

STOCHASTIC APPROXIMATION FOR REGULATING CIRCADIAN CYCLES, A PRECISION MEDICINE VIEWPOINT

Alexey Nikolaev

Felisa J. Vázquez-Abad

Department of Computer Science

Hunter College and Graduate Center, City University of New York

CUNY Institute for Computer Simulation, Stochastic Modeling and Optimization

695 Park ave NYC, 10065 NY, USA

ABSTRACT

Circadian cycles and other self-regulatory biological processes are the result of complex interactions between gene expression and molecular interactions. In this paper we study a Petri net model of the circadian clock and use gradient estimation methods for finding optimal input rates. The significance of our research is the potential early identification of pathologies caused by aberrant cycles, and the discovery of those rates that are of main importance for the control of the cycles, enabling specific cures for people, in accordance with personalized (or precision) medicine. We use SPSA to drive the simulation to the optimal rates that result in a desired period, then propose a surrogate model for gradient estimation that evaluates the exact gradient for an "aggregate" system described by ODEs. Our hybrid model for gradient estimation addresses the high-dimensionality problem and can potentially increase the efficiency of the optimization method by at least one order of magnitude.

1 INTRODUCTION

The recent announcement from the White House to allocate a budget of \$215 million in 2016 for *precision medicine* acknowledges the need to increase research efforts to address the specific causes of disease, which may vary from person to person, rather than develop cures for the "average" person (MSNBC 2015).

Most of the research in this area falls within the scope of genomics and health informatics. In this work we focus on a specific problem, related to defects in self-regulatory mechanisms that keep the cells' biological clocks in order. Because of the daily periodic exposure to sunlight, cells must adjust to 24-hour patterns to function correctly. These are called *circadian cycles* (meaning "close to a day"). In mammals there is a master clock located in the hypothalamus that can function autonomously (without need of light exposure). The master clock can regulate other peripheral clocks via humoral signals and other chemicals. Peripheral circadian clocks are present in the liver, heart, lungs and kidneys, among other organs. It has been recently discovered (Sahar and Sassone-Corsi 2009) that aberrant circadian rhythm may cause tumorigenesis (when otherwise normal cells become cancer cells). Other mono-cellular mechanisms that rely on keeping good biological clocks are present in insulin regulation, sleeping patterns, etc.

The circadian mechanisms usually exhibit periodic patterns for the amount of certain molecules in the cell exhibiting peaks every 24 hours. Usually these patterns alternate for different molecules. As explained below, such patterns can be explained through the interaction of activation and repression mechanisms that affect these two types of molecules.

The mathematical model that we use in this paper describes a complete functioning circadian self-regulated mechanism for two molecules in a single cell. We use a stochastic process to model the number of molecules present in the cell. Various processes determine rates of production, destruction, binding and

unbinding of various molecules in order to achieve the periodic behavior. Importantly, our model requires a set of 15 *rates* in order to reproduce the simplest circadian self-regulatory mechanism. The question that we address is the following. *When the rhythm is aberrant, what are the optimal rates that are required in order to restore the 24 hour cycle?*

We pose this problem as a stochastic approximation problem and solve it first using SPSA (Fu and Hill 1997). Then we propose a hybrid model to increase the efficiency in gradient estimation with potential increase in efficiency.

Establishing the optimal rates may provide indication for plausible treatment. Our model permits the use of stochastic simulation to first mimic the situation observed for the real biological system under study, and then to determine which rates should be increased or decreased to promote resetting of the clocks. Importantly, gradient estimation can provide insight as to which of the firing mechanisms is more important to restore the 24 hour cycle. How the rates can be controlled is outside the scope of our research for now and would require laboratory analysis of single cell systems. As an example, in (Seo, Park, Lim, Kim, Lee, Baldwin, and Park 2012) it was determined that both light and temperature may affect the circadian clock in plant cells. These are examples of *exogenous* stimuli that may help to modify the rates.

The paper is organized as follows. Section 2 presents a model for a self regulated clock. The chemical equations can be described in a fluid approximation as a set of differential equations or ODE's (Vilar, Kueh, Barkai, and Leibler 2002) that can provide information on an aggregate level. However when the number of molecules is small the fluid approximation is not justified. Section 3 explains the discrete event (Petri net) model that we use, following the ideas proposed by (Gillespie 1977). Section 4 describes the hybrid model for estimating the period of the clock. Section 5 states the optimization problem and the hybrid gradient estimation technique. The last section 6 contains the simulation results using finite differences and SPSA.

2 BIOCHEMICAL MODEL FOR SELF-REGULATED CLOCK

We first review the basic mechanisms for gene expression and regulation. A *gene* is a functional segment of a DNA molecule that encodes the structure of a certain biomolecule: usually, the product is a protein, however it can be a functional RNA as well. The process of constructing a protein from its corresponding gene is called *gene expression*.

The first stage of this process is called *transcription*: at this stage the DNA segment that encodes the protein is copied to a new relatively short RNA molecule that is called messenger RNA (or mRNA for short). After that, at the second stage called *translation*, the mRNA is translated into the protein. To sum up, for making a protein molecule, first, a mRNA is constructed, which becomes the blueprint for making the protein itself.

The rates at which the transcription process may increase or decrease when certain proteins (called *transcription factors*) bind to the gene. Two main types of gene regulation are **activation** and **repression**.

The gene regulatory network (GRN) consists of two genes D_A and D_R , the former is encoding a transcription factor A (“activator”), and the latter is encoding another protein that we call R (“repressor”). The genes are first transcribed into mRNA (we call them m_A and m_R) at corresponding rates α_A and α_R . The messenger RNAs m_A and m_R are then translated into corresponding proteins with the rates β_A and β_R .

