OPEN ACCESS Freely available online

PLOS | COMPUTATIONAL BIOLOGY

Uncoupled Analysis of Stochastic Reaction Networks in Fluctuating Environments

Christoph Zechner[1¤], **Heinz Koeppl**[2*]

1 Department of Information Technology and Electrical Engineering, ETH Zurich, Zurich, Switzerland, **2** Department of Electrical Engineering and Information Technology, Technische Universität Darmstadt, Darmstadt, Germany

Abstract

The dynamics of stochastic reaction networks within cells are inevitably modulated by factors considered extrinsic to the network such as, for instance, the fluctuations in ribosome copy numbers for a gene regulatory network. While several recent studies demonstrate the importance of accounting for such extrinsic components, the resulting models are typically hard to analyze. In this work we develop a general mathematical framework that allows to uncouple the network from its dynamic environment by incorporating only the environment's effect onto the network into a new model. More technically, we show how such fluctuating extrinsic components (e.g., chemical species) can be marginalized in order to obtain this decoupled model. We derive its corresponding process- and master equations and show how stochastic simulations can be performed. Using several case studies, we demonstrate the significance of the approach.

Citation: Zechner C, Koeppl H (2014) Uncoupled Analysis of Stochastic Reaction Networks in Fluctuating Environments. PLoS Comput Biol 10(12): e1003942. doi:10.1371/journal.pcbi.1003942

Editor: Bard Ermentrout, University of Pittsburgh, United States of America

Received May 20, 2014; **Accepted** September 28, 2014; **Published** December 4, 2014

Data Availability: The authors confirm that all data underlying the findings are fully available without restriction. All relevant data are within the paper and its Supporting Information files.

Funding: CZ and HK acknowledge the support from the Swiss National Science Foundation (www.snf.ch), grant no. PP00P2 128503. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

* Email: heinz.koeppl@bcs.tu-darmstadt.de

¤ Current address: Department of Biosystems Science and Engineering, ETH Zurich, Zurich, Switzerland

Introduction

Biochemical systems involving low-copy molecules demand for mathematical models that account for the intrinsic stochasticity [1]. In recent years, however, realization has grown that intrinsic noise alone cannot account for the observed substantial phenotypic variability among isogenic cells. That is, fluctuations in the intracellular environment, commonly termed extrinsic noise, represent an additional source of variability [2,3,4].

Several recent studies focus on separating intrinsic and extrinsic fluctuations through dual-reporter measurements [5,6]. Other approaches model extrinsic noise through certain parameters (i.e., the translation rate) of a kinetic model which is calibrated subsequently using flow-cytometry [7,8] or time-lapse microscopy data [9,10]. All of those approaches have in common that they consider the biochemical process under study – i.e., the expression of a gene – as a small subpart that is embedded into a larger dynamical system. Accordingly, they rely on augmenting the original kinetic model by certain environmental components which are assumed to be fixed but random [11,10,8,9] or fluctuating over time [5,12,13]. In fact, such models agree very well with the variability that is observed experimentally, but on their downside, suffer from the increased dimensionality - somehow defeating the original purpose of tractable dedicated models.

A natural question arising in that context is whether we can find a proper dynamical description of just the system of interest as if it was still embedded into its stochastically modulating environment. In other words, we aim to find a "self-contained" stochastic model that summarizes all system behaviors attainable under all possible realizations of the extrinsic fluctuations. Such models could then be used to perform an uncoupled analysis of a reaction network subject to extrinsic noise. The mathematical correct answer that we provide in this work is the marginalization of the system dynamics with respect to those extrinsic fluctuations. Interestingly it turns out that the resulting model exploits its own stochasticity to emulate the effect of extrinsic noise, leading to a self-exciting process. A simple instance of such self-excitation is the Polya urn scheme: at each draw from the Polya urn with balls of two colors the drawn ball and a fixed number of new balls of the same color as the draw are placed in the urn. This scheme is known to be equivalent to Bernoulli trials marginalized over random (and here correspondingly extrinsic) success rates [14]. Intuitively, due to its self-excitation, the number of draws of the same color for a Polya urn over repeated draws displays a much richer and dispersed dynamics than the number of successful draws in a Bernoulli trial with fixed success rate.

For the purpose of inference we recently proposed a first attempt of such marginalization for the special case of fixed but random environmental conditions [9]. In this work we develop a general mathematical framework from which the uncoupled dynamics can be constructed in a principled manner, regardless whether the environment is constant or dynamically changing.

Author Summary

Cellular reaction networks show a substantial degree of complexity, and many of the regulatory mechanisms are yet to be discovered. Therefore, building detailed mathematical models of complete networks is not only unrewarding but most often impossible. Instead, focus is put on small subnetworks such as the transcriptional circuitry associated with a few genes. Practically, those subnetworks are still embedded in the cellular environment and are consequently modulated by it. The mathematical framework that we present in this work allows to build models of such subnetworks only, while still incorporating the impact of the environment. Such models are instrumental not only for advancing our biological understanding through them, but also for the reliable forward engineering of new subnetworks in synthetic biology.

