



Published in final edited form as:

*Int J Uncertain Quantif.* 2020 ; 10(6): 515–542. doi:10.1615/int.j.uncertaintyquantification.2020033241.

## BAYESIAN INFERENCE OF STOCHASTIC REACTION NETWORKS USING MULTIFIDELITY SEQUENTIAL TEMPERED MARKOV CHAIN MONTE CARLO

Thomas A. Catanach<sup>1,\*</sup>, Huy D. Vo<sup>2</sup>, Brian Munsky<sup>2</sup>

<sup>1</sup>Sandia National Laboratories, Livermore, CA, 94550

<sup>2</sup>Dept. of Chemical and Biological Engr., Colorado State University, Fort Collins, CO, 80521

### Abstract

Stochastic reaction network models are often used to explain and predict the dynamics of gene regulation in single cells. These models usually involve several parameters, such as the kinetic rates of chemical reactions, that are not directly measurable and must be inferred from experimental data. Bayesian inference provides a rigorous probabilistic framework for identifying these parameters by finding a posterior parameter distribution that captures their uncertainty. Traditional computational methods for solving inference problems such as Markov Chain Monte Carlo methods based on classical Metropolis-Hastings algorithm involve numerous serial evaluations of the likelihood function, which in turn requires expensive forward solutions of the chemical master equation (CME). We propose an alternate approach based on a multifidelity extension of the Sequential Tempered Markov Chain Monte Carlo (ST-MCMC) sampler. This algorithm is built upon Sequential Monte Carlo and solves the Bayesian inference problem by decomposing it into a sequence of efficiently solved subproblems that gradually increase both model fidelity and the influence of the observed data. We reformulate the finite state projection (FSP) algorithm, a well-known method for solving the CME, to produce a hierarchy of surrogate master equations to be used in this multifidelity scheme. To determine the appropriate fidelity, we introduce a novel information-theoretic criteria that seeks to extract the most information about the ultimate Bayesian posterior from each model in the hierarchy without inducing significant bias. This novel sampling scheme is tested with high performance computing resources using biologically relevant problems.

### Keywords

Bayesian inference; Stochastic modeling; Systems Biology; UQ; MCMC; SMC

## 1. INTRODUCTION

A distinguishing feature of biology is the diversity manifested by living things across different scales, from the readily observed multitude of species to the differences between individuals of the same species. At the microscopic level, a population of cells with the same

\*Address all correspondence to: Thomas A. Catanach, tacatan@sandia.gov.

genetic code, growing under the same lab conditions, could still display phenotypic variability in gene products [1–5]. Phenotypic variability has been observed in an increasing volume of data obtained from single-cell, single-molecule measurements enabled by recent progresses in chemical labeling and imaging techniques [6–8].

Much of the variability in gene expression is attributed to the stochasticity of vital cellular processes (e.g., transcription, translation) that are subjected to the randomness of molecular interactions. Stochastic reaction networks (SRN) represent a class of models that have been widely used to capture temporal and spatial fluctuations in single-cell gene expression [9]. SRN models treat the copy numbers of biochemical species, i.e. the number of molecules of a given type within a cell, as states in a discrete-space, continuous-time Markov process, where chemical reactions are represented by transitions between states. Given an SRN model, the probabilities of gene expression states within a cell can be computed by solving the chemical master equation (CME). This is a dynamical system in an infinite-dimensional space that describes the evolution of the probability distribution of all states. The finite state projection (FSP) is a well-known approximation method to obtain high-fidelity solutions of the CME [10]. This method reduces the intractable state space of the original SRN into a finite subset chosen based on a proven error bound, turning the infinite-dimensional CME system into a finite problem of linear differential equations.

The present work is concerned with the selection, parameter estimation, and uncertainty propagation of these reaction network models within the Bayesian framework. Bayesian methods are a powerful tool for system identification for SRN models because they provide rigorous uncertainty quantification by identifying a probability distribution over plausible model parameters instead of selecting a single model that may fit the data well [11–15]. This distribution over the models given the data is called the posterior distribution. Quantifying model parameter uncertainty is critical because it is difficult to model the full complexity of the biological system that may exhibit experimental context dependence [16]. Further, once model parameter uncertainty has been quantified, further experiments can be designed to provide new information about the system [17–21].

In this paper, we focus on data obtained by the experimental technique of single-molecule fluorescent in situ hybridization (smFISH). These datasets consist of independent single-cell measurements, each of which measures the copy number of biochemical species at a single time point. The standard approach to sample from the posterior distribution implied by this data is to use Markov Chain Monte Carlo (MCMC) algorithms such as the random walk Metropolis-Hastings MCMC sampler [11,22]. With high-fidelity CME solutions enabled by the FSP, one can compute the likelihood of observing these single-cell data and then perform Bayesian inference for model parameters. However, this approach suffers from two drawbacks. The first drawback is that the MCMC is inherently serial, preventing it from utilizing the massively parallel processing capability provided by modern high performance computing clusters. Therefore, if standard MCMC techniques are used, several tens to hundreds of thousands of sequential model evaluations may be needed to adequately sample the posterior distribution. The second drawback is that the FSP solutions required for the likelihood function are expensive, often requiring several minutes per model evaluation for a moderately sized problem. Typically, the number of differential equations that the FSP

algorithm needs to solve grows exponentially with the number of species in the network, so the size of the state space and transition matrix quickly grows intractable. This has motivated Approximate Bayesian Computation (ABC) approaches that replace the computationally expensive single-cell likelihood function with less expensive model-data discrepancy functions [23–25]. Samples produced by ABC are, in general, not distributed according to the true posterior distribution, and a careful choice of summary statistics is critical for the performance and reliability of ABC samplers. Another prominent direction to circumvent the cost of solving the CME is to use moment-based likelihood functions [26,27]. When the SRN consists only of linear first-order reactions, the equations describing the first centralized moments (such as mean and variance) can be solved exactly at very low computational cost, which is a large advantage over FSP-based approaches. However, for general reaction networks, only approximations to the moments based on moment closure techniques are available and the appropriate choice of closure method is essential for making reliable inference [27]. There has also been recent progress on non-Bayesian approaches to fitting SRN models to smFISH data, such as by replacing the expensive CME-based likelihood with the Wasserstein distance that could be much easier to estimate for certain networks [28].

In this work, we propose a different approach that uses a parallel and multifidelity MCMC computational framework to produce samples from the true posterior distribution. Many approaches to parallel MCMC methods have been proposed either based on parallel proposals or parallel Markov chains [29–32]. One family of popular parallel MCMC methods are those based upon sequential Monte Carlo (SMC) samplers [31,33–35]. For this work, we replace the standard MCMC methods with the Sequentially Tempered Markov Chain Monte Carlo (ST-MCMC) [35], which is a massively parallel sampling scheme based on SMC. This method transports a population of model parameter samples through a series of intermediate annealing levels that reflect the gradual increase in the influence of the data likelihood. At each level, MCMC is used to explore the intermediate distribution and re-balance the distribution of the samples. Since this method is population based, these MCMC steps can be done in parallel. Further, this method can effectively adapt to the target posterior distribution to speed up sampling.

Similarly, approaches to multifidelity MCMC have been explored in the literature such as multifidelity delayed acceptance schemes, Multilevel Markov Chain Monte Carlo, and multifidelity approaches to SMC [36–41]. Multifidelity delayed acceptance schemes have been applied to Bayesian inference for the CME before [22,42,43]. Within these methods, a fast surrogate of the expensive likelihood function is used to pre-screen proposed samples within MCMC before they are accepted or rejected based on the expensive CME likelihood. This method still requires many sequential full model evaluations in order to sample the posterior. Multilevel MCMC [40] uses a hierarchy of models, such as different discretization grids of a PDE, to design an estimator for a specific quantity of interest. This uses parallel Markov chains at different model fidelities to estimate a correction to the quantity of interest estimate incurred by refining the model fidelity. Multifidelity SMC methods like the Multilevel Sequential<sup>2</sup> Monte Carlo sampler [41] use an embarrassingly parallel approach and a hierarchy of multifidelity models. Therefore, only a few sequential full model evaluations may be needed. We take this approach to develop a multifidelity form of ST-

MCMC for solving the CME. Within our method, instead of only considering a series of annealing levels, we also consider a hierarchy of model fidelity. Thus, the solution of the full inference problem is broken down by steering the samples from a distribution that reflects a low fidelity model with little influence from the likelihood to a distribution that reflects the high-fidelity model with the full influence of the likelihood. By performing the early updates using fast models, the sampler can quickly converge to the most important regions of parameter space, where a higher fidelity model can then be used to better assess which regions of this high-likelihood space are most likely. The key challenge for applying multifidelity methods in the SMC context is deciding which annealing factor and model fidelity is appropriate at a given level. Latz et al. suggest an approach based on the effective sample size of the population [41]. We take a different approach by leveraging a limited number of high fidelity model solves to estimate the information gained about the ultimate posterior given a current model fidelity and annealing factor. This information-theoretic criteria can effectively identify when the lower fidelity model is overly biasing the solution and should be discarded in favor of a higher fidelity model.

We demonstrate the efficiency and accuracy of our novel scheme when solving parameter estimation, model selection, and uncertainty propagation for stochastic chemical kinetic models in the Bayesian framework. This new approach to multifidelity ST-MCMC using fast surrogates from reduced order models based on a novel reformulation of the FSP significantly reduces the number of expensive likelihood function required to sample the posterior. The example problems are based on models from the system and synthetic biology literature. These include a three-dimensional repressilator gene circuit, a spatial bursting gene expression network, and a stochastic transcription network for the inflammation response gene IL1beta [44]. As such the primary contributions of this paper are:


1. Development of multifidelity ST-MCMC for Bayesian inference of SRNs
2. Introduction of an information-theoretic criteria for assessing the appropriate model fidelity within multifidelity ST-MCMC
3. Description of a novel surrogate model of the CME to be used within multifidelity inference problems

This paper is organized as follows: Section 2 describes general background for stochastic reaction network modeling, finite state projection, Bayesian inference, and MCMC methods. Section 3 describes the Multifidelity Sequential Tempered MCMC algorithm and the information-theoretic criteria for adapting model fidelity. Section 4 describes a novel surrogate models for the CME. Section 5 describes three experiments to identify model parameters in SRNs using Multifidelity ST-MCMC. Finally, Section 6 concludes.

## 2. BACKGROUND

### 2.1 Stochastic reaction networks for modeling gene expression

A reaction network consists of  $N$  different chemical species  $S_1, \dots, S_N$  that are interacting via the following  $M$  chemical reactions



We are interested in keeping track of the integral vectors  $\mathbf{x} \equiv (x_1, \dots, x_N)^T$ , where  $x_i$  is the population of the  $i$ th species. Assuming constant temperature and volume, the time-evolution of this system can be modeled by a continuous-time, discrete-space Markov process [9]. The  $j$ th reaction channel is associated with a stoichiometric vector

$\nu_j = (\nu_{1,j}^{prod} - \nu_{1,j}^{react}, \dots, \nu_{N,j}^{prod} - \nu_{N,j}^{react})^T$  ( $j = 1, \dots, M$ ) such that, if the system is in state  $\mathbf{x}$  and reaction  $j$  occurs, the system transitions to state  $\mathbf{x} + \nu_j$ . Given  $\mathbf{x}(t) = \mathbf{x}$ , the propensity  $\alpha_j(\mathbf{x}; \theta) dt$  determines the probability that reaction  $j$  occurs in the next infinitesimal time interval  $[t, t + dt)$ , where  $\theta$  is the vector of model parameters. In other words,

$$\text{Prob}(\mathbf{x}(t + dt) = \mathbf{x} + \nu_j | \mathbf{x}(t) = \mathbf{x}) = \alpha_j(\mathbf{x}; \theta) dt.$$

An important case of reaction networks are those that follow mass-action kinetics, whose propensity functions take the form

$$\alpha_j(\mathbf{x}; \theta) = c_j(\theta) \binom{x_1}{\nu_{1,j}^{react}} \cdot \dots \cdot \binom{x_N}{\nu_{N,j}^{react}}. \quad (2)$$

In this formulation,  $c_j(\theta)$  is the probability per unit time that a combination of specific reactant molecules can react via reaction  $j$ , and the remaining factor is the number of ways the existing molecules can be combined to form the left side of the chemical equation (1).

The time-evolution of the probability distribution of this Markov process is the solution of the linear system of differential equations known as the chemical master equation (CME)

$$\begin{cases} \frac{d}{dt} \mathbf{p}(t) = \mathbf{A}(\theta) \mathbf{p}(t), t \in [0, t_f] \\ \mathbf{p}(0) = \mathbf{p}_0 \end{cases}, \quad (3)$$

where  $\mathbf{p}(t)$  is the time-dependent probability distribution of all states,  $p(\mathbf{x}(t) = \mathbf{x}_j | \mathbf{x}(0))$ . The initial distribution  $\mathbf{p}_0$  is assumed to be given, and  $\mathbf{A}(\theta)$  is the infinitesimal generator of the Markov process, defined entrywise as

$$\mathbf{A}(\mathbf{y}, \mathbf{x}; \theta) = \begin{cases} \alpha_j(\mathbf{x}; \theta) \text{ if } \mathbf{y} = \mathbf{x} + \nu_j \\ -\sum_{j=1}^M \alpha_j(\mathbf{x}; \theta) \text{ if } \mathbf{y} = \mathbf{x} \\ 0 \text{ otherwise} \end{cases}. \quad (4)$$

Here, we have made explicit the dependence of  $\mathbf{A}$  on the model parameter vector  $\theta$ , which we need to infer from experimental data.

## 2.2 The finite state projection

Typically, reaction networks model open biochemical systems, where the set of all possible molecular states is unbounded. This makes the CME an infinite-dimensional linear system

of ODEs. The finite state projection (FSP) is a well-known strategy to systematically reduce this linear system into a finite surrogate model with a strict error bound.

The FSP can be thought of as a special class of projection-based model reduction applied to the CME. Specifically, let  $\Omega$  be a finite subset of the CME state space. The projection of the CME operator  $\mathbf{A}$  onto the subspace spanned by the point-mass measures  $\{\delta_{\mathbf{x}} | \mathbf{x} \in \Omega\}$  is given by

$$\tilde{\mathbf{A}}_{\Omega}(\mathbf{y}, \mathbf{x}) = \begin{cases} \mathbf{A}(\mathbf{y}, \mathbf{x}) & \text{if } \mathbf{x}, \mathbf{y} \in \Omega \\ 0 & \text{otherwise} \end{cases}. \quad (5)$$

We can then define a reduced model of the dynamical system (3) based on this projection as

$$\frac{d}{dt} \tilde{\mathbf{p}}_{\Omega}(t) = \tilde{\mathbf{A}}_{\Omega} \tilde{\mathbf{p}}_{\Omega}(t), \quad t \in [0, t_f]. \quad (6)$$

Clearly, to solve (6) we only need to keep track of the equations corresponding to states in the finite set  $\Omega$ , which is amenable to numerical treatments.

In contrast to generic projection methods, the gap between the reduced-order model and the true CME can be computed for the FSP. Indeed, Munsky and Khammash [45, Theorem 2] proved that the truncation error can be quantified in  $\ell^1$  norm as

$$\|\mathbf{p}(t) - \tilde{\mathbf{p}}_{\Omega}(t)\|_1 = 1 - \sum_{\mathbf{x} \in \Omega} \tilde{\mathbf{p}}_{\Omega}(t, \mathbf{x}). \quad (7)$$

Clearly, the right hand side can be readily computed from the solution of the reduced system (6). From this precise error quantification, we have effective iterative method for solving the CME. Choosing an error tolerance  $\varepsilon > 0$ , starting from any initial set  $\Omega := \Omega_0$ , we solve system (6) and check that the right hand side of (7) is less than  $\varepsilon$ . If this fails, we add more states to  $\Omega_0$  to get a strictly larger set  $\Omega_1$  and repeat the procedure until we find an approximation that satisfies our error tolerance.

As the sequence of subsets  $\Omega_j$  grows until it eventually covers the whole state space, we might expect that the finite-time solution of (6) will likewise converge to the true solution. This is indeed the case for all models in practice, with only a few theoretical counterexamples in which the Markov chain is explosive [10]. Sufficient conditions for the convergence of the FSP can be checked based on the form of the propensity function [46,47]. In practice, reaction networks tend to have only reactions between two or fewer molecules, with propensity functions in mass-action form (2), and these are guaranteed to be approximable with the FSP [46].