The protein A is able to bind to the genes D_A and D_R , which is why we call it a transcription factor. Binding increases the rates of production to $\alpha'_A > \alpha_A$ and $\alpha'_R \gg \alpha_R$ of their mRNAs, and the effect is particularly strong for the gene D_R . We use the notation D_A^b and D_R^b to represent the gene under binding.

The protein R , on the other hand, is not a transcription factor. It binds with the protein A making a complex C , which exists until the protein A in the complex degenerates. Figure 1 shows a diagram with the basic interactions that occur in the process.

The following chemical equations state the various processes (Vilar, Kueh, Barkai, and Leibler 2002). We have also included actual values for the rates that we use later on in our simulation experiments.

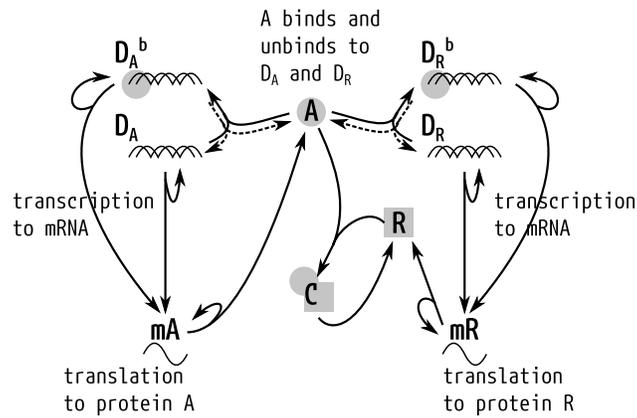
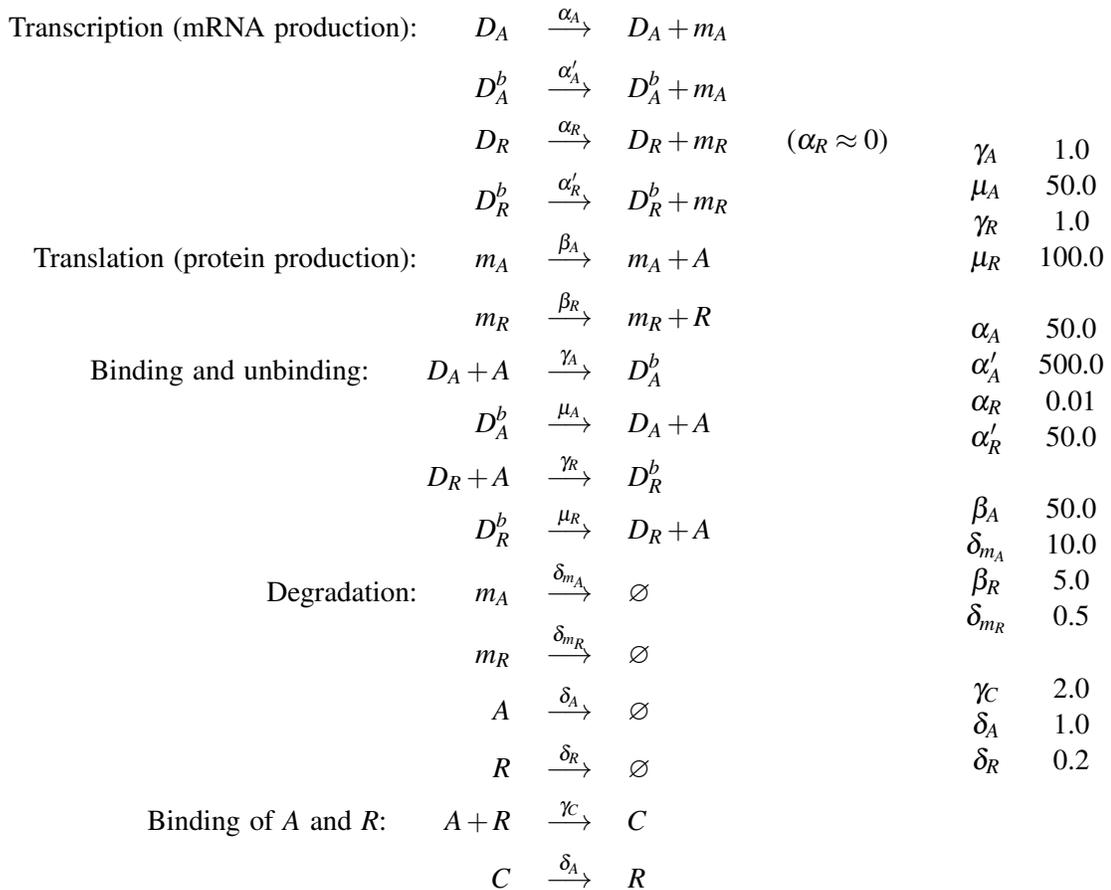


Figure 1: A diagram of the circadian clock model.



The vector of rates will be denoted by $\theta \in (\mathbb{R}^+)^{16}$. These rates control the dynamics of the process. We use the common notation $[X]$ for the concentration of protein X in the cell.

Figure 4 shows the concentrations of m_A , m_R , A , R , and C for one period of the circadian clock. The dynamics of the system can be described as follows. The system starts with two genes D_A and D_R in the unbound state with the concentrations of the other chemicals equal to zero (or close to zero). D_A immediately starts producing mRNA m_A , which get translated into A . After producing a few molecules

of A , they quickly bind to the genes, which switch to the bound state, D_A^b and D_R^b , and stay in this state almost always as long as the molecules of the protein A are still present in the system. In this bound state, the production rate of m_A is constant, and the concentration $[m_A]$ increases until the production and decay balance each other. The ODE approximation $\frac{\partial [m_A]}{\partial t} = \alpha'_A - \delta_{m_A}[m_A]$, yields the steady state concentration of $[m_A] = \alpha'_A / \delta_{m_A} \approx 50$.