Results

Mathematical framework

We describe the time-evolution of a stochastic reaction network by a continous-time Markov chain (CTMC) X with M chemical species and N reaction channels. The system state at time t is denoted $X(t)$ and we write its random path on time intervals $[0,a]$ as $\mathbf{X}_a$. Throughout we will follow the usual convention to refer to upper-case and lower-case versions of a symbol as a random variable and its realization, respectively. Furthermore, we assume that X depends on another multivariate Markov process Z through its hazard functions in the form

$$h_i(x,z) = c_i(z)g_i(x), \tag{1}$$

with c_i some positive function and g_i a polynomial determined by the law of mass-action, for instance. For reactions independent of Z, we thus have $c_i(z) \equiv c_i$. Typically, Z is another jump or diffusion process corresponding to a set of modulating *environmental* species or conditions that are considered extrinsic to the system of interest, whereas the species in X represent the actual system of interest. For example, Z could be the fluctuating ribosome copy numbers affecting the kinetics of a gene regulatory network represented by X. Although a more general treatment is possible, we assume a feed-forward structure between Z and X, which means that Z modulates X but not vice-versa. Consequently, the dynamics of the joint system $Y(t) = (Z(t), X(t))$ can be described by a Markov process Z together with a conditional Markov chain $X|Z$.

Uncoupled dynamics. Mathematical descriptions of the joint system $Y(t)$ are readily obtained using available techniques for modeling Markovian dynamics [5,15,12]. For complexity reasons, however, we aim for models that can properly describe *only* the interesting components $X(t)$. In order to see that marginalization over Z yields the desired model, let us first consider two dependent random variables A and B described by a joint probability distribution $p(a,b) = p(a|b)p(b)$. If we are interested in analyzing A under all possible values of B, we need to average the probability at $A = a$ over all possible values of B, i.e.,

$$p(a) = \int p(a,b)\mathrm{d}b = \mathbb{E}[p(a|B)].$$

Note that as a consequence of averaging probabilities, any value of A possible under the joint distribution $p(a,b)$ is possible under the marginal $p(a)$, while this does not necessarily apply to $p(a|b)$ for any choice of B. The resulting *marginal* distribution $p(a)$ is an exact mathematical description of A only, or in other words, it allows to analyze A uncoupled from B.

In case of the coupled processes $Z(t)$ and $X(t)$, we analogously *marginalize* the joint Markov chain Y with respect to the environmental process Z. While such a marginalization involves several difficulties, the idea remains the same: we try to construct an uncoupled process X which directly admits the marginal path distribution $p(\mathbf{x}_t) = \mathbb{E}[p(\mathbf{x}_t|\mathbf{Z}_t)]$. As a result, we obtain a jump process, which - in contrast to the conditional process $X|Z$ - no longer depends on the environment Z. We remark that a straightforward marginalization of the joint master equation of Z and X generally leads to intractable propensities [16,5]. Based on the innovation theorem [17] we demonstrate in section S.1 in S1 Text that the hazard functions of the uncoupled process (later referred to as the *marginal* hazard functions) can be generally written as

$$\lambda_i(\mathbf{X}_t) = \mathbb{E}[c_i(Z(t))|\mathbf{X}_t]g_i(X(t)), \tag{2}$$

where the expectation is taken with respect to the conditional distribution $\pi(z,t|\mathbf{X}_t)$. The latter describes the conditional probability of the environmental process $Z(t)$ given the entire history (or filtration) of process X until time t. Using the expected value of that distribution, the feed-forward influence of Z on the hazard functions of X can be replaced by a deterministic function of X, which no longer depends on the actual state of Z. Instead, the uncoupled process X becomes *self-exciting*, meaning that it exerts a feedback on itself. Hence, given that we can evaluate Eq.2, we have a means to simulate X while bypassing the need to draw realizations of Z. This has for instance been exploited for the exact simulation of diffusion-driven Poisson processes [18]. Note that the uncoupled process X is no longer Markovian, since the conditional expectation - and hence the hazard function - possibly depends on

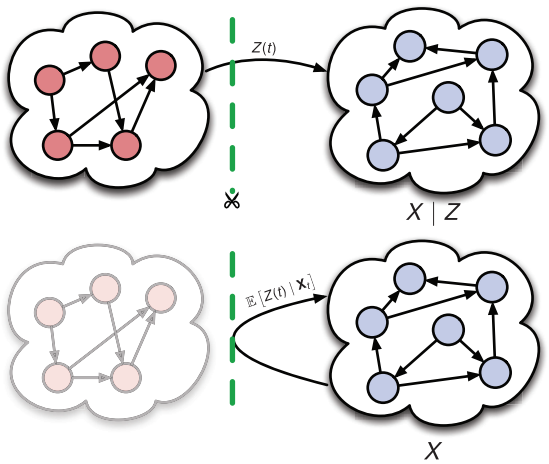

Figure 1. Uncoupled stochastic dynamics. The environmental process Z modulates the dynamics of the process under study X, e.g., through one of its hazard functions. Marginalization with respect to Z yields the uncoupled dynamics of X, whereas the original dependency on the environment Z is replaced by its optimal estimator given the history of X. Consequently, the marginal process X is self-exciting, i.e., it exerts a feedback on itself.
doi:10.1371/journal.pcbi.1003942.g001

the full process history $\mathbf{X}_t$. A schematic illustration of that uncoupling is given in Fig. 1.

