface parameters. Borrero and Escobedo⁴⁹ presented two techniques to minimize the FFS error for a fixed cost either by optimizing the number of trial runs at each interface or by iteratively refining the placement of the interfaces. Kratzer *et al.*⁵⁰ extended this idea with an algorithm to automatically define the interface number and position on-the-fly by constraining the number of successful trials to be the same for each interface. However, none of these previous works included a systematic treatment of the error arising from the estimate of the flux out of the initial state and they all assume that computational cost per unit distance along the order parameter is fixed. Jiang *et al.*⁵¹ later showed that interface placement along a complex transition landscape, e.g., one with a metastable intermediate state, cannot be correctly optimized by methods that assume cost is proportional to the interface distance.

Here, we present a method to optimally perform an FFS simulation of a nonequilibrium stationary process at a given margin of error and confidence interval. The key advance of our method is to estimate all of the interface probabilities and costs from a short pilot simulation and then use those parameters to optimize the configuration of a full FFS simulation. We minimize the total computational cost to achieve a user specified statistical error, which greatly reduces the number of choices that need to be made by the user. Our method includes a new formal treatment of the sampling error arising in the calculation of the flux out of the initial state, which is shown to have a significant influence on the total cost. Unlike previous efforts, the cost in our method varies by interface, which enables it to account for changes in computational efficiency along the order parameter. We evaluate the capability of our method to control sampling error on three different models exhibiting rare event dynamics. In two one-dimensional models and a multidimensional model, sampling error dominates and is controlled precisely. We then investigate a multidimensional model, where memory effects become significant and require oversampling. Finally, we show that the performance advantage of forward flux pilot sampling (FFPilot) is substantial compared to direct sampling (DS) and previous FFS methods.

The remainder of this work is organized as follows: in Sec. II, we introduce our theoretical framework for estimating the sampling error in a FFS simulation. In Secs. III A and III B, we derive the optimizing equation used by our method to control sampling error and then describe the FFPilot algorithm. In Secs. III C–III E, we evaluate the accuracy and performance of our method using three increasingly complex models. In Sec. III F, we analyze sources of error outside of the sampling error. Finally, in Sec. III G, we evaluate the error introduced due to memory in terms of the covariance between the simulation phases.

II. THEORY AND METHODS

A. Simulation of rare events in stochastic processes

The standard stochastic simulation protocol, here referred to as direct sampling (DS), proceeds in two iterated steps: (1) increment the system time with the time of the next stochastic event and (2) update the system state according to the event. These two steps are repeated until some predetermined stopping condition

(simulation steps, simulation time, etc.) is reached, at which point the simulation is terminated. Repeated DS simulation produces a dataset from which ensemble average quantities can be estimated by straightforward averaging. Running more DS replicates leads to increased accuracy.

As discussed above, DS has the disadvantage that it requires a long simulation time to sample each rare event. Therefore, calculating rare event statistics is computationally expensive. Enhanced sampling (ES) methods use a combination of constraints on simulation trajectories and statistical unbiasing methods in order to enrich the sampling of rare events. Forward flux sampling (FFS) is a popular ES method that was initially proposed by Allen and co-workers.³³ We implemented the FFS algorithm as follows:

1. Find the steady states of the model system, here termed $\mathcal{A}$ and $\mathcal{B}$. These will serve as an initial and a final state for the simulation.
2. Choose a one dimensional parameter $\mathcal{O}$ for which $\mathcal{O}[\mathcal{A}] < \mathcal{O}[\mathcal{B}]$.
3. Choose a set $\{\lambda_0, \dots, \lambda_N\}$ of interface values of $\mathcal{O}$, such that $\mathcal{O}[\mathcal{A}] < \lambda_i < \mathcal{O}[\mathcal{B}]$ for every λ_i . We say that a trajectory has fluxed forward with respect to a λ_i when it crosses the interface traveling in the direction of increasing $\mathcal{O}$, and that it has fluxed backward when it crosses traveling in the direction of decreasing $\mathcal{O}$.
4. Begin FFS phase 0:
 - (a) Execute a DS trajectory initialized at $\mathcal{A}$ with a reflecting barrier at the midpoint between $\mathcal{O}[\mathcal{A}]$ and $\mathcal{O}[\mathcal{B}]$. Multiple phase 0 trajectories can be executed in parallel.
 - (b) Each time the trajectory fluxes forward across λ_0 , record the elapsed simulation time τ since the previous crossing event, and also record the state $\mathcal{X}$.
 - (c) Once n_0 samples of τ have been collected, terminate the trajectories and move to the next phase.
5. Begin FFS phase $i > 0$:
 - (a) Randomly choose a state $\mathcal{X}$ from the collection of states at which any trajectory in the previous phase crossed λ_{i-1} .
 - (b) Execute a DS trajectory initialized at $\mathcal{X}$. Allow the trajectory to run until it either crosses λ_0 back into $\mathcal{A}$ or moves forward across λ_i . Terminate the trajectory and record a trajectory outcome value, either 0 or 1, depending on whether the trajectory moved backward or forward, respectively. If the trajectory moved forward, add the endpoint, which will lie along λ_i , to the set of states that will be used to initialize trajectories during the next phase.
 - (c) Repeat steps 5a and 5b.
 - (d) Once n_i trajectory outcomes have been collected, terminate the phase. Every phase i trajectory can be executed in parallel.
6. The procedure for FFS phase i is then repeated for phase $i + 1$ (during which trajectories are launched from λ_i) until the final phase, phase N , is reached. Trajectories in phase N begin at λ_{N-1} and may move forward across λ_N and into $\mathcal{B}$.

The overall aim of FFS is to ratchet a simulation from $\mathcal{A} \rightarrow \mathcal{B}$ across the state space. The immediate goal of each phase is to estimate a phase weight. During phase i , samples are taken from an observable $\mathbf{w}_i$ that behaves as a random variable. $\mathbf{w}_0$ is the waiting time in between forward flux events across λ_0 , and each $\mathbf{w}_{i>0}$ is the

final outcome (0 for fall back or 1 for flux forward) of each trajectory launched in that i . The phase weight w_i is defined as the true expected value of the random variable $\mathbf{w}_i$,

$$w_i = E[\mathbf{w}_i] = \begin{cases} \tau_{\mathcal{A}} & \text{if } i = 0 \\ P(\lambda_i | \lambda_{i-1}) & \text{otherwise,} \end{cases} \quad (1)$$

where $\tau_{\mathcal{A}}$ is the expected waiting time in between λ_0 crossing events, and $P(\lambda_i | \lambda_{i-1})$ is the probability that a trajectory launched from λ_{i-1} (i.e., launched during phase i) crosses forward past λ_i before it falls back behind the starting interface λ_0 . The phase weights can be used to reweight the output of an FFS simulation so as to calculate unbiased statistics. Note that $\widehat{w}_0$ has different units than $\widehat{w}_{i>0}$, but we choose to discuss them all in terms of a single phase weight formalism as it leads to simpler expressions in our derivations.

B. Estimating values of interest from DS and ES stochastic simulations

As mentioned in Sec. II A, estimating values of interest from DS simulations is straightforward. One valid estimator of any ensemble average quantity is simply the arithmetic mean taken across an appropriate set of direct observations. For example, in order to calculate the *MFPT* of the switching process that takes some multistate system from $\mathcal{A} \rightarrow \mathcal{B}$, n DS simulations are initialized in $\mathcal{A}$. Each of these n replicate simulations is allowed to run until the first time it enters $\mathcal{B}$, at which point it is terminated. The time that each simulation i ran for is then a single sample FPT_i of the first passage time of the switching process. The mean of these first passage time samples is an estimate of *MFPT*,⁵²

$$\widehat{MFPT}_{\text{ds}} = \frac{1}{n} \sum_{i=1}^n FPT_i, \quad (2)$$

where $\widehat{MFPT}_{\text{ds}}$ is the estimator (i.e., estimation function) of *MFPT* specific to DS simulation. Throughout this paper, we place a $\widehat{}$ above a symbol to indicate that we are referring to an estimator of a value rather than to the value itself.

Equivalent ensemble average estimators can be calculated using FFS. Although the results of FFS simulations are biased and partitioned, statistical protocols allow FFS results to be recombined and reweighted so as to recapitulate the results of the unbiased ensemble. In order to perform these reweightings, we need estimates of the phase weights. During phase i of FFS, n_i samples $\{\mathbf{w}_{i1}, \dots, \mathbf{w}_{in_i}\}$ of an underlying random process $\mathbf{w}_i$ are taken. The mean of this sample can be used as an estimator $\widehat{w}_i$ of phase weight i ,

$$\widehat{w}_i = \frac{\sum_{j=1}^{n_i} \mathbf{w}_{ij}}{n_i} = \begin{cases} \frac{\sum_{j=1}^{n_0} \tau_j}{n_0} & \text{if } i = 0 \\ \frac{n_i^s}{n_i} & \text{otherwise,} \end{cases} \quad (3)$$

where each τ_i is a sample of the waiting time in between phase 0 forward flux events and n_i^s is the total count of trajectories that successfully fluxed forward during phase $i > 0$.

In an idealized situation, in which it were possible to know the exact values of the phase weights w_i , the exact value of the *MFPT* could be found via a simple combination W of the phase weights,

$$MFPT = \frac{\tau_A}{\prod_{i>0} P(\lambda_i | \lambda_{i-1})} = \frac{w_0}{\prod_{i>0} w_i} = W. \quad (4)$$

In reality, we only know the phase weight estimators $\widehat{w}_i$, so we must instead estimate the value of the $MFPT$ by way of the combination of estimators $\widehat{W}$,

$$\widehat{MFPT}_{\text{ffs}} = \frac{\widehat{w}_0}{\prod_{i>0} \widehat{w}_i} = \widehat{W}, \quad (5)$$

where $\widehat{MFPT}_{\text{ffs}}$ is the estimator of $MFPT$ specific to FFS simulation. Other quantities of interest, such as the stationary probability density function (PDF), can also be calculated using the phase weights.³⁵

C. Predicting simulation error in terms of margin of error

Stochastic simulation results are not single valued, but are instead estimators, random variables with associated distributions. For example, when attempting to calculate $MFPT$, each complete round of simulations can be thought of as performing a single draw from the distribution of possible $\widehat{MFPT}$ values. A prediction about the likely level of error in any given set of simulations can be made based on the characteristics of this estimate distribution. One way to quantify confidence in this prediction is in terms of the margin of error. The margin of error of a random variable X is defined as the ratio of half the width of a specified confidence interval and its expected value,

$$\zeta[X, \alpha] = \frac{ub_\alpha[X] - lb_\alpha[X]}{2E[X]}, \quad (6)$$

where $\zeta[X, \alpha]$ is the margin of error of X at confidence level α , $E[X]$ is the expected value of X , and $lb_\alpha[X]$ and $ub_\alpha[X]$ are the lower and upper confidence bounds, respectively.

For many values of interest, $\zeta[X, \alpha]$ is dependent on simulation parameters that are set by the user. For example, when calculating $MFPT$, the margin of error $\zeta[\widehat{MFPT}, \alpha]$ is dependent upon the total number of replicate simulations (when using DS), or upon the count of trajectories run in each phase (when using FFS).

D. Determining margin of error from simulation parameters

It is straightforward to determine the margin of error of an estimator that depends on a single underlying observable. For example, consider the margin of error of an $MFPT$ estimate as determined via DS simulation, $\zeta[\widehat{MFPT}_{\text{ds}}, \alpha]$. It has been shown that observations of first passage times between two well separated states follow an exponential distribution during DS simulations.²¹ Given this distribution of first passage times, the central limit theorem⁵³ gives the distribution of $MFPT$ estimates in the limit of large sample size,

$$\frac{\sqrt{n}(\widehat{MFPT}_{\text{ds}} - E[FPT])}{\sqrt{V[FPT]}} \xrightarrow{D} \mathcal{N}(0, 1), \quad (7)$$

where n is the count of first passage time observations, $E[FPT] = MFPT$ is the expected value of the first passage time,

$V[FPT]$ is the variance, $\mathcal{N}(0, 1)$ is the standard normal distribution (i.e., mean 0 and variance 1), and $\xrightarrow{D}$ signifies that the distribution on the left converges to the one on the right. Equation (7) implies⁵³ that an estimate of $MFPT$ calculated using n observations will follow a normal distribution such that

$$\begin{aligned} E[\widehat{MFPT}_{\text{ds}}] &= MFPT, \\ V[\widehat{MFPT}_{\text{ds}}] &= \frac{V[FPT]}{n} = \frac{MFPT^2}{n}, \end{aligned} \quad (8)$$

where the fact that $V = E^2$ for an exponential distribution was used to factor out $V[FPT]$. The lower and upper bounds of the confidence interval of any normally distributed random variable X can be found using a standard formula⁵⁴ that depends on the first two moments of X ,

$$\begin{aligned} lb_\alpha[X] &= E[X] - z_\alpha \sqrt{V[X]}, \\ ub_\alpha[X] &= E[X] + z_\alpha \sqrt{V[X]}, \end{aligned} \quad (9)$$

where z_α is the z score associated with the confidence level α (e.g., $z_{0.95} \approx 1.96$).⁵⁵ Plugging Eqs. (8) and (9) into Eq. (6) yields the margin of error of the DS $MFPT$ estimator,

$$\zeta[\widehat{MFPT}_{\text{ds}}, \alpha] = \frac{z_\alpha}{\sqrt{n}}. \quad (10)$$

Thus, if say, 10^5 replicate trajectories are produced during a DS simulation, the resultant estimate of $MFPT$ will have no more than $\frac{z_{0.95}}{\sqrt{10^5}} \approx \frac{1.96}{316} = 0.62\%$ difference from the true value 95% of the time. A similar calculation can be performed for any observable.