Soon after, the production of mR catches up, and its concentration reaches the steady state level at $[m_R] = \alpha'_R / \delta_{m_R} \approx 100$. The concentrations of m_A and mR remain at this level until eventually all proteins A deplete, and the genes switch to the unbound state (thus decreasing the production rates of m_A and mR).

While the concentrations of m_A , and mR stay close to their steady state, the proteins A and R start getting produced. However, since the production rate of A is higher than the production rate of R , there is more A present in the system, and the quick complexation process make almost all R bind with A , producing the complex C . So, we get concentrations $[A]$ and $[C]$ growing, while $[R]$ stays close to zero. However, since C itself degenerates into R , eventually the complex C will “eat up” all free proteins A , thus depleting their concentration.

After this event, the genes switch to the unbound state, the production of m_A and mR drops. With the huge supply of C , which degenerate into R , and still a considerable amount of mR , which also keep producing R , the amount of $[R]$ increases. A few proteins A that get produced at this stage quickly join with R and eventually degenerate. So, this second phase of the oscillation is primarily governed by the production of R followed by their exponential decay. When R get completely depleted, the production of m_A and A restarts, and the cycle repeats itself.

3 SIMULATION MODEL PETRI NET

The ODE model for chemical reactions can be a good approximation of an aggregate behavior when the number of molecules is high, thereby describing the dynamics of the concentration of material. However, in biochemical settings where the number of molecules may be very few, it is not always appropriate to use the concentration as a state description. An alternative model was introduced by (Gillespie 1977) assuming that particles are created or destructed individually, according to the given rates. When the state of the system (measured in occupancy or number of molecules of each of the various components) is X , there are a number of possible *events* that can happen either producing new particles, destroying particles or binding. Each event corresponds to one chemical equation. The model assumes a Markovian structure where all residual times are independent exponential random variables.

Let $Y(t) = (X(t), D(t)) \in \mathbb{R}^5 \times \{0, 1\}^2$ be the process that counts the number of molecules of each type. Specifically, let $X_A(t) = [A](t)$, $X_{m_A}(t) = [m_A](t)$, $X_{m_R}(t) = [mR](t)$, $X_R(t) = [R](t)$, $X_C(t) = [C](t)$, $D_1(t) = [D_A](t) \in \{0, 1\}$, $D_2(t) = [D_R](t) \in \{0, 1\}$. In particular, $(D_1(t), D_2(t)) = (1, 1)$ when no protein A is bound to the genes, and it is equal to $(0, 0)$ when A are bound to both genes. Notice that the number of molecules in each component of X does not have an upper bound and their dynamics follows a general multidimensional birth and death process. The model is a hybrid model because the rates are dependent on the *regime* dictated by the DNA component $D(t)$. Notice also that this component does not behave as a Markov Hidden Model, but it is dependent on the state $Y(t)$.

Let k label the possible *events* each associated with a different chemical equation. We define the increment vector $v_{i,k}$ as the number of molecules of the type i that are added or removed from the system when the event k happen:

$$v_{ik} = \begin{cases} +1 & \text{if event } k \text{ increases component } Y_i \text{ by one} \\ -1 & \text{if event } k \text{ decreases component } Y_i \text{ by one} \\ 0 & \text{otherwise,} \end{cases}$$

The *propensity* $a_k(Y, \theta)$ is the corresponding event rate (i.e. inverse expected time) at which event k occurs given state Y and the reaction rates θ . The propensity of a chemical reaction is proportional to

its rate and the concentrations of the reactants. For example, if k is the event of the production of A , the propensity $a_k(X, D, \theta) = \beta_A X_{m_A}$ and $v_{A,k} = +1$, while for k' the event of degradation of A , we'll have $a_{k'}(X, D, \theta) = \delta_A X_A$ and $v_{A,k'} = -1$. We use the notation \mathfrak{F}_t for the natural filtration of the process, that is, $\mathfrak{F}_t = \sigma(X(s), D(s); s \leq t)$.

The simulation model is a particular case of the standard clock simulation model. Assuming exponential residual times for each chemical equation, the time for the next event has exponential distribution with rate $\mathbf{a}(Y, \theta) = \sum_k a_k(Y, \theta)$, and the probability that the next event is event j is $a_j(Y, \theta) / \mathbf{a}(Y, \theta)$.

This model leads to a Petri net model for simulation shown in Figure 2 with corresponding algorithm:

- Initialize the number of molecules Y .
- Loop:
 - Compute the propensities $a_k(Y, \theta)$ of each reaction (event).
 - Sample and fire the next reaction. $\mathbb{P}(\text{event } j) = a_j(Y, \theta) / \mathbf{a}(Y, \theta)$.
 - Update time $t \leftarrow t + \Delta$, where $\Delta \sim \text{Exp}(\mathbf{a}(Y, \theta))$.
 - Update the state $Y := Y + v_k$.