Associated filtering problem. Although the construction of the uncoupled dynamics is general, any practical implementation thereof will depend on an explicit computation of the conditional expectation in Eq. 2. This expectation estimates the environmental state $Z(t)$ given the full history of the uncoupled process $\mathbf{X}_t$ and therefore, can be understood as the solution to a *stochastic filtering* problem [19]. Filtering techniques deal with the problem of optimally reconstructing a hidden stochastic process at time t from noisy observations of that process up to time t. In the situation considered here, the hidden process corresponds to the environment $Z(t)$, which gets reconstructed from the "observed" history $\mathbf{X}_t$ through the conditional mean in Eq. 2.

We assume that the environment $Z(t)$ admits a probability distribution $p(z,t)$ described by a Kolmogorov-forward equation of the form

$$\frac{\partial}{\partial t}p(z,t)=\mathcal{A}p(z,t), \tag{3}$$

where $\mathcal{A}$ represents the temporal change of $p(z,t)$, i.e., is the infinitesimal generator of Z. For instance, if Z is a diffusion process, $\mathcal{A}$ corresponds to the Fokker-Planck operator, while in case of a CTMC, $\mathcal{A}$ is given by the difference operator of the chemical master equation (CME). In terms of filtering, Eq. 3 corresponds to the process model of Z. Furthermore, we know that at a given time t, the solution of X can be written as a sum of independent but time-transformed Poisson processes [20], each of them corresponding to a particular reaction channel. Consequently, the observation model is given by a set of Poisson counting observations with the hazard functions given in Eq.1. This is closely related to Markov-modulated Poisson processes [21] and their corresponding optimal filtering [22].

While a more general treatment is provided in S1 Text, we assume in the following that a one-dimensional process Z is modulating X through its k-th reaction of order zero. We further restrict ourselves to the case where c_k is a linear function of z, i.e., $c_k(z)=c_k z$. Under those assumptions, it can be shown that the conditional process $Z(t)|\mathbf{X}_t$ follows a filtering distribution $\pi(z,t|\mathbf{X}_t)$ with

$$\begin{aligned}&\mathrm{d}\pi(z,t)=\\&[\mathcal{A}\pi(z,t)-c_k(z-M_1(t))\pi(z,t)]\mathrm{d}t+\left[\frac{z-M_1(t)}{M_1(t)}\pi(z,t)\right]\mathrm{d}R_k(t),\end{aligned} \tag{4}$$

with $R_k(t)$ as the number of reactions of type k up to time t in $\mathbf{X}_t$ and $M_1(t)=\mathbb{E}[Z(t)|\mathbf{X}_t]$ (for a derivation see S.2 in S1 Text). Since Eq. 4 shows an implicit dependency on its own mean, it is complicated to handle numerically. A simpler equation can be obtained for an unnormalized variant $\tilde{\pi}(z,t)$ of the filtering distribution [23], i.e.,

$$\mathrm{d}\tilde{\pi}(z,t)=[\mathcal{A}\tilde{\pi}(z,t)-c_k z\tilde{\pi}(z,t)]\mathrm{d}t+[z-1]\tilde{\pi}(z,t)\mathrm{d}R_k(t), \tag{5}$$

with $\pi(z,t|\mathbf{X}_t)\equiv\xi(t)\tilde{\pi}(z,t)$ and $\xi(t)$ a time-dependent normalizing factor independent of z. Thus, Eq.5 describes a scaled version of the normalized filtering distribution from Eq.4. However, once $\tilde{\pi}$ has been numerically solved for, it can be easily rescaled such that it integrates (or sums up) to one for all t.

Note that Eq. 5 is a stochastic partial differential equation (SPDE) in case Z describes a diffusion process or a stochastic difference-differential equation (SDDE) if Z is a CTMC. In the latter case, the solution of Eq.5 can be compactly written as

$$\tilde{\mathbf{\Pi}}(t)=e^{\left(Q-c_k\Lambda\right)\left(t-t_{R_k(t)}\right)}\left(\prod_{l=1}^{R_k(t)}\Lambda e^{(Q-c_k\Lambda)(t_l-t_{l-1})}\right)\Pi_0 \tag{6}$$

with $\tilde{\mathbf{\Pi}}(t)=(\tilde{\pi}(0,t),\ldots,\tilde{\pi}(L-1,t))^T$, L the number of reachable states of Z, $\Lambda=\mathrm{diag}(0,\ldots,L-1)$, $\Pi_0\in\mathbb{R}^L$ the initial distribution over Z and $Q\in\mathbb{R}^{L\times L}$ the generator matrix of Z. Note that we define Q to be a left stochastic matrix, i.e., its rows sum up to zero.