E. Minimizing the computational cost required to achieve a desired error goal

We define the computational cost $\mathcal{C}$ of a simulation to be equivalent to the average in-simulation time required to complete it (alternatively, one may use the count of simulation steps). When estimating $MFPT$ using DS simulation, the computational cost is

$$\mathcal{C}_{\text{ds}} = n \cdot MFPT, \quad (11)$$

where n is the total number of replicate trajectories. Equation (10) illustrates the direct tradeoff between computational cost and simulation accuracy. In order to find the minimum number of replicate trajectories that are required to achieve a particular error goal in DS simulations, Eq. (10) can be solved for n ,

$$n = \left(\frac{z_\alpha}{\zeta} \right)^2. \quad (12)$$

III. RESULTS

A. Derivation of the FFPilot optimizing equation

In this subsection, we find the number of trajectories to launch in each FFS phase (which we will also refer to as sample count or n_i) that will minimize the simulation run time while fixing the margin of error ζ of the estimator of $MFPT$. We will refer to the $MFPT$ estimator in this subsection as $\widehat{W}$. The derivation of the optimal choice of n_i starts with the characterization of the moments and distribution of $\widehat{W}$. We then use the moments and distribution of $\widehat{W}$

to find ζ as a function of the per-phase sample counts n_i . We then use the method of Lagrange multipliers to find the optimal choice of n_i that will keep ζ fixed while minimizing the simulation run time.

1. The moments and distribution of the FFS MFPT estimator $\widehat{W}$

Each individual FFS phase weight w_i can be thought of as the first moment (i.e., the mean) of an observable $\mathbf{w}_i$ of a random process that is specific to phase i . By the end of every FFS phase i , n_i samples have been drawn from $\mathbf{w}_i$ (see Sec. II A for a more concrete description), at which point the phase weight w_i is estimated as

$$\widehat{w}_i = \frac{1}{n_i} \sum_{j=1}^{n_i} \mathbf{w}_{ij}.$$

Given the above form of $\widehat{w}_i$, and given that the observable $\mathbf{w}_i$ meets certain regularity conditions,⁵⁶ the asymptotic distribution of the individual $\widehat{w}_i$ terms can be determined from the central limit theorem,

$$\frac{\sqrt{n_i}(\widehat{w}_i - w_i)}{\sqrt{V[\mathbf{w}_i]}} \xrightarrow{D} \mathcal{N}(0, 1). \quad (13)$$

The moments of each $\widehat{w}_i$ can be determined by the appropriate interpretation of Eq. (13),

$$\begin{aligned} E[\widehat{w}_i] &= w_i, \\ V[\widehat{w}_i] &= \frac{V[\mathbf{w}_i]}{n_i}. \end{aligned} \quad (14)$$

An estimator with this type of convergence behavior is said to be $\sqrt{n}$ -consistent.

Given that the MFPT estimator $\widehat{W}$ is defined in Eq. (5) as a function $g[\widehat{\mathbf{w}}]$ of a vector of $\sqrt{n}$ -consistent estimators $\widehat{\mathbf{w}} = \{\widehat{w}_0, \widehat{w}_1, \dots, \widehat{w}_N\}$, we can apply the multivariate delta method⁵⁷ to determine the moments and distribution of $\widehat{W}$. From the multivariate delta method, we know that

$$\frac{\widehat{W} - w_0 \prod_{i>0} w_i^{-1}}{\sqrt{\nabla_{\mathbf{w}}^T \Sigma \nabla_{\mathbf{w}}}} \xrightarrow{D} \mathcal{N}(0, 1), \quad (15)$$

where $\nabla_{\mathbf{w}}$ is the gradient of $g(\widehat{\mathbf{w}})$ evaluated at $\mathbf{w}$, $\nabla_{\mathbf{w}}^T$ is the transpose of $\nabla_{\mathbf{w}}$, and Σ is the covariance matrix of the phase weight estimators $\widehat{w}_i$. We can determine the moments of $\widehat{W}$ by interpreting Eq. (15),

$$\begin{aligned} E[\widehat{W}] &= \frac{w_0}{\prod_{i>0} w_i} = MFPT, \\ V[\widehat{W}] &= \nabla_{\mathbf{w}}^T \Sigma \nabla_{\mathbf{w}}. \end{aligned} \quad (16)$$

A simpler form of $V[\widehat{W}]$ can be found. To begin with, we find $\nabla_{\mathbf{w}}$. In column vector form, it is

$$\nabla_{\mathbf{w}} = \begin{pmatrix} w_0^{-1} w_0 \prod_{j>0} w_j^{-1} \\ -w_1^{-1} w_0 \prod_{j>0} w_j^{-1} \\ \vdots \\ -w_N^{-1} w_0 \prod_{j>0} w_j^{-1} \end{pmatrix} = \begin{pmatrix} s_0 \\ -s_1 \\ \vdots \\ -s_N \end{pmatrix}, \quad (17)$$

where we have substituted s_i for $w_i^{-1} w_0 \prod_{j>0} w_j^{-1}$ for the sake of brevity. If the covariance of $\widehat{w}_i$ and $\widehat{w}_j$ is written as σ_{ij} , then the covariance matrix Σ is

$$\Sigma = \begin{pmatrix} \sigma_{00} & \sigma_{01} & \cdots & \sigma_{0N} \\ \sigma_{10} & \sigma_{11} & \cdots & \sigma_{1N} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{N0} & \sigma_{N1} & \cdots & \sigma_{NN} \end{pmatrix},$$

and the variance of $\widehat{W}$ can be written as

$$V[\widehat{W}] = (s_0 - s_1 \cdots - s_N) \begin{pmatrix} \sigma_{00} & \sigma_{01} & \cdots & \sigma_{0N} \\ \sigma_{10} & \sigma_{11} & \cdots & \sigma_{1N} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{N0} & \sigma_{N1} & \cdots & \sigma_{NN} \end{pmatrix} \begin{pmatrix} s_0 \\ -s_1 \\ \vdots \\ -s_N \end{pmatrix}.$$

If we impose the assumption of independence on all of the phase weight estimators $\widehat{w}_i$, then only the diagonal elements of the covariance are nonzero,

$$V[\widehat{W}] = (s_0 - s_1 \cdots - s_N) \begin{pmatrix} \sigma_{00} & 0 & \cdots & 0 \\ 0 & \sigma_{11} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sigma_{NN} \end{pmatrix} \begin{pmatrix} s_0 \\ -s_1 \\ \vdots \\ -s_N \end{pmatrix}.$$

Under this condition of independence, the form of $V[\widehat{W}]$ can be simplified considerably,

$$\begin{aligned} V[\widehat{W}] &= \sum_i s_i^2 \sigma_{ii}, \\ V[\widehat{W}] &= \sum_i w_i^{-2} \sigma_{ii} \frac{w_0^2}{\prod_{j>0} w_j^2}, \end{aligned}$$

and since $\sigma_{ii} = V[\widehat{w}_i]$,

$$V[\widehat{W}] = \frac{w_0^2}{\prod_{j>0} w_j^2} \sum_i \frac{V[\widehat{w}_i]}{w_i^2}.$$

Finally, plugging in substitutions from Eq. (4) and from Eq. (14), the variance of $\widehat{W}$ is

$$V[\widehat{W}] = MFPT^2 \sum_i \frac{V[\mathbf{w}_i]}{w_i^2 n_i}. \quad (18)$$

Alternatively, the variance of $\widehat{W}$ can be derived from the formula for the variance of a product of random variables (see the [supplementary material](#)). The expressions derived from each technique agree in the high sample count limit.

The result in Eq. (18) agrees with and is similar to the established result of Allen *et al.*⁴⁸ concerning the variance of estimates produced by a complete FFS simulation. Unlike earlier work, however, we have imposed no particular form on $\mathbf{w}_i$, the random process underlying each phase i , i.e., we do not assume here that it is a Bernoulli process. As will be seen in Sec. III A 4, this generalization allows us to study the contributions of phase 0 to the overall error of an FFS simulation for the first time.

2. Margin of error of $\widehat{W}$

Now, we derive a formula for $\zeta[\widehat{W}, \alpha]$, the margin of error of the MFPT estimator $\widehat{W}$. From Eq. (15), we know that $\widehat{W}$ follows a

normal distribution. The lower and upper confidence bounds of $\widehat{W}$ are found by plugging the moments of $\widehat{W}$ [given by Eqs. (16) and (18)] into the bound formulas for a normally distributed random variable [given by Eq. (9)],

$$\begin{aligned} ub_{\alpha}[\widehat{W}] &= MFPT\left(1 + z_{\alpha}\sqrt{\sum_i \frac{V[\mathbf{w}_i]}{w_i^2 n_i}}\right), \\ lb_{\alpha}[\widehat{W}] &= MFPT\left(1 - z_{\alpha}\sqrt{\sum_i \frac{V[\mathbf{w}_i]}{w_i^2 n_i}}\right). \end{aligned} \quad (19)$$

Plugging Eq. (19) into the margin of error definition [given by Eq. (6)] yields the desired margin of error,

$$\zeta[\widehat{W}, \alpha] = z_{\alpha}\sqrt{\sum_i \frac{V[\mathbf{w}_i]}{w_i^2 n_i}}. \quad (20)$$

3. Derivation of the general optimizing equation

As we can see from Eq. (20), there are many different choices of n_i that will give the same value of $\zeta[\widehat{W}, \alpha]$. What we really want to find is the optimal choice of n_i that will minimize the simulation run time while keeping $\zeta[\widehat{W}, \alpha]$ fixed. We can find a formula for this optimal choice using the method of Lagrange multipliers.⁵⁸

For the method of Lagrange multipliers, we need a function to minimize the target function $f[x]$ and a function to hold constant, the constraint equation $g[x]$. We use the total computational cost $\mathcal{C}$ as the target function, which for FFS is

$$f[n_i] = \mathcal{C}_{\text{ffs}} = \sum_{i=0}^N n_i c_i, \quad (21)$$

where c_i is the average computational cost per sample. For the constraint equation, we square both sides of Eq. (20) and set it equal to zero,

$$g[n_i] = \sum_i \frac{k_i}{n_i} - \frac{\zeta^2}{z_{\alpha}^2} = 0, \quad (22)$$

where

$$k_i = \frac{V[\mathbf{w}_i]}{w_i^2}.$$

Now that we have chosen a target and a constraint function, the next step of the method is to combine Eqs. (21) and (22) in order to write out the Lagrangian,

$$\mathcal{L} = \sum_i n_i c_i + \lambda \left(\sum_i \frac{k_i}{n_i} - \frac{\zeta^2}{z_{\alpha}^2} \right), \quad (23)$$

and then find its gradient $\nabla \mathcal{L}$,

$$\nabla \mathcal{L} = \{\partial_{\lambda} \mathcal{L}, \partial_{n_0} \mathcal{L}, \dots, \partial_{n_N} \mathcal{L}\}.$$

We find the values of each of the partial derivatives individually. Calculating $\partial_{\lambda} \mathcal{L}$ is trivial, and when finding each separate $\partial_{n_i} \mathcal{L}$, we can eliminate all but one term from both sums, since terms that do not depend on n_i will vanish,

$$\begin{aligned} \partial_{\lambda} \mathcal{L} &= \sum_i \frac{k_i}{n_i} - \frac{\zeta^2}{z_{\alpha}^2}, \\ \partial_{n_i} \mathcal{L} &= c_i - \frac{\lambda k_i}{n_i^2}. \end{aligned}$$

Now, we can write out the actual gradient

$$\nabla \mathcal{L} = \left\{ \sum_i \frac{k_i}{n_i} - \frac{\zeta^2}{z_{\alpha}^2}, \dots, c_i - \frac{\lambda k_i}{n_i^2}, \dots \right\}. \quad (24)$$

The next step is to set each component of the gradient Eq. (24) equal to zero and solve the resulting set of $N + 2$ equations for λ and each n_i . We begin by solving $\partial_{n_i} \mathcal{L} = 0$ for n_i in terms of λ ,

$$\begin{aligned} c_i - \frac{\lambda k_i}{n_i^2} &= 0, \\ n_i &= \left\{ -\frac{\sqrt{\lambda} \sqrt{k_i}}{\sqrt{c_i}}, \frac{\sqrt{\lambda} \sqrt{k_i}}{\sqrt{c_i}} \right\}. \end{aligned} \quad (25)$$

Next, we solve $\partial_{\lambda} \mathcal{L} = 0$ for $\sqrt{\lambda}$ by substituting in the positive expression for n_i found in Eq. (25),

$$\begin{aligned} \sum_i \frac{k_i}{\frac{\sqrt{\lambda} \sqrt{k_i}}{\sqrt{c_i}}} - \frac{\zeta^2}{z_{\alpha}^2} &= 0, \\ \sqrt{\lambda} &= \frac{z_{\alpha}^2}{\zeta^2} \sum_i \sqrt{c_i k_i}. \end{aligned} \quad (26)$$

Now, we eliminate λ from our expression for n_i in Eq. (25) by substituting in the expression for $\sqrt{\lambda}$ we found in Eq. (26),

$$n_i = \frac{z_{\alpha}^2}{\zeta^2} \sqrt{\frac{k_i}{c_i}} \sum_j \sqrt{c_j k_j}. \quad (27)$$

4. The optimizing equation for FFS

A form of the general optimizing equation given in Eq. (27) that is more specific to FFS can be found by considering the properties of the observable $\mathbf{w}_i$ in each phase. During a phase $i > 0$, each trajectory launched and finished is equivalent to a single sample taken from $\mathbf{w}_{i>0}$. The j th trajectory of a phase will either succeed (i.e., cross forward to the next interface) with probability p_{ij} , or it will fail (i.e., fall back into its starting basin), with probability $1 - p_{ij}$. The sample taken from $\mathbf{w}_{i>0}$ is 1, if the trajectory succeeds, and 0 otherwise. Thus, the outcome of the j th trajectory of phase i is a Bernoulli random variable with probability parameter p_{ij} .