For the rest of the paper, we only concern ourselves with non-explosive SRNs where the FSP converges. Given such models, any exhaustive sequence of subsets  $\{\Omega_j\}$  suffices to guarantee that the FSP solution eventually satisfies any prespecified error tolerance. However, some choices are more efficient than other, and it is also more advantageous to

partition the time interval into smaller timesteps and use a smaller  $\Omega_j$  on each step [48]. We will return to these observations in section 4.3.

A final point to make about the FSP is that, if the sequence of reduced sets  $\Omega_j$  are increasing, that is,  $\Omega_j \subset \Omega_{j+1}, j = 1, 2, \dots$ , the resulting truncation error has been shown to decrease monotonically [45, Theorem 1] (also see [49, Theorem 2.5]). This gives us a natural way to form a hierarchy of reduced models for the CME, and we will return to this in section 4.1. With the computed distribution of the SRN in our hands, we can now match them directly to experimental data using a straightforward likelihood function.

### 2.3 Bayesian inference of SRN models from discrete single-cell measurements

In this work, we focus on inferring the parameters for the reaction network from discrete, single-cell datasets [3,6–8,50] that consist of several snapshots of many independent cells taken at discrete times  $t_1, \dots, t_T$ . The snapshot at time  $t_j$  records gene expression in  $n_j$  independent cells, each of which can be collected in the data vector  $\mathbf{c}_{j,p}, p = 1, \dots, n_j$  of molecular populations in cell  $j$  at time  $t_j$ . Since smFISH experiments [6,50] measure a new batch of cells at every time point, there is no correlation between observations at different time points.

Assume that a model class  $\mathcal{M} = \{M(\theta)\}_{\theta \in \Theta}$  of stochastic reaction networks has been chosen to model the data consisting of a fixed set of reactions with unknown reaction parameters  $\theta$ . Let  $p(t, \mathbf{x} | M(\theta))$  denote the entry of the CME solution corresponding to state  $\mathbf{x}$  at time  $t$ , given by SRN model  $M(\theta)$ . The log-likelihood of the dataset  $\mathcal{D}$  given  $M(\theta)$  is given by

$$L(\mathcal{D} | \theta) = \sum_{i=1}^T \sum_{j=1}^{n_i} \log p(t_i, \mathbf{c}_{j,i} | M(\theta)). \quad (8)$$

A common approach to fitting this model is to find model parameters  $\theta$  that maximize the log-likelihood function. However, there are potentially many other model parameters that could fit the data approximately as well as the maximum likelihood model.

In contrast, the Bayesian approach quantifies the range of parameter uncertainty. Bayesian inference is rooted in the Bayesian philosophy in which our uncertainty is modeled using probability distributions [51]. Inference begins with a prior distribution,  $p(\theta)$ , that captures our initial beliefs. After data has been observed, the likelihood of the data given a model class and associated parameters is found as  $p(\mathcal{D} | \theta) = \exp(L(\mathcal{D} | \theta))$ . Then by applying Bayes' Theorem, we construct the posterior distribution on model parameters that reflects our updated beliefs:

$$p(\theta | \mathcal{D}) = \frac{p(\mathcal{D} | \theta)p(\theta)}{p(\mathcal{D})} \quad (9)$$

Here,  $p(\mathcal{D})$  is a normalization constant known as the model evidence. Sampling from the posterior distribution on parameters allows us to quantify our uncertainty regarding the parameter fit and predictions. Therefore, the Bayesian approach allows us to be confident in

the inference results, understand the influence of uncertainty on predictions, and design experiments to reduce the parameter uncertainty. However, it can be computationally challenging.

The Bayesian framework also provides a criteria to select or weigh different model classes as data become available. Suppose, instead of a single model class, we are given  $K$  possible network structures that could potentially explain the observations. Let

$\mathcal{M}^k = \{M^k(\theta^k)\}_{\theta^k \in \Theta^k}$  denote the  $k$ -th class, where the parameter domains  $\Theta^k$  need not have the same dimensionality. Each model class is associated with a prior weight  $P(\mathcal{M}^k)$  that represents the prior level of belief in each class. If  $\mathcal{D}$  denotes the dataset as before, we can compute the model evidence of  $\mathcal{M}^k$  as

$$P(\mathcal{D} | \mathcal{M}^k) = \int_{\theta \in \Theta^k} p(\mathcal{D} | \theta, \mathcal{M}^k) P(\theta | \mathcal{M}^k) d\theta. \quad (10)$$

The posterior probability of each model class can then be computed by applying Bayes' Theorem:

$$P(\mathcal{M}^k | \mathcal{D}) = \frac{P(\mathcal{D} | \mathcal{M}^k) P(\mathcal{M}^k)}{\sum_{j=1}^K P(\mathcal{D} | \mathcal{M}^j) P(\mathcal{M}^j)}. \quad (11)$$

These probabilities reflect the posterior weighting of the different model classes and can be used to make average predictions over the models. The drawback of Bayesian model selection in the context of stochastic gene expression is the computational cost of computing the model evidences. We will return to this issue in section 3 where we show how the Multifidelity ST-MCMC framework provides an efficient way to estimate model evidence.

## 2.4 Markov Chain Monte Carlo samplers

Markov Chain Monte Carlo algorithms are widely used for sampling the posterior distribution of a Bayesian inference problem. These methods design a Markov chain whose stationary distribution is the target posterior distribution,  $p(\theta | \mathcal{D})$ . Therefore, by simulating the evolution of samples,  $\theta$ , according to the Markov chain, correlated samples are asymptotically drawn from the posterior distribution. A common MCMC method is the Metropolis-Hastings algorithm. The algorithm begins by initializing the parameter state to some  $\theta_0$ . Then at a step  $i + 1$  in the evolution, a candidate sample  $\theta'$  is drawn according to a proposal distribution  $Q(\theta' | \theta_i)$ . This candidate is then accepted or rejected with probability  $\alpha(\theta' | \theta_i)$  given by

$$\alpha(\theta' | \theta_i) = \min\left(1, \frac{p(\mathcal{D} | \theta') p(\theta') Q(\theta_i | \theta')}{p(\mathcal{D} | \theta_i) p(\theta_i) Q(\theta' | \theta_i)}\right) \quad (12)$$

This acceptance probability is independent of the normalization constant  $p(\mathcal{D})$  and therefore computationally tractable. If the candidate is accepted  $\theta_{i+1} = \theta'$ ; otherwise,  $\theta_{i+1} = \theta_i$ . This



algorithm iterates until a sufficient number of posterior samples have been generated to accurately represent the posterior.

Effective sample size (ESS) provides a metric to judge whether sufficient samples have been generated. The ESS of a  $N$  sample population  $\theta_{i=1 \dots N}$  corresponds to the number of independent samples,  $N_{ESS}$ , which would estimate a quantity of interest,  $\hat{q} = E[q(\theta)]$ , with the same variance as the estimate from the  $N$  samples. Therefore, when designing a sampler we want to maximize  $N_{ESS}$  to attain the highest possible sampling efficiency. For MCMC this corresponds to minimizing sample correlation. However, even very effective samplers often need tens of thousands of sequential model evaluations to generate sufficient effective samples, making this form of MCMC challenging for computationally expensive models.

## 2.5 Sequential Monte Carlo samplers

To overcome many of the challenges associated with a standard Metropolis-Hastings based MCMC method, parallel methods, like Sequential Monte Carlo (SMC), have been introduced to better leverage high performance computing resources. SMC methods for Bayesian inference transport a sample population, initially distributed so that it can approximate expectations with respect to the prior, to one which can approximate posterior expectations [31,33–35]. Typically, this means a population of samples initially distributed according to the prior being transformed into a population of samples approximately distributed according to the posterior. For Sequential Tempered MCMC (ST-MCMC) [35], we break down the inference problem into a series of annealing levels  $i$  defined by an annealing factor  $\beta_i \in [0, 1]$ . Each level defines an intermediate distribution,  $\pi_{\beta_i}(\theta)$ , which we would like to use to generate samples and to compute expectations. These intermediate distributions take the form of

$$\pi_{\beta_i}(\theta) = p(\theta \mid \mathcal{D}, \beta_i) = \frac{p(\mathcal{D} \mid \theta)^{\beta_i} p(\theta)}{\int p(\mathcal{D} \mid \theta)^{\beta_i} p(\theta) d\theta} \quad (13)$$

This annealing approach is common to many SMC methods used for Bayesian inference and can be thought of as gradually integrating the influence of the data into the solution. Tempering provides several benefits namely: 1) robust handling of potentially multimodal or unidentifiable posteriors, 2) smoother evolution of the parallel sample population to avoid different rates of convergence to the posterior, 3) online adaptation of the MCMC sampler, and 4) estimation of the model evidence for model selection through thermodynamic integration.

To simplify the problem of transporting samples from the prior,  $\pi_0(\theta)$ , to posterior,  $\pi_1(\theta)$ , we transport samples sequentially through each level in the sequence, i.e.  $\pi_{\beta_i}(\theta)$  to  $\pi_{\beta_{i+1}}(\theta)$  with  $\beta_{i+1} > \beta_i$ . Because we can control the size of the jump  $\beta = \beta_{i+1} - \beta_i$  we can ensure that this change is not too drastic as to cause poor approximation of the true distribution i.e. too drastic a decrease in the ESS. Transporting samples is done in three steps:

1. Reweight the previous sample population, distributed according to  $\pi_{\beta_i}(\theta)$ , with unnormalized weights  $w_i = p(\mathcal{D} \mid \theta_i)^{\Delta\beta}$  to reflect expectations with respect to the new distribution  $\pi_{\beta_{i+1}}(\theta)$ .
2. Resample the population according to the weights so that the samples now reflect  $\pi_{\beta_{i+1}}(\theta)$ .
3. Seed a Markov chain starting at each sample and then use MCMC to explore  $\pi_{\beta_{i+1}}(\theta)$ .

The MCMC step is essential to ensure the sample population does not degenerate since the reweighting and resampling steps reduce the ESS of the population. MCMC increases the ESS because it decorrelates the seeds and explores the target distribution, causing samples to better reflect it. Typically,  $\beta$  is chosen adaptively to not decrease the ESS too much during the update. This is achieved by finding a  $\beta$  such that the coefficient of variation (COV) of the sample weights equals a target  $\kappa$ . The COV approximates the ESS by  $N_{ESS} \approx \frac{N}{1 + \kappa^2}$ .

Therefore, we find a  $\beta > 0$  that solves the equation

$$\kappa = \frac{\sqrt{\frac{1}{N} \sum_{i=1}^N (w_i(\Delta\beta) - \hat{w}(\Delta\beta))^2}}{\hat{w}(\Delta\beta)}. \quad (14)$$

Here,  $w_i(\Delta\beta) = p(\mathcal{D} \mid \theta_i)^{\Delta\beta}$  and  $\hat{w}(\Delta\beta) = \frac{1}{N} \sum_{i=1}^N w_i(\Delta\beta)$ . Typically, we choose  $\kappa = 1$ , which corresponds to a target ESS of  $N/2$ . With this method for finding  $\beta$ , we then sequentially move through all the adaptively tuned annealing levels until we reach the final posterior reflected by  $\pi_1(\theta)$ . For more details about this algorithm, see [35].

### 3. MULTIFIDELITY ST-MCMC

For expensive models, ST-MCMC and similar SMC-based methods may still be computationally prohibitive. One approach to overcome this computational burden is to utilize a multifidelity model hierarchy that can speed up sampling. The key idea is that for early levels of the ST-MCMC algorithm a low fidelity, but computationally cheap model, may be sufficiently informative to guide the samples towards the ultimate posterior distribution. This is because at early levels, the annealing factor causes the contribution of the likelihood to be damped, so perturbations in the likelihood caused by the decrease in model fidelity are less important. Intuitively, a lower fidelity model may be useful when the bias it introduces in the likelihood function is less than the variance of the likelihood at the annealing level. We consider different strategies for rigorously defining this intuition in the rest of this section.

We extend ST-MCMC to multifidelity ST-MCMC by defining intermediate levels both in terms of their annealing factor  $\beta$  and the choice of model fidelity  $M$ . We consider a hierarchy of models  $\mathcal{M} = \{M_j: j = 1 \dots K\}$  with increasing fidelity and computational cost. This algorithm is described in Algorithm 1. The key challenge is determining the best strategy for

choosing  $\beta_l$  and  $m_l$  at Step 2 since the rest of the algorithm proceeds like standard ST-MCMC.

---

**Algorithm 1:** Multifidelity ST-MCMC

---

**Inputs :** Prior distribution  $p(\theta)$   
 Model fidelity hierarchy  $M = \{M_j : j = 1, \dots, K\}$   
 Likelihood function  $p(\mathcal{D} | \theta, M_j)$   
 Number of samples  $N$

**Output :** Posterior samples  $\theta_{1:N}$

---

```

begin
  Initialize: Set level counter  $l = 0$ 
  Set annealing factor  $\beta_0 = 0$ 
  Set model level  $m_0 = 1$ 
  Define the first intermediate distribution  $\pi_0(\theta) = p(\theta)$ 
  Draw initial samples  $\theta_{1:N}^0 \sim \pi_0(\theta)$ 
  while  $\pi_{\beta_l}(\theta | M_{m_l}) \neq p(\theta | \mathcal{D}, M_{m_l})$  do
    1) Increment level counter  $l = l + 1$ 
    2) Choose the next  $\beta_l$  and  $m_l$  based on the previous level sample population  $\theta_{1:N}^{l-1}$ 
    3) Define the next intermediate distribution  $\pi_{\beta_l}(\theta | M_{m_l}) \propto p(\mathcal{D} | \theta, M_{m_l})^{\beta_l} p(\theta)$ 
    4) Compute the unnormalized importance weights for the population as  $w_i = \frac{p(\mathcal{D} | \theta_i^{l-1}, M_{m_l})}{p(\mathcal{D} | \theta_i^{l-1}, M_{m_{l-1}})}$ 
    5) Resample the population according to the normalized importance weights to initialize  $\theta_{1:N}^l$ 
    6) Evolve the samples  $\theta_{1:N}^l$  using MCMC with stationary distribution  $\pi_{\beta_l}(\theta | M_{m_l})$ 
  end
  return  $\theta_{1:N} = \theta_{1:N}^l$ 
end

```

---

### 3.1 Tempering and bridging using an Effective Sample Size criteria

One approach to choosing the appropriate annealing factor and model fidelity is a combined likelihood tempering and model bridging scheme discussed in Latz et al. [41]. This scheme is based upon the ESS statistic discussed in Section 2.5. Within Multifidelity ST-MCMC, at every level  $l$  of Algorithm 1, we choose whether to temper by changing  $\beta_l = \beta_{l-1} + \beta$  or to bridge by changing the model fidelity,  $m_l = m_{l-1} + 1$ . This choice is made by measuring the ESS of the sample population with respect to the next change in the model fidelity by computing the unnormalized weights as if we were to bridge:

$$w_i = \frac{\pi_{\beta_l - 1}(\theta_i^{l-1} | M_{m_{l-1} + 1})}{\pi_{\beta_{l-1}}(\theta_i^{l-1} | M_{m_{l-1}})} = \frac{p(\mathcal{D} | \theta_i^{l-1}, M_{m_{l-1} + 1})^{\beta_l - 1}}{p(\mathcal{D} | \theta_i^{l-1}, M_{m_{l-1}})^{\beta_{l-1}}}. \quad (15)$$

We can then compute the coefficient of variation of the weights,  $w_i$ , to determine if it exceeds a target  $\kappa$ . If it does, we choose to bridge to the next model fidelity because the sample population is beginning to degenerate, so it no longer has sufficient ESS to approximate the next level intermediate posterior. If the COV is less than  $\kappa$ , we choose to keep the current model but instead increase  $\beta$ . The next beta is chosen using the same strategy as before by solving Equation 14.

### 3.2 Information-theoretic criteria for model fidelity adaptation

We introduce a new criteria for model fidelity selection based on information theory. This criteria is motivated by the fact that the ESS-based strategy described above only decides to change fidelity based on the next model in the hierarchy. This means the sampler may continue to use a low fidelity model because it still meets the ESS criteria with respect to the next model, even when it is drifting away from the high fidelity posterior. Instead, we introduce a method that utilizes a limited number of full fidelity model evaluations to help us better decide when to bridge model fidelity. Depending on the computational cost of the full fidelity simulations, the improved bridging strategy and the improved robustness of this method may outweigh the cost of these full fidelity solutions.