Conditional moment dynamics. In order to evaluate Eq. 2, we only require the mean (i.e., the first moment) of the filtering distribution, i.e., $M_1(t)$. In general, however, the mean also depends on the second-order moment, which in turn depends on the third-order moment and so forth. Generally, the i-th order non-central moment is found by multiplying both sides of Eq.4 with z^i and summing (or integrating) over all $z\in\mathcal{Z}$, i.e.,

$$\begin{aligned}\sum_{z\in\mathcal{Z}}z^i\mathrm{d}\pi(z,t)=&\sum_{z\in\mathcal{Z}}z^i(\mathcal{A}\pi(z,t)-c_k[z-M_1(t)]\pi(z,t))\mathrm{d}t\\&+\sum_{z\in\mathcal{Z}}z^i\left[\frac{z-M_1(t)}{M_1(t)}\right]\pi(z,t)\mathrm{d}R_k(t)\end{aligned} \tag{7}$$

such that the filtering moment dynamics up to order i can be generally written as

$$\begin{aligned}&\mathrm{d}M_1(t)=\\&[\mathcal{D}_1(t)-c_k(M_2(t)-M_1(t)M_1(t))]\mathrm{d}t+\tfrac{M_2(t)-M_1(t)M_1(t)}{M_1(t)}\mathrm{d}R_k(t)\\&\vdots\\&\mathrm{d}M_i(t)=\\&[\mathcal{D}_i(t)-c_k(M_{i+1}(t)-M_1(t)M_i(t))]\mathrm{d}t+\tfrac{M_{i+1}(t)-M_1(t)M_i(t)}{M_1(t)}\mathrm{d}R_k(t),\end{aligned} \tag{8}$$

with $\mathcal{D}_i(t)=\sum_{z\in\mathcal{Z}}z^i\mathcal{A}\pi(z,t)$ as the unconditional moment dynamics. The computation of moments in case of multivariate environments is performed analogously. Although Eq.8 is generally infinite-dimensional, there are several relevant scenarios, for which the moment dynamics are *closed*, i.e., only depend on higher-order moments up to a certain order. This is for instance the case, if $Z(t)$ is a Cox-Ingersoll-Ross (CIR) process [24] or any finite state Markov chain. On the other hand, if the moment dynamics are infinite-dimensional, suitable assumptions on the filtering distribution π can be imposed to yield a closed moment-dynamics (see [25] and S.3 in S1 Text). An important closure is found by analyzing Eq. 6: especially for large c_k we have that $e^{(Q-c_k\Lambda)(t_l-t_{l-1})}\approx e^{-c_k\Lambda(t_l-t_{l-1})}$ and furthermore,

$$\tilde{\mathbf{\Pi}}(t)\approx\left(\prod_{l=1}^{R_k(t)}\Lambda e^{-c_k\Lambda\left(t_l-t_{l-1}\right)}\right)\Pi_0=\Lambda^{R_k(t)}e^{-c_k\Lambda t}\Pi_0 \tag{9}$$

suggesting that Eq. 6 can be well approximated by a Gamma-distribution. We note that the Gamma-distribution is fully characterized by two parameters – or equivalently – its first two moments $M_1(t)$ and $M_2(t)$. As a consequence, we may express the third order moment as a function of the first two moments, i.e., $M_3(t)=-M_1(t)M_2(t)+2M_2^2(t)/M_1(t)$, such that the second conditional moment closes as

$$\mathrm{d}M_2(t)=$$

$$\left[\mathcal{D}_2(t)-2c_k\frac{M_2(t)}{M_1(t)}\left(M_2(t)-M_1^2(t)\right)\right]\mathrm{d}t+2\left[\frac{M_2^2(t)}{M_1^2(t)}-M_2(t)\right]\mathrm{d}R_k(t).$$

Further discussion on moment-closure is provided in section S.3 in S1 Text. In the following, we demonstrate the uncoupled dynamics using several numerical and analytical case studies.