For multidimensional systems, p_{ij} is dependent upon the starting point of a trajectory. Since the starting point of each trajectory is chosen at random, p_{ij} in general varies from trajectory to trajectory. Thus, each $\mathbf{w}_{i>0}$ is technically a mixture of Bernoulli random variables. However, for the purposes of the optimizing equation, it can be shown that no accuracy is lost if each $\mathbf{w}_{i>0}$ is treated as a single Bernoulli random variable (see Appendix A) with moments,

$$\begin{aligned} E[\mathbf{w}_{i>0}] &= w_i = p_i, \\ V[\mathbf{w}_{i>0}] &= p_i(1 - p_i), \end{aligned} \quad (28)$$

where p_i is the crossing probability in phase i [i.e., $P(\lambda_i | \lambda_{i-1})$].

The k_i terms in Eq. (27) can be expanded to yield

$$n_i = \frac{z_\alpha^2}{\zeta^2} \sqrt{\frac{V[\mathbf{w}_i]}{w_i^2 c_i}} \sum_j \sqrt{\frac{V[\mathbf{w}_j] c_j}{w_j^2}}. \quad (29)$$

The moments from Eq. (28) can then be plugged into Eq. (29) to yield the FFS specific form of the optimizing equation,

$$n_i = \begin{cases} \frac{z_\alpha^2}{\zeta^2} \sqrt{\frac{V[\mathbf{w}_0]}{w_0^2 c_0}} \left(\sqrt{\frac{V[\mathbf{w}_0] c_0}{w_0^2}} + \sum_{j=1}^N \sqrt{\frac{(1-p_j) c_j}{p_j}} \right) & \text{if } i = 0 \\ \frac{z_\alpha^2}{\zeta^2} \sqrt{\frac{1-p_i}{p_i c_i}} \left(\sqrt{\frac{V[\mathbf{w}_0] c_0}{w_0^2}} + \sum_{j=1}^N \sqrt{\frac{(1-p_j) c_j}{p_j}} \right) & \text{otherwise.} \end{cases} \quad (30)$$

We call this form the FFPilot optimizing equation.

In regards to phase 0, the precise forms of $E[\mathbf{w}_0] = w_0^2$ and $V[\mathbf{w}_0]$ are unknown. We have found that $\mathbf{w}_0$, the waiting time in between phase 0 forward flux events, does not in general follow an exponential distribution (see [supplementary material](#), Fig. S2). In fact, the distribution of $\mathbf{w}_0$ seems to be highly model dependent. Thus, in order to avoid any assumptions about $\mathbf{w}_0$, we leave the ratio $\frac{V[\mathbf{w}_0]}{w_0^2}$ unexpanded in Eq. (30).

Of the assumptions made in deriving Eq. (30), two of the most significant are the assumption of large sample size and the assumption of the uncorrelatedness of the phases during an FFS simulation. The large sample size assumption underlies the validity of Eqs. (14) and (15). In general, this assumption can be satisfied by setting a minimum floor on the number of samples n_i taken in each phase (an implementation of this sample size floor is discussed in Sec. III B).

Ensuring that the uncorrelatedness assumption is satisfied is altogether trickier. Most of the pre-existing FFS literature takes the uncorrelatedness of the phases as a given,^{47–49} but in practice we have found this to not always be the case. This issue is discussed in detail in Sec. III F. In brief, systems with complex, high dimensional state spaces can have correlations between the outcomes of trajectories across the different phases. This results in nonzero covariance between the different phase weights, which is effectively like adding an extra term to Eq. (18). In other words, when the phases are correlated, our approach will somewhat underestimate the actual variance, and Eq. (30) will somewhat underestimate the number of samples required to achieve a particular error goal.

B. FFPilot: A sampling algorithm designed to take advantage of the optimization equation

We wanted to be able to apply the optimizing equation to real biochemical networks, but to do so we need some prior knowledge of the system under study. As shown in Eq. (30), two global and $2 \cdot (N + 1)$ phase-specific parameters are required in order to apply the optimization equation and thereby calculate the optimal value of n_i . The two global parameters, the margin of error ζ and the confidence interval z score z_α , are independent of the model being studied and are set according to the desired error goal. The other parameters, the ratio $\frac{V[\mathbf{w}_0]}{w_0^2}$ (from phase 0), the successful crossing probabilities p_i (from phases $i > 0$), and the per-sample computational costs c_i , are

model dependent and vary for each different combination of model, order parameter, and interface placement.

In general, the exact values of the model-dependent parameters are unknown and estimates must be used instead. Rather than simplifying using assumptions, such as constant cost,^{48,49} we produce rough but conservative estimates of the necessary parameters using a pilot simulation. By conservative, we mean that the estimates, when plugged into the optimization equation, will be likely to give values of n_i that are at least as large as the true values. This condition ensures that simulations run with n_i trajectories per phase will produce results that are at least as accurate as the specified error goal.

We call our new enhanced sampling protocol FFPilot. The basic concept of FFPilot is to break an FFS simulation up into two stages. First, a pilot stage is executed (see [supplementary material](#), Fig. S1), from which the parameters required for the optimization equation are estimated. Then, based on the results of the optimization equation, a production stage is planned and executed, from which the actual simulation output is calculated.

The FFPilot algorithm proceeds as follows:

- Specify an error goal in terms of a target margin of error. Optionally, the confidence interval (which defaults to 95%) can be specified as well. As with standard FFS, the user must specify an order parameter and interface placements.
- Begin the FFPilot pilot stage:
 - Set the pilot stage sample count n_{pilot} to a single fixed value. Throughout this paper, we used a value of $n_{\text{pilot}} = 10^3$.
 - Run a complete FFS simulation, following the algorithm described in Sec. II A. Unlike standard FFS, the number of samples to collect in each phase is determined by a blind optimization method (see [Appendix B](#)).
 - Phase 0 proceeds the same as in standard FFS, using $n_0 = n_{\text{pilot}}$ as the sample count.
 - In phases $i > 0$, trajectories are run until n_{pilot} successful forward flux events are observed.
- When the pilot stage is finished, estimate the values required for the optimization equation from the results of the pilot simulation. Use confidence intervals to form conservative estimates that, when plugged into the optimization equation, are likely to yield values of n_i that are as large as or larger than those required for the error goal.
- Begin the FFPilot production stage:
 - Determine n_i , the number of samples to collect in each phase, based on the error goal and Eq. (30), the FFPilot optimizing equation (as parameterized in step 3).
 - Run another complete FFS simulation using the values of n_i calculated in step 4a.
- Collect results from the production stage simulation and use/analyze them in the same way as would be done for the standard FFS simulation. The results from the pilot stage are ignored for the purposes of calculating the final simulation results as sampling differences in states at the interfaces would introduce additional error.

In terms of the error in the final simulation results, the optimization equation is inaccurate in the low sample count limit. Therefore, we enforce a minimum floor (10^3) on the count of samples taken in each phase of the production stage.

The effectiveness of the pilot stage blind optimization method can be related to earlier findings of Glasserman *et al.*⁵⁹ and Borrero *et al.*⁴⁹ They showed that a fixed quantity of computational effort is optimally spent during an FFS simulation when the interfaces and the trajectory counts per phase are arranged in such a way that each interface encounters an equal flux of trajectories crossing them. Although we do not constrain the computational effort spent during our blind optimization, our approach produces equal flux across each interface as well.

C. Rare event model

We began our testing of FFPilot with a toy model of a barrier crossing process, which we refer to as the rare event model (REM). REM models a particle in a discrete potential field, in which there are two metastable states, $\mathcal{A}$ and $\mathcal{B}$, connected by a transition path (see top of Fig. 1). Particles in $\mathcal{A}$ have a constant propensity to initiate

a transition by entering the transition path. Because of the constant propensity, the waiting times in between transition attempts are exponentially distributed.

The transition path itself is composed of a sequence of N barriers; particles enter the path before barrier 1. At each successive barrier, a particle will instantaneously either proceed to the next barrier with probability p_i , or fall back into $\mathcal{A}$ with probability $1 - p_i$. If the particle successfully passes the final barrier, it enters $\mathcal{B}$. In effect, the particle's fate once it enters the transition path can be thought of as the outcome of N weighted coin flips. If all N coins land heads up, the particle completes the transition $\mathcal{A} \rightarrow \mathcal{B}$. Otherwise, the particle falls back into $\mathcal{A}$.

We designed REM to map precisely onto FFPilot sampling in order to directly test the validity of the assumptions and simplifications that were made in the derivation of the FFPilot optimizing equation [Eq. (27)]. Simulations of REM can be carried out using either DS or FFPilot. To simulate a single replicate of a particle starting in $\mathcal{A}$ using DS, first the time until the particle leaves $\mathcal{A}$ is randomly selected from an exponential random variable according to the propensity of entering the transition path. The particle's behavior at each barrier is then randomly chosen, either falling back to $\mathcal{A}$ or proceeding to the next barrier according to the appropriate p_i . If the particle falls back into $\mathcal{A}$, the process is repeated until it successfully passes to $\mathcal{B}$. The total time the particle took to transition to $\mathcal{B}$ is the $\mathcal{A} \rightarrow \mathcal{B}$ first passage time for that trajectory.

To simulate a single replicate of a particle starting in $\mathcal{A}$ using FFPilot, first an order parameter and the interface positions must be chosen (see section Sec. II A). We chose the barrier number i as the order parameter, and placed the interfaces between each barrier i . The phase 0 weight is calculated by first drawing many samples of the $\mathcal{A}$, leaving time according to the propensity, and then taking the mean of those samples. The remaining phase weights are determined by repeatedly starting a particle at barrier i , randomly selecting if it continues on to the next barrier according to p_i , and then calculating the average probability of success from the observations. The pilot stage of FFPilot is accomplished by first running the phase 0 weight calculation n_{pilot} times, then running each phase $i > 0$ weight calculation until n_{pilot} success events are observed. The outcome of the pilot stage is then fed into the FFPilot optimizing equation (30), and the results are used to determine how many samples to take during the FFPilot production stage. For the purposes of parameterizing Eq. (30), the phase 0 relative variance and the per-phase costs c_i are all set equal to 1 (see Table I for complete parameters). MFPT is then estimated as the product of the production stage phase weights.

We first studied the distribution of the MFPT estimators. Taken together, Eqs. (15), (16), and (18) describe the normal distribution that repeated estimation of MFPT is expected to produce. In our derivation, we have assumed that we are working in the high sampling limit for values of w_i and n_i of interest. To test this assumption, we performed $1.6 \cdot 10^5$ independent simulations of REM using both DS and FFPilot using a 1% error goal. Figure 2 shows the distributions from our simulations. The black line in each figure is the normal distribution with mean and variance given by Eqs. (16) and (18). The binned MFPT estimates from both the DS and the FFPilot REM simulations are in excellent agreement with the predicted normal distribution.

FIG. 1. Schematics of the model systems investigated. (Top) The Rare Event Model (REM, see Sec. III C and Table I for complete details). (Middle) Self-regulating Gene (SRG, see Sec. III D and Tables II and III). (Bottom) Genetic Toggle Switch (GTS^{EX}, see Sec. III E and supplementary material, Table S1).

TABLE I. Parameterization and other simulation data from our simplified Rare Event Model (REM). The phase weights w_i were copied from those of $GTS_{\theta=1}^{EX}$ (see Table V). All sample costs c were set to 1.

