A Scheme for Molecular Computation of Maximum Likelihood Estimators for Log-Linear Models

Manoj Gopalkrishnan

Tata Institute of Fundamental Research, Mumbai 400 005, India
`manoj.gopalkrishnan@gmail.com`, `http://www.tcs.tifr.res.in/~manoj`

Abstract. We propose a novel molecular computing scheme for statistical inference. We focus on the much-studied statistical inference problem of computing maximum likelihood estimators for log-linear models. Our scheme takes log-linear models to reaction systems, and the observed data to initial conditions, so that the corresponding equilibrium of each reaction system encodes the corresponding maximum likelihood estimator. The main idea is to exploit the coincidence between thermodynamic entropy and statistical entropy. We map a Maximum Entropy characterization of the maximum likelihood estimator onto a Maximum Entropy characterization of the equilibrium concentrations for the reaction system. This allows for an efficient encoding of the problem, and reveals that reaction networks are superbly suited to statistical inference tasks. Such a scheme may also provide a template to understanding how *in vivo* biochemical signaling pathways integrate extensive information about their environment and history.

1 Introduction

The sophisticated behavior of cells emerges from the computations that are being performed by the underlying biochemical reaction networks. These biochemical pathways have been studied in a "top-down" manner, by looking for recurring motifs, and signs of modularity [18]. There is also an opportunity to study these pathways in a "bottom-up" manner by proposing primitive building blocks which can be composed to create interesting and technologically valuable behavior. This "bottom-up" approach connects with work in the Molecular Computation community whose goal is to generate sophisticated behavior using DNA hybridization reactions [23,24,19,31,27,6,7,22,3,25] and other Artificial Chemistry approaches [5,10].

We propose a new building block for molecular computation. We show that the mathematical structure of reaction networks is particularly well adapted to compute Maximum Likelihood Estimators for log-linear models, allowing a pithy encoding of such computations by reactions. According to [12]:

> Log-linear models are arguably the most popular and important statistical models for the analysis of categorical data; see, for example, Bishop, Fienberg and Holland (1975) [4], Christensen (1997) [8]. These powerful models, which include as special cases graphical models [see, e.g., Lauritzen (1996) [16]] as well as many logit models [see, e.g., Agresti (2002) [1], Bishop, Fienberg and Holland (1975) [4]], have applications in many scientific areas, ranging from social and biological sciences, to privacy and disclosure limitation problems, medicine, data mining, language processing and genetics. Their popularity has greatly increased in the last decades...

In order to respond in a manner that maximizes fitness, a cell has to correctly estimate the overall state of its environment. Receptors that sit on cell walls collect a large amount of information about the cellular environment. Processing and integration of this spatially and temporally extensive and diverse information is carried out in the biochemical reaction pathways. We propose that this processing and integration may be advantageously viewed from the lens of machine learning.

Our proposal entails that *schemes for statistical inference by reaction networks are of biological significance, and are deserving of as thorough and extensive a study as schemes for statistical inference by neural networks.* In particular, machine learning is not just a tool for the analysis of biochemical data, but theoretical and technological insights from machine learning could provide a deep and fundamental way, and perhaps "the" correct way, to think about biochemical networks. We view the scheme we present here as a promising first step in this program of applying machine learning insights to biochemical networks.

The problem: We illustrate the main ideas of our scheme with an example. Following [21], consider the **log-linear model** (also known as toric model) described by the **design matrix** $A = \begin{pmatrix} 2 & 1 & 0 \\ 0 & 1 & 2 \end{pmatrix}$. This means that we are observing an event with three possible mutually exclusive outcomes, call them $X_1, X_2,$ and X_3, which represent respectively the columns of A. The rows of A represent "hidden variables" θ_1 and θ_2 respectively which parametrize the statistics of the outcomes in the following way specified by the columns of A:

$$P[X_1 \mid \theta_1, \theta_2] \propto \theta_1^2$$
$$P[X_2 \mid \theta_1, \theta_2] \propto \theta_1 \theta_2$$
$$P[X_3 \mid \theta_1, \theta_2] \propto \theta_2^2$$

where the constant of proportionality normalizes the probabilities so they sum to 1. [1]

Suppose several independent trials are carried out, and the outcome X_1 is observed $x_1 \in (0,1)$ fraction of the time, the outcome X_2 is observed $x_2 \in (0, 1 - x_1)$ fraction of the time, and the outcome X_3 is observed $x_3 = 1 - x_1 - x_2$ fraction of the time. We wish to find the maximum likelihood estimator $(\hat{\theta}_1, \hat{\theta}_2) \in \mathbb{R}_{>0}^2$ of the parameter (θ_1, θ_2), i.e., that value of θ which maximizes the likelihood of the observed data.

Our contribution: We describe a scheme that takes the design matrix A to a reaction network that solves the maximum likelihood estimation problem. In Definition 8, we describe our scheme for every matrix A over the integers with all column sums equal. All our results hold in this generality.