Within multifidelity ST-MCMC, if the algorithm is at annealing level  $l$  with annealing factor  $\beta_l \in [0, 1]$  and has been sampling the intermediate posterior defined by a model,  $M_{m_l}$ , in a

model hierarchy  $\mathcal{M} = \{M_j: j = 1 \dots K\}$ , we would like to know whether  $M_{m_l}$  still provides meaningful information about the ultimate posterior once we move to level  $l+1$  with annealing factor  $\beta_{l+1}$ . Here we assume the ultimate posterior is  $p(\theta | \mathcal{D}, M_K)$ , where  $M_K$  is the highest fidelity model. Therefore, unlike the previous ESS-based method, we begin by proposing a tempering step under the assumption that the current model fidelity is valid. We find the proposed  $\beta_{l+1}$  by solving Equation 14.

If  $M_{m_l}$  no longer provides meaningful information at the next level, we use the next highest fidelity model in the algorithm,  $M_{m_l+1}$ . This criteria can be formulated using a generalization of information theory [52], where the information gained about the full posterior,  $p(\theta | \mathcal{D}, M_K)$ , by moving from level  $l$  to  $l+1$  with model  $M_{m_l}$  is:

$$\begin{aligned} & \mathcal{I}_{p(\theta | \mathcal{D}, M_K)}[p(\theta | \mathcal{D}, M_{m_l}, \beta_{l+1}) \| p(\theta | \mathcal{D}, M_{m_l}, \beta_l)] \\ &= D_{\text{KL}}[p(\theta | \mathcal{D}, M_K) \| p(\theta | \mathcal{D}, M_{m_l}, \beta_l)] - \\ & D_{\text{KL}}[p(\theta | \mathcal{D}, M_K) \| p(\theta | \mathcal{D}, M_{m_l}, \beta_{l+1})] \\ &= \int p(\theta | \mathcal{D}, M_K) \log \frac{p(\theta | \mathcal{D}, M_{m_l}, \beta_{l+1})}{p(\theta | \mathcal{D}, M_{m_l}, \beta_l)} d\theta. \end{aligned} \quad (16)$$

If this quantity is positive, then the intermediate posterior defined by  $\beta_{l+1}$  and  $m_l$  is closer to the ultimate posterior than the previous level, so we choose  $m_{l+1} = m_l$ . However if this is negative, this update is driving the distribution away from the ultimate posterior, so we should use a higher fidelity model for the next update, thus  $m_{l+1} = m_l + 1$ .

If we choose to update the model fidelity, we consider two strategies for choosing  $\beta_{l+1}$  for the next level. In the first strategy, keeping with the ESS-based tempering and bridging framework from above, is to set  $\beta_{l+1} = \beta_l$ . The second strategy is to tune  $\beta_{l+1}$  to try to attain an ESS target. The first strategy is often more computationally efficient, but may not be as robust if changing model fidelity introduces significant variations in the likelihood. To tune  $\beta_{l+1}$ , we first define the importance weight for transitioning from a level defined by  $\beta_j$  and  $m_j$  to a level defined by  $\beta_{j+1}$  and  $m_{j+1} = m_j + 1$  as

$$w_i = \frac{\pi_{\beta_{l+1}}(\theta_i^l | M_{m_l+1})}{\pi_{\beta_l}(\theta_i^l | M_{m_l})} = \frac{p(\mathcal{D} | \theta_i^l, M_{m_l+1})^{\beta_{l+1}}}{p(\mathcal{D} | \theta_i^l, M_{m_l})^{\beta_l}} \quad (17)$$

Using the same approach as before, we can then tune  $\beta_{l+1}$  to meet some ESS target based upon the COV of the weights. However, unlike in previous problems, this might not be achievable. If Equation (14) has a solution, we chose the largest  $\beta_{l+1}$  such that the COV target is met. If Equation (14) does not have a solution, we find the  $\beta_{l+1}$  that minimizes the COV and thus maximizes the ESS.

### 3.3 Computing the information-theoretic criteria

Since computing the information in Equation 16 requires marginalizing over the posterior, it can be challenging. However, this computation can be approximated using the samples from ST-MCMC. The first step is to recognize the connection between computing this criteria and estimating the model evidence:

$$\begin{aligned}
 & \mathcal{J}_{p(\theta | \mathcal{D}, M_K)} [p(\theta | \mathcal{D}, M_{m_l}, \beta_{l+1}) || p(\theta | \mathcal{D}, M_{m_l}, \beta_l)] \\
 &= \int p(\theta | \mathcal{D}, M_K) \log \frac{p(\theta | \mathcal{D}, M_{m_l}, \beta_{l+1})}{p(\theta | \mathcal{D}, M_{m_l}, \beta_l)} d\theta \\
 &= \int p(\theta | \mathcal{D}, M_K) \log p(\mathcal{D} | \theta, M_{m_l})^{4\beta} \frac{p(\mathcal{D} | M_{m_l}, \beta_l)}{p(\mathcal{D} | M_{m_l}, \beta_{l+1})} d\theta
 \end{aligned} \tag{18}$$

Here,  $p(\mathcal{D} | M, \beta)$  is the model evidence, i.e. the normalization, for the likelihood defined by the model  $M$  with an annealing factor  $\beta$ :

$$p(\mathcal{D} | M, \beta) = \int p(\mathcal{D} | \theta, M)^\beta p(\theta) d\theta \tag{19}$$

By noting the relationship to model evidence, the ratio of the evidences can be expressed as:

$$\begin{aligned}
 \frac{p(\mathcal{D} | M_{m_l}, \beta_l)}{p(\mathcal{D} | M_{m_l}, \beta_{l+1})} &= \frac{p(\mathcal{D} | M_{m_l}, \beta_l)}{\int p(\mathcal{D} | \theta, M_{m_l})^{\beta_{l+1}} p(\theta) d\theta} \\
 &= \frac{1}{\int p(\mathcal{D} | \theta, M_{m_l})^{4\beta} \frac{p(\mathcal{D} | \theta, M_{m_l})^{\beta_l} p(\theta)}{p(\mathcal{D} | M_{m_l}, \beta_l)} d\theta} \\
 &= \frac{1}{\mathbb{E}_{\theta | \mathcal{D}, M_{m_l}, \beta_l} [p(\mathcal{D} | \theta, M_{m_l})^{4\beta}]}.
 \end{aligned} \tag{20}$$

$\mathbb{E}_{\theta | \mathcal{D}, M_{m_l}, \beta_l}$  indicates that the expectation is taken over the distribution  $p(\theta | \mathcal{D}, M_{m_l}, \beta_l)$ .

Therefore,

$$\begin{aligned}
 & \int p(\theta | \mathcal{D}, M_K) \log p(\mathcal{D} | \theta, M_{m_l})^{4\beta} \frac{p(\mathcal{D} | M_{m_l}, \beta_l)}{p(\mathcal{D} | M_{m_l}, \beta_{l+1})} d\theta \\
 &= \int p(\theta | \mathcal{D}, M_K) \log \frac{p(\mathcal{D} | \theta, M_{m_l})^{4\beta}}{\mathbb{E}_{\theta | \mathcal{D}, M_{m_l}, \beta_l} [p(\mathcal{D} | \theta, M_{m_l})^{4\beta}]} d\theta.
 \end{aligned} \tag{21}$$

Since we cannot yet sample  $p(\theta | \mathcal{D}, M_K)$  we use importance sampling to express this integral in terms of the level  $l$  distribution, which we have samples for:

$$\begin{aligned}
 & \int p(\theta \mid \mathcal{D}, M_K) \log \frac{p(\mathcal{D} \mid \theta, M_{m_l})^{4\beta}}{E_{\theta \mid \mathcal{D}, M_{m_l}, \beta_l} [p(\mathcal{D} \mid \theta, M_{m_l})^{4\beta}]} d\theta \\
 &= \int p(\theta \mid \mathcal{D}, M_{m_l}, \beta_l) \frac{p(\theta \mid \mathcal{D}, M_K)}{p(\theta \mid \mathcal{D}, M_{m_l}, \beta_l)} \log \frac{p(\mathcal{D} \mid \theta, M_{m_l})^{4\beta}}{E_{\theta \mid \mathcal{D}, M_{m_l}, \beta_l} [p(\mathcal{D} \mid \theta, M_{m_l})^{4\beta}]} d\theta \quad (22) \\
 &\propto \int p(\theta \mid \mathcal{D}, M_{m_l}, \beta_l) \frac{p(\mathcal{D} \mid \theta, M_K)}{p(\mathcal{D} \mid \theta, M_{m_l})^{\beta_l}} \log \frac{p(\mathcal{D} \mid \theta, M_{m_l})^{4\beta}}{E_{\theta \mid \mathcal{D}, M_{m_l}, \beta_l} [p(\mathcal{D} \mid \theta, M_{m_l})^{4\beta}]} d\theta \\
 &= \mathcal{J}_{criteria}.
 \end{aligned}$$

This integral only needs to be known up to a constant of proportionality since we only need to assess whether it is positive. We can then express it in terms of expectations as:

$$\begin{aligned}
 \mathcal{J}_{criteria} &= E_{\theta \mid \mathcal{D}, M_{m_l}, \beta_l} \left[ \frac{p(\mathcal{D} \mid \theta, M_K)}{p(\mathcal{D} \mid \theta, M_{m_l})^{\beta_l}} \log p(\mathcal{D} \mid \theta, M_{m_l})^{4\beta} \right] - \\
 &E_{\theta \mid \mathcal{D}, M_{m_l}, \beta_l} \left[ \frac{p(\mathcal{D} \mid \theta, M_K)}{p(\mathcal{D} \mid \theta, M_{m_l})^{\beta_l}} \right] \log E_{\theta \mid \mathcal{D}, M_{m_l}, \beta_l} [p(\mathcal{D} \mid \theta, M_{m_l})^{4\beta}]. \quad (23)
 \end{aligned}$$

We can now estimate whether  $\mathcal{J}_{criteria}$  is positive or negative to determine if information is gained or lost by this next update. To approximate these expectations we use the NST-MCMC samples at level  $l$ , where  $\theta_i^l: i = 1 \dots N$ , which are approximately distributed according to  $p(\theta \mid \mathcal{D}, M_{m_l}, \beta_l)$ . We also use the evaluation of the full fidelity model likelihood at these points:

$$\begin{aligned}
 & E_{\theta \mid \mathcal{D}, M_{m_l}, \beta_l} \left[ \frac{p(\mathcal{D} \mid \theta, M_K)}{p(\mathcal{D} \mid \theta, M_{m_l})^{\beta_l}} \log p(\mathcal{D} \mid \theta, M_{m_l})^{4\beta} \right] \\
 &\approx \sum_{i=1}^N \frac{p(\mathcal{D} \mid \theta_i^l, M_K)}{p(\mathcal{D} \mid \theta_i^l, M_{m_l})^{\beta_l}} \log p(\mathcal{D} \mid \theta_i^l, M_{m_l})^{4\beta}. \quad (24)
 \end{aligned}$$

$$E_{\theta \mid \mathcal{D}, M_{m_l}, \beta_l} \left[ \frac{p(\mathcal{D} \mid \theta, M_K)}{p(\mathcal{D} \mid \theta, M_{m_l})^{\beta_l}} \right] \approx \sum_{i=1}^N \frac{p(\mathcal{D} \mid \theta_i^l, M_K)}{p(\mathcal{D} \mid \theta_i^l, M_{m_l})^{\beta_l}}. \quad (25)$$

$$E_{\theta \mid \mathcal{D}, M_{m_l}, \beta_l} [p(\mathcal{D} \mid \theta, M_{m_l})^{4\beta}] \approx \sum_{i=1}^N p(\mathcal{D} \mid \theta_i^l, M_{m_l})^{4\beta}. \quad (26)$$

### 3.4 Multifidelity ST-MCMC and Bayesian model selection

SMC and ST-MCMC methods not only enable robust solutions of Bayesian inference problems for parameter calibration, but also enable Bayesian model selection by providing asymptotically unbiased estimates of the model evidence. Model evidence estimates are generally highly computationally expensive since they require estimating the normalization constant,

$$p(\mathcal{D} | M) = \int p(\mathcal{D} | \theta, M) p(\theta | M) d\theta, \quad (27)$$

which consists of marginalizing the likelihood over the prior distribution. If the high probability content of the prior differs significantly from the most likely parameters according to the likelihood, it is difficult to estimate this integral using Monte Carlo samples. Instead, SMC type methods break down this estimate into a series of Monte Carlo approximations over the intermediate distribution levels previously discussed. As such, a hierarchy of multifidelity models can also be used to accelerate this estimate within the Multifidelity ST-MCMC framework. Using the methods described in [31,53] a SMC based sampler, like Multifidelity ST-MCMC, can estimate the model evidence of the highest fidelity model,  $M_K$ , by estimating the product:

$$p(\mathcal{D} | M_K) = \prod_{l=1}^L \frac{p(\mathcal{D} | m_l, \beta_l)}{p(\mathcal{D} | m_{l-1}, \beta_{l-1})} = \prod_{l=1}^L c_l. \quad (28)$$

where  $p(\mathcal{D} | m, \beta) = \int p(\mathcal{D} | \theta, M_m)^\beta p(\theta) d\theta$ , and  $L$  is the final level of ST-MCMC. The ratio  $c_l$  can be written as

$$\begin{aligned} c_l &= \frac{\int p(\mathcal{D} | \theta, M_{m_l})^{\beta_l} p(\theta) d\theta}{p(\mathcal{D} | m_{l-1}, \beta_{l-1})} \\ &= \int \frac{p(\mathcal{D} | \theta, M_{m_l})^{\beta_l}}{p(\mathcal{D} | \theta, M_{m_{l-1}})^{\beta_{l-1}}} \frac{p(\mathcal{D} | \theta, M_{m_{l-1}})^{\beta_{l-1}} p(\theta)}{p(\mathcal{D} | m_{l-1}, \beta_{l-1})} d\theta \\ &= E_{\theta | \mathcal{D}, M_{m_{l-1}}, \beta_{l-1}} \left[ \frac{p(\mathcal{D} | \theta, M_{m_l})^{\beta_l}}{p(\mathcal{D} | \theta, M_{m_{l-1}})^{\beta_{l-1}}} \right] \\ &\approx \frac{1}{N} \sum_{i=1}^N w_i^l. \end{aligned} \quad (29)$$

where  $w_i^l$  are the unnormalized resampling weights at level  $l$  for the sample population  $\theta_{i=1}^{l-1} \dots N$ . Therefore, using the weights we already computed as part of Multifidelity ST-MCMC, we are able to compute an estimate of the model evidence. The error in the estimate of the model evidence is controlled by the ESS of the Monte Carlo samples used to estimate  $c_l$  and the choice of annealing factors  $\beta_l$ . For a comprehensive analysis of the statistical properties of the estimator see [31,53].

## 4. MULTIFIDELITY REDUCED MODELS OF THE CHEMICAL MASTER EQUATION

The FSP algorithm introduced in section 2 is commonly used to compute the likelihood of observed data measurements. When used within MCMC sampling, the FSP is usually implemented in one of the following two ways.

The first is to fix a single, large, subset of states for all parameter samples [11]. Since the probability distribution of the CME changes significantly as the MCMC explores the parameter space, it is very difficult to specify a finite state set that accurately captures a significant portion of the probability mass for all times and all parameters. One can end up choosing a static FSP that is either inaccurate or inefficient. This scenario is similar to when a static discretization scheme (e.g., finite element) is employed in the simulation of parametric partial differential equation models, in which the manually chosen grid size may turn out to be too coarse for some parameter regimes and excessive for others.

This drawback motivates the second approach that instead uses adaptive CME solvers. There have been many adaptive formulations of the FSP [48,54–57] in which an approximation to the full-fidelity CME is sought within a user-specified error tolerance by expanding the state set iteratively. There could be regions in the parameter space where the adaptive state set has to be expanded to an enormous size to accurately approximate the CME solution. Yet, most of these parameter combinations fit poorly to the data, and therefore the large computational effort for their forward solutions does not provide useful information about the posterior. On the other hand, since reaction networks usually comprise of nonlinear and unpredictable interactions, it is difficult to know a priori which parameter values would give rise to such difficult (but meaningless) forward solutions of the full-fidelity CME. Simple techniques to regularize the cost of the forward solutions by restricting either the computational time or the number of time steps may run the risk of mistakenly ignoring genuinely informative parameter candidates whose evaluation just happens to require high computational cost.