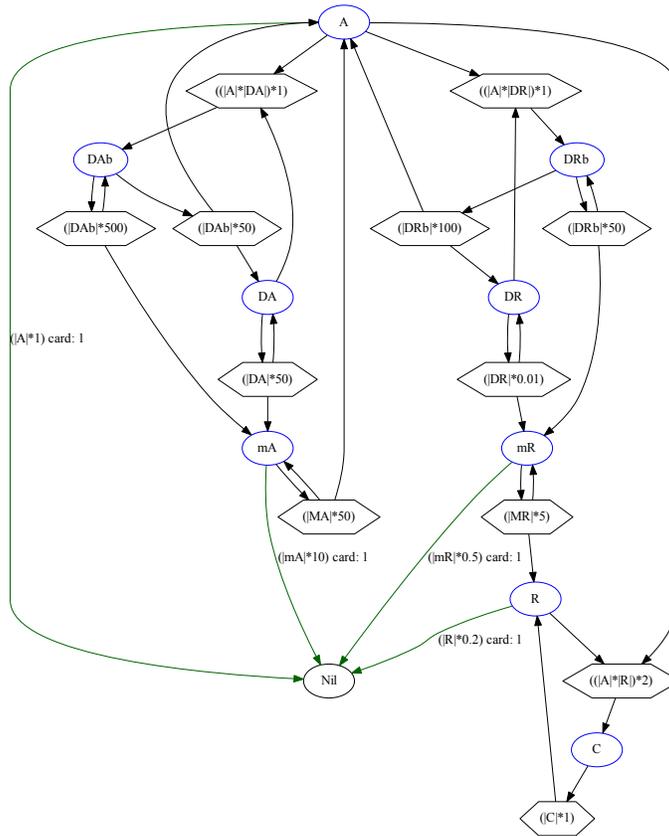


Figure 2: A Petri net of the circadian clock model generated by Beta Workbench.

To estimate the period $P(\theta)$ of the oscillations, it turns out that the simulation can run for exactly one period, starting at the time $t = 0$ with the initial concentrations equal to the concentrations of the chemicals that are normal at that moment (the choice of the initial conditions is discussed below, when we talk about

the systems of the ODEs). The simulation runs until the concentration of R depletes ($[R] = 0$), and the production of A starts again (in particular, we chose the condition $[A] > 10$ for this).

4 HYBRID MODEL FOR ESTIMATION

In order to improve the efficiency of the method we propose to use a hybrid model that approximates the discrete event Markov chain process with its expected or aggregate behavior dominated by ODE's only during certain subintervals within each cycle, avoiding the problem of "fractional" number of molecules that the full ODE model has.

Theorem 1 For any time T , conditioning on the event of no regime changes: $\{D(t+s) = D; s \leq T\}$ the process X satisfies:

$$\mathbb{E}(X_i(t+s) - X_i(t) | \mathfrak{F}_t) = \int_t^{t+s} \sum_k v_{i,k} \mathbb{E}(a_k(X(u), D, \theta) | \mathfrak{F}_t) du, \tag{1}$$

for all $i \in I = \{A, R, mA, mR, C\}$ and for any $s \leq T$.

Proof. Let $h > 0$ be an infinitesimal quantity. It follows from the exponential distribution and merging of Poisson processes that the probability of having two or more events within $[t, t+s]$ is $\mathcal{O}(h^2)$. The probability of no events is $1 - \mathcal{O}(h)$, and in this case $X(t+h) - X(t) = 0$. When there is only one event happening then $X_i(t+h) - X_i(t) = v_{ik}$ with a corresponding probability $h a_k(X(t), D, \theta) + \mathcal{O}(h^2)$, thus

$$\frac{\mathbb{E}(X_i(t+h) - X_i(t) | \mathfrak{F}_t)}{h} = \sum_k v_{i,k} a_k(X(t), D, \theta) + \mathcal{O}(h).$$

Given a constant regime D on $[t, t+T)$ the propensities are continuous functions of the state X . For given $h > 0$, let $t_0 = 0, t_n = t + nh$. Then using a telescopic sum,

$$\begin{aligned} \mathbb{E}(X_i(t+s) - X_i(t) | \mathfrak{F}_t) &= \\ &= \mathbb{E} \left(\sum_{n=1}^{\lfloor s/h \rfloor} (X_i(t_n) - X_i(t_{n-1})) \middle| \mathfrak{F}_t \right) + \mathcal{O}(h) = \mathbb{E} \left(\sum_{n=1}^{\lfloor s/h \rfloor} h \frac{\mathbb{E}(X_i(t_n) - X_i(t_{n-1}) | \mathfrak{F}_{t_{n-1}})}{h} \middle| \mathfrak{F}_t \right) + \mathcal{O}(h) \\ &= \mathbb{E} \left(\sum_{n=1}^{\lfloor s/h \rfloor} h \left(\sum_k v_{i,k} a_k(X(t_{n-1}), D, \theta) \right) \right) + \mathcal{O}(h) \xrightarrow{h \rightarrow 0} \mathbb{E} \left(\int_t^{t+s} \sum_k v_{i,k} a_k(X(u), D, \theta) du \middle| \mathfrak{F}_t \right). \end{aligned}$$

□

It is interesting to note that if, and only if, the propensities are linear in the state X , then the so-called averaged process $\langle X(t) \rangle$ satisfies the ODE

$$\frac{dx(t)}{dt} = \sum_k v_{i,k} a_k(x(t), D, \theta). \tag{2}$$

Gillespie (Gillespie 2000) mentions that this ODE cannot be accurate unless one looks at the limit when the number of molecules is large enough that one can approximate it by a continuous vector. We remark here that the ODE may be accurate for describing the *expected* behavior of the processes, provided that the propensities are linear, and that the regime is constant. For our particular model of the activator and repressor proteins, there is a chemical equation for which $a_k(X(t), D, \theta) = \gamma_C X_A(t) X_R(t)$ is the propensity for the increase of X_C and decrease in both X_A and X_R . Because of the correlations between components, in this case $\mathbb{E}(a_k(X(t), D, \theta)) \neq a_k(\mathbb{E}(X(t), D, \theta))$.