Uncoupled dynamics for network simulation

The marginal simulation algorithm (MSA). Although the uncoupled dynamics of X are non-Markovian, the Markov property can be enforced by virtually extending the state space by the first moment of Eq. 8, summarizing the history of X. As a result, one can simulate sample paths of the uncoupled process using standard methods that can account for the explicit time-dependency of the hazard functions [26]. In general, such algorithms rely on the generation of random waiting times for each of the reaction channels. All reactions that are independent of $Z(t)$ will retain their exponentially distributed waiting times. In contrast, the time τ_k that passes until a reaction of type k happens is distributed according to

$$P_k(\tau_k<s|\mathbf{X}_t)=1-e^{-c_k\int_0^s M_1(t+T)\mathrm{d}T}. \tag{10}$$

We note that as long as no reaction of type k happens, $\mathrm{d}R_k(t)$ is zero and hence, $M_1(t)$ is found by solving a set ordinary differential equations (ODEs). Since that solution is not generally known in closed form, we cannot directly sample from Eq. 10. However, several efficient solutions to that problem have been developed in the context of inhomogeneous Poisson processes, e.g., such as the method of *thinning* [27]. Once a reaction has fired, the filtering moments need to be updated by the terms multiplying the firing process $\mathrm{d}R_k(t)$ in Eq. 8 (i.e., they exhibit a discontinuity). The following lines describe a possible implementation of the MSA.

Algorithm 1 (Marginal simulation algorithm) *The algorithm simulates the uncoupled dynamics of a reaction network with N reactions associated with stoichiometric change vectors* $v_1,\ldots,v_N$. *The k-th reaction is assumed to be driven by the environmental network Z. The algorithm requires a real-valued constant* $\hat{\lambda}\geq c_k M_1(t)$ *to be given as an input.*

1 Initialize variables $t\leftarrow 0$ *and* $x\leftarrow x_0$.
while $t<T$ ***do***
for $i=1,\ldots,N$ ***do***
  if $i=k$ *(marginal hazard)* ***then***
    Initialize $\tau_i\leftarrow 0$ *and* $u\leftarrow\infty$.
    while $u>c_k M_1(t+\tau_i)/\hat{\lambda}$ ***do***
    Simulate $\hat{\tau}\sim\mathrm{Exp}(\hat{\lambda})$.
    Set $\tau_i\leftarrow\tau_i+\hat{\tau}$.
    Simulate $u\sim\mathcal{U}(0,1)$.
    end while
  else
    Simulate exponential waiting-time $\tau_i\sim\mathrm{Exp}(c_i g_i(x))$.
  end if
end for
Choose reaction associated with the minimal waiting-time $j\leftarrow\underset{i=1,\ldots,N}{\mathrm{argmax}}\,\tau_i$.
Update state $x\leftarrow x+v_j$.
Update time $t\leftarrow t+\tau_j$.
if $j=k$ ***then***
  Update $M_1(t)$ *by the terms accompanying* $\mathrm{d}R_k(t)$ *in Eq. 8.*
end if
Output t and x.
end while

Evidently, simulation from Eq. 10 (i.e., through the thinning method in lines 5-10 of Algorithm 1) comes at higher cost than simulating from an exponential distribution, since in general, it relies on the numerical integration of an ODE. However, reactions associated with the environmental part no longer need to be simulated, which yields a significant reduction in computational effort as soon as the environmental network is large and expensive to simulate due to high propensity reactions, for instance.

An illustrative example. We will now instantiate the proposed framework using a simple telegraph model of transcription, i.e.,

$$Z_0 \underset{c_2}{\overset{c_1}{\rightleftharpoons}} Z_1$$

$$Z_1 \xrightarrow{c_3} Z_1+X$$

with propensities given by the law of mass action. The two-state promoter Z stochastically switches between its inactive state Z_0 and its active state Z_1, whereas only the latter allows transcription of mRNA (denoted by X). For the sake of illustration, this model is now understood as a process X which is driven by an environmental perturbation Z. We thus aim to find the uncoupled dynamics of X, where fluctuations in Z have been marginalized.

In order to evaluate the marginal hazard function, we need to compute the conditional expectation $M_1(t)=\pi(0,t)0+\pi(1,t)1=\pi(1,t)$. It is straightforward to show that the unconditional probability distribution over the promoter states satisfies the CME

$$\tfrac{\mathrm{d}}{\mathrm{d}t}P(t)=QP(t)=\begin{pmatrix}-c_1 & c_2\\ c_1 & -c_2\end{pmatrix}P(t), \tag{11}$$

with $P(t)=[P(0,t),P(1,t)]^T$. From equation Eq. 5 we further know that the unnormalized conditional probability $\tilde{\mathbf{\Pi}}(t)=[\tilde{\pi}(0,t),\tilde{\pi}(1,t)]^T$ admits the linear stochastic differential equation

$$\begin{aligned}\mathrm{d}\tilde{\mathbf{\Pi}}(t)&=(Q-c_3\Lambda)\tilde{\mathbf{\Pi}}(t)\mathrm{d}t+(\Lambda-I)\tilde{\mathbf{\Pi}}(t)\mathrm{d}R_3(t)\\&=\begin{pmatrix}-c_1 & c_2\\ c_1 & -(c_2+c_3)\end{pmatrix}\tilde{\mathbf{\Pi}}(t)\mathrm{d}t+\begin{pmatrix}-1 & 0\\ 0 & 0\end{pmatrix}\tilde{\mathbf{\Pi}}(t)\mathrm{d}R_3(t).\end{aligned} \tag{12}$$