The framework of the Multifidelity ST-MCMC sampler allows us to conceive of a compromise. In particular, we introduce a series of surrogate CMEs whose solution complexity is uniformly bounded across all parameters. An adaptive FSP method with strict error tolerance is applied only to these surrogate master equations. In other words, the solution to each surrogate CME can be computed accurately by the FSP with bounded cost, where coarser, yet cheaper, surrogate CMEs are used to guide the sampling process at the early annealing level, while the more expensive solutions of the high-fidelity CME are only required at the last annealing steps. This allows us to avoid overcommitting to parameters that have low posterior probabilities during the early annealing levels, yet still guarantee accurate computation of the likelihood at later sampling stages.

### 4.1 Implicitly defined finite state projection for constructing surrogate CME models

Surrogate models can be derived by adding restrictive assumptions to the physics of the original model. In particular, consider a hypothetical physical biological surrogate of the original cells modeled by the full CME in which all cellular processes ‘freeze’ when



molecular copy numbers reach a certain set of thresholds. As we increase these thresholds, the surrogate cells behave more freely and closer to the original cells and the master equation describing their behavior becomes closer to the original CME, illustrated in Fig. 1.

Let  $b_1, \dots, b_N$  be bounds on the copy number of species 1 through  $N$ . We define an approximate SRN whose propensities are surrogates of the original SRN propensities and are given by

$$\hat{\alpha}_j(\mathbf{x}) = \alpha_j(\mathbf{x}) \prod_i [x_i \leq b_i], \quad (30)$$

where  $[E]$  takes value 1 if expression  $E$  is true and zero otherwise. Since there are no further transitions once the process enters a state that exceeds the bounds, the state space of the surrogate chemical master equation is effectively reduced to the hyper-rectangle  $H(\mathbf{b}) = \times_{i=1}^N \{0, \dots, b_i\}$ . Thus, the infinite-dimensional system of differential equations (3) is replaced by the finite-dimensional surrogate dynamical system

$$M(\mathbf{b}) : \frac{d}{dt} \hat{\mathbf{p}}_H(t) = \hat{\mathbf{A}}_H \hat{\mathbf{p}}_H(t), \quad \hat{\mathbf{p}}_H(0) = \hat{\mathbf{p}}_0|_H, \quad (31)$$

where the truncated infinitesimal generator  $\hat{\mathbf{A}}_H$  is defined similar to eq. (4) but with the exact propensities replaced by the surrogate propensities given in eq. (30).

We note that eq. (31) is equivalent to eq. (6) with  $\Omega = H(\mathbf{b})$ . Thus, our surrogate propensities implicitly define a finite state projection of the original CME. We also note that the surrogate CMEs need not be constrained within a hyper-rectangle as considered here (see, e.g., [56,57] for examples of non-rectangular FSP). It may be beneficial to derive a sequence of transformations to the original propensities, with the approximations chosen in such a way that alleviate the computational burden of solving the original model by, e.g., making the lower-fidelity dynamical system less stiff than the high-fidelity one. We leave this more general strategy to future work.

For the present choice hyper-rectangular state spaces, we recall the important result mentioned in section 2.2 that, as the entries of  $\mathbf{b}$  increase monotonically, the state space  $H(\mathbf{b})$  includes more states and the truncation error, measured as the  $\ell_1$ -distance between  $\hat{\mathbf{p}}_H$  and the true CME solution  $\mathbf{p}(t)$  decreases *monotonically*. This provides us with a straightforward and natural way to form a hierarchy of surrogate models within the Multifidelity ST-MCMC framework.

## 4.2 Using the hierarchy of surrogate CMEs within the ST-MCMC framework

With the surrogate CME models formulated, a strict hierarchy of surrogate models can be defined by a sequence of bounding vectors  $\mathbf{b}^{(1)} \ \mathbf{b}^{(2)} \ \dots \ \mathbf{b}^{(K)}$  where the “ ” sign applies element-wise. The corresponding surrogate models  $M_l := M(\mathbf{b}^{(l)})$  are then defined as in eq. (31). As mentioned earlier, the error in the surrogate CMEs decrease monotonically as we increase the bounds. Therefore,  $\{M_l\}$  forms a hierarchy in which each level attains more fidelity than its predecessor.

At the  $l$ -th level, the log-likelihood function in eq. (8) is approximated by [58]

$$L(\mathcal{D} | \theta) \approx L_{M_l}(\theta) = \sum_{i=1}^T \sum_{j=1}^{n_i} \log p(t_i, \min(c_{j,i}, \mathbf{b}^{(l)}) | M_l(\theta)). \quad (32)$$

In the surrogate log-likelihood, the data is projected onto the finite state space  $H(\mathbf{b}^{(l)})$ , and the probabilities of the data at different time points are computed from the surrogate Markov model  $M_l$ . Clearly, as  $l$  increases, the surrogate function  $L_{M_l}(\theta)$  becomes a more accurate approximation to the true log-likelihood  $L(\mathcal{D} | \theta)$ . In the ideal situation where the hierarchy is allowed to have infinite depth, these surrogates are guaranteed to form a sequence that asymptotically converges pointwise to the true log-likelihood from below. This is shown formally in the following proposition.

**Proposition 4.1.**—Let the sequence of bounds  $\{\mathbf{b}^{(l)}\}_{l=1}^{\infty} \subset \mathbb{N}^N$ , where  $\mathbf{b}^{(l)} := (b_1^{(l)}, \dots, b_N^{(l)})$  be chosen such that  $\mathbf{b}^{(l)} \leq \mathbf{b}^{(l+1)}$  elementwise (i.e.,  $b_j^{(l)} \leq b_j^{(l+1)}$ ). Assume that the continuous-time Markov chain underlying the SRN is non-explosive for all parameters and that the initial distribution of the CME (3) has finite support. For each fixed value of the parameter  $\theta$ , we have the following:

1.  $L_{M_l}(\theta) \rightarrow L(\mathcal{D} | \theta)$  as  $l \rightarrow \infty$ .
2. There exists a subsequence  $l_i$  such that  $L_{M_{l_i}} \uparrow L(\mathcal{D} | \theta)$ .

Here, the log-likelihood function  $L(\mathcal{D} | \theta)$  is defined as in (8) and the surrogate  $L_{M_{l_i}}(\theta)$  is defined as in (32).

**Proof.:** Without loss of generality, we assume that the initial distribution is concentrated at a single state  $\mathbf{x}_0$ . Let  $R$  be the number of reactions that occur during the finite time interval  $[0, t_T]$  given the starting state  $\mathbf{x}_0$ . If there exists  $\varepsilon > 0$  for which  $\mathbb{P}(R > R_\varepsilon) \geq \varepsilon$  for every choice of  $R_\varepsilon$ , then we have  $\mathbb{P}(R = \infty) \geq \varepsilon$ , violating the assumption of non-explosion. Thus, for every  $\varepsilon > 0$ , there exists  $R_\varepsilon$  such that  $\mathbb{P}(R > R_\varepsilon) < \varepsilon$ . Furthermore, we can find  $l_\varepsilon$  such that  $H(\mathbf{b}^{(l_\varepsilon)})$  contains all states that are reachable from  $\mathbf{x}_0$  via  $R_\varepsilon$  reactions or fewer. Thus, the probability for a sample path within  $[0, t_T]$  of the surrogate CME to ever exceed  $H(\mathbf{b}^{(l)})$  is less than  $\varepsilon$ , and the corresponding solution of the surrogate CME is guaranteed to be less than  $\varepsilon$  away (in one-norm) from the true CME solution. This proves (i).

We have  $\min(c_{j,i}, \mathbf{b}^{(l)}) = c_{j,i}$  for sufficiently large  $l$ . Entry-wise, the FSP approximations increase monotonically [10, Theorem 2.2], so  $p(t_i, c_{j,i} | M_l(\theta))$  increases monotonically. We can then choose the subsequence  $\{l_i\}$  from  $\{l\}$  by simply truncating the leading elements until  $\mathbf{b}^{(l)}$  disappears from the  $\min(\cdot, \cdot)$  function in eq. (32). This proves (ii).  $\square$

In summary, the FSP scheme allows us to define a hierarchy of surrogate master equations that approach the true CME as the surrogate state space enlarges. From this, we can define a sequence of surrogate log-likelihood functions that converge to the true log-likelihood from

below. These surrogates could be used within the Multifidelity ST-MCMC framework introduced in section 3. Before we do so, however, we must first ensure that an accurate solution to the system (31) could be computed efficiently.

### 4.3 Fast and accurate solution of the surrogate master equation

Although the surrogate master equation (31) is a significant reduction from the infinite-dimensional CME, the number of states included in the truncated state space  $H(b)$  still grows as  $\mathcal{O}(b_1 \cdot \dots \cdot b_N)$  and the surrogate CME can quickly become expensive as we increase the entries of  $b$ . However, in practice, the probability mass of the solution vector  $\tilde{p}_H(t)$  tends to concentrate at a much smaller subset of states. It is therefore advantageous to approximate  $\hat{p}_H(t)$  with a more compactly supported distribution. More precisely, if we let  $\varepsilon > 0$  be an error tolerance, we can use a distribution  $\tilde{p}_\Omega$  supported on  $\Omega \subset H$  such that  $\|\hat{p}_H - \tilde{p}_\Omega\| \leq \varepsilon$ . Here, the FSP error bound (8) plays a critical role in choosing the appropriate support set  $\Omega$ . We note that this error bound was recently utilized by Fox et al. [58] to compute rigorous lower and upper bounds for the true log-likelihood function (8), from which comparison between certain models could be done even at a low-fidelity FSP solution. We do not pursue this direction in the present work.

To efficiently compute the solution of the surrogate CMEs using the principles just mentioned, we employ a new FSP implementation recently developed by Vo and Munsky [59]. This solver divides the time interval of interest  $[0, t_f]$  into subintervals  $I_j := [t_j, t_{j+1})$ ,  $j = 0, \dots, n_{step} - 1$  with  $0 := t_0 < t_1 < \dots < t_{n_{step}} := t_f$ . On each time subinterval  $I_j$ , the dense tensor  $\hat{p}_H(t)$  that is the solution of the surrogate CME (31) is approximated by a sparse tensor  $\tilde{p}_{\Omega_j}(t)$  supported on  $\Omega_j \subset H$ , obtained from solving

$$\frac{d}{dt} \tilde{p}_{\Omega_j}(t) = \tilde{A}_{\Omega_j} p_{\Omega_j}(t), t \in [t_j, t_{j+1}), \quad (33)$$

where

$$\tilde{A}_{\Omega_j}(y, x) = \begin{cases} \hat{A}_H(y, x) & \text{if } x, y \in \Omega_j \\ 0 & \text{otherwise} \end{cases}$$

Clearly, in solving (33), we only need to keep track of the equations corresponding to states in  $\Omega_j$  and that reduces the computational cost significantly.

From the FSP error bound (7), we derive an error-control criteria of the form

$$g_j(t) = 1 - \mathbf{1}^T \tilde{p}_{\Omega_j}(t) \leq \frac{t}{t_f} \varepsilon. \quad (34)$$

If at some  $t \in [t_j, t_{j+1})$  we find that the inequality is not satisfied, more states are added to  $\Omega_j$  and the integration starts again from  $t_j$  until the criteria is satisfied over the whole interval. The determination of the time steps  $t_j$  is left to the ODE integrator employed for solving eq. (33). The advantage of using such an adaptive FSP approach, compared to the classic FSP as formulated in [10], is that we need not start with the full state space at the beginning of the

integration. Rather, a small projection can be used to advance the solution over the early time points, while more states are included at the later time points only when needed using the time-dependent FSP error bound (34).

The state sets  $\Omega_j$  are chosen as integral solutions of a set of inequality constraints. In particular, they have the form

$$\Omega_j = \{ \mathbf{x} \in H(b) \mid f_i(\mathbf{x}) \leq c_i^{(j)} \}, \quad (35)$$

where  $f_i$  are functions that are chosen a priori, and  $c_i > 0$  are positive scalars. To expand  $\Omega_j$ , we simply increase  $c_i^{(j)}$  and run a breadth-first-search routine to explore all reachable states that satisfy the relaxed inequality constraints.

Implementation-wise, the approximate solution  $\tilde{\mathbf{p}}_{\Omega_j}$  is stored in the coordinate format similar to that used for sparse tensors [60]. The list of tensor indices is managed with the Distribute Dictionary data structure in the software package Zoltan [61,62]. We also make use of parallel objects from the PETSc library [63–65].

These MPI-based libraries allow our implementation to scale into multiple computing nodes, details of the parallel version of our CME solver is communicated elsewhere [59]. However, it is worth pointing out that, as the number of cores increase, there is a diminishing return on the speedup gained from parallelizing the forward solution of the CME due to the communication cost inherent in numerical operations such as matrix-vector multiplications. Therefore, simply plugging a parallel forward solution code on an increasing number of nodes into a serial MCMC sampler such as Metropolis-Hastings will have diminishing benefits. The ST-MCMC, in contrast, allows us to achieve better utilization of the computing resources, since it is embarrassingly parallel. Doubling the number of processors simply enables us to simultaneously sample twice as many parameter samples in about the same computational time.

We also note that using an adaptive solver such as one we present here incurs some numerical error in the surrogate likelihood function. However, we expect this error to be negligible with a conservative choice of error tolerance. In particular, the error threshold  $\epsilon$  is always fixed at  $10^{-8}$  in our numerical tests.

In the next section, we will confirm the accuracy and efficiency of our combined multifidelity sampler and adaptive model reduction scheme when applied to two biologically inspired problems and one on a real experimental dataset.

## 5. NUMERICAL EXAMPLES

In the following tests we compare the four variants of the ST-MCMC described above: the Full-fidelity, ESS-Bridge, IT-Bridge, and Tuned IT-Bridge schemes. The full-fidelity scheme is the classic ST-MCMC with every likelihood evaluation using the highest model fidelity. The remaining three schemes are Multifidelity ST-MCMC variants in which the bridging between fidelity levels are determined based on the ESS, the new information-theoretic

criteria with or without  $\beta$ -tuning (see eq. (17) and the preceding discussion in section 3.2). In each of these Multifidelity schemes, the surrogate likelihood is formulated as described in section 4.1. When all propensities are time-invariant as in the first two examples, the reduced system of ODEs in eq. (33) is solved by computing the action of the matrix exponential operator using the Krylov approximation with Incomplete Orthogonalization Procedure [66–68], with the Krylov basis size fixed at 30. In the case of time-varying propensities in the third example, we use the Four stage third order L-stable Rosenbrock-W scheme [69] implemented in the TS module of PETSc [70]. All these ODEs solvers are set with conservative absolute tolerance of  $10^{-14}$ , and relative tolerance of  $10^{-4}$ . All the codes for the numerical experiments are available at the public Github repository at <https://github.com/voduchuy/StMcmcCme>.

### 5.1 Parameter inference for repressilator gene circuit

We first consider a three-species model inspired by the well-known repressilator gene circuit [1]. This model consists of three species, TetR,  $\lambda$ cI and LacI, which constitute a negative feedback network (Table 1). We simulate a dataset that consists of five measurement times 2, 4, 6, 8, and 10 minutes, with 1000 cells measured at each time point. These numbers of single-cell measurements are typical of smFISH experiments [6,44]. We assume that all cells start at the state  $x_0 = (\text{TetR}, \lambda\text{cI}, \text{LacI}) = (0, 0, 0)$ , so that at the initial time where there are no gene products.

The hierarchy of surrogate CMEs (cf. (31)) is defined by the bounds

$$b^{(l)} = \begin{bmatrix} b_{\text{TetR}}^{(l)} \\ b_{\lambda\text{cI}}^{(l)} \\ b_{\text{LacI}}^{(l)} \end{bmatrix} = \begin{bmatrix} \left\lceil c_1 + (l-1) \frac{d_1 - c_1}{L_{\max} + 1} \right\rceil \\ \left\lceil c_2 + (l-1) \frac{d_2 - c_2}{L_{\max} + 1} \right\rceil \\ \left\lceil c_3 + (l-1) \frac{d_3 - c_3}{L_{\max} + 1} \right\rceil \end{bmatrix}$$

where  $(c_1, c_2, c_3) = (20, 40, 40)$  and  $(d_1, d_2, d_3) = (50, 100, 100)$ , with  $L_{\max} = 10$ . Therefore, the multifidelity ST-MCMC will transit through ten levels, with the highest-fidelity model having a state space of size  $51 \times 101 \times 101$ .