Corollary 2 For any finite $\theta > 0$ the process $\{Y(t); t > 0\}$ is tight (or, equivalently: “bounded in probability”).

Proof. The proof of the result is based on a stochastic Lyapounov argument. The key observation is that all decaying propensities are proportional to the amount of material, but increasing rates are not. We need consider only the X component of the state because there are only two DNA molecules, so $D(\cdot)$ is bounded w.p.1. Given any regime with constant D , X_{mA} (and similarly for X_{mR}) satisfy

$$\mathbb{E}(X_{mA}(t+s) | \mathfrak{F}_t) = X_{mA}(t) + \int_t^{t+s} (\theta_{k_1} - \theta_{k_2} X_{mA}(u)) du,$$

for some indices k_1, k_2 . Suppose that $[t, t+s)$ is a time interval where $X_{mA}(\cdot) > \theta_{k_1}/\theta_{k_2}$. Then the expected drift is negative: the process behaves as a strict supermartingale over all such time intervals. A non-negative supermartingale converges with probability one to a finite value, thus

$$\lim_{K \rightarrow \infty} \mathbb{P} \left(\sup_{t \geq 0} (X_{mR}(t)) > K \right) = 0,$$

which implies that $X_{mA}(\cdot) \leq \theta_{k_1}/\theta_{k_2}$ infinitely often w.p.1., proving tightness for the RNA molecules.

Consider now the aggregate count of protein molecules $Z(t) = X_A(t) + X_R(t) + X_C(t)$. Because each component is non-negative, tightness of Z implies that of all three components. From the chemical equations

$$\mathbb{E}(Z(t+s) | \mathfrak{F}_t) = Z(t) + \int_t^{t+s} (\beta_A X_{mA}(u) + \beta_R X_{mR}(u) - \delta_A X_A(u) - \delta_R X_R(u) - \gamma_C X_A(u) X_R(u)) du.$$

Given the condition $C(K) = \sup(\beta_A X_{mA}(u) + \beta_R X_{mR}(u)) \leq K; u \in [t, t+s)$ there is a bounded region in the plane (x_A, x_R) such that outside that region the integrand above is strictly negative. If the process stays outside this bounded region for an interval of time within $[t, t+s)$ then Z is a non-negative supermartingale and converges towards the bounded region w.p.1. Because $\mathbb{P}(C(K)) \rightarrow 0$ as $K \rightarrow \infty$ this implies that $\{Z(t); t \geq 0\}$ is tight. \square

A full cycle can be divided into three distinct periods, each of which corresponds approximately to a different regime. The first part of the cycle describes the dynamics of A , assuming that $D = (0, 0)$ during this period. The second part of the cycle describes the joint dynamics of R and C when there is no activator present, assuming $D = (1, 1)$ until R depletes to the level of C molecules. The third part is a hybrid process when all molecules exist in relatively small quantities until R depletes and a few molecules of A are produced (in simulations, we use the condition $X_A \geq 10$).

Regime $D = (0, 0)$. In this regime A is bound to the DNAs, there is large supply of A , i.e. $x_A \gg 1$, and $\gamma_C x_A \gg \delta_R$, which implies that all supply of R almost immediately get transformed into the complex C , and we can assume the decay rate of R to be close to zero, $\delta_R \approx 0$.

If these assumptions hold, there is very little amount of R in the system. Practically all R immediately joins with A producing the complex C . There are two main sources of R : translation from mR and the decay of C , this gives the rate $\beta_R x_{mR}(t) + \delta_A x_C(t)$. And because R almost immediately becomes C this is the production rate of C . At the same times, because the production of C requires molecules A , the same rate contributes to the depletion of A .

Thus we obtain the following (approximate) system of ODE's:

$$\begin{aligned}\frac{dx_{mA}(t)}{dt} &= \alpha'_A - \delta_{mA}x_{mA}(t), \\ \frac{dx_{mR}(t)}{dt} &= \alpha'_R - \delta_{mR}x_{mR}(t), \\ \frac{dx_A(t)}{dt} &= \beta_Ax_{mA}(t) - (\beta_Rx_{mR}(t) + \delta_Ax_C(t)) - \delta_Ax_A(t), \\ \frac{dx_C(t)}{dt} &= (\beta_Rx_{mR}(t) + \underbrace{\delta_Ax_C(t)}_{=0}) - \delta_Ax_C(t).\end{aligned}$$

Regime $D = (1, 1)$. The protein A molecules are not bound to the DNAs and $x_A \approx 0$, $x_R \gg x_A$, and $x_C \gg x_A$. Because $\gamma_{CxR} \gg \delta_A$, all molecules of the protein A immediately get transformed into C . At this stage, x_A is at its steady state value α_A/δ_{mA} . Therefore, the production of A is equal to $\beta_A(\alpha_A/\delta_{mA})$, and since all A are almost immediately transformed into C , this value is the production rate of C , and simultaneously it contributes to the depletion of R .