As long as no transcription event takes place (i.e., $\mathrm{d}R_3(t)=0$), the conditional distribution from Eq.12 is just given by a linear homogeneous ODE, whose solution is known in closed form. We refer to this distribution as $\tilde{\mathbf{\Pi}}^-(t)$. In case a reaction happens at time t, the distribution will be updated according to

$$\tilde{\mathbf{\Pi}}^+(t)=\tilde{\mathbf{\Pi}}^-(t)+\begin{pmatrix}-1 & 0\\ 0 & 0\end{pmatrix}\tilde{\mathbf{\Pi}}^-(t)=\begin{pmatrix}0\\ \tilde{\pi}^-(0,t)\end{pmatrix}. \tag{13}$$

Upon normalization of the distribution $\tilde{\mathbf{\Pi}}^-(t)$ we may express its expectation between two consecutive firing times t_1 and t_0 with

$t_1 > t_0$ as

$$M_1^-(t) = \frac{(c_1 - c_2 - c_3)\sinh\left(\frac{1}{2}(t-t_0)\rho\right) + \rho\cosh\left(\frac{1}{2}(t-t_0)\rho\right)}{(c_1 + c_2 - c_3)\sinh\left(\frac{1}{2}(t-t_0)\rho\right) + \rho\cosh\left(\frac{1}{2}(t-t_0)\rho\right)}, \quad (14)$$

with $t \in [t_0, t_1)$ and $\rho = \sqrt{(c_1 + c_2 + c_3)^2 - 4c_1c_3}$. Furthermore, normalization of $\tilde{\mathbf{\Pi}}^+(t)$ shows that when the next reaction at time t_1 fires, the conditional expectation instantaneously changes to one. This is consistent with the fact that for a transcription event to happen at t_1, the promoter must be in its active state at least until t_1. Fig. 2 illustrates the computation of the marginal hazard function used during MSA.

Another important quantity to describe the uncoupled model is its Kolmogorov-forward- (or master-) equation. We remark that since the uncoupled dynamics are non-Markovian, they do not satisfy a conventional master equation. Instead, such processes are described by *generalized master equations* (GME) that can account for memory effects in the dynamics (see S.5 in S1 Text for further discussion). In particular, there exists a fundamental relation between a GME's memory and the corresponding network's waiting-time distributions [28]. For the uncoupled telegraph model, the waiting-time distribution $f_3(\tau)$ of the transcription events is fully tractable using Eq. 10 (for an explicit expression see S1 Mathematica Notebook). It is therefore straightforward to develop a Chapman-Kolmogorov type of equation for the transcriptional dynamics $X(t) = R_3(t)$, i.e.,

$$P(x,t) = \int_0^t f_3(\tau) P(x-1, t-\tau) \mathrm{d}\tau \quad (15)$$

for $x > 0$ and $P(0,t) = 1 - F_3(t)$ with $F_3(t) = \int_0^t f_3(\tau)\mathrm{d}\tau$ as the cumulative waiting-time distribution. Following [28], Eq. 15 can be transformed into a gain-loss-type of master equation, i.e.,

$$\frac{\mathrm{d}}{\mathrm{d}t} P(x,t) = \int_0^t \phi_3(\tau)[P(x-1,t-\tau) - P(x,t-\tau)]\mathrm{d}\tau, \quad (16)$$

where the memory kernel $\phi_3(\tau)$ is related to $f_3(\tau)$ through the *Montroll-Weiss* equation (see S.5 in S1 Text). Eq. 16 represents the desired marginal master equation, from which further analysis could be deduced.

Comparison of MSA to existing approaches. The impact of environmental fluctuations on a dynamical system of interest is as diverse as the timescale on which they operate. For instance, extrinsic noise in the context of gene expression might be slowly varying (e.g., correlates well with the cell-cycle [29,30]), while fluctuations in transcription factor abundance might be significantly faster than the expression kinetics downstream. From a technical point of view, timescales range from constant environmental conditions that are random but fixed [31] to regimes where the fluctuations are very fast, such that quasi-steady-state (QSS) assumptions become applicable [16,32]. A QSS-based approach for simulating a system X in the presence of extrinsic noise Z corresponds to simulating the conditional CTMC $X|Z$, where Z is replaced by the mean of Z. Alternatively, one may try to replace a fluctuating environment Z through a random but fixed environment of same variance but this leads to an overestimation of the process variance in X [5], as discussed in a later section. To investigate the two above simplifying assumptions and compare them to the exact solution obtained via SSA and MSA, we performed a simulation study on a linear three-stage birth-death model given in Fig. 3a. In this case only species C is considered of interest whose uncoupled dynamics are obtained by marginalizing A and B. The results from Fig. 3b and Fig. 3c show that MSA facilitates accurate and fast approximations also under intermediate environmental time-scales where QSS- and static environmental assumptions break down.