We conduct parameter inference in  $\log_{10}$ -transformed space. The prior for the parameters is chosen to be a multivariate normal distribution (in  $\log_{10}$  space) with a diagonal covariance matrix (see table 2). We ran both the ST-MCMC with the highest-fidelity surrogate CME and the multifidelity ST-MCMC on 29 nodes, with 36 cores per node with 1044 parallel chains. For each level, samples were evolved using Metropolis-Hastings MCMC until a correlation target of 0.6 was reached. The proposal distribution was adaptively tuned as part of the algorithm. Details on the ST-MCMC sampler and its tuning can be found in [35]. Fig. 2 shows the time taken for each sampling scheme to reach a certain annealing level, with the multifidelity schemes with our proposed IT-based criteria outperforming the state-of-the-art fixed-fidelity ST-MCMC and ESS Bridging schemes. Specifically, while the full-fidelity ST-MCMC took over 25 hours to finish, the Multifidelity ST-MCMC with ESS, Information-theoretic, and Tuned Information-theoretic Bridging took respectively 7.2, 4.1 and 5.3 hours,

resulting in speedup factors of about 3.5, 6.2 and 4.8. The novel Information-theoretic (IT) schemes are clearly faster than the ESS-based scheme in this example, with the untuned IT scheme almost twice as fast as the ESS-based scheme. We observe that for the early levels of the algorithms, when  $\beta$  is small, the lowest fidelity model is sufficiently informative. Further, at these early levels the ESS-based scheme is slightly faster than the others since it does not require any full model evaluations. However, after  $\beta$  gets larger, the IT-based methods start to outperform the ESS-based method since they use the full model evaluations to judge that they do not need to bridge to the higher fidelity models as quickly as the ESS-based scheme does. Although the prior assigns a probability density of only about  $8.766 \times 10^{-20}$  to the true parameter vector, all samplers were able to bring the particles close to the true parameters (Fig. 4). There is no notable difference in the shapes of the posterior distributions constructed from the samples of these two schemes (Fig.3 and table 2).

## 5.2 Bayesian comparison of compartmental models of gene expression

We next explore the application of multifidelity ST-MCMC to the problem of model selection. We consider a class of compartmental multi-state gene expression models based on the model considered in [71]. The model separates biomolecules into the nuclear and cytoplasmic compartments. The reaction network consists of a gene that could switch between an inactivated state  $G_0$  and several activated states  $G_i$ ,  $i = 1, \dots, n_G - 1$ . When activated, these gene can be transcribed into RNA molecules within the nucleus at the rate of  $r_i$  molecule/minute on average. These nuclear mRNA molecules are then transported into the cytoplasm at a rate of  $k_{trans}$  molecule/min, where they degrade at the probabilistic rate  $\gamma$  molecule/minute. Overall, the model consists of  $n_G + 2$  species: genes that are at different states, nuclear mRNA and cytoplasmic mRNA. These molecular species that can go through  $3n_G + 1$  reaction channels (Table 3). Only the copy numbers of the nuclear and cytoplasmic mRNA species are observable in experiments. We want to use model selection to decide the number  $n_G$  of gene states that best explain the observed data.

We simulate a ground truth dataset based on the model with  $n_G := 3$ , which consists of 1000 single-cell measurements for each time point  $t \in \{2, 4, 6, 8, 10\}$  (with hour as time unit). We then use the multifidelity ST-MCMC using the information-theoretic criteria with  $\beta$ -tuning to estimate the model evidence for three classes of reaction networks that consist of two, three, and four gene states and compare these results with the model evidence founding using the full model. We choose the information-theoretic criteria with  $\beta$ -tuning over the other multifidelity approaches because it is the most robust and uses the  $\beta$ -tuning to avoid sampling degeneracy when the model fidelity changes. This is very important for computing model evidence because the model evidence estimate error is related to the KL-divergence between the intermediate distributions. The 2 and 3 state model were run using 7 nodes with 36 cores, while the 4 state model was run on 14 nodes also with 36 cores each. Each ST-MCMC used 1008 parallel samples. For each level, chains were run using MCMC until a correlation target of 0.6 was reached or 100 iterations exceeded. These results are summarized in Table 4. The model evidence ranges are computed using the approach described in [53].

From our results, we observe that not only does the Tuned-IT multifidelity ST-MCMC provide consistent estimate of the model evidence compared to the full-model based ST-MCMC, it actually predicts less error. All the while taking less time with speed up factors of about 1.6, 3.8 and 3.2 for the 2, 3, and 4 gene state model respectively. The improved estimate of the multifidelity approach is likely due to the fact that it uses more intermediate levels so it has a finer discretionary of the thermodynamic integration used to estimate the evidence.

The computed evidence indicates that the present data does not significantly favor one model choice over others. Clearly more experiments are required to provide conclusive evidence for model selection, and the multifidelity framework allows us to realize the insufficiency of data faster than the full-fidelity scheme. This is an important advantage in practice, as the faster assessment of current experimental data will likely reduce the lag time between consecutive batches of experiments. Further it may be possible to integrate the same multifidelity ST-MCMC based approach into Bayesian experimental design to speed up estimating the expected information gain from various experimental setups in order to design experiments to better discriminate between the models.

### 5.3 Stochastic transcription of the inflammation response gene IL1beta

Having explored the performance of the Multifidelity ST-MCMC schemes with FSP on theoretical examples with simulated datasets, we apply our method on modeling real datasets. We consider the expression of the IL1beta gene in response to LPS stimulation that was studied in Kalb et al. [44]. The dataset consists of mRNA counts for IL1beta measured right before applying LPS stimulation, as well as those at [0.5, 1, 2, 4] hours after. We consider a three-state gene expression model with a time-varying deactivation rate. We assume the initial state (2, 0, 0, 0). The observed mRNA counts are fit to the solutions of the CME at times  $T_0 + \{0, 0.5, 1, 2, 4\}$  hour, where the time offset  $T_0$  is to be estimated. The influence of LPS-induced signaling molecules is modeled by the function of the form

$$S(t) = \max\{0, \exp(-r_1(t - T_0))(1 - \exp(-r_2(t - T_0)))\}. \quad (36)$$

This signal affects the rate by which the gene turns off,

$$k_{1,0}(t) = \max\{0, a_{10} - b_{10}S(t)\}.$$

This results in a chemical reaction network with time-varying propensities and eleven uncertain parameters (Table 5). Similar to the previous example, the gene state is hidden and the data only contains measurements of the mRNA copy numbers.

Fig. 5 summarizes the performance of the four sampling schemes. The full-fidelity model considered in this example has a state space of only 18,432 states, and could be solved quickly without any reduction scheme. Yet, we still observe significant speedup from the Multifidelity schemes. Specifically, the Multifidelity ST-MCMC with ESS, Information-theoretic and Tuned Information-theoretic Bridging took respectively 3153, 2384, and 2703 seconds to finish, with speedup factors of 1.7, 2.2, and 2.0 over the full-fidelity scheme that takes over 5468 seconds. In Figure 6, we see the evolution of the model parameters for the

different methods, and we can use this to better understand differences from Figure 5. It took significantly longer for the ESS-based method to bridge to higher fidelity models than the IT-based methods. However, when it did bridge it jumped straight to the highest model fidelity. For several parameters, we see that their evolution under the ESS-based scheme accumulated significant bias at lower  $\beta$  levels before being corrected when the ESS sampler started bridging much later on (e.g. parameters  $r_1$ ,  $k_{01}$ , and  $T_0$ ). Furthermore, once bridging occurred the distributions were very far apart from each other so the sample population degenerated. This can be seen in the bias that occurs in parameters  $r_2$ ,  $b_{10}$  and  $\alpha_1$  immediately after bridging. The fact that the ESS-based method does not use any information from the full posterior explains this delay and degeneracy as it is unable to detect the emergence of bias. In contrast, the IT-based methods use the guidance of full model evaluations to better recognize the emergence of bias, so they correct it quicker. Therefore, the IT-based schemes have a smoother evolution and as a result take less time. However, we do observe bias occur in  $r_2$  after one bridging step for the IT-Bridge without  $\beta$  tuning, indicating some sampling degeneracy. This is not the case for IT-Bridge with  $\beta$  tuning since it is specifically designed to avoid degeneracy.

All sampling schemes essentially arrived at the same posterior estimates for the model parameters (Table 6). Despite significant posterior variance for some parameters, the Bayesian prediction for the distributions of RNA copy number has negligible uncertainties, and they appear to correspond reasonably well with the experimental data at the beginning and the end of the measurement time period (Fig. 7). We notice that this is not necessarily the only model structure that could explain the data, and there may yet be other models that could fit and predict single-cell behavior more accurately. The speedup enabled by the Multifidelity framework will allow the researcher the ability to more rapidly propose, assess, and choose between different alternative models.

## 6. CONCLUSION

Rapid advancements in experimental techniques are allowing biologists to collect quantitative data about cellular processes at ever smaller scales with increasing detail [6–8]. Mathematical models have become an indispensable part in the process of learning and making predictions from this data. Stochastic reaction networks (SRNs) form a powerful class of models that have found widespread use within the quantitative biology community [7]. Identifying these models from the data, however, is a challenging task due to the computational cost of solving the chemical master equation (CME). This has prevented a fully Bayesian statistical framework from being adopted widely in real biological studies. In this paper, we seek to address the challenge of applying the Bayesian philosophy to analyzing stochastic gene expression data by proposing an efficient computational framework for Bayesian parameter calibration and model selection for SRNs. This framework combines novel multifidelity formulations of the massively parallel ST-MCMC sampler with surrogate models of the CME. Numerical tests demonstrate that this combined approach leads to significant savings in comparison to a state-of-the-art method that uses solely the high-fidelity models. Further, we also propose a new criteria for tuning model fidelity within multifidelity SMC type methods based on information theory that compares favorably to effective sample size based techniques.



The research reported here may potentially lead to fruitful future directions. With respect to surrogate models, the approach proposed here for the efficient solution of the surrogate master equations is only one among various alternatives that have been proposed over the years since the introduction of the FSP algorithm [10]. Another attractive option for constructing multifidelity models is to utilize a low-rank tensor format such as the quantized tensor train that has been proposed for the forward solution of the CME [72–77]. It is also possible to exploit bounds on the log-likelihood function as done in Fox et al. [58]. While the present paper focuses on inference from smFISH data, our approach may also be adapted to analyze other types of single-cell data, such as time-course fluorescence measurements whose likelihood is also amenable to the FSP approach [78]. The improved efficiency may lead to more widespread adoptions of the Bayesian approach in answering biological questions. We refer to Catanach et al. [16] for an example of a Bayesian approach to studying the phenomenon of context dependence in synthetic gene circuits using Bayesian model selection, which required significant computational resources.

There are also many directions for improving Multifidelity ST-MCMC in general. We expect estimating the information gain criteria could be significantly improved. One possibility is using a more advanced sampling scheme that leverages model evaluations from across the multifidelity hierarchy. This could even further reduce the number of full model evaluations needed at each level. Another avenue of research is designing Multifidelity ST-MCMC specifically for estimating a quantity of interest to a given accuracy as is done in Multilevel MCMC. If we have a design object, ST-MCMC may not need to progress through the full model hierarchy or all annealing levels in order to provide enough information to estimate the quantity of interest to the desired accuracy. By further reducing the computational cost of Bayesian methods like ST-MCMC, engineers and scientists will be better able to integrate uncertainty quantification into their workflow. Therefore, as high performance computing resources are becoming increasingly accessible, we expect the Multifidelity ST-MCMC framework to provide a useful tool for researchers who are interested in model calibration and uncertainty propagation for complex models.

## ACKNOWLEDGEMENT

HV and BM were supported by National Institutes of Health (R35 GM124747). We thank James Werner and Daniel Kalb for kindly sharing with us the data from their smFISH experiment. The cited work [44] was performed, in part, at the Center for Integrated Nanotechnologies, an Office of Science User Facility operated for the U.S. Department of Energy (DOE) Office of Science by Los Alamos National Laboratory (Contract 89233218CNA000001) and Sandia National Laboratories (Contract DE-NA-0003525). The work presented here was also funded in part by the Department of Energy Office of Advanced Scientific Computing Research through the John von Neumann Fellowship. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

We also thank Ania-Ariadna Baetica for providing constructive comments on the manuscript and Jed Duersch for discussions regarding information theory.

## REFERENCES

1. Elowitz MB and Leibler S, A synthetic oscillatory network of transcriptional regulators., *Nature*, 403(6767):335–338, 2000. [PubMed: 10659856]
2. Munsky B, Trinh B, and Khammash M, Listening to the noise: random fluctuations reveal gene network parameters., *Mol. Syst. Biol.* 5(318):318, 2009. 10.1038/msb.2009.75 [PubMed: 19888213]
3. Neuert G, Munsky B, Tan RZ, Teytelman L, Khammash M, and Van Oudenaarden A, Systematic identification of signal-activated stochastic gene regulation, *Science (80-.)*, 339(6119):584–587, 2013. [www.sciencemag.org/cgi/content/full/339/6119/584/DC1](http://www.sciencemag.org/cgi/content/full/339/6119/584/DC1)
4. Singh A and Dennehy JJ, Stochastic holin expression can account for lysis time variation in the bacteriophage  $\lambda$ , *J. R. Soc. Interface*, 11(95), 2014. 10.1098/rsif.2014.0140<http://rsif.royalsocietypublishing.org>.
5. Van Boxtel C, Van Heerden JH, Nordholt N, Schmidt P, and Bruggeman FJ, Taking chances and making mistakes: Non-genetic phenotypic heterogeneity and its consequences for surviving in dynamic environments, *J. R. Soc. Interface*, 14(132), 2017. <https://royalsocietypublishing.org/doi/pdf/10.1098/rsif.2017.0141>
6. Raj A, van den Bogaard P, Rifkin SA, van Oudenaarden A, and Tyagi S, Imaging individual mRNA molecules using multiple singly labeled probes, *Nat. Methods*, 5:877, 9 2008. 10.1038/nmeth.1253<https://www.nature.com/articles/nmeth.1253/#supplementary-information> [PubMed: 18806792]
7. Munsky B, Fox Z, and Neuert G, Integrating single-molecule experiments and discrete stochastic models to understand heterogeneous gene transcription dynamics, *Methods*, 85:12–21, 9 2015. <http://www.ncbi.nlm.nih.gov/pubmed/26079925><http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC4537808> [PubMed: 26079925]
8. Li G and Neuert G, Multiplex RNA single molecule FISH of inducible mRNAs in single yeast cells, *Sci. data*, 6(1):94, 6 2019. [PubMed: 31209217]
9. Gillespie DT, A rigorous derivation of the chemical master equation, *Phys. A Stat. Mech. its Appl*, 188(1–3):404–425, 9 1992. <http://linkinghub.elsevier.com/retrieve/pii/037843719290283V>
10. Munsky B and Khammash M, The finite state projection algorithm for the solution of the chemical master equation, *J. Chem. Phys.* 124(4):44104, 2006. 10.1063/1.2145882<http://aip.scitation.org/toc/jcp/124/4><http://jcp.aip.org/resource/1/jcpsa6/v124/i4/p044104{s1>
11. Gomez-Schiavon M, Chen LF, West AE, and Buchler NE, BayFish: Bayesian inference of transcription dynamics from population snapshots of single-molecule RNA FISH in single cells, *Genome Biol*, 18(1):164, 2017. <https://genomebiology.biomedcentral.com/track/pdf/10.1186/s13059-017-1297-9> [PubMed: 28870226]
12. Schnoerr D, Sanguinetti G, and Grima R, Approximation and inference methods for stochastic biochemical kinetics—a tutorial review, *J. Phys. A*, 50(9):093001, 1 2017. <https://doi.org/10.1088%2F1751-8121%2Faa54d9>
13. Tiberi S, Walsh M, Cavallaro M, Hebenstreit D, and Finkenstädt B, Bayesian inference on stochastic gene transcription from flow cytometry data, In *Bioinformatics*, Vol. 34, pp. i647–i655, 2018. [PubMed: 30423089]
14. Weber L, Raymond W, and Munsky B, Identification of gene regulation models from single-cell data, *Phys. Biol.* 15(5):055001, 5 2018. <https://doi.org/10.1088%2F1478-3975%2Faabc31> [PubMed: 29624181]
15. Lin YT and Buchler NE, Exact and efficient hybrid monte carlo algorithm for accelerated bayesian inference of gene expression models from snapshots of single-cell transcripts, *J. Chem. Phys.* 151(2):024106, 2019. 10.1063/1.5110503 [PubMed: 31301707]
16. Catanach TA, McCardell R, Baetica AA, and Murray RM, Context dependence of biological circuits, *bioRxiv*, 2018. <https://www.biorxiv.org/content/early/2018/07/03/360040>
17. Lindley DV, On a Measure of the Information Provided by an Experiment, *Ann. Math. Stat.* 27(4):986–1005, 2007. [https://projecteuclid.org/download/pdf/\\_1/euclid.aoms/1177728069](https://projecteuclid.org/download/pdf/_1/euclid.aoms/1177728069)
18. Huan X and Marzouk YM, Simulation-based optimal Bayesian experimental design for nonlinear systems, *J. Comput. Phys.* 232(1):288–317, 2013. 10.1016/j.jcp.2012.08.013