Phase	w_i^a	c_i^{fb}	c_i^{sc}	c_i^d	$n_i \times 10^{-3e}$
0	45.3	...	...	1	1834
1	0.092	1	1	1	2020
2	0.274	1	1	1	1047
3	0.136	1	1	1	1621
4	0.150	1	1	1	1527
5	0.242	1	1	1	1138
6	0.636	1	1	1	486
7	0.711	1	1	1	410
8	0.858	1	1	1	261
9	0.913	1	1	1	199
10	0.973	1	1	1	107
11	0.989	1	1	1	67
12	0.998	1	1	1	30
MFPT	$1.074 \cdot 10^6$				

^aThe phase weight of phase i .
^bThe expected cost of a failed phase i trajectory.
^cThe expected cost of a successful phase i trajectory.
^dThe expected cost of a phase i sample.
^eThe phase i sample count (in thousands) for a 1% error goal with 95% confidence.

Next, we tested how well the FFPilot approach was able to control sampling error in simulations of REM. If the method works as expected, 95% of simulations should have errors at or below the FFPilot error goal. We executed 1000 FFPilot simulations of REM at 3 different error goals (10%, 3.2%, and 1%). We used the full FFPilot algorithm to determine how many trajectories to start at each interface.

The percent errors of the $MFPT$ calculated in each of these simulations are shown in Fig. 3. The percent errors were calculated relative to the analytically determined $MFPT$. As can be seen, the 95th error percentiles (marked by the red lines) are located along $x = y$, indicating that the overall error in the $MFPT$ estimates was constrained to the error goal. REM has no source of error aside from

FIG. 2. $MFPT$ estimates for the rare event model (REM) calculated using (left) DS and (right) FFPilot. Shown are (blue bars) binned $MFPT$ estimates from $1.6 \cdot 10^5$ simulations with a 1% error goal and (black lines) $MFPT$ estimate distributions from Eq. (7) and Eq. (15), respectively.

FIG. 3. Actual error vs error goal of FFPilot simulations of the REM. Each strip shows the $MFPT$ estimation error from (blue dots) 1000 independent FFPilot simulations and (red lines) the 95th percentiles of the errors. Jitter was added to the x -position of the dots for visualization.

the sampling error, and under these conditions FFPilot precisely controls the total simulation error.

D. Selfregulating gene model

We next tested FFPilot with a relatively simple biochemical network, the selfregulating gene (SRG) model.⁶ SRG models the expression of a single protein A . A is produced through autocatalysis, and decays via a first order process (see Fig. 1).

In the deterministic formulation of SRG, the rate of change in the quantity of protein A is

$$\frac{dA}{dt} = k_{low} + (k_{high} - k_{low}) \frac{A^h}{k_{50}^h + A^h} - A. \quad (31)$$

For a given set of parameters, the fixed points of the state space of SRG can be found by setting Eq. (31) equal to 0 and solving for A . For all of the parameters we used in our simulations, there are three fixed points, two stable and one unstable. One of the stable fixed points corresponds to a state with a low count of A , and the other corresponds to a state with a high count of A .

To formulate SRG as a stochastic model, we use the chemical master equation (CME) (see Table II for complete reaction list). The CME models the probability for the system to be in any state. Additionally, fluctuations due to population noise can cause the system to transition back and forth between the low and high states. The

TABLE II. Reaction scheme for our Selfregulating Gene (SRG) models.

Reactions	Rates	Description
$\emptyset \rightarrow A$	$k_{low} + (k_{high} - k_{low}) \frac{A^h}{k_{50}^h + A^h}$	Expression
$A \rightarrow \emptyset$	βA	Decay

MFPT between the states is related to the entropic barrier separating them.

We wanted to study how the height of the barrier between the low *A* and high *A* states affects the accuracy of FFPilot. Towards this end, we parameterized 3 different variants of SRG with different barrier heights, and thus different *MFPT* values. We call these three variants SRG_{h=2.4}, SRG_{h=2.3}, and SRG_{h=2.2}, after the Hill coefficient used in the protein *A* production rate law. We tuned the other parameters in the model in order to approximately balance the occupancy of the low *A* and high *A* states in each of the variants (see Table III for parameter values).

For all FFPilot simulations of SRG, we used the count of protein *A* as the order parameter. We determined the positions of the interfaces by first placing λ_0 a quarter of the distance (in terms of the order parameter) from the lower fixed point to the intermediate fixed point, then λ_N three quarters of the distance from the lower fixed point to the upper fixed point. We then placed 11 more interfaces spaced evenly between λ_0 and λ_N .

Unlike REM, trajectories in the low *A* state do not cross λ_0 and enter the transition pathway with a fixed propensity. Instead, the propensity changes dynamically with the system state, giving rise to a complex distribution of waiting times in between crossing events. In deriving the FFPilot optimizing equation [more specifically, when deriving Eq. (18)], we assumed that the first two central moments of the phase 0 waiting time distribution (w_0 and $V[w_0]$) exist. In order to test this assumption, we executed simulations, in which we only performed phase 0, collecting 10^6 crossing events for λ_0 .

Figure 4 shows the phase 0 interevent time distributions. The tail of each distribution is fit well by a single exponential distribution, but the distribution near 0 is not. We estimated the value of the relative variance, $\frac{V[w_0]}{w_0^2}$, used in the phase 0 terms of the FFPilot optimizing equation [Eq. (30)] to be 6.80, 6.43, and 5.96 for SRG_{h=2.4}, SRG_{h=2.3}, and SRG_{h=2.2}, respectively.

Next, we examined how well the FFPilot approach was able to control the sampling error with respect to *MFPT*. We executed 1000 FFPilot simulations of SRG_{h=2.4}, SRG_{h=2.3}, and SRG_{h=2.2} using error

TABLE III. Parameterizations and other simulation data from our SRG models. Per-phase weights, costs, and sample counts were estimated from the average results of 1000 FFPilot simulations (1% error goal). *MFPT* values were estimated from the results of 10^5 DS simulations, plus or minus the 95% confidence interval bounds.

Parameter	SRG _{h=2.4}						SRG _{h=2.3}						SRG _{h=2.2}					
<i>h</i>	2.4						2.3						2.2					
<i>k_{low}</i>	10						10						10					
<i>k_{high}</i>	200						200						200					
<i>k₅₀</i>	91						92.5						94					
β	1						1						1					
$\mathcal{A}$	12						12						11					
$\mathcal{B}$	155						156						157					
Phase	λ_i^a	w_i^b	c_i^c	c_i^{sd}	c_i^e	$n_i \times 10^{-3f}$	λ_i	w_i	c_i^f	c_i^s	c_i	$n_i \times 10^{-3}$	λ_i	w_i	c_i^f	c_i^s	c_i	$n_i \times 10^{-3}$
0	23.0	9.66	...	...	9.66	641	23.0	6.89	...	...	6.89	703	23.0	4.78	...	...	4.78	762
1	33.6	0.007	0.130	0.913	0.135	26 395	34.1	0.010	0.137	0.936	0.150	18 975	34.6	0.018	0.138	0.846	0.151	12 067
2	44.2	0.071	1.56	1.21	1.53	2 383	45.2	0.112	1.65	1.20	1.60	1620	46.2	0.140	1.64	1.30	1.60	1 253
3	54.8	0.277	2.89	1.14	2.41	850	56.2	0.296	2.99	1.24	2.47	714	57.8	0.354	3.05	1.23	2.41	556
4	65.3	0.565	4.22	1.30	2.57	447	67.3	0.580	4.35	1.28	2.57	387	69.3	0.572	4.45	1.39	2.70	337
5	75.9	0.855	5.45	1.12	1.75	254	78.4	0.825	5.63	1.23	2.00	237	80.9	0.814	5.78	1.23	2.08	212
6	86.5	0.962	6.48	1.03	1.24	147	89.5	0.946	6.75	1.08	1.39	147	92.5	0.924	7.01	1.24	1.67	141
7	97.1	0.993	7.35	0.844	0.890	73	101	0.986	7.71	0.926	1.02	85	104	0.977	8.12	1.09	1.25	88
8	108	0.999	8.11	0.700	0.709	34	112	0.997	8.56	0.836	0.862	46	116	0.993	9.07	0.905	0.959	53
9	118	1.00	8.83	0.767	0.769	16	123	0.999	9.42	0.825	0.833	25	127	0.997	10.0	0.981	1.00	32
10	129	1.00	...	0.779	0.779	8	134	1.00	10.4	0.894	0.897	14	139	0.999	11.0	0.955	0.965	20
11	139	1.00	...	1.05	1.05	5	145	1.00	...	1.07	1.07	9	150	0.999	12.3	1.24	1.25	14
12	150	1.00	...	1.33	1.33	3	156	1.00	...	1.44	1.44	6	162	1.00	13.8	1.48	1.49	10
<i>MFPT</i>	$(1.521 \pm 0.0094) \cdot 10^5$						$(4.661 \pm 0.029) \cdot 10^4$						$(1.271 \pm 0.0079) \cdot 10^4$					

^aThe position of interface *i* in terms of order parameter *A*.

^bThe phase weight of phase *i*.

^cThe expected cost of a failed phase *i* trajectory.

^dThe expected cost of a successful phase *i* trajectory.

^eThe expected cost of a phase *i* sample.

^fThe phase *i* sample count (in thousands) for a 1% error goal with 95% confidence.

FIG. 4. Waiting times between forward crossing events during phase 0 of forward flux simulations of the selfregulating gene (SRG) models. Each line shows a distribution calculated from 10^6 samples from a single phase 0 trajectory of (blue) $\text{SRG}_{h=2.4}$, (orange) $\text{SRG}_{h=2.3}$, and (green) $\text{SRG}_{h=2.2}$.

goals of 1%, 3.2%, and 10%. We estimated *MFPT* from each simulation and then found the percent error relative to the value estimated from a DS simulation executed with an error goal of 0.62%.

The errors are shown in Fig. 5. The 95th error percentiles are again located precisely along $x = y$. The accuracy of the estimates shows that FFPilot is able to control error in both phase 0 (regardless of the exact distribution of the interevent times) and the remaining phases for SRG.

We also looked at the contribution of phase 0 to the overall cost of the pilot stage. Applying Eq. (21) to values from Table III, we found that, for all variants of SRG, phase 0 required around ~35% of the total simulation time. This finding is in contrast to the long-standing assumption in the FFS literature that phase 0 does not significantly contribute to the cost of a simulation and should therefore be extensively sampled.⁴⁸

E. Genetic toggle switch model

The last model we investigated using FFPilot was a more complex gene regulatory network, one of a family of extensively studied systems commonly referred to as a genetic toggle switch (GTS).^{60–67} Our GTS has eight species that interact with one another via 24 reactions, all of which are first or second order (see Table IV). GTS consists of a single piece of operator DNA, *O*. When *O* is not bound to anything, it can produce either of two proteins, *A* or *B*. *A* and *B* can both form homodimers and those dimers can both bind back to *O*. When *O* is bound to a dimer of a single type of protein, it can only produce more of that same protein. If it is bound with dimers of both types of proteins, it is inactive (see Fig. 1). Proteins can decay from any state. This model is similar to the general toggle switch, e.g., as described by Warren and ten Wolde.⁶¹

The combination of positive feedback (of monomer production on dimer/operator binding) and negative feedback (of

FIG. 5. Actual error vs error goal of FFPilot simulations for (top) $\text{SRG}_{h=2.4}$, (middle) $\text{SRG}_{h=2.3}$, and (bottom) $\text{SRG}_{h=2.2}$. Each strip shows (blue dots) 1000 independent FFPilot simulations and (red lines) the 95th percentiles of the errors.

dimer/operator binding on production of the competing monomer) gives GTS bistable dynamics. The system as a whole switches between a state with high levels of *A* and low levels of *B*, and a state with low levels of *A* and high levels of *B*.

For GTS, we defined two order parameters. One, which we called Δ , is the difference between the total count of protein *B* and the total count of protein *A*. The other, which we called Σ , is the sum of the total count of protein *A* and the total count of protein *B*. In

TABLE IV. Reaction scheme for our Genetic Toggle Switch (GTS^{EX}) models.

Reactions		Rate constants				Description
		General		Exclusive		
Protein A	Protein B	Forward	Reverse	Forward	Reverse	
$2A \rightleftharpoons A_2$	$2B \rightleftharpoons B_2$	10	10	5	5	Dimerization
$O + A_2 \rightleftharpoons OA_2$	$O + B_2 \rightleftharpoons OB_2$	50	10	5	1	Operator binding
$OB_2 + A_2 \rightleftharpoons OA_2B_2$	$OA_2 + B_2 \rightleftharpoons OA_2B_2$	50	10	...	...	Operator binding
$O \longrightarrow O + A$	$O \longrightarrow O + B$	α		θ		Expression
$OA_2 \longrightarrow OA_2 + A$	$OB_2 \longrightarrow OB_2 + B$	α		θ		Bound expression
$A \longrightarrow \emptyset$	$B \longrightarrow \emptyset$	1		$\theta/4$		Decay
$A_2 \longrightarrow \emptyset$	$B_2 \longrightarrow \emptyset$	1/2		...		Decay
$OA_2 \longrightarrow O$	$OB_2 \longrightarrow O$	1/2		...		Decay
$OA_2B_2 \longrightarrow OB_2$	$OA_2B_2 \longrightarrow OA_2$	1/2		...		Decay

terms of the underlying species counts, the order parameters can be written as

$$\Delta = B + 2B_2 + 2OB_2 - (A + 2A_2 + 2OA_2),$$
$$\Sigma = A + 2A_2 + 2OA_2 + B + 2B_2 + 2OB_2 + 4OA_2B_2,$$

where A and B are the monomer counts, A_2 and B_2 are the dimer counts, and OA_2 , OB_2 , and OA_2B_2 are the dimer-operator complex counts.