Varying environmental time-scales. Simulation of the joint system (X,Z) becomes particularly challenging if the environmental fluctuations are fast, while with Eq. 8 the complexity of the marginal process simulation is invariant with respect to the time-scale of the environment. To demonstrate this effect, we performed a simulation study using a ten-dimensional, non-linear environmental network (Fig. 3d). Different time-scales were realized by multiplying the vector of environmental rate constants by a constant factor $q \in \{1,10,20\}$, effectively changing the number of reactions that have to be simulated on average. The results from Fig. 3e demonstrate that the computation time of the SSA simulation strongly increases with q. In contrast, the MSA's efficiency appears largely invariant with respect to the environmental time-scale, possibly leading to a substantial reduction in computational effort. We remark that for all three time-scales, the MSA algorithm achieved very high accuracies (i.e., estimated Kolmogorov distances below 3.5%). However, If the environmental time-scale is comparably slow (i.e., if $q = 1$), the extra effort needed for computing the marginal hazard function dominates, in which case SSA appears favorable.

Bistable environmental fluctuations. We further analyzed the case where the environmental network exhibits a more complex dynamics. In particular, we considered a bistable

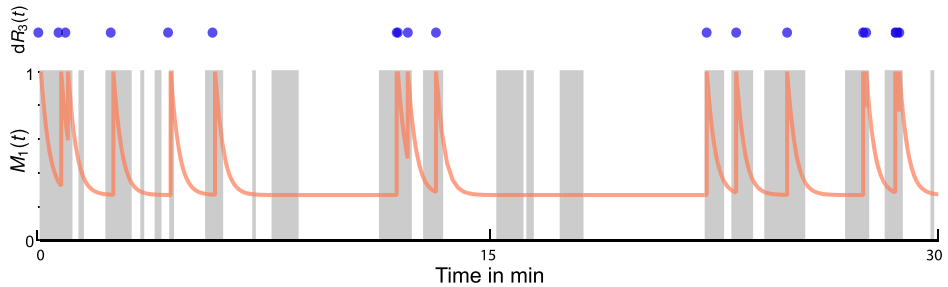

Figure 2. Schematic illustration of the marginal simulation algorithm. The red line shows the computation of the marginal hazard function $M_1(t)$ of the telegraph model (i.e., using Eq.14). The blue dots indicate the corresponding transcription events. Also shown is a possible corresponding realization of the promoter state (gray bars indicate the time the promoter has been active). The figure illustrates that using the MSA algorithm, only the events associated with X need to be simulated (in this case the transcription of mRNA). In contrast, SSA requires explicit simulation of all environmental states, which – depending on the time-scale – may become computationally prohibitive.
doi:10.1371/journal.pcbi.1003942.g002

Models

All models and simulations were implemented in MATLAB (The MathWorks, Natick, MA). Source codes for reproducing the numerical results have been attached as S1 Database.

Supporting Information

S1 Text Supplementary theory and derivations.
(PDF)

S1 Mathematica Notebook Mathematica notebook for the telegraph model.
(NB)

S1 Database MATLAB source codes used for simulations.
(ZIP)

Acknowledgments

We thank Jakob Ruess for sharing code for generating moments of chemical reaction networks.

Author Contributions

Conceived and designed the experiments: CZ HK. Performed the experiments: CZ. Analyzed the data: CZ. Contributed reagents/materials/analysis tools: CZ HK. Wrote the paper: CZ HK.