19. Ruess J, Miliás-Argeitis A, and Lygeros J, Designing experiments to understand the variability in biochemical reaction networks., *J. R. Soc. Interface*, 10(88):20130588, 2013. <http://www.ncbi.nlm.nih.gov/pubmed/23985733>{%37}5Cn<http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3785824> [PubMed: 23985733]
20. Fox ZR and Munsy B, The finite state projection based Fisher information matrix approach to estimate information and optimize single-cell experiments, *PLOS Comput. Biol.*, 15(1):e1006365, 1 2019. <http://dx.plos.org/10.1371/journal.pcbi.1006365> [PubMed: 30645589]
21. Busetto AG, Hauser A, Krummenacher G, Sunnåker M, Dimopoulos S, Ong CS, Stelling J, and Buhmann JM, Near-optimal experimental design for model selection in systems biology, *Bioinformatics*, 29(20):2625–2632, 07 2013. 10.1093/bioinformatics/btt436 [PubMed: 23900189]
22. Vo HD, Fox Z, Baetica A, and Munsy B, Bayesian Estimation for Stochastic Gene Expression Using Multifidelity Models, *J. Phys. Chem. B*, 123(10):2217–2234, 3 2019. <http://pubs.acs.org/doi/10.1021/acs.jpcc.8b10946> [PubMed: 30777763]
23. Wu Q, Smith-Miles K, and Tian T, Approximate Bayesian computation schemes for parameter inference of discrete stochastic models using simulated likelihood density, *BMC Bioinformatics*, 15(12), 11 2014.
24. Prescott TP and Baker RE, Multifidelity Approximate Bayesian Computation, 2018. <https://github.com/tpprescott/mf-abc>.<http://arxiv.org/abs/1811.09550>
25. Warne DJ, Baker RE, and Simpson MJ, Rapid bayesian inference for expensive stochastic models, 2019.
26. Zechner C, Ruess J, Krenn P, Pelet S, Peter M, Lygeros J, and Koepl H, Moment-based inference predicts bimodality in transient gene expression, *Proc. Natl. Acad. Sci. U. S. A.*, 109(21):8340–8345, 2012. [PubMed: 22566653]
27. Cao Z and Grima R, Accuracy of parameter estimation for auto-regulatory transcriptional feedback loops from noisy data, *J. R. Soc. Interface*, 16(153), 2019.
28. Öcal K, Grima R, and Sanguinetti G, Parameter estimation for biochemical reaction networks using Wasserstein distances, 53(3):034002, 2020.
29. Rosenthal JS, Parallel computing and monte carlo algorithms, *Far east journal of theoretical statistics*, 4(2):207–236, 2000.
30. Liu JS, Liang F, and Wong WH, The multiple-try method and local optimization in metropolis sampling, *Journal of the American Statistical Association*, 95(449):121–134, 2000. <https://amstat.tandfonline.com/doi/abs/10.1080/01621459.2000.10473908>
31. Ching J and Chen YC, Transitional markov chain monte carlo method for bayesian model updating, model class selection, and model averaging, *Journal of Engineering Mechanics*, 133(7):816–832, 2007. <https://ascelibrary.org/doi/abs/10.1061/%28ASCE%290733-9399%282007%29133%3A7%28816%29>
32. Wu S, Angelikopoulos P, Papadimitriou C, and Koumoutsakos P, Bayesian Annealed Sequential Importance Sampling: An Unbiased Version of Transitional Markov Chain Monte Carlo, *ASCE-ASME J Risk and Uncert in Engrg Sys Part B Mech Engrg*, 4(1), 09 2017, 011008. 10.1115/1.4037450
33. Moral PD, Doucet A, and Jasra A, Sequential Monte Carlo samplers, *J. R. Stat. Soc. B*, 68:411–436, 2006. <https://www.stats.ox.ac.uk/~doucet/delmoral{ }doucet{ }jasra{ }sequentialmontecarlosamplersJRSpdf>
34. Kantas N, Beskos A, and Jasra A, Sequential monte carlo methods for high-dimensional inverse problems: A case study for the navier–stokes equations, *SIAM/ASA Journal on Uncertainty Quantification*, 2(1):464–489, 2014. 10.1137/130930364
35. Catanach TA and Beck JL Bayesian updating and uncertainty quantification using sequential tempered mcmc with the rank-one modified metropolis algorithm, 2018.
36. Andrés Christen J and Fox C, Markov chain Monte Carlo using an approximation, *J. Comput. Graph. Stat.*, 14(4):795–810, 2005.
37. Efendiev Y, Hou T, and Luo W, Preconditioning markov chain monte carlo simulations using coarse-scale models, *SIAM Journal on Scientific Computing*, 28(2):776–803, 2006. 10.1137/050628568

38. Cui T, Marzouk YM, and Willcox KE, Data-driven model reduction for the Bayesian solution of inverse problems, *Int. J. Numer. Methods Eng*, 102(5):966–990, 2015. <https://arxiv.org/pdf/1403.4290.pdf>
39. Koutsourelakis PS, A multi-resolution, non-parametric, Bayesian framework for identification of spatially-varying model parameters, *J. Comput. Phys*, 228(17):6184–6211, 9 2009.
40. Dodwell TJ, Ketelsen C, Scheichl R, and Teckentrup AL, A hierarchical multilevel markov chain monte carlo algorithm with applications to uncertainty quantification in subsurface flow, *SIAM/ASA Journal on Uncertainty Quantification*, 3(1):1075–1108, 2015. 10.1137/130915005
41. Latz J, Papaioannou I, and Ullmann E, Multilevel Sequential Monte Carlo for Bayesian inverse problems, *J. Comput. Phys*, 368:154–178, 2018. <https://arxiv.org/pdf/1709.09763.pdf>
42. Golightly A, Henderson DA, and Sherlock C, Delayed acceptance particle MCMC for exact inference in stochastic kinetic models, *Stat. Comput*, 25(5):1039–1055, 2015. <https://link.springer.com/content/pdf/10.1007/s11222-014-9469-x.pdf>
43. Sherlock C, Golightly A, and Henderson DA, Adaptive, Delayed-Acceptance MCMC for Targets With Expensive Likelihoods, *J. Comput. Graph. Stat*, 26(2):434–444, 2017. <https://arxiv.org/pdf/1509.00172.pdf>
44. Kalb DM, Adikari SH, Hong-Geller E, and Werner JH, Single-cell correlations of mRNA and protein content in a human monocytic cell line after LPS stimulation, *PLOS ONE*, 14(4):1–16, 2019. 10.1371/journal.pone.0215602
45. Munsky B and Khammash M, Transient analysis of stochastic switches and trajectories with applications to gene regulatory networks, *IET Systems Biology*, 2(5):323–333, 9. 2008. [PubMed: 19045827]
46. Gauckler L and Yserentant H, Regularity and approximability of the solutions to the chemical master equation, *ESAIM. Math. Model. Numer. Anal*, 48:1757–1775, 2014. <http://search.proquest.com/openview/600d089885e46d52e2df571119fd91a8/1?pq-origsite=gscholar>
47. Gupta A, Briat C, and Khammash M, A Scalable Computational Framework for Establishing Long-Term Behavior of Stochastic Reaction Networks, *PLoS Comput. Biol*, 10(6), 2014. <http://journals.plos.org/ploscompbiol/article/file?id=10.1371/journal.pcbi.1003669&type=printable>
48. Munsky B and Khammash M, A multiple time interval finite state projection algorithm for the solution to the chemical master equation, *J. Comput. Phys*, 226(1):818–835, 9 2007.
49. Kuntz J, Thomas P, Stan GB, and Barahona M, The exit time finite state projection scheme: Bounding exit distributions and occupation measures of continuous-time markov chains, *SIAM J. Sci. Comput*, 41(2):A748–A769, 2019. 10.1137/18M1168261
50. Femino AM, Fay FS, Fogarty K, and Singer RH, Visualization of single RNA transcripts in situ, *Science (80-.)*, 280(5363):585–590, 4 1998.
51. Jaynes ET, *Probability theory: The logic of science*, Cambridge university press, 2003.
52. Duersch JA and Catanach TA, Generalizing information to the evolution of rational belief, *Entropy*, 22(1):108, 2020.
53. Calderhead B and Girolami M, Estimating bayes factors via thermodynamic integration and population mcmc, *Computational Statistics & Data Analysis*, 53(12):4028–4045, 2009. <http://www.sciencedirect.com/science/article/pii/S0167947309002722>
54. Burrage K, Hegland M, Macnamara S, and Sidje RB, A Krylov-based finite state projection algorithm for solving the chemical master equation arising in the discrete modelling of biological systems, In: Langeville A and Stewart WJ (Eds.), *Proc. 150th Markov Anniv. Meet*, pp. 1–18, Charleston, SC, USA, 2006. Bosen books.
55. Wolf V, Goel R, Mateescu M, and Henzinger TA, Solving the chemical master equation using sliding windows., *BMC Syst. Biol*, 4:42, 2010. [PubMed: 20377904]
56. Sidje RB and Vo HD, Solving the chemical master equation by a fast adaptive finite state projection based on the stochastic simulation algorithm, *Math. Biosci*, 269:10–16, 2015. [PubMed: 26319118]
57. Cao Y, Terebus A, and Liang J, State Space Truncation with Quantified Errors for Accurate Solutions to Discrete Chemical Master Equation, *Bull. Math. Biol*, 78(4):617–661, 2016. <https://link.springer.com/10.1007/s11538-016-0149-1> [PubMed: 27105653]

58. Fox Z, Neuert G, and Munsky B, Finite state projection based bounds to compare chemical master equation models using single-cell data, *J. Chem. Phys.*, 145(7):074101, 2016. 10.1063/1.4960505 [PubMed: 27544081]
59. Vo HD and Munsky B, High performance solver for the chemical master equation, in preparation, 2019.
60. Bader BW and Kolda TG, Efficient matlab computations with sparse and factored tensors, *SIAM Journal on Scientific Computing*, 30(1):205–231, 2008.
61. Devine K, Boman E, Heaphy R, Hendrickson B, and Vaughan C, Zoltan data management services for parallel dynamic applications, *Comput. Sci. Eng.*, 4(2):90–97, 2002.
62. Boman EG, Catalyurek UV, Chevalier C, and Devine KD, The Zoltan and Isorropia parallel toolkits for combinatorial scientific computing: Partitioning, ordering, and coloring, *Sci. Prog.*, 20(2):129–150, 2012.
63. Balay S, Gropp WD, McInnes LC, and Smith BF, Efficient management of parallelism in object oriented numerical software libraries, In: Arge E, Bruaset AM, and Langtangen HP (Eds.), *Modern Software Tools in Scientific Computing*, pp. 163–202. Birkhäuser Press, 1997.
64. Balay S, Abhyankar S, Adams MF, Brown J, Brune P, Buschelman K, Dalcin L, Eijkhout V, Gropp WD, Kaushik D, Knepley MG, McInnes LC, Rupp K, Smith BF, Zampini S, and Zhang H, PETSc users manual, Tech. Rep. ANL-95/11 - Revision 3.6, Argonne National Laboratory, 2015.
65. Balay S, Abhyankar S, Adams MF, Brown J, Brune P, Buschelman K, Dalcin L, Eijkhout V, Gropp WD, Kaushik D, Knepley MG, McInnes LC, Rupp K, Smith BF, Zampini S, and Zhang H PETSc Web page. <http://www.mcs.anl.gov/petsc>, 2015.
66. Koskela A, Approximating the matrix exponential of an advection-diffusion operator using the incomplete orthogonalization method, *Lect. Notes Comput. Sci. Eng.*, 103:345–353, 2015.
67. Vo HD and Sidje RB, Approximating the large sparse matrix exponential using incomplete orthogonalization and Krylov subspaces of variable dimension, *Numer. Linear Algebr. with Appl.*, 24(3), 2017.
68. Gaudreault S, Rainwater G, and Tokman M, Kiops: A fast adaptive krylov subspace solver for exponential integrators, *J. Comput. Phys.*, 372:236–255, 2018. <http://www.sciencedirect.com/science/article/pii/S0021999118304042>
69. Rang J and Angermann L, New rosenbrock w-methods of order 3 for partial differential algebraic equations of index 1, *BIT Numer. Math.*, 45(4):761–787, 12 2005. 10.1007/s10543-005-0035-y
70. Abhyankar S, Brown J, Constantinescu EM, Ghosh D, Smith BF, and Zhang H, PETSc/TS: A Modern Scalable ODE/DAE Solver Library, 6 2018. <http://arxiv.org/abs/1806.01437>
71. Munsky B, Li G, Fox ZR, Shepherd DP, and Neuert G, Distribution shapes govern the discovery of predictive models for gene regulation, *Proc. Natl. Acad. Sci.*, 115(29):7533–7538, 7 2018. <http://www.ncbi.nlm.nih.gov/pubmed/29959206><http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC6055173> [PubMed: 29959206]
72. Kazeev V, Khammash M, Nip M, and Schwab C, Direct Solution of the Chemical Master Equation using Quantized Tensor Trains, *PLoS Comput Biol.*, 10(3), 2014.
73. Kazeev V and Schwab C, Tensor Approximation of Stationary Distributions of Chemical Reaction Networks, *SIAM J. Matrix Anal. Appl.*, 36(3):1221–1247, 2015. <http://epubs.siam.org/doi/10.1137/130927218>
74. Dolgov SV and Savostyanov D, Alternating minimal energy methods for linear systems in higher dimensions, *SIAM J. Sci. Comput.*, 36(5):A2248–A2271, 2014.
75. Dolgov S and Khoromskij B, Simultaneous state-time approximation of the chemical master equation using tensor product formats, *Numerical Linear Algebra with Applications*, 22(2):197–219, 2015. <https://onlinelibrary.wiley.com/doi/abs/10.1002/nla.1942>
76. Vo HD and Sidje RB, An adaptive solution to the chemical master equation using tensors, *J. Chem. Phys.*, 147(147), 2017. 10.1063/1.4994917<http://aip.scitation.org/toc/jcp/147/4>
77. Dolgov Sergey V, A Tensor Decomposition Algorithm for Large ODEs with Conservation Laws, *cmam*, 19(1):23, 2018. <https://www.degruyter.com/view/j/cmam.2019.19.issue-1/cmam-2018-0023/cmam-2018-0023.xml>

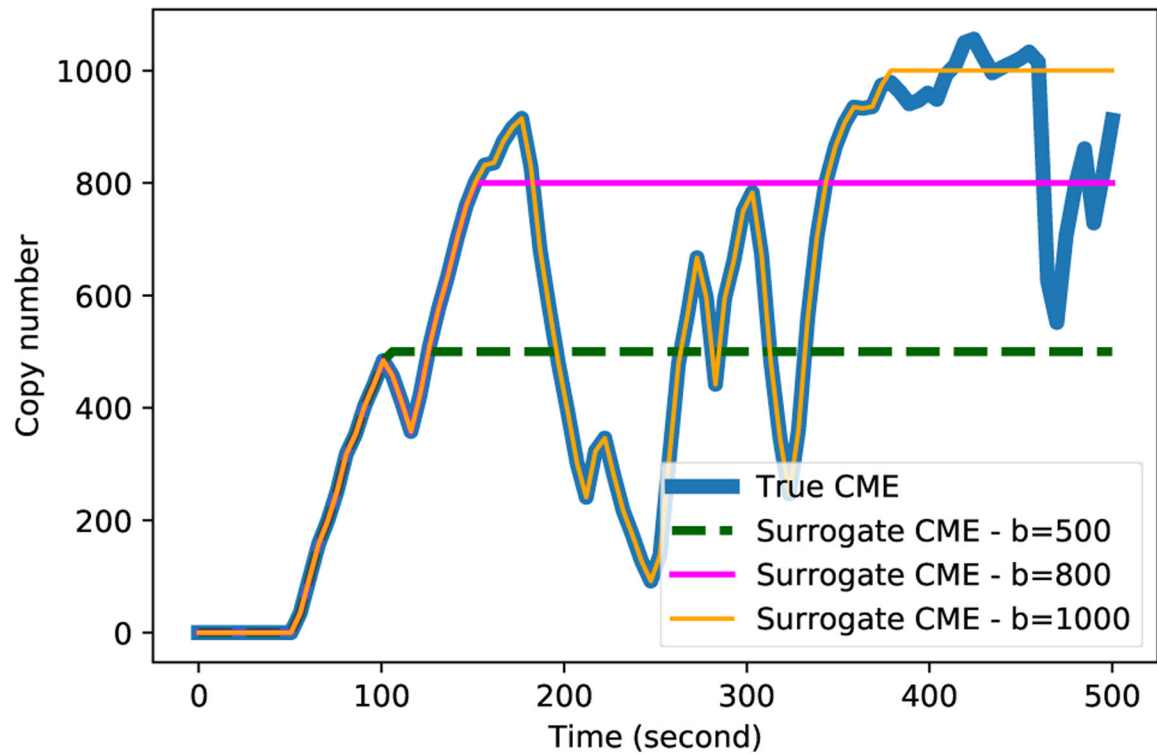
78. Andreychenko A, Mikeev L, Spieler D, and Wolf V, Parameter identification for Markov models of biochemical reactions, In Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics), Vol. 6806 LNCS, pp. 83–98. Springer, Berlin, Heidelberg, 2011.