Therefore, the (approximate) system of ODE's is

$$\begin{aligned}\frac{dx_{mR}(t)}{dt} &= \alpha_R - \delta_{mR}x_{mR}(t), \\ \frac{dx_R(t)}{dt} &= \beta_Rx_{mR}(t) - \beta_A(\alpha_A/\delta_{mA}) - \delta_Rx_R(t) + \delta_Ax_C(t), \\ \frac{dx_C(t)}{dt} &= \beta_A(\alpha_A/\delta_{mA}) - \delta_Ax_C(t)\end{aligned}$$

Initial Conditions for ODE's. Because the regimes alternate, the initial conditions for the regime $D = (0, 0)$ are the steady state concentrations of x_C , x_{mR} , and x_{mA} in the regime $D = (1, 1)$. These concentrations can be determined exactly by setting $dx_C/dt = 0$, $dx_{mR}/dt = 0$ and $dx_{mA}/dt = 0$, which yield:

$$x_{mA}(0) = \frac{\alpha_A}{\delta_{mA}} = 5, \quad x_{mR}(0) = \frac{\alpha_R}{\delta_{mR}} \approx 0, \quad x_A(0) = 0, \quad x_C(0) = \frac{\beta_A \alpha_A}{\delta_{mA} \delta_A} = 250.$$

And the concentrations of the chemicals at the end of this first regime $D = (1, 1)$ evaluated numerically, give us the initial conditions the second regime $D = (0, 0)$.

We use *Mathematica* to integrate the first system of ODEs. With this solution we find the moment when A depletes, that is:

$$P_A(\theta) = \inf(t > 0: x_A(t) \leq 0), \tag{3}$$

which we do numerically. The concentrations are used as initial condition for the second part of the period. We use *Mathematica* to integrate this second system of ODEs. Then find the moment of time when the concentrations of R and C become equal, that is:

$$P_R(\theta) = \inf(t > 0: x_R(t) \leq \kappa), \tag{4}$$

where κ is the steady state value of C in the second regime.

The last part of the cycle is from the point where $x_R(t) = x_C(t)$ until the point when R depletes and the production of A starts again. This part of the cycle has more complex dynamics, with non-linear rates and small amounts of molecules and regime changes that trigger the start of a new cycle. On this part we do not use the aggregate behavior, but the actual simulation of the process $\{X(t), D(t)\}$. The estimation of the period P can then be done using $\hat{P}(\theta) = P_A(\theta) + P_{RC}(\theta) + \hat{P}_C(\theta)$, where the last part is estimated with simulation.

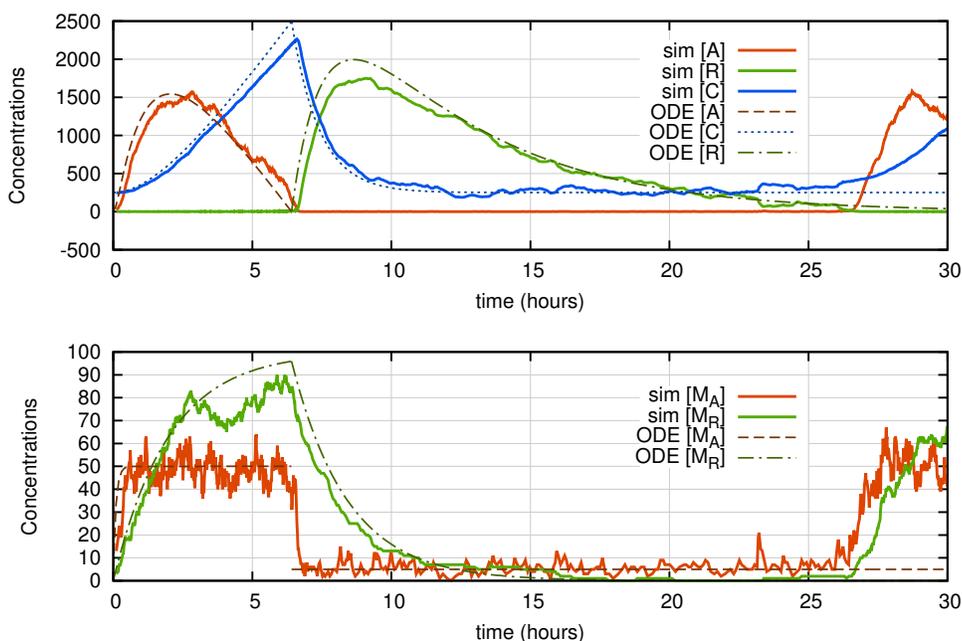


Figure 3: A comparison between the simulation results and the systems of ODEs.

Figure 4 shows a simulation trajectory using the Gillespie simulation model against the hybrid process using ODE approximation in the first two parts.

For the simulations we used the rates specified in the chemical equations, and simulated the process for 100,000 periods. The hybrid estimation method uses the ODE's for the first two parts and simulations for the last part, while the naive Monte Carlo estimation uses simulation all along. The hybrid estimator of the period is $\hat{P}^H(\theta) = 6.41037 + 14.3311 + (5.06924 \pm 0.044) = 25.811 \pm 0.044$ and it took 1.5 minutes of CPU time on a laptop computer using Linux. The naive Monte Carlo estimator is $\hat{P}^N(\theta) = 25.6343 \pm 0.052$ and it took 18.4 minutes of CPU time. **Thus, hybrid estimation increases efficiency 20 fold.**

5 GRADIENT ESTIMATION

5.1 Motivation

In this section we seek to determine the exact rates that result in a desired period, for instance 24 hour cycles. The input parameters or *control variables* are the 15 different rates that control the chemical reactions. We call θ the vector of rates. Notice that the propensities a_k are themselves functions of θ and of the current state of the process. The state describes the number of molecules of each kind as time evolves. In order to achieve a desired period we pose the optimization problem as a tracking problem, with

$$J(\theta) = \frac{1}{2}(P(\theta) - \pi)^2,$$

where π is the desired cycle time (such as 24 hours). The problem is then to minimize $J(\theta)$. Many tracking problems have a monotonic structure, where $P(\theta)$ is either decreasing or increasing in each of its arguments. In that situation it is common to use the Robbins Monro procedure directly with

$$\theta_{n+1} = \theta_n - M\varepsilon_n(\hat{P}(\theta_n) - \pi)$$

where $\{\varepsilon_n\}$ is a suitable step size sequence, and M is a constant diagonal matrix with ± 1 values on the diagonal, depending on the monotonicity of the function. However in our case the influence of rates in the

increase or decrease of concentration of proteins is very complex and not necessarily monotonic. Thus it is necessary to use the information on the gradient of the period with respect to the rates.