References

1. McAdams HH, Arkin A (1997) Stochastic mechanisms in gene expression. Proc Natl Acad Sci USA 94: 814–819.
2. Elowitz MB, Levine AJ, Siggia ED, Swain PS (2002) Stochastic gene expression in a single cell. Science 297: 1183–6.
3. Colman-Lerner A, Gordon A, Serra E, Chin T, Resnekov O, et al. (2005) Regulated cell-to-cell variation in a cell-fate decision system. Nature 437: 699–706.
4. Raser JM, O'Shea EK (2004) Control of stochasticity in eukaryotic gene expression. Science 304: 1811–1814.
5. Hilfinger A, Paulsson J (2011) Separating intrinsic from extrinsic fluctuations in dynamic biological systems. Proc Natl Acad Sci USA 108: 12167–12172.
6. Swain PS, Elowitz MB, Siggia ED (2002) Intrinsic and extrinsic contributions to stochasticity in gene expression. Proceedings of the National Academy of Sciences 99: 12795–12800.
7. Ruess J, Milias-Argeitis A, Lygeros J (2013) Designing experiments to understand the variability in biochemical reaction networks. Journal of The Royal Society Interface 10.
8. Zechner C, Ruess J, Krenn P, Pelet S, Peter M, et al. (2012) Moment-based inference predicts bimodality in transient gene expression. Proc Natl Acad Sci USA 109: 8340–8345.
9. Zechner C, Unger M, Pelet S, M P, Koeppl H (2014) Scalable inference of heterogeneous reaction kinetics from pooled single-cell recordings. Nat Methods 11: 197–202.
10. Finkenstädt B, Woodcock DJ, Komorowski M, Harper CV, Davis JR, et al. (2013) Quantifying intrinsic and extrinsic noise in gene transcription using the linear noise approximation: An application to single cell data. The Annals of Applied Statistics 7: 1960–1982.
11. Hasenauer J, Waldherr S, Doszczak M, Radde N, Scheurich P, et al. (2011) Identification of models of heterogeneous cell populations from population snapshot data. BMC Bioinformatics 12: 125.
12. Shahrezaei V, Ollivier JF, Swain PS (2008) Colored extrinsic fluctuations and stochastic gene expression. Mol Syst Biol 4: 196.
13. Newby JM (2012) Isolating intrinsic noise sources in a stochastic genetic switch. Physical biology 9: 026002.
14. Johnson N, Kotz S (1977) Urn Models and Their Application. Wiley & Sons, New York.
15. Koeppl H, Zechner C, Ganguly A, Pelet S, Peter M (2012) Accounting for extrinsic variability in the estimation of stochastic rate constants. Int J Robust Nonlin 22: 1103–1119.
16. Rao CV, Arkin AP (2003) Stochastic chemical kinetics and the quasi-steady-state assumption: Application to the gillespie algorithm. The Journal of Chemical Physics 118: 4999–5010.
17. Aalen OO, Borgan, Gjessing HK (2008) Survival and event history analysis: a process point of view. Springer Verlag.
18. Giesecke K, Kakavand H, Mousavi M (2011) Exact simulation of point processes with stochastic intensities. Operations research 59: 1233–1245.
19. Bain A, Crisan D (2009) Fundamentals of stochastic filtering. Springer, New York.
20. Anderson DF, Kurtz TG (2011) Continuous time Markov chain models for chemical reaction networks. In: Design and Analysis of Biomolecular Circuits, Springer. pp. 3–42.
21. Snyder DL, Miller MI (1975) Random Point Processes in Time and Space. Wiley & Sons, New York.
22. Elliott RJ, Malcolm WP (2005) General smoothing formulas for markov-modulated poisson observations. Automatic Control, IEEE Transactions on 50: 1123–1134.
23. Eden UT, Brown EN (2008) Continuous-time filters for state estimation from point process models of neural data. Statistica Sinica 18: 1293.
24. Leijdekker V, Spreij P (2011) Explicit computations for a filtering problem with point process observations with applications to credit risk. Probability in the Engineering and Informational Sciences 25: 393–418.
25. Hespanha J (2008) Moment closure for biochemical networks. In: Communications, Control and Signal Processing, 2008. ISCCSP 2008. 3rd International Symposium on. IEEE, pp. 142–147.
26. Anderson DF (2007) A modified next reaction method for simulating chemical systems with time dependent propensities and delays. J Chem Phys 127: 214107.
27. Lewis PAW, Shedler GS (1979) Simulation of nonhomogeneous Poisson processes by thinning. Naval Research Logistics Quarterly 26: 403–413.
28. Montroll EW, Weiss GH (1965) Random walks on lattices. ii. Journal of Mathematical Physics 6.
29. Rosenfeld N, Young JW, Alon U, Swain PS, Elowitz MB (2005) Gene regulation at the single-cell level. Science 307: 1962–1965.
30. Volfson D, Marciniak J, Blake WJ, Ostroff N, Tsimring LS, et al. (2006) Origins of extrinsic variability in eukaryotic gene expression. Nature 439: 861–864.
31. Zechner C, Deb S, Koeppl H (2013) Marginal dynamics of stochastic biochemical networks in random environments. 2013 European Control Conference (ECC): 4269–4274.
32. Kepler TB, Elston TC (2001) Stochasticity in transcriptional regulation: origins, consequences, and mathematical representations. Biophysical journal 81: 3116–3136.
33. Schlögl F (1972) Chemical reaction models for non-equilibrium phase transitions. Zeitschrift für Physik 253: 147–161.
34. Bowsher CG, Voliotis M, Swain PS (2013) The fidelity of dynamic signaling by noisy biomolecular networks. Plos Comp Biol 9: e1002965.
35. Jayanthi S, Nilgiriwala KS, Del Vecchio D (2013) Retroactivity controls the temporal dynamics of gene transcription. ACS Synthetic Biology 2: 431–441.
36. Lestas I, Vinnicombe G, Paulsson J (2010) Fundamental limits on the suppression of molecular fluctuations. Nature 467: 174–178.
37. Munsky B, Khammash M (2006) The finite state projection algorithm for the solution of the chemical master equation. J Chem Phys 124: 044104.