Author Manuscript

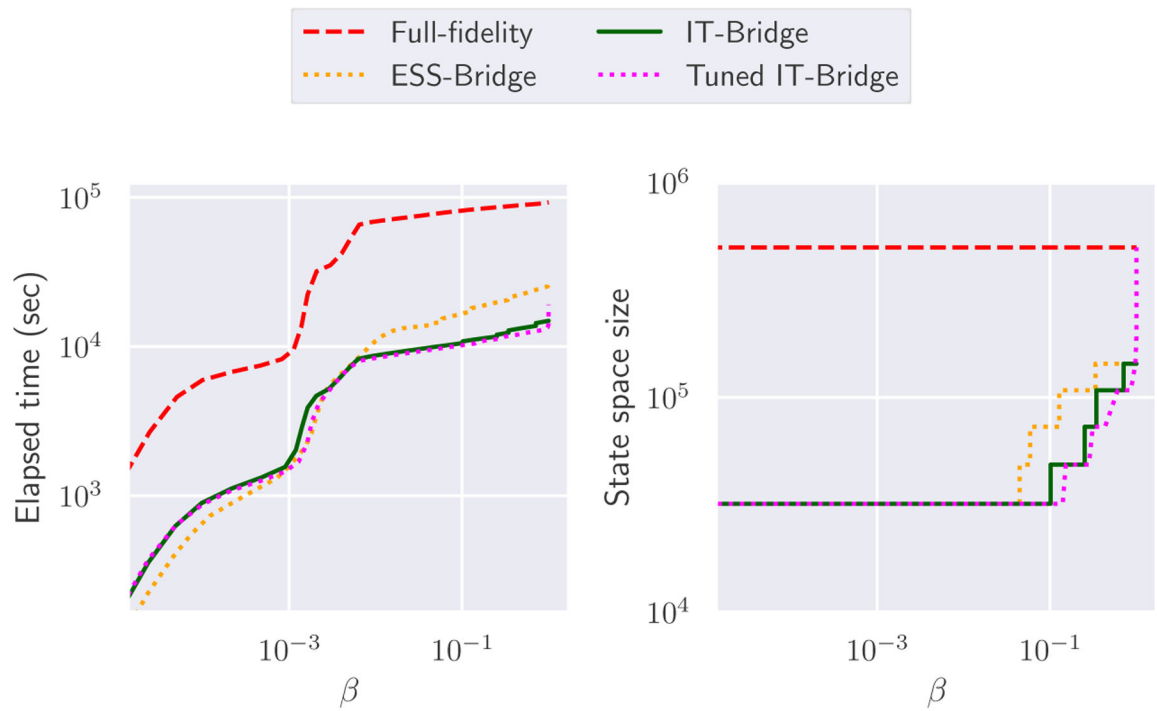
Author Manuscript

Author Manuscript

Author Manuscript

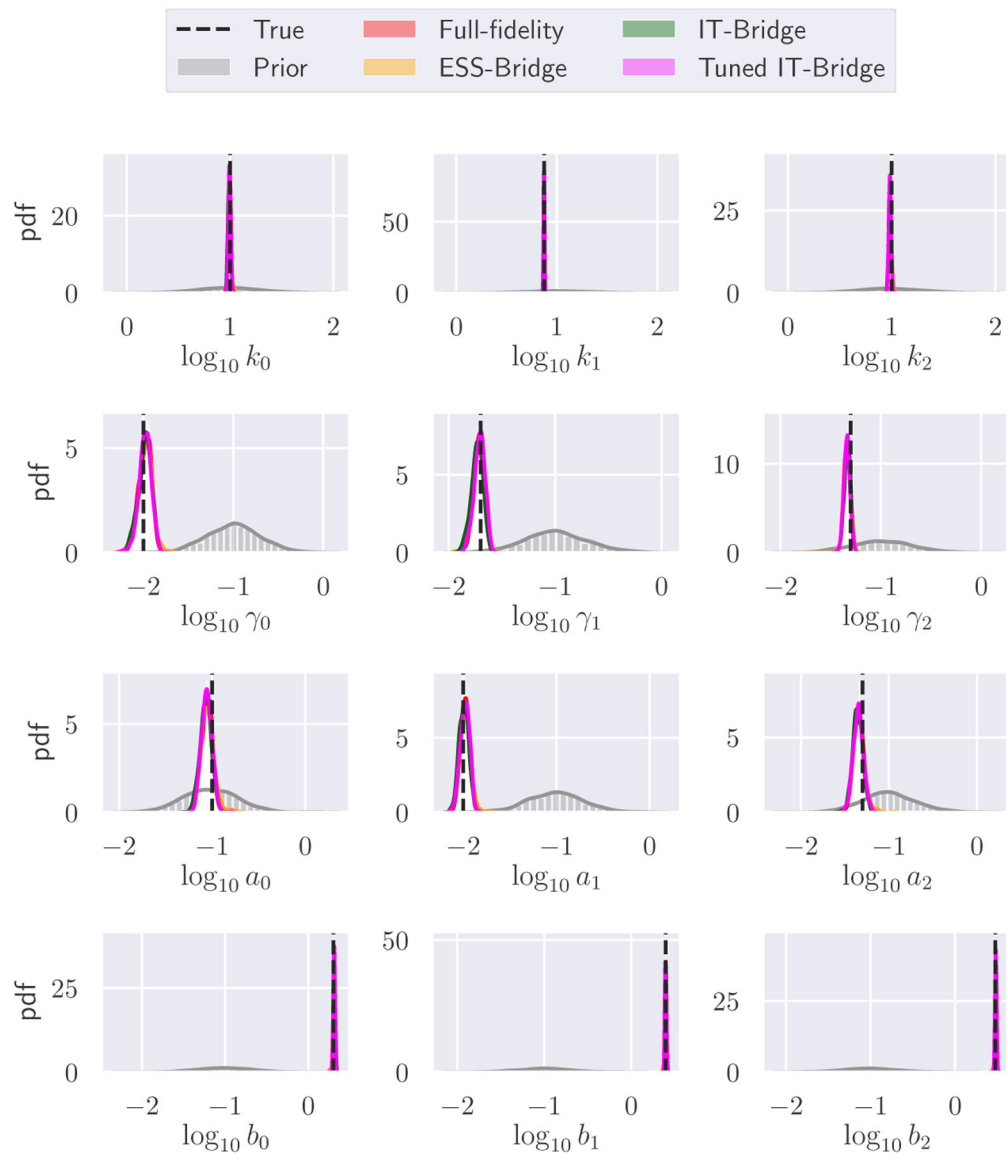
**FIG. 1:**

Illustrative realizations of the full and surrogate CMEs for a simple system with mass-action propensity. Given the same random seed, the simulated trajectory of the surrogate CME will be identical to that of the true CME until the state reaches a threshold,  $b$ , where the surrogate trajectory freezes. Increasing the threshold reduces the chance that the surrogate trajectories hit the bounds and consequently more realizations of the true CME are captured by the surrogate model.

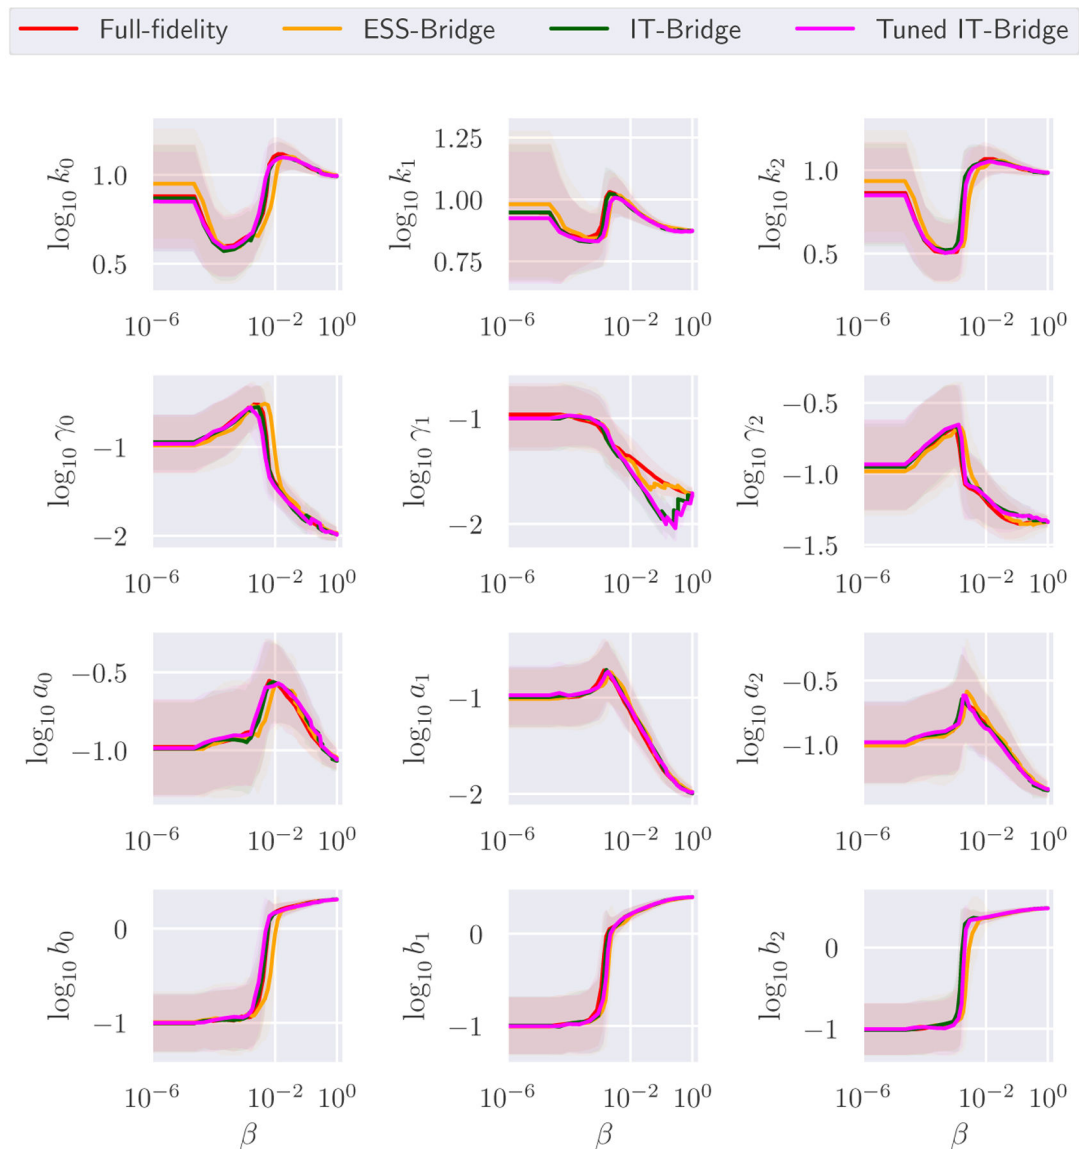


**FIG. 2:** Performance of ST-MCMC samplers on the repressilator example. The horizontal axis represent the annealing factor, i.e. inverse temperature. Significant speed up is observed for the multifidelity schemes.