Although all of our GTS simulations are based on a stochastic master equation formulation of the system, it is helpful to consider the more straightforward deterministic formulation when trying to understand the system's overall behavior (see [supplementary material](#), Table S1 for the deterministic rate equations). For a given set of parameters, the fixed points of the deterministic formulation can be found. For all of the parameter sets we used in our simulations, there are three fixed points in terms of Δ , two stable fixed points and one unstable fixed point. One of the stable fixed points corresponds to the state with a high level of A and a low level of B , and the other stable fixed point corresponds to the state with a low level of A and a high level of B . We refer to these two states as $\mathcal{A}$ and $\mathcal{B}$, respectively.

We parameterized the protein birth rate as α , with higher birth rates leading to higher steady state values of A and B . Higher steady states result in longer MFPT for switching. We used three

different variants of GTS in our simulations, $GTS_{\alpha=150}$, $GTS_{\alpha=400}$, and $GTS_{\alpha=650}$. [Figure 6](#) shows the stationary probability density function (PDF) for the system, note that the position of the steady states $\mathcal{A}$ and $\mathcal{B}$ increases as does α .

For FFPilot simulations of GTS, we used Δ as the order parameter. We tiled the state space in terms of Δ by placing interface λ_0 three standard deviations below the mean of $\mathcal{A}$ at $(\mathcal{A} - 3\sqrt{\mathcal{A}})$, interface λ_N at $(\mathcal{B} - 3\sqrt{\mathcal{B}})$, and then filling the gap with interfaces evenly spaced about 20 units of Δ apart.

We next ran a series of simulations to examine how well the full FFPilot protocol was able to control the sampling error in estimations of MFPT of the $\mathcal{A} \rightarrow \mathcal{B}$ switching process. We executed FFPilot simulations of $GTS_{\alpha=150}$, $GTS_{\alpha=400}$, and $GTS_{\alpha=650}$ using error goals of 1%, 3.2%, and 10%. We estimated MFPT for each simulation and then found the percent errors relative to the results from high accuracy DS simulations of equivalent models, which were run with a 0.62% error goal. The percent errors of the MFPT estimates are shown in [Fig. 7](#). Each 95th error percentile (marked by the red lines) falls exactly along the $x = y$ line, indicating that FFPilot is precisely controlling error in these simulations of a multidimensional model.

As the dynamics along the transition path are complicated in the multidimensional GTS model, we wanted to know how the cost of launching trajectories varied among the phases.

FIG. 6. Stationary probability distributions for (left) $GTS_{\alpha=150}$, (center) $GTS_{\alpha=400}$, and (right) $GTS_{\alpha=650}$.

FIG. 7. Actual error vs error goal of FFPilot simulations for (top) $GTS_{\alpha=150}$, (middle) $GTS_{\alpha=400}$, and (bottom) $GTS_{\alpha=650}$. Each strip shows (blue dots) independent FFPilot simulations and (red lines) the 95th percentiles of the errors.

Figure 8 shows the cost of launching a trajectory at each simulation phase. We again found that phase 0 contributed significantly to the cost of FFPilot simulations. Furthermore, we found that early phases are cheap to run and the cost increases dramatically near the transition barrier. Indeed, we saw that the cost of a trajectory varied by two orders of magnitude in some cases. As discussed above, FFPilot optimizes the total cost of the simulation accounting for differences in trajectory costs along the transition path. It runs more cheap trajectories and fewer expensive trajectories to optimally balance cost and accuracy. Other methods for

FIG. 8. Cost in simulation steps of calculating a trajectory by phase for (top) $GTS_{\alpha=150}$, (middle) $GTS_{\alpha=400}$, and (bottom) $GTS_{\alpha=650}$.

optimizing FFS that do not account for variable costs may exhibit slower performance.

To study the efficiency gains of FFPilot, we performed a performance comparison of FFPilot both against traditional FFS with a matched margin of error and FFS with constant flux, which in the case of equal cost has been shown to be the optimal strategy.^{49,50} We also compared against DS. It is straightforward to use Eq. (11) to calculate the number of DS trajectories needed to make a comparison against FFPilot. For FFS, however, the comparison is more difficult. There is no general way (apart for our optimizing equation) to determine the number of trajectories to launch at each phase to achieve a specific error goal. We settled for a FFS scenario in which we executed an identical number of trajectories N at each phase and then

calculated the margin of error with 95% confidence vs N empirically. We ran a large number of FFS simulations at several values of N and then extrapolated this relationship to the desired error goal (see Fig. S3). We did the same for constant flux FFS with the number of successful trajectories.

We ran simulations at 1%, 3.2%, and 10% error for each of our GTS models using all four methods. We then calculated the total number of simulation steps that were required for each, which is a fair readout of simulation cost. Figure 9 shows the results of the comparison for 1% error and Fig. S4 for all three error goals. FFPilot ranges from 30X to 300X faster than DS, with the speedup increasing exponentially with the barrier height, as expected.

Compared to traditional FFS, FFPilot has a performance improvement of 60%–120%. The constant flux method is better than traditional FFS, but FFPilot still outperforms it, being 15%–50% faster. These speedups mean that by increasing the number of cheap trajectories and decreasing the number of expensive trajectories to

be launched, FFPilot can offer a significant improvement in runtime while maintaining the same error.

In Fig. S5, we show the cumulative cost of each simulation type broken down by phase. FFPilot is faster than traditional FFS throughout the entire simulation. Compared to constant flux FFS, though, FFPilot allocates more of the total simulation time in the earlier simulation phases where trajectories are cheap to execute. At the transition point, trajectories are the most expensive and FFPilot crosses over constant flux in terms of performance by limiting the number of these expensive trajectories that need to be run.

However, this analysis overlooks the fact that without running a sample set of simulations to calculate the optimal N , it is impossible for a user to know how many trajectories to specify for a traditional FFS or constant flux FFS simulation. Thus, the utility of FFPilot over traditional FFS is not only the runtime improvement, but the fact that it automatically calculates the best way for a FFS simulation to be configured to achieve a desired level of statistical error.

F. Correlation error in the exclusive genetic toggle switch

Finally, it is well known that in high dimensional models there can be non-Markovian memory effects that introduce correlations across FFS interfaces.⁴⁸ Trajectory reweighting⁶⁸ and flat histogram methods, such as flatPERM,⁶⁹ have been used to account for similar correlations in various enhanced sampling contexts. A detailed study of how these effects influence error in FFS simulation of a nonequilibrium biochemical networks has not been presented.

Therefore, we studied the performance of FFPilot in the case where memory effects are significant. There is a variant of the GTS, known as the exclusive genetic toggle switch (GTS^{EX}), in which only one protein dimer can bind to the operator at a time. In this case, switching is driven also by the rare dissociation event of the dimer, not only the location of the system in phase space. We formulated a version of GTS^{EX} by removing the multiple binding reactions from our GTS model and also allowing protein decay only from the monomer form, to be consistent with previous work⁶¹ (see Table IV).

We wanted to be able to tune the rarity of the $A \rightarrow B$ event without changing the steady state of GTS^{EX} . To do so, we added a relative protein turnover parameter θ . The rate constants of all of the expression and decay reactions for both A and B are multiplied by θ . Since θ does not affect the birth/death ratio of each protein, the steady state levels of both A and B are constant with respect to θ . However, θ does have a large effect on the rate of $A \rightarrow B$ switching. We used three different variants of GTS^{EX} in our simulations, $GTS_{\theta=1}^{EX}$, $GTS_{\theta=10}^{EX}$, and $GTS_{\theta=100}^{EX}$. For FFPilot simulations, we used Δ as the order parameter and tiled the state space in terms of Δ by placing λ_0 at $\Delta = -27$, λ_N at $\Delta = 27$, and then placing 11 more interfaces evenly spaced in between. See Table V for complete parameters.

We next ran a test to examine how well the full FFPilot protocol was able to control the sampling error in estimations of $MFPT$ of the $A \rightarrow B$ switching process. We executed 1000 FFPilot simulations of $GTS_{\theta=1}^{EX}$, $GTS_{\theta=10}^{EX}$, and $GTS_{\theta=100}^{EX}$ using error goals of 1%, 3.2%, and 10%. We estimated $MFPT$ for each simulation and then

FIG. 9. (Top) Simulation steps required to achieve 1% margin of error with 95% confidence vs $MFPT$ for three GTS models. The mean numbers of steps are shown for (red) DS, (green) traditional FFS, (magenta) constant flux FFS, and (blue) FFPilot. (Bottom) The speedup of FFPilot relative to (red) DS, (green) traditional FFS, and (magenta) constant flux FFS. Connecting lines are added for illustration.

TABLE V. Parameterizations and other simulation data from our GTS^{EX} models. Per-phase weights, costs, and sample counts were estimated from the average results of 1000 FFPilot simulations (1% error goal). MFPT values were estimated from the results of 10⁵ DS simulations, plus or minus the 95% confidence interval bounds.

Parameter	GTS ^{EX} _{θ=1}						GTS ^{EX} _{θ=1}						GTS ^{EX} _{θ=10}					
θ	0.1						1						10					
$\mathcal{A}$	$O = 0, OA_2 = 1, OB_2 = 0, A = 4, B = 0, A_2 = 16, B_2 = 0$																	
$\mathcal{B}$	$O = 0, OA_2 = 0, OB_2 = 1, A = 0, B = 4, A_2 = 0, B_2 = 16$																	
Phase	λ_i^a	w_i^b	c_i^c	c_i^d	c_i^e	$n_i \times 10^{-3f}$	w_i	c_i^f	c_i^g	c_i	$n_i \times 10^{-3}$	w_i	c_i^f	c_i^g	c_i	$n_i \times 10^{-3}$		
0	−27.0	453	...	...	453	1664	45.3	...	...	45.3	1337	4.52	...	...	4.52	985		
1	−22.5	0.092	19.7	63.5	23.7	8010	0.092	1.97	6.35	2.37	6436	0.090	0.197	0.643	0.238	4698		
2	−18.0	0.271	148	113	138	1728	0.274	14.8	11.2	13.8	1382	0.281	1.51	1.12	1.40	977		
3	−13.5	0.128	307	176	290	1897	0.136	30.7	17.1	28.8	1482	0.170	3.13	1.46	2.85	946		
4	−9.0	0.111	442	120	406	1738	0.150	44.0	10.5	38.9	1201	0.339	4.46	0.597	3.15	568		
5	−4.5	0.082	536	131	503	1849	0.242	52.6	10.1	42.3	859	0.676	4.98	0.529	1.97	356		
6	0.0	0.437	659	118	423	684	0.636	61.4	8.93	28.0	451	0.872	5.14	0.463	1.06	268		
7	4.5	0.632	830	167	411	467	0.711	73.6	12.7	30.3	365	0.873	5.83	0.657	1.32	240		
8	9.0	0.861	994	140	259	310	0.858	85.1	11.1	21.6	276	0.918	6.43	0.599	1.08	208		
9	13.5	0.940	1153	165	225	209	0.913	97.3	14.4	21.6	210	0.925	7.08	0.861	1.33	178		
10	18.0	0.988	1283	122	135	116	0.973	108	11.5	14.1	140	0.960	7.58	0.809	1.08	141		
11	22.5	0.997	1397	156	159	53	0.989	117	15.4	16.5	81	0.972	8.14	1.27	1.46	102		
12	27.0	1.00	1493	150	150	19	0.998	126	14.9	15.1	37	0.989	8.61	1.36	1.44	63		
MFPT	$(7.015 \pm 0.043) \cdot 10^7$						$(1.074 \pm 0.0067) \cdot 10^6$						$7.701 \pm 0.048) \cdot 10^3$					

^aThe position of interface i in terms of order parameter Δ .^bThe phase weight of phase i .^cThe expected cost of a failed phase i trajectory.^dThe expected cost of a successful phase i trajectory.^eThe expected cost of a phase i sample.^fThe phase i sample count (in thousands) predicted by FFPilot to achieve a 1% error goal, 95% confidence.

found the percent errors relative to the results from high accuracy DS simulations.

The percent errors of the MFPT estimates are shown in Fig. 10. As can be seen in the figure, the 95th error percentiles (marked by the red lines) are located somewhat above $x = y$, indicating that the overall errors in the estimated MFPT values are above the desired errors. The 95th percentile lines do decrease along with error goal, implying that FFPilot partially but not completely controls error in simulations of GTS^{EX}. The anomalous dispersion decreases as the height of the barrier between $\mathcal{A}$ and $\mathcal{B}$ in probability space increases. This implies that the extra error is caused by a system dependent property and is not directly related to undersampling.

The phase weights from the FFPilot simulation can also be used to reconstruct the stationary PDF of the system. We used the output of a 10% error goal simulation from each of GTS^{EX}_{θ=1}, GTS^{EX}_{θ=1}, and GTS^{EX}_{θ=10} to generate the PDFs shown in Fig. 11 and Fig. S21. As can be seen from a comparison with the PDFs generated from extensive direct sampling, FFPilot reproduces a highly accurate PDF, especially in the low probability transition region. Thus, we conclude that the excess error in the MFPT estimate does not come from

incorrect sampling of the stationary PDF. We also see the two stable fixed points we expect and a transition path between them with a barrier that increases in height as θ decreases. Thus, the extra error also does not come from additional stable states, as other authors have observed for GTS variants.^{62,63,67}

We sought to further quantify the cause of the extra error in the GTS^{EX} simulations. We chose the condition with the largest anomalous errors, GTS^{EX}_{θ=10} executed with an error goal of 10%, and examined the phase weight estimates produced by each of the 1000 replicate simulations we had run with that condition. These phase weights are shown in the top half of Fig. 12. The phase weights estimated by an equivalent FFPilot simulation run with an error goal of 0.1% are shown as dashed lines, and serve as a point of reference (there is no exact method for extracting the phase weights from a DS simulation). The dispersion of phase weight estimates around the reference weight is much greater in certain phases, especially phases 4 and 5. By itself, this is not an indication that FFPilot is failing to correctly estimate sampling counts for these phases. By design, FFPilot allows for different levels of dispersion in different phases when it is determining the optimal simulation plan.

FIG. 10. Actual error vs error goal of FFPilot simulations for (top) $\text{GTS}_{\theta=1}^{\text{EX}}$, (middle) $\text{GTS}_{\theta=10}^{\text{EX}}$, and (bottom) $\text{GTS}_{\theta=10}^{\text{EX}}$. Each strip shows (blue dots) 1000 independent FFPilot simulations and (red lines) the 95th percentiles of the errors.

In order to determine how much of the phase weight dispersion represents FFPilot functioning as intended and how much of the dispersion is truly anomalous, we calculated an optimal set of error goal targets for each simulation phase. From the FFPilot optimizing equation [Eq. (30)], we derived analytic expressions for the

per-phase error goals,

$$\begin{aligned}\zeta_{i=0} &= \zeta \sqrt{\frac{\sqrt{\frac{V[\mathbf{w}_0]c_0}{w_i^2}}}{\sqrt{\frac{V[\mathbf{w}_0]c_0}{w_i^2}} + \sum_{j=1}^N \sqrt{\frac{(1-p_j)c_j}{p_j}}}}, \\ \zeta_{i>0} &= \zeta \sqrt{\frac{\sqrt{\frac{(1-p_i)c_i}{p_i}}}{\sqrt{\frac{V[\mathbf{w}_0]c_0}{w_i^2}} + \sum_{j=1}^N \sqrt{\frac{(1-p_j)c_j}{p_j}}}}.\end{aligned}\quad (32)$$

Just as with the overall error goal, in any given phase $100 \cdot \alpha\%$ percent of simulations will have a level of sampling error in the phase weight estimate at or below the relevant per-phase error goal ζ_i . We parameterized Eq. (32) using the phase weight and phase cost estimates from the very high accuracy (0.1% error goal) FFPilot simulation of $\text{GTS}_{\theta=10}^{\text{EX}}$ mentioned above.

The error goal targets we calculated are shown as the dashed lines in the bottom half of Fig. 12. The phase weight percent errors (calculated relative to the estimates from the 0.1% error goal simulation) are shown as dots, and the red lines mark the 95th percentiles of the errors. If the red line and the dashed line overlap for a particular phase, it means that the error in this phase is dominated by the sampling error, which FFPilot is able to completely account for. If the red line is above the dashed line, then there is more error occurring in that phase than was predicted by FFPilot, and the magnitude of the separation of the two lines represents the magnitude of the anomalous (as opposed to the predicted) error. Interestingly, FFPilot estimates the majority of phase weights to within the desired error goal. The extra error in the MFPT estimate appears to be due primarily to extra error in only three of the phase weight estimates, those from phases 3–5. Furthermore, the bulk of the extra error is concentrated in just two of those phase weight estimates, those from phases 4 and 5. Interfaces λ_4 and λ_5 also happen to be the interfaces immediately preceding the transition midpoint.

We hypothesized that there must be some particular feature of the state space landscape of GTS^{EX} that the simulations are exploring during phases 4 and 5 that is responsible for the anomalous error. We further reasoned that the same features that are responsible for what previous authors have called memory correlations and landscape variance⁴⁸ could be the source of the excess error. Errors of these types are due to two factors: (1) misrepresentation of some regions of the state space in the set of trajectory starting points at λ_i and (2) significant differences in $P(\lambda_i|\lambda_{i-1})$ as a function of trajectory starting state. The total probability factor $P(\lambda_i|\lambda_{i-1})$ that is measured in each phase $i > 0$ can be thought of as a mixture of many independent probabilities, one for each state along λ_{i-1} , weighted by the (normalized) count of times the state is used as a starting point for a phase i trajectory. If either of factors (1) or (2) occurs alone, $P(\lambda_i|\lambda_{i-1})$ will still be correctly estimated. However, if the factors occur together, they can lead to significant errors. In other words, if the landscapes assembled at λ_3 and λ_4 are heterogeneous across replicate simulations, and if differences in those landscapes can lead to differences in the effective value of $P(\lambda_i|\lambda_{i-1})$, then landscape error may be the cause of the anomalous simulation error we observe in our simulations of GTS^{EX} . Of particular importance is the fact that

FIG. 11. Stationary probability distributions for (left) $GTS_{\theta=1}^{EX}$, (center) $GTS_{\theta=1}^{EX}$, and (right) $GTS_{\theta=10}^{EX}$. The top row was calculated using FFPilot and the bottom row using direct sampling.

FIG. 12. $GTS_{\theta=10}^{EX}$ phase weight estimates and errors. (top) The phase weights as estimated by (blue dots) 1000 independent FFPilot simulations run to a 10% error goal. The dashed black lines show the phase weights from an FFPilot simulation with a 0.1% error goal. (bottom) The blue dots show the percent error of each phase weight estimate given in the top subplot, relative to the 0.1% error goal simulation. Also shown are (red lines) the observed 95th error percentiles, and (dashed black lines) the expected 95th error percentiles.

the landscape error in phase i is due to errors in the landscape assembled from the endpoints of successful trajectories launched during phase $i - 1$. Thus, no amount of extra sampling performed during phase i alone can completely abolish the landscape error.

To facilitate further analysis, we defined an additional order parameter, Ω , which takes a value from $[-1, 1]$ based solely on the state of the single operator: $\Omega = -OA_2 + OB_2$. Equivalently, Ω can be said to have one of three categorical values: $\Omega \rightarrow \{OA_2, O, OB_2\}$.

We wanted to test if the conditions for landscape error were present in FFPilot simulations of $GTS_{\theta=10}^{EX}$. To this end, we chose two simulations, which we will call replicate 6 and replicate 8, that had a large divergence in their $MFPT$ estimates. Looking at the phase weight estimates produced by these two simulations (Fig. 13), the divergence in the $MFPT$ estimates can be seen to be mostly due

FIG. 13. The phase weights of two replicates: (blue x) replicate 6 and (gold +) replicate 8, from a $GTS_{\theta=10}^{EX}$ FFPilot simulation run to an error goal of 10%. Circles show weights for phases 4 and 5 calculated via an alternative method using Eq. (33).

to divergence in the phase weight 4 and 5 estimates (as expected). Exploring phase 5 in greater depth, we calculated the landscape occupancy along λ_4 in terms of the orthogonal order parameter Σ . The λ_4 occupancies for replicates 6 and 8 $P(\Sigma|\Omega, \lambda_4)$ (which are binned by Σ and operator state Ω) are shown in the lower left-hand subplots of Fig. 14 as lines colored blue or gold, respectively.

FIG. 14. Breakout of the calculation of phase weight 5 for replicates 6 and 8. Each group of three plots shows data from a slice of genetic toggle switch (GTS) state space with a different Ω , the state of the operator DNA. Each subplot within a group shows different probability data. (dots) The probability of a trajectory launched from a point on λ_4 fluxing forward to λ_5 vs the orthogonal order parameter Σ . (lines) The occupancy along Σ at interface λ_4 for (blue) replicate 6 and (gold) replicate 8. (bars) The cumulative probability of fluxing forward across λ_5 when starting at the specified Ω for (blue) replicate 6 and (gold) replicate 8.

Replicate 6 has somewhat higher occupancy in the OA_2 and O states, and replicate 8 has higher occupancy in the OB_2 states. Thus, condition (1) for landscape error in phase 5, heterogeneous occupancies along λ_4 , is indeed satisfied.

Next, we launched 10^6 independent trajectories from each starting state along λ_4 that had nonzero occupancy in either replicate 6 or 8. Just as in a normal FFPilot simulation, we stopped each trajectory when it either reached the next interface or fell back into the initial basin, and we took note of the stopping states. This gave us a highly accurate estimate of $P(\lambda_5|\lambda_4)$ as a function of trajectory starting state. These state-dependent probability values $P(\lambda_5|\Sigma, \Omega, \lambda_4)$ (which are also binned by Σ and operator state Ω) are shown as filled circles in Fig. 14. $P(\lambda_5|\Sigma, \Omega, \lambda_4)$ varies a great deal across both Σ and Ω , meaning that condition (2) is also satisfied for our simulations of $GTS_{\theta=10}^{EX}$, and that landscape error is indeed a possible explanation for the anomalous error.

In order to test if landscape error alone is a sufficient explanation for the observed anomalous error, we recalculated the phase 4 and 5 weights of replicates 6 and 8 from their landscapes alone, according to

$$P(\lambda_5|\lambda_4) = \sum_{\Sigma} \sum_{\Omega} [P(\lambda_5|\Sigma, \Omega, \lambda_4) \times P(\Sigma, \Omega|\lambda_4)]. \quad (33)$$

If landscape error is indeed the sole cause of the anomalous error, we expected that the phase weights derived from Eq. (33) would closely match those originally estimated by replicates 6 and 8, even though these original estimates vary greatly between the replicates. In this view, the phase $i > 0$ weight estimate produced by each simulation converges (with increasing trajectory count) to a unique value of $P(\lambda_i|\lambda_{i-1})$, as determined by their heterogeneous samples of the landscape along λ_{i-1} . The recalculated phase weight values are plotted as empty circles in Fig. 13, and they do indeed closely agree with the original estimates. Thus, we conclude that sampling error is indeed being handled correctly by FFPilot, and the anomalous error in our $GTS_{\theta=10}^{EX}$ simulation results is due to landscape error.

C. Relationship between correlation error and phase weight covariance

To further investigate the landscape error, we considered whether covariance between the phase weight estimators showed the extra error. Recall that in deriving Eq. (18), we assumed that the phase weight estimators $\hat{w}_i$ were independent and that the off-diagonal elements of the covariance matrix Σ were zero. When this assumption does not hold, there will be additional unaccounted for variance in the MFPT estimator $\hat{W}$.

To check this assumption for our $GTS_{\theta=10}^{EX}$ model, we numerically estimated the covariance matrix for the case where $\theta = 10$. We performed 100 independent FFPilot and FFS simulations at a variety of accuracy levels. For FFPilot simulations, we used error goals of 1%, 3.2%, and 10% and for FFS simulations, a constant number of trajectories per interface from 1000–25 000. Figures S13 and S14 show the distribution of phase weight estimators obtained during the FFS runs. We next calculated the variance of each $\hat{w}_i$ and for $i > 0$ compared to the expected variance obtained from Eqs. (14) and (28),

$$V[\hat{w}_{i>0}] = \frac{p_i(1-p_i)}{n_i}. \quad (34)$$

FIG. 15. (a) Comparison of the phase weight estimator variances for the indicated phases. Shown are (x) the measured variance from 100 independent FFPilot simulations and (line) the expected variance from Eq. (34). The top row shows the variances with no oversampling and the bottom row shows the variances with the optimal oversampling: 20X, 10X, 10X, 10X, and 10X for phases 0–4, respectively. Note that phase 5 was not oversampled. (b) The covariance matrices for the phase weight estimators at 10% error goal (top) without and (bottom) with oversampling.

Figure 15(a) and Figs. S15 and S16 show the comparison. For phases 3–5, the variance is much greater than expected, which agrees with our earlier analysis, revealing which phases had the most uncontrolled error.

We then calculated the covariance matrix from the simulations. Figure 15(b) and Fig. S17 show the results. A clear cluster of high covariance is apparent between phases 3–5 and, additionally, phase 0 has high covariance with most of the other phases. However, phase 0 has higher absolute variance than the other phases and if one instead looks at the Pearson correlation coefficient (see Fig. S18) that correlation between phase 0 disappears while the high correlation between phases 3–5 remains.

Given that the phases with high covariance correspond to the phases with high error, we wanted to know how oversampling these phases would affect the covariance matrix. We oversampled various phases and recalculated the covariances. The bottom row in Fig. 15 shows the covariance results with oversampling. When we applied oversampling in phases 0–4, the variances were near their expected levels and the magnitude of the covariances decreases substantially. It is interesting to note in Fig. 15 how the variance for phase 5 moves to the expected value, even though no additional sampling was done in phase 5. Increasing the sampling in earlier phases can reduce the error in later phases through the landscape correlations.

We also tried various combinations of oversampling. Figure S19 shows the results of two different experiments. First, we oversampled sequential phases one at a time, starting with phase 0. In this case, the covariance gradually diminishes as more phases are oversampled. Second, we started with oversampling only phase 4 and then sequentially added prior phases, down to phase 0. In this case, the covariance does not gradually change. It remains roughly the same until phase 0 is added to the oversampling and then it dramatically jumps lower. Thus, we conclude that oversampling starting in phase 0 is necessary to reduce underrepresentation of trajectories likely to cross the barrier that might relax slowly during the crossing. It is this slow relaxation that drives the landscape correlation errors in our GTS model. Further discussion of oversampling is given in [supplementary material](#), S2.

IV. DISCUSSION

Above, we presented our FFPilot approach for automatically parameterizing a FFS simulation to both minimize the simulation time and constrain the sampling error at a user specified margin of error and confidence interval. By performing an inexpensive pilot simulation, we obtain estimates of the weights and computational costs at each FFS interface, which are used to optimize a more thorough production simulation. Our pilot simulation approach also provides the advantage that individual interface placement, weights, and costs can be quickly evaluated before a potentially costly simulation is performed, if desired.

Unlike previous FFS optimization techniques, our method accounts for the statistics of phase 0 and also for varying computational cost along the order parameter. Both of these features are important for ensuring that error is controlled while minimizing the simulation time. Particularly, optimizing the time spent in phase 0 is important as our results indicate that these calculations can consume a significant fraction of the total simulation time.

Our results show that FFPilot correctly and efficiently controls the sampling error in FFS simulations. For both one and multi-dimensional systems that relax quickly, the *MFPT* estimates fall precisely within the specified error bounds. For systems with correlations along the transition path, our testing reveals that while sampling error is well controlled by FFPilot, error due to the system dependent landscape pushes the total error outside of the specified bounds. Errors in the estimation of early crossing distributions lock-in and cannot relax during later phases.

Since many multidimensional systems of significant complexity might relax more quickly in some degrees of freedom than others, a general approach is needed to control for correlation error in FFS while minimizing the simulation time. Our analysis of the covariance matrix suggests that it may be possible to control for the landscape error in a revised formulation of our optimizing equation (30) that includes terms for the covariance. However, it is not straightforward to estimate the covariances from the pilot simulation due to the lack of direct correspondence between individual trajectories in the various phases. Using the approach we followed here,

calculating the covariances from many independent simulations, would be computationally prohibitive. At present, we recommend first calculating the covariance matrix with a set of inexpensive low-accuracy FFPilot simulations and then using it to detect the presence of memory effects and, if necessary, determine an oversampling strategy for a more accurate production simulation.

Studying stochastic biochemical systems with metastable states will become increasingly important as more regulatory and developmental networks are elucidated in sufficient detail to permit quantitative modeling. Our findings raise the important issue of what other obstacles exist for efficient rare event sampling of these systems using FFS. For example, how to control landscape error in systems with many metastable states, with multiple transition paths, or with metastable intermediates? One possibility would be to study each transition path separately, using a branching approach, and then recombine the results based on the branching probabilities.

In summary, our FFPilot method provides a significant speedup compared to direct simulation of systems with rare event dynamics. Even with oversampling to control landscape error, speedups on the order of 100X for systems with long first passage times can be expected. FFPilot also provides a performance improvement compared to previous forward flux sampling methods, by accurately minimizing computational cost. The automatic optimization of simulation parameters to achieve a desired level of sampling error through the use of our optimization equation makes FFS simulations more robust and efficient.

SUPPLEMENTARY MATERIAL

See the [supplementary material](#) for additional derivations, Figs. S1–S20, and Table S1.

ACKNOWLEDGMENTS

The authors thank the members of Roberts lab for discussions. This work was supported by the National Science Foundation under Grant No. PHY-1707961 and the National Institutes of Health under Grant No. T32 GM008403.

Code implementing the FFPilot algorithm is available as part of our LMES software package for high-performance simulation of stochastic biological models (<https://www.robertslabjhu.info/home/software/lmes/>). We have also created a tutorial describing in detail how to use LMES to execute FFPilot simulations (<https://www.robertslabjhu.info/home/tutorials/tutorials/#ffpilot>).

APPENDIX A: THE UPPER BOUND OF THE VARIANCE OF A SUM OF BERNOULLI DISTRIBUTIONS

The overall process of launching trajectories in order to determine the phase weight during each phase of FFS simulation can be thought of as a single draw from a sum of many independent Bernoulli distributions, also called a Poisson binomial distribution (PBD). This viewpoint emphasizes the statistical equivalence between the outcome of the j th trajectory in an FFS phase, which will either fall back into its starting basin or flux forward to the next interface, and the outcome of the j th Bernoulli random variable in a PBD, which will take on a value of either 0 or 1. In both cases “success” (fluxing forward, drawing a 1) occurs with some probability

p_j inherent to the individual process, while “failure” (returning to basin, drawing a 0) occurs with probability $(1 - p_j)$. The expected value and the variance of a draw from a PBD of n terms (i.e., the sum over an independent draw from each of its constituent Bernoulli random variables) is

$$\begin{aligned}\mu &= \sum_{j=1}^n p_j, \\ V &= \sum_{j=1}^n (1 - p_j)p_j.\end{aligned}\quad (\text{A1})$$

Equivalently, Eq. (A1) is also the mean and variance of the total count of successful trajectories n_i^s in FFS phase $i > 0$, given that $n = n_i$ trajectories were run.

The variance of a PBD is maximized when the probability parameters of each of its Bernoulli random variables are all the same p ,

$$p_1 = p_2 = \dots = p_n = p,$$

and so the upper bound on the variance of a sum of Bernoulli random variables is

$$V \leq np(1 - p). \quad (\text{A2})$$

Proof. For a given value of μ , the method of Lagrange Multipliers can be used to maximize V with respect to the choice of particular values of the p_i terms. Appropriate constraint and target equations can be taken from Eq. (A1),

$$\begin{aligned}g[p_j] &= \sum_{j=1}^n p_j - \mu = 0, \\ f[p_j] &= V = \sum_{j=1}^n (1 - p_j)p_j.\end{aligned}\quad (\text{A3})$$

A Lagrangian can be formed from Eq. (A3),

$$\mathcal{L} = \sum_{j=1}^n ((1 - p_j)p_j) - \lambda \left(\mu - \sum_{k=1}^n p_k \right).$$

Next, we find the gradient of the Lagrangian,

$$\begin{aligned}\partial_\lambda \mathcal{L} &= \sum_{j=1}^n p_j - \mu, \\ \partial_{p_j} \mathcal{L} &= -2p_j + \lambda + 1.\end{aligned}$$

Now, we set each part of the gradient to zero and solve the resulting set of equations. We start by solving for p_i in terms of λ ,

$$p_j = \frac{\lambda + 1}{2}.$$

Next, we solve for λ alone,

$$\lambda = \frac{2\mu}{n} - 1.$$

Finally, we plug the solution for λ into the gradient of p_i and solve for p_i alone,

$$p_j = \frac{\mu}{n} = p. \quad (\text{A4})$$

Thus, the spot at which every p_i is equal to $\frac{\mu}{n}$ is a critical point for the variance of a Bernoulli mixture distribution. It can further be shown that the above critical point is a maximum for the constrained variance using the Bordered Hessian variation of the classical second derivative test. The Bordered Hessian⁷⁰ of a Lagrangian can be defined as

$$H \equiv \begin{pmatrix} \frac{\partial \mathcal{L}}{\partial \lambda^2} & \frac{\partial \mathcal{L}}{\partial \lambda p_1} & \frac{\partial \mathcal{L}}{\partial \lambda p_2} & \cdots & \frac{\partial \mathcal{L}}{\partial \lambda p_N} \\ \frac{\partial \mathcal{L}}{\partial \lambda p_1} & \frac{\partial \mathcal{L}}{\partial p_1^2} & \frac{\partial \mathcal{L}}{\partial p_1 p_2} & \cdots & \frac{\partial \mathcal{L}}{\partial p_1 p_N} \\ \frac{\partial \mathcal{L}}{\partial \lambda p_2} & \frac{\partial \mathcal{L}}{\partial p_1 p_2} & \frac{\partial \mathcal{L}}{\partial p_2^2} & \cdots & \frac{\partial \mathcal{L}}{\partial p_2 p_N} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \mathcal{L}}{\partial \lambda p_N} & \frac{\partial \mathcal{L}}{\partial p_1 p_N} & \frac{\partial \mathcal{L}}{\partial p_2 p_N} & \cdots & \frac{\partial \mathcal{L}}{\partial p_N^2} \end{pmatrix},$$

which for our Lagrangian works out to

$$H = \begin{pmatrix} 0 & 1 & 1 & \cdots & 1 \\ 1 & -2 & 0 & \cdots & 0 \\ 1 & 0 & -2 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & 0 & \cdots & -2 \end{pmatrix}. \quad (\text{A5})$$

$\frac{\mu}{n}$ is a maximum if and only if H is negative definite. The negative definiteness of H can be demonstrated by showing that the signs of its leading principal minors demonstrate the appropriate alternating pattern.⁷⁰ For any finite value of n , this Hessian can be diagonalized using the Gaussian elimination technique. This makes it easy to calculate the determinants of the various $x \times x$ upper-left submatrices and to show that the signs of the said determinants do indeed follow the pattern of $(-1)^{x-1}$ for $x \geq 3$, satisfying the condition for negative definiteness. This conclusion can be generalized to arbitrary values of n using a proof by induction (the details of which are omitted for brevity). Thus, H is always a negative definite matrix, and so setting each $p_j = \frac{\mu}{n} = p$ does indeed maximize V for a given μ .

From Eq. (A2), the upper bound on the variance of the count of successful trajectories n_i^s in FFS phase $i > 0$ is

$$V[n_i^s] \leq n_i p_i (1 - p_i).$$

From Eq. (3), the phase weight estimator $\widehat{w}_{i>0}$ can be given in terms of the success count n_i^s ,

$$\widehat{w}_{i>0} = \frac{n_i^s}{n_i}.$$

The variance of a quotient $V[x/y]$ is $V[x]/y^2$, given that y is a constant.⁵⁸ Thus, the upper bound on the variance of $\widehat{w}_{i>0}$ is

$$V[\widehat{w}_{i>0}] = V\left[\frac{n_i^s}{n_i}\right] = \frac{V[n_i^s]}{n_i^2} \leq \frac{1}{n_i} p_i (1 - p_i).$$

Finally, Eq. (14) can be used to derive an upper bound on $V[\mathbf{w}_{i>0}]$ from the bound on $V[\widehat{w}_{i>0}]$,

$$V[\mathbf{w}_{i>0}] = n_i V[\widehat{w}_{i>0}] \leq p_i (1 - p_i). \quad (\text{A6})$$

Parameterizing the FFPilot optimizing equation [Eq. (30)] with a value of $V[\mathbf{w}_{i>0}]$ that is at least as large as the true value helps to ensure that the calculated n_i is at least sufficient to achieve the given error goal. The upper bound on the variance given in Eq. (A6) is equivalent to the variance of a single Bernoulli random variable with parameter p_i . Thus, for the purposes of parameterizing the FFPilot optimizing equation, we treat each $\mathbf{w}_{i>0}$ as a single Bernoulli random variable without any loss in accuracy.

APPENDIX B: BLIND OPTIMIZATION METHOD

Taken as a whole, the FFPilot approach to optimizing simulations can function reliably only if the results of the pilot stage are highly accurate (at least in terms of the individual parameter estimates) and computationally inexpensive (relative to the production stage). No prior knowledge of the system under study is used in the setup of the pilot stage. Thus, a “blind” optimization method, one that uses no information about the current phase or any other, must be used during this initial stage.

The blind optimization method that we use during the FFPilot pilot stage works by altering the conditions under which a simulation phase is terminated. During a standard FFS simulation, simulation phase $i > 0$ is terminated once a fixed number of trajectories n_i have launched from λ_{i-1} and have ended, regardless of where (in state space) those trajectories have ended. During a pilot stage, we instead terminate phase $i > 0$ only once a fixed number $n_i^s = n_{\text{pilot}}$ of successful trajectories (i.e., the ones that reach λ_i) have been observed.

The advantage of using our blind optimization method is that it is able to produce estimates of the phase weight $w_{i>0} = p_i$ with constrained maximum error. The margin of error for a single phase $i > 0$ can be calculated from Eqs. (20) and (28),

$$\zeta[\widehat{p}_i] = z_\alpha \sqrt{\frac{1 - p_i}{p_i n_i}}.$$

The total count of trajectories n_i required to produce n_{pilot} successful trajectories in phase $i > 0$ converges to $\frac{n_{\text{pilot}}}{p_i}$. Thus, with respect to the pilot stage, the above equation can be rewritten as

$$\zeta[\widehat{p}_i] = z_\alpha \sqrt{\frac{1 - p_i}{n_{\text{pilot}}}}. \quad (\text{B1})$$

The error increases as p_i becomes smaller, but it remains within a finite bound even as p_i goes to zero (see [supplementary material](#), Fig S10). By default and throughout this paper, we use the fixed values of $n_{\text{pilot}} = 10^4$ and $z_\alpha = z_{.95} \approx 1.96$ for all phases of the pilot stage. These values of n_{pilot} and z_α give a maximum $\widehat{p}_i$ margin of error of 2%.

During phase 0, the distribution of samples taken from the underlying random variable $\mathbf{w}_0$ (i.e., the set of observed waiting times in between λ_0 forward crossing events) is model dependent. This means that no formulation equivalent to Eq. (B1) is possible for the phase weight $w_0 = \tau_A$. However, the estimator $\widehat{\tau}_A$ can be in general assumed to be $\sqrt{n}$ -consistent (as described in Secs. II D and III A 1). Plugging Eqs. (9) and (14) into Eq. (6) gives the asymptotic

margin of error for phase 0 of the pilot stage,

$$\zeta[\hat{\tau}_A] = \frac{z_\alpha}{\sqrt{n_{\text{pilot}}}} \frac{\sqrt{V[\mathbf{w}_0]}}{\tau_A}. \quad (\text{B2})$$

It can be ensured that $\tau_A \approx \sqrt{V[\mathbf{w}_0]}$ through appropriate placement of $\mathbf{w}_0$ (i.e., away from a basin of attraction). Given that τ_A and $V[\mathbf{w}_0]$ are appropriately matched, the margin of error of phase 0 will be roughly proportional to $\frac{z_\alpha}{\sqrt{n_{\text{pilot}}}}$. For $n_{\text{pilot}} = 10^4$ and $z_\alpha \approx 1.96$, this also works out to a $\hat{\tau}_A$ margin of error of about 2%. Thus, the blind optimization approach used in the FFPilot pilot stage controls error in $\hat{\tau}_A$, though not in such a conveniently bounded fashion as $\hat{p}_i$.

REFERENCES

- ¹F. Harold, *The Way of the Cell: Molecules, Organisms, and the Order of Life* (Oxford University Press, USA, 2003).
- ²J. X. Zhou, M. D. S. Aliyu, E. Aurell, and S. Huang, *J. R. Soc., Interface* **9**, 3539 (2012).
- ³C. H. Waddington and H. Kacser, *The Strategy of the Genes: A Discussion of Some Aspects of Theoretical Biology* (Allen & Unwin, London, 1957).
- ⁴R. J. Nichols, S. Sen, Y. J. Choo, P. Beltrao, M. Zietek, R. Chaba, S. Lee, K. M. Kazmierczak, K. J. Lee, A. Wong, M. Shales, S. Lovett, M. E. Winkler, N. J. Krogan, A. Typas, and C. A. Gross, *Cell* **144**, 143 (2011).
- ⁵H. Qian and L. M. Bishop, *Int. J. Mol. Sci.* **11**, 3472 (2010).
- ⁶E. Roberts, S. Be'er, C. Bohrer, R. Sharma, and M. Assaf, *Phys. Rev. E* **92**, 062717 (2015).
- ⁷W. Wu and J. Wang, *J. Chem. Phys.* **139**, 121920 (2013).
- ⁸D. T. Gillespie, *J. Comput. Phys.* **22**, 403 (1976).
- ⁹D. T. Gillespie, *Annu. Rev. Phys. Chem.* **58**, 35 (2007).
- ¹⁰D. Bratsun, D. Volfson, L. S. Tsimring, and J. Hasty, *Proc. Natl. Acad. Sci. U. S. A.* **102**, 14593 (2005).
- ¹¹Y. Cao, H. Li, and L. Petzold, *J. Chem. Phys.* **121**, 4059 (2004).
- ¹²M. A. Gibson and J. Bruck, *J. Phys. Chem. A* **104**, 1876 (2000).
- ¹³Y. Cao, H.-M. Lu, and J. Liang, *Proc. Natl. Acad. Sci. U. S. A.* **107**, 18445 (2010).
- ¹⁴E. Roberts, A. Magis, J. O. Ortiz, W. Baumeister, and Z. Luthey-Schulten, *PLoS Comput. Biol.* **7**, e1002010 (2011).
- ¹⁵C. Li and J. Wang, *Proc. Natl. Acad. Sci. U. S. A.* **111**, 14130 (2014).
- ¹⁶B. Zhang and P. G. Wolynes, *Proc. Natl. Acad. Sci. U. S. A.* **111**, 10185 (2014).
- ¹⁷M. C. Zwier, A. J. Pratt, J. L. Adelman, J. W. Kaus, D. M. Zuckerman, and L. T. Chong, *J. Phys. Chem. Lett.* **7**, 3440 (2016).
- ¹⁸T. D. Smith, M. J. Tse, E. L. Read, and W. F. Liu, *Integr. Biol.* **8**, 946 (2016).
- ¹⁹R. Sharma and E. Roberts, *Phys. Biol.* **13**, 036003 (2016).
- ²⁰P. Baron, *Reaction Rate Theory and Rare Events*, 1st ed. (Elsevier B.V., 2017).
- ²¹N. B. Becker and P. R. ten Wolde, *J. Chem. Phys.* **136**, 174119 (2012).
- ²²H. Kahn and T. E. Harris, *Natl. Bur. Stand. Appl. Math. Ser.* **12**, 27 (1951).
- ²³G. M. Torrie and J. P. Valleau, *J. Comput. Phys.* **23**, 187 (1977).
- ²⁴M. Souaille and B. Roux, *Comput. Phys. Commun.* **135**, 40 (2001).
- ²⁵A. Laio and M. Parrinello, *Proc. Natl. Acad. Sci. U. S. A.* **99**, 12562 (2002).
- ²⁶J. F. Dama, M. Parrinello, and G. A. Voth, *Phys. Rev. Lett.* **112**, 240602 (2014).
- ²⁷C. Dellago, P. G. Bolhuis, F. S. Csajka, and D. Chandler, *J. Chem. Phys.* **108**, 1964 (1998).
- ²⁸G. E. Crooks and D. Chandler, *Phys. Rev. E* **64**, 026109 (2001).
- ²⁹N. Guttenberg, A. R. Dinner, and J. Weare, *J. Chem. Phys.* **136**, 234103 (2012).
- ³⁰H. Jung, K.-i. Okazaki, and G. Hummer, *J. Chem. Phys.* **147**, 152716 (2017).
- ³¹G. A. Huber and S. Kim, *Biophys. J.* **70**, 97 (1996).
- ³²D. M. Zuckerman and L. T. Chong, *Annu. Rev. Biophys.* **46**, 43 (2017).
- ³³R. J. Allen, P. B. Warren, and P. R. ten Wolde, *Phys. Rev. Lett.* **94**, 018104 (2005).
- ³⁴R. J. Allen, D. Frenkel, and P. R. ten Wolde, *J. Chem. Phys.* **124**, 024102 (2006).
- ³⁵C. Valeriani, R. J. Allen, M. J. Morelli, D. Frenkel, and P. Rein ten Wolde, *J. Chem. Phys.* **127**, 114109 (2007).
- ³⁶R. J. Allen, C. Valeriani, and P. Rein ten Wolde, *J. Phys.: Condens. Matter* **21**, 463102 (2009).
- ³⁷A. Warmflash, P. Bhimalapuram, and A. R. Dinner, *J. Chem. Phys.* **127**, 154112 (2007).
- ³⁸A. Dickson, A. Warmflash, and A. R. Dinner, *J. Chem. Phys.* **130**, 074104 (2009).
- ³⁹A. Dickson, A. Warmflash, and A. R. Dinner, *J. Chem. Phys.* **131**, 154104 (2009).
- ⁴⁰B. W. Zhang, D. Jasnow, and D. M. Zuckerman, *J. Chem. Phys.* **132**, 054107 (2010).
- ⁴¹D. Bhatt, B. W. Zhang, and D. M. Zuckerman, *J. Chem. Phys.* **133**, 014110 (2010).
- ⁴²J. L. Adelman and M. Grabe, *J. Chem. Phys.* **138**, 044105 (2013).
- ⁴³R. M. Donovan, A. J. Sedgewick, J. R. Faeder, and D. M. Zuckerman, *J. Chem. Phys.* **139**, 115105 (2013).
- ⁴⁴R. M. Donovan, J.-J. Tapia, D. P. Sullivan, J. R. Faeder, R. F. Murphy, M. Dittich, and D. M. Zuckerman, *PLoS Comput. Biol.* **12**, e1004611 (2016).
- ⁴⁵A. Dickson and A. R. Dinner, *Annu. Rev. Phys. Chem.* **61**, 441 (2010).
- ⁴⁶F. A. Escobedo, E. E. Borrero, and J. C. Araque, *J. Phys.: Condens. Matter* **21**, 333101 (2009).
- ⁴⁷N. B. Becker, R. J. Allen, and P. R. ten Wolde, *J. Chem. Phys.* **136**, 174118 (2012).
- ⁴⁸R. J. Allen, D. Frenkel, and P. R. ten Wolde, *J. Chem. Phys.* **124**, 194111 (2006).
- ⁴⁹E. E. Borrero and F. A. Escobedo, *J. Chem. Phys.* **129**, 024115 (2008).
- ⁵⁰K. Kratz, A. Arnold, and R. J. Allen, *J. Chem. Phys.* **138**, 164112 (2013).
- ⁵¹H. Jiang, M. Pu, and Z. Hou, *Sci. China: Chem.* **57**, 165 (2013).
- ⁵²D. T. Gillespie, *J. Chem. Phys.* **74**, 5295 (1981).
- ⁵³D. Olive, *Statistical Theory and Inference* (Springer, Cham, 2014).
- ⁵⁴S. M. Ross, *Introductory Statistics* (Academic Press, 2010).
- ⁵⁵The z score z_α for any confidence level α can be calculated as $z_\alpha = \sqrt{2}\text{Erfinv}(\alpha)$, where Erfinv is the Inverse Error Function, and α is expressed as a fraction. Some authors write the z score as $z_{(1-\alpha)/2}$ instead of z_α .
- ⁵⁶G. W. Oehlert, *Am. Stat.* **46**, 27 (1992).
- ⁵⁷L. Wasserman, *All of Statistics, A Concise Course in Statistical Inference* (Springer Science & Business Media, New York, NY, 2013).
- ⁵⁸K. F. Riley, M. P. Hobson, and S. J. Bence, *Mathematical Methods for Physics and Engineering, A Comprehensive Guide* (Cambridge University Press, 2006).
- ⁵⁹P. Glasserman, P. Heidelberger, P. Shahabuddin, and T. Zajic, *Oper. Res.* **47**, 585 (1999).
- ⁶⁰T. S. Gardner, C. R. Cantor, and J. J. Collins, *Nature* **403**, 339 (2000).
- ⁶¹P. B. Warren and P. R. ten Wolde, *J. Phys. Chem. B* **109**, 6812 (2005).
- ⁶²D. Schultz, A. M. Walczak, J. N. Onuchic, and P. G. Wolynes, *Proc. Natl. Acad. Sci. U. S. A.* **105**, 19165 (2008).
- ⁶³H. Feng and J. Wang, *Sci. Rep.* **2**, 550 (2012).
- ⁶⁴M. Lu, J. Onuchic, and E. Ben Jacob, *Phys. Rev. Lett.* **113**, 078102 (2014).
- ⁶⁵T. Biancalani and M. Assaf, *Phys. Rev. Lett.* **115**, 208101 (2015).
- ⁶⁶M. J. Tse, B. K. Chu, M. Roy, and E. L. Read, *Biophys. J.* **109**, 1746 (2015).
- ⁶⁷B. K. Chu, M. J. Tse, R. R. Sato, and E. L. Read, *BMC Syst. Biol.* **11**, 14 (2017).
- ⁶⁸P. B. Warren and R. J. Allen, *Mol. Phys.* **116**, 3104 (2018).
- ⁶⁹T. Prellberg and J. Krawczyk, *Phys. Rev. Lett.* **92**, 120602 (2004).
- ⁷⁰J. R. Magnus and H. Neudecker, *Matrix Differential Calculus with Applications in Statistics and Econometrics* (University of Texas Press, 1999).