Gradient estimation is also important in that the *sensitivity* to each of the rates θ provides information about the relative impact that each chemical process has. Rates with no statistically significant impact on the period should not be targeted for development of possible treatment.

5.2 Hybrid Gradient Estimation

The ODE's that we use to compute $P_A(\theta)$ and $P_R(\theta)$ can also be used to compute $\nabla_{\theta}P_A(\theta)$ and $\nabla_{\theta}P_R(\theta)$, as follows. In general notation, the differential equations for $x_A(t)$ in the first part, and for $x_R(t)$ in the second are of the form (2). Because the propensities are affine functions of θ the *derivative processes* $\xi_j(t) = \partial x(t)/\partial \theta_j; j = 1, \dots, 16$ are well defined and they satisfy the companion system of ODE's:

$$\frac{d\xi_{ij}(t)}{dt} = \sum_k v_{ik} \left(\nabla_x a(x(t), D, \theta)^T \xi_j(t) + \frac{\partial a(x(t), D, \theta)}{\partial \theta_j} \right); \quad j = 1, \dots, 15; \quad i \in I.$$

The general form of PA and PB in equations (3) and (4) is $P_{\alpha}(\theta) = \inf(t > 0: x_i(t) \leq \alpha)$, therefore, using the Implicit Function Theorem

$$\frac{\partial P_{\alpha}(\theta)}{\partial \theta_j} = \left(\frac{dx_i}{dt}(P_{\alpha}(\theta)) \right)^{-1} \xi_{ij}(P_{\alpha}(\theta)).$$

We used *Mathematica* to integrate the derivative processes along with the original processes $x(t)$. In our example the equations simplify because many chemical equations are independent of each rate θ_j , and the partial derivatives of the propensities are either constant or linear in the state space.

The gradients for the third part of the period, for which we don't have an ODE representation, were estimated in using finite differences.

Before applying the hybrid estimation technique, we first estimated the gradient using finite differences, running Monte-Carlo simulation for the entire period. The results are shown in the table below:

Table 1: Components of the gradient of the expected period $\mathbb{E}(P(\theta))$. Naive Monte-Carlo method.

	γ_A	μ_A	γ_R	μ_R	α_A	α'_A	α'_R	α''_R
Mean	-2.7*	0.048	0.1	-0.014	-0.048	0.0107	14	-0.044
StdErr	0.4	0.008	0.4	0.004	0.007	0.0008	37	0.008
	β_A	δ_{mA}	β_R	δ_{mR}	γ_C	δ_A	δ_R	
Mean	0.051	-0.36	-0.30	-3.7*	1.09*	-8.5*	-78.2*	
StdErr	0.007	0.04	0.08	0.8	0.19	0.4	1.9	

Then, for the most important components (the derivatives with respect to δ_R and δ_A) we computed the hybrid gradient estimation:

$$\begin{aligned} \frac{\partial \mathbb{E}(P(\theta))}{\partial \delta_R} &= 0 + (-62.5) + (-15.3 \pm 1.6) = -77.8 \pm 1.6, \\ \frac{\partial \mathbb{E}(P(\theta))}{\partial \delta_A} &= -4.667 + (-0.017) + (-3.8 \pm 0.3) = -8.5 \pm 0.3. \end{aligned}$$

The numbers match perfectly with the results of the simpler method, while the running time for the hybrid method has decreased dramatically (approximately by the factor of 15 – 20).

Also we would like to note that the numbers in the table agree with our assessment that very few rates have a significant effect on the period. The rates that have statistically significant non-zero derivative, while having a relatively large magnitude of the derivative are signaled by an asterisk.

6 STOCHASTIC OPTIMIZATION VIA SIMULATION

The optimization approach that we followed as benchmark is an application of SPSA (He, Fu, and Marcus 2003), which ensures almost sure convergence to θ^* when minimizing $J(\theta)$, provided that this is convex. At each iteration n , we sample N random perturbations $(\Delta_n)_1 \dots (\Delta_n)_N$ uniformly from the set of two possible values $\{+1, -1\}$ and we simulate the biological system for $T = 500$ hours, at rate $\theta_n + c_n \Delta_n$ and in parallel we also simulate $T = 500$ hours at rate $\theta_n - c_n \Delta_n$. The gradient is then approximated by the symmetric finite difference and the rates are updated with:

$$(\theta_{n+1})_i = (\theta_n)_i - \varepsilon_n \frac{\hat{J}(\theta_n + c_n \Delta_n) - \hat{J}(\theta_n - c_n \Delta_n)}{2c_n (\Delta_n)_i}.$$

The parameters $c_n = \text{const}/(n+1)^{0.101}$, $\varepsilon_n = \text{const}/(n+2)^{0.602}$ are shown to be optimal in (Bhatnagar, Prasad, and Prashanth 2013).

Figure 4 shows the result of the SPSA for a target cycle of $\pi = 48$ hours. Observe that the most important rate turns out to be δ_R , the degradation rate of the protein R , which is consistent with the model: The decay of R has a longer tail now, and so the period increases.