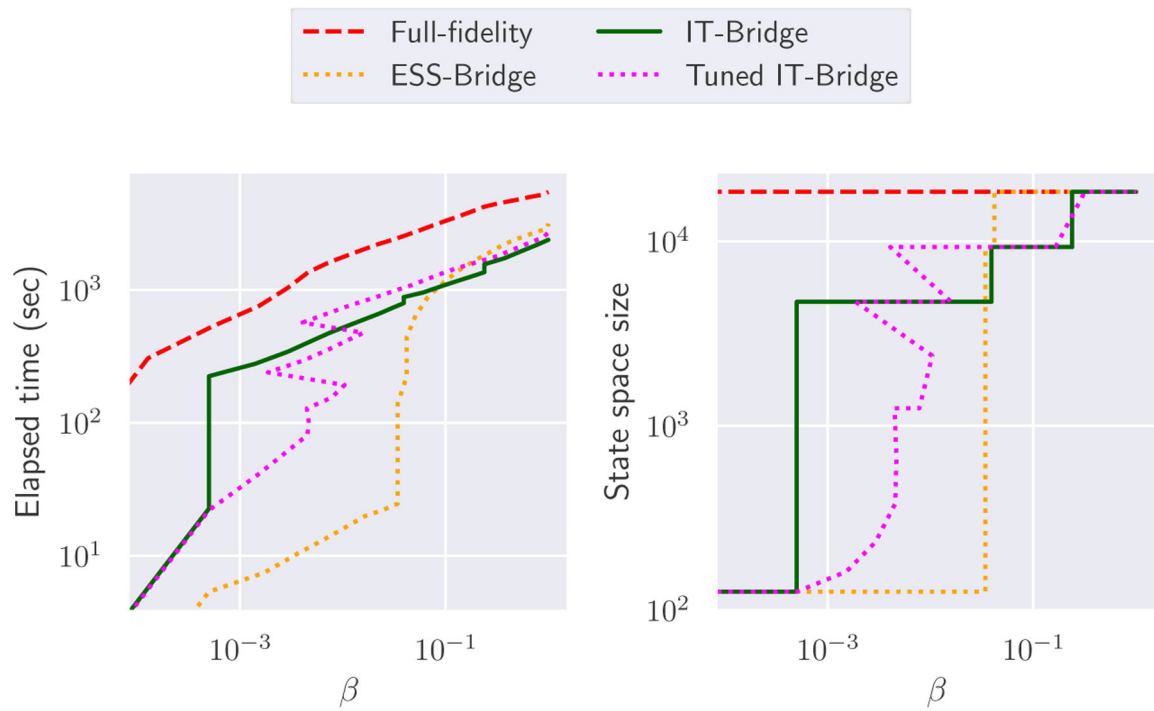


**FIG. 3:**

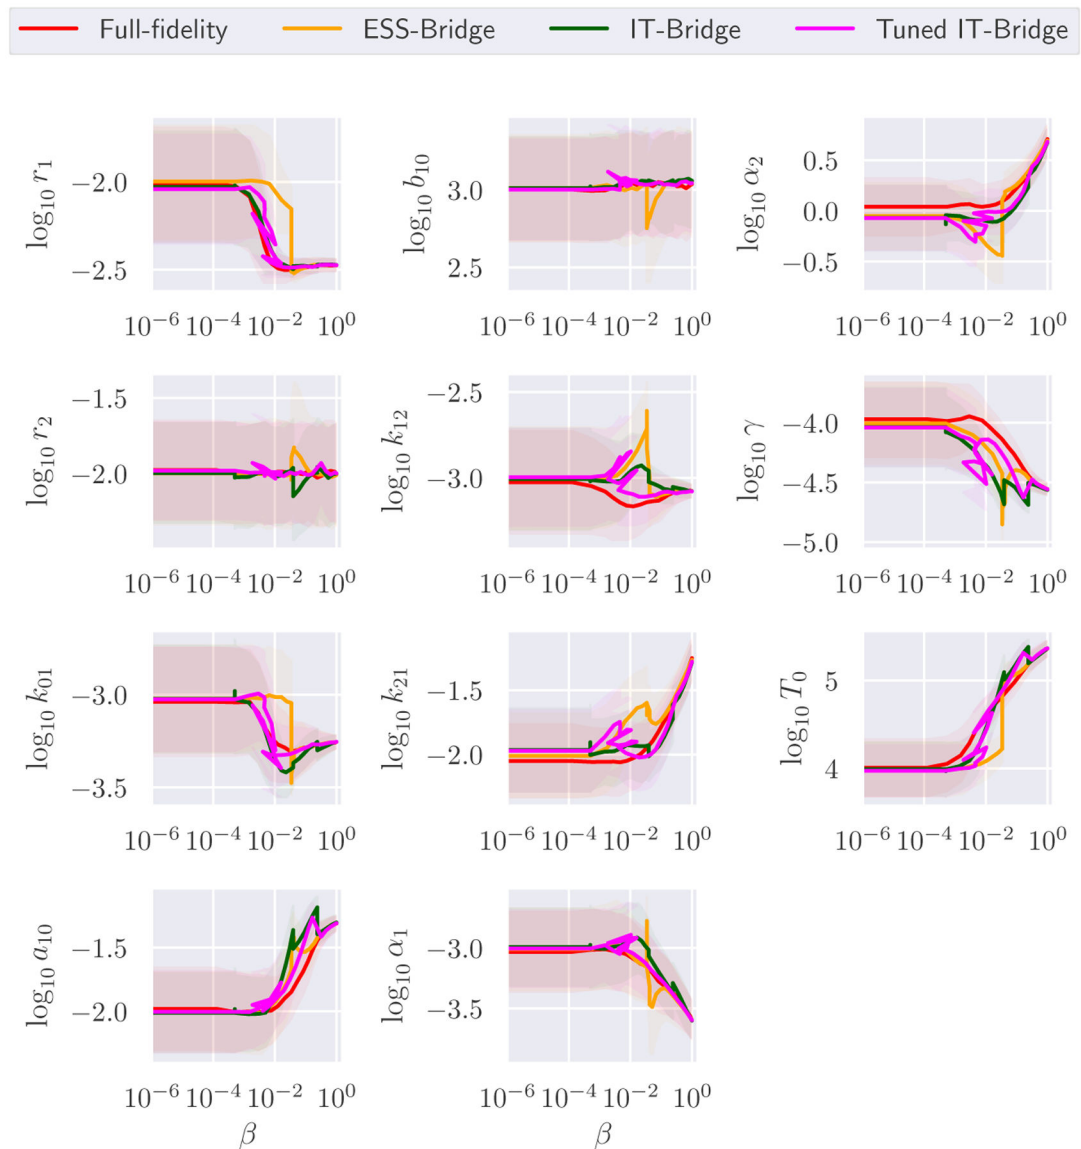
Prior and posterior densities in the repressilator example. See table 2 for the numerical values of the estimated means and standard deviations of these posterior distributions. It is evident that all methods converge to virtually the same distribution.

**FIG. 4:**

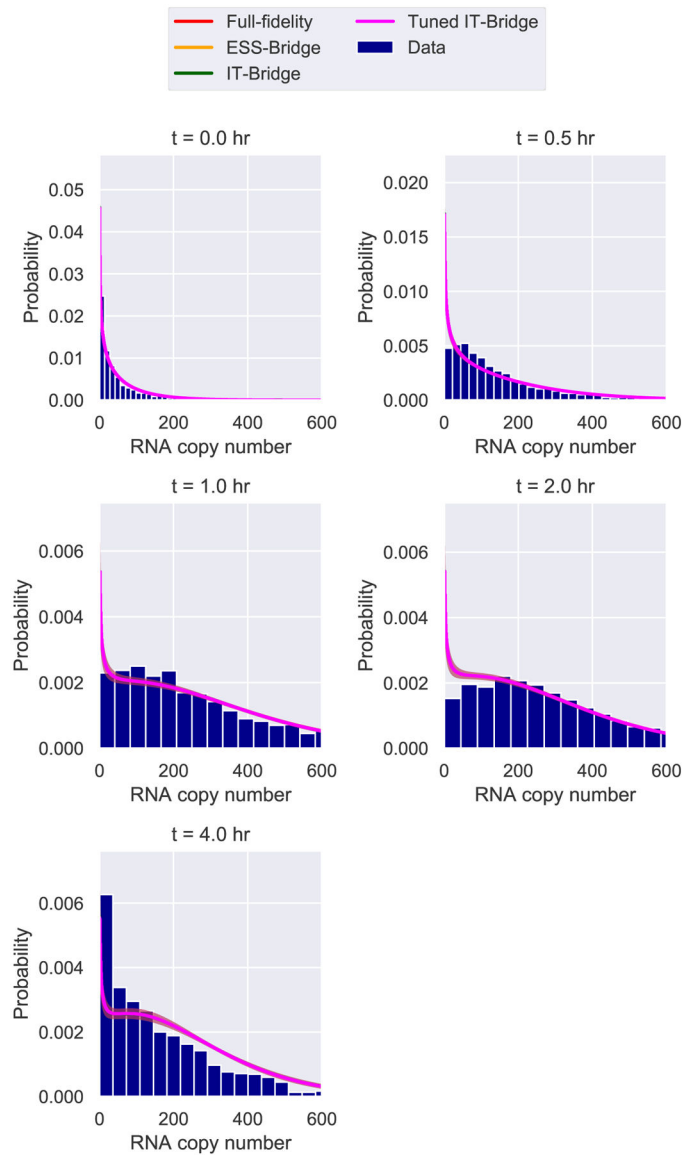
Evolution of the population of samples for the repressilator model parameters using four different ST-MCMC variants: full-fidelity, multifidelity strategies with bridging based on ESS, Information-theoretic Criteria and Tuned Information-theoretic Criteria. The solid lines represent the history of the sample means. The area of the mean  $\pm$  standard deviation is presented in the shaded region. Notice in the  $\gamma_1$  parameter that bias starts to accumulate for the IT-based methods. This bias is corrected when the sampler starts bridging since the bias began to exceed the natural parameter variability.



**FIG. 5:** Performance of ST-MCMC samplers on the IL1beta example. The horizontal axis represents the inverse temperature.



**FIG. 6:** Evolution of the population of samples for the IL1beta model parameters using four different ST-MCMC variants: full-fidelity, multifidelity strategies with bridging based on ESS, Information-theoretic Criteria and Tuned Information-theoretic Criteria. The solid lines represent the history of the sample means. The area of the mean  $\pm$  standard deviation is presented in the shaded region.



**FIG. 7:** Comparison of data and the posterior mRNA distribution predictions for the IL1beta transcription model at zero and four hour after LPS induction. The mean Bayesian prediction for the mRNA probability distribution is computed by averaging the solution of the CME over all posterior samples. The area of one standard deviation around the mean is shown in shade. Visually speaking, samples from different ST-MCMC formulations yield identical predictions.

**TABLE 1:**

Reactions and propensities in the repressilator model. ( $[X]$  is the number of copies of the species X.)

	reaction	propensity
1.	$\emptyset \rightarrow \text{TetR}$	$k_0 / (1 + a_0[\text{LacI}]^{b_0})$
2.	$\text{TetR} \rightarrow \emptyset$	$\gamma_0 [\text{TetR}]$
3.	$\emptyset \rightarrow \lambda\text{cI}$	$k_1 / (1 + a_1[\text{TetR}]^{b_1})$
4.	$\lambda\text{cI} \rightarrow \emptyset$	$\gamma_1 [\lambda\text{cI}]$
5.	$\emptyset \rightarrow \text{LacI}$	$k_2 / (1 + a_2[\lambda\text{cI}]^{b_2})$
6.	$\text{LacI} \rightarrow \emptyset$	$\gamma_2 [\text{LacI}]$

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**TABLE 2:**

Model parameters in the repressilator example. The second column presents the parameters of the prior distribution, where we use a Gaussian prior in the  $\log_{10}$ -transformed parameter space with a diagonal covariance matrix. The last four columns present the posterior mean and standard deviation of model parameters estimated using four methods: fixed-fidelity ST-MCMC (Fixed), Multifidelity ST-MCMC with ESS-Bridging, Multifidelity ST-MCMC with IT-Bridging, and Multifidelity ST-MCMC with  $\beta$ -tuning and IT-Bridging.

Parameter	True	Prior	Posterior			
			Full-fidelity	ESS-Bridge	IT-Bridge	Tuned IT-Bridge
$\log_{10}(k_0)$	1.00	$10.00 \pm 0.3$	$1.00 \pm 0.01$	$1.00 \pm 0.02$	$0.99 \pm 0.01$	$1.00 \pm 0.01$
$\log_{10}(\gamma_0)$	-2.00	$0.10 \pm 0.3$	$-1.98 \pm 0.07$	$-1.98 \pm 0.08$	$-1.98 \pm 0.07$	$-1.98 \pm 0.07$
$\log_{10}(a_0)$	-1.00	$0.10 \pm 0.3$	$-1.05 \pm 0.06$	$-1.05 \pm 0.07$	$-1.07 \pm 0.06$	$-1.06 \pm 0.06$
$\log_{10}(b_0)$	0.30	$0.10 \pm 0.3$	$0.31 \pm 0.01$	$0.31 \pm 0.01$	$0.31 \pm 0.01$	$0.31 \pm 0.01$
$\log_{10}(k_1)$	0.88	$10.00 \pm 0.3$	$0.87 \pm 0.00$	$0.87 \pm 0.01$	$0.87 \pm 0.00$	$0.87 \pm 0.00$
$\log_{10}(\gamma_1)$	-1.70	$0.10 \pm 0.3$	$-1.71 \pm 0.05$	$-1.73 \pm 0.06$	$-1.73 \pm 0.05$	$-1.71 \pm 0.05$
$\log_{10}(a_1)$	-2.00	$0.10 \pm 0.3$	$-1.98 \pm 0.05$	$-1.98 \pm 0.06$	$-1.99 \pm 0.05$	$-1.98 \pm 0.05$
$\log_{10}(b_1)$	0.40	$0.10 \pm 0.3$	$0.40 \pm 0.01$	$0.40 \pm 0.01$	$0.40 \pm 0.01$	$0.40 \pm 0.01$
$\log_{10}(k_2)$	1.00	$10.00 \pm 0.3$	$0.98 \pm 0.01$	$0.99 \pm 0.01$	$0.98 \pm 0.01$	$0.98 \pm 0.01$
$\log_{10}(\gamma_2)$	-1.30	$0.10 \pm 0.3$	$-1.34 \pm 0.03$	$-1.34 \pm 0.04$	$-1.34 \pm 0.03$	$-1.34 \pm 0.03$
$\log_{10}(a_2)$	-1.30	$0.10 \pm 0.3$	$-1.35 \pm 0.06$	$-1.34 \pm 0.07$	$-1.36 \pm 0.06$	$-1.35 \pm 0.06$
$\log_{10}(b_2)$	0.48	$0.10 \pm 0.3$	$0.48 \pm 0.01$	$0.48 \pm 0.01$	$0.48 \pm 0.01$	$0.48 \pm 0.01$

**TABLE 3:**

Reactions and propensities in the compartmental gene expression model.

reaction index	reaction	propensity
$1, \dots, n_G$	$G_{i-1} \rightarrow G_i, i = 1, \dots, n_G - 1$	$k_i^+ [G_{i-1}]$
$n_G + 1, \dots, 2n_G$	$G_i \rightarrow G_{i-1}, i = 1, \dots, n_G - 1$	$k_i^- [G_i]$
$2n_G + 1, \dots, 3n_G - 1$	$G_i \rightarrow G_i + \text{RNA}_{nuc}, i = 1, \dots, n_G - 1$	$r_i [G_i]$
$3n_G$	$\text{RNA}_{nuc} \rightarrow \text{RNA}_{cyt}$	$k_{trans} [\text{RNA}_{nuc}]$
$3n_G + 1$	$\text{RNA}_{cyt} \rightarrow \emptyset$	$\gamma [\text{RNA}_{cyt}]$

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript



**TABLE 4:**

Comparison of the model evidence computation for the 2,3, and 4 state gene expression model using ST-MCMC with the full fidelity and Multifidelity ST-MCMC with  $\beta$ -tuning. The evidence estimates from the full-fidelity and multifidelity methods are consistent, but the multifidelity scheme is significantly faster.

Model	Full-fidelity		Tuned IT Bridge	
	Log Evidence	Time (Sec)	Log Evidence	Time (Sec)
2 Gene	$-20108.8 \pm 5.4$	1244	$-20112.6 \pm 2.0$	758
3 Gene	$-20111.9 \pm 5.6$	67496	$-20115.7 \pm 2.0$	17511
4 Gene	$-20113.0 \pm 5.7$	76546	$-20117.5 \pm 1.8$	23777

**TABLE 5:**

Reactions and propensities in the IL1beta model.

	reaction	propensity
1.	$G_0 \rightarrow G_1$	$k_{01} [G_0]$
2.	$G_1 \rightarrow G_2$	$k_{12} [G_1]$
3.	$G_2 \rightarrow G_1$	$k_{21} [G_2]$
4.	$G_1 \rightarrow G_0$	$k_{10}(t) = \max\{0, a_{10} - b_{10}S(t)\}$ , see eq.(36)
5.	$\emptyset \rightarrow \text{RNA}$	$a_1 [G_1] + a_2 [G_2]$
6.	$\text{RNA} \rightarrow \emptyset$	$\gamma[\text{RNA}]$

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**TABLE 6:**

Model parameters in the IL1beta example. The second column presents the parameters of the prior distribution, where we use a Gaussian prior in the  $\log_{10}$ -transformed parameter space with a diagonal covariance matrix. The last four columns present the posterior mean and standard deviation of model parameters estimated using the ST-MCMC with full-fidelity model and the Multifidelity ST-MCMC with three different bridging strategies.

Parameter	Prior	Posterior			
		Full-fidelity	ESS-Bridge	IT-Bridge	Tuned IT-Bridge
$\log_{10}(r_1)$	$-2.00 \pm 0.33$	$-2.48 \pm 0.03$	$-2.47 \pm 0.03$	$-2.47 \pm 0.03$	$-2.47 \pm 0.03$
$\log_{10}(r_2)$	$-2.00 \pm 0.33$	$-2.00 \pm 0.34$	$-2.00 \pm 0.33$	$-2.00 \pm 0.33$	$-2.00 \pm 0.32$
$\log_{10}(k_{01})$	$-3.00 \pm 0.33$	$-3.26 \pm 0.03$	$-3.25 \pm 0.04$	$-3.25 \pm 0.03$	$-3.25 \pm 0.03$
$\log_{10}(a_{10})$	$-2.00 \pm 0.33$	$-1.31 \pm 0.06$	$-1.31 \pm 0.06$	$-1.30 \pm 0.06$	$-1.31 \pm 0.06$
$\log_{10}(k_{10})$	$3.00 \pm 0.33$	$3.04 \pm 0.34$	$3.05 \pm 0.34$	$3.06 \pm 0.30$	$3.04 \pm 0.31$
$\log_{10}(k_{12})$	$-3.00 \pm 0.33$	$-3.08 \pm 0.03$	$-3.08 \pm 0.04$	$-3.08 \pm 0.04$	$-3.08 \pm 0.03$
$\log_{10}(k_{21})$	$-2.00 \pm 0.33$	$-1.25 \pm 0.15$	$-1.26 \pm 0.14$	$-1.28 \pm 0.14$	$-1.28 \pm 0.13$
$\log_{10}(a_1)$	$-3.00 \pm 0.33$	$-3.60 \pm 0.15$	$-3.60 \pm 0.16$	$-3.60 \pm 0.16$	$-3.60 \pm 0.16$
$\log_{10}(a_2)$	$0.00 \pm 0.33$	$0.71 \pm 0.15$	$0.70 \pm 0.14$	$0.68 \pm 0.13$	$0.68 \pm 0.13$
$\log_{10}(\gamma)$	$-4.00 \pm 0.33$	$-4.56 \pm 0.05$	$-4.56 \pm 0.05$	$-4.56 \pm 0.05$	$-4.56 \pm 0.05$
$\log_{10}(T_0)$	$4.00 \pm 0.33$	$5.36 \pm 0.09$	$5.36 \pm 0.09$	$5.37 \pm 0.08$	$5.37 \pm 0.09$