– In Definition 8.5, we show how to obtain from the matrix A, a reaction network that computes the maximum likelihood distribution. Specialized to our example, note that the kernel of the matrix A is spanned by the vector $(1, -2, 1)^T$. We encode this by the reversible reaction

$$X_1 + X_3 \underset{1}{\overset{1}{\rightleftharpoons}} 2X_2$$
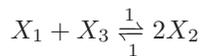
– In Theorem 5, we show that if this reversible reaction is started at initial concentrations $X_1(0) = x_1, X_2(0) = x_2, X_3(0) = x_3$, and the dynamics proceeds according to the law of mass action with all specific rates set to 1:

$$\dot{X}_1(t) = \dot{X}_3(t) = -X_1(t)X_3(t) + X_2^2(t), \qquad \dot{X}_2(t) = -2X_2^2(t) + 2X_1(t)X_3(t)$$

---

[1] It is more common in statistics and statistical mechanics literature to write $\theta_1 = e^{-E_1}$ and $\theta_2 = e^{-E_2}$ in terms of "energies" E_1, E_2 so that $P[X_2 \mid E_1, E_2] \propto e^{-E_1 - E_2}$ for example.

then the reaction reaches equilibrium $(\hat{x}_1, \hat{x}_2, \hat{x}_3)$ where $\hat{x}_1 + \hat{x}_2 + \hat{x}_3 = 1$ and $\hat{x}_1 \propto \hat{\theta}_1^2$, $\hat{x}_2 \propto \hat{\theta}_1\hat{\theta}_2$, and $\hat{x}_3 \propto \hat{\theta}_2^2$, so that $(\hat{x}_1, \hat{x}_2, \hat{x}_3)$ represents the probability distribution over the outcomes X_1, X_2, X_3 at the maximum likelihood $\hat{\theta}_1, \hat{\theta}_2$.

– This part of our scheme involves only reversible reactions, and requires no catalysis (see [13, Theorem 5.2] and Lemma 2). One difficulty with implementing such schemes has been that empirical control over kinetics is rather poor. Exquisitely setting the specific rates of individual reactions to desired values is very tricky, and requires a detailed understanding of molecular dynamics. Our scheme avoids this problem since any choice of specific rates that leads to the same equilibrium will do. Hence we can freely set the specific rates so long as the equilibrium constants (ratio of forward and backward specific rates) have value 1. This is an equilibrium thermodynamic condition that is much easier to ensure in vitro. This combination of reversible reactions, no catalysis, and robustness to the values of the specific rates may make this scheme particularly easy and efficient to implement.

– In Definition 8.2, we show how to obtain from the matrix A a reaction network that computes the maximum likelihood estimator. Specialized to our example, we obtain the reaction network with 5 species $X_1, X_2, X_3, \theta_1, \theta_2$ and the 5 reactions:

$$X_1 + X_3 \rightleftharpoons 2X_2, \qquad 2\theta_1 \to 0, \qquad X_1 \to X_1 + 2\theta_1,$$
$$\theta_1 + \theta_2 \to 0, \qquad X_2 \to X_2 + \theta_1 + \theta_2.$$
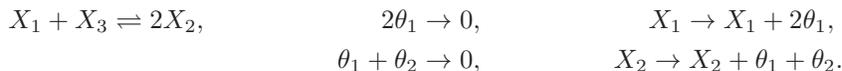
The number of species equals the number of rows plus the number of columns of A. The reactions are not uniquely determined by the problem, but become so once we choose a basis for the kernel of A and a maximal linearly independent set of columns. Here we have chosen columns 1 and 2. Each column of A determines a pair of irreversible reactions.

– Theorem 6 implies that if this reaction system is launched at initial concentrations $X_1(0) = x_1, X_2(0) = x_2, X_3(0) = x_3$ and arbitrary concentrations of $\theta_1(0)$ and $\theta_2(0)$, and the dynamics proceeds according to the law of mass action with all specific rates set to 1:

$$\dot{X}_1(t) = \dot{X}_3(t) = -X_1(t)X_3(t) + X_2^2(t), \quad \dot{\theta}_1(t) = -2\theta_1^2(t) + 2X_1(t) - \theta_1(t)\theta_2(t) + X_2(t),$$
$$\dot{X}_2(t) = -2X_2^2(t) + 2X_1(t)X_3(t), \qquad\qquad \dot{\theta}_2(t) = -\theta_1\theta_2(t) + X_2(t),$$

then the reaction reaches equilibrium $(\hat{x}_1, \hat{x}_2, \hat{x}_3, \hat{\theta}_1, \hat{\theta}_2)$ where $(\hat{\theta}_1, \hat{\theta}_2)$ is the maximum likelihood estimator for the data frequency vector (x_1, x_2, x_3) and $(\hat{x}_1, \hat{x}_2, \hat{x}_3)$ represents the probability distribution over the outcomes X_1, X_2, X_3 at the maximum likelihood. We prove global convergence: our dynamical system provably converges to the desired equilibrium. Global convergence results are known to be notoriously hard to prove in reaction network theory [14].

– A number of schemes have been proposed for translating reaction networks into DNA strand displacement reactions [27,22,6,7]. Adapting these schemes to our setting should allow molecular implementation of our MLE-solving reaction networks with DNA molecules.

2   Maximum Likelihood Estimation in toric models

The definitions and results in this section mostly follow [21]. Because we require a slightly stronger statement, and Theorem 1 allows a short, easy, and insightful proof, we give the proof here for completeness.

In statistics, a **parametric model** consists of a family of probability distributions, one for each value of the parameters. This can be described as a map from a manifold of parameters into a manifold of probability distributions. If this map can be described by monomials as below, then the parametric statistical model is called a **toric** or **log-linear** model, as we now describe.

Definition 1 (Toric Model). *Let m, n be positive integers. The probability simplex and its relative interior are:*

$$\Delta^n := \{(x_1, x_2, \ldots