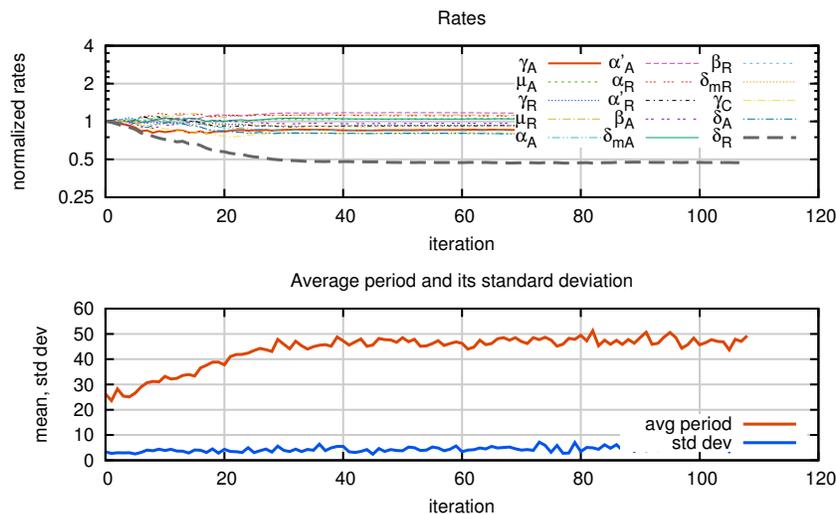


Figure 4: SPSA iteration procedure is converging to oscillations with the period of 48 hours.

The hybrid gradients are estimated using *Mathematica* code for two parts of the cycle and simulation using our own Gillespie algorithm simulator written in C for the third part. We are currently consolidating our code to perform the full hybrid gradient estimation at each iteration of the stochastic approximation. From our experiments on gradient estimation we expect a significant increase in the efficiency of the algorithm. In addition, from the insight gained above, we will incorporate a learning stage to identify which derivatives are close to zero, and stop evaluating those as the stochastic approximation evolves.

7 CONCLUDING REMARKS

We are currently working on the problem of applying IPA to the derivative estimation of the third part of the cycle. The simulation model as stated here does not satisfy the Lipschitz continuity assumptions required for unbiasedness. However under a different representation it is possible to show unbiasedness of a Filtered Monte Carlo gradient estimator. The results will be reported elsewhere. In future we will focus on online versions of the algorithms that can achieve monitoring and change detection, with the goal of triggering early treatment to restore normal cycle operation.

REFERENCES

- Bhatnagar, S., H. Prasad, and L. Prashanth. 2013. “Stochastic Recursive Algorithms for Optimization”. *Lecture Notes in Control and Information Sciences* 434.
- Fu, M. C., and D. Hill. 1997. “Optimization of Discrete Event Systems via Simultaneous Perturbation Stochastic Approximation”. *IEEE Transactions* 29 (3): 233–243.
- Gillespie, D. T. 1977. “Exact Stochastic Simulation of Coupled Chemical Reactions”. *The Journal of Physical Chemistry* 81 (25): 2340–2361.
- Gillespie, D. T. 2000. “The Chemical Langevin Equation”. *The Journal of Chemical Physics* 113 (1): 297–306.
- He, Y., M. C. Fu, and S. I. Marcus. 2003. “Convergence of Simultaneous Perturbation Stochastic Approximation for Nondifferentiable Optimization”. *IEEE Transactions on Automatic Control* 48 (8): 1459–1463.
- MSNBC, M. Richinick 2015. “Obama Seeks \$215 Million for Personalized Medicine”. <http://www.msnbc.com/msnbc/obama-seeks-215-million-personalized-medicine>. Accessed: 2015-04-08.
- Sahar, S., and P. Sassone-Corsi. 2009. “Metabolism and Cancer: the Circadian Clock Connection”. *Nature Reviews Cancer* 9 (12): 886–896.
- Seo, P. J., M.-J. Park, M.-H. Lim, S.-G. Kim, M. Lee, I. T. Baldwin, and C.-M. Park. 2012. “A Self-regulatory Circuit of CIRCADIAN CLOCK-ASSOCIATED1 Underlies the Circadian Clock Regulation of Temperature Responses in Arabidopsis”. *The Plant Cell Online* 24 (6): 2427–2442.
- Vilar, J. M., H. Y. Kueh, N. Barkai, and S. Leibler. 2002. “Mechanisms of Noise-resistance in Genetic Oscillators”. *Proceedings of the National Academy of Sciences* 99 (9): 5988–5992.

AUTHOR BIOGRAPHIES

ALEXEY NIKOLAEV is a PhD candidate at the Graduate Center of the City University of New York (CUNY). He received his undergraduate degree in Physics from Moscow State University in 2008 and his Master’s degree in Physics at Clarkson University in 2011. His interests are in bioinformatics, systems biology, algorithms, and numerical methods. His email address is anikolaev@gradcenter.cuny.edu

FELISA VÁZQUEZ-ABAD is Professor of Computer Science at Hunter College and Executive Director of the CUNY Institute CoSSMO. She obtained a B.Sc. in Physics in 1983 and a M.Sc. in Statistics and Operations Research from the Universidad Nacional Autónoma de México. In 1989 she obtained a Ph.D. in Applied Mathematics from Brown University. She became a professor at the University of Montreal, Canada in 1993, where she remained until 2004 when she became a professor at the University of Melbourne, Australia, until she moved to New York in 2009. In 2014 she founded the CUNY Institute CoSSMO. Her interests focus on the optimization of complex systems under uncertainty, primarily to build efficient self-regulated learning systems. Her email address is fvazquez@hunter.cuny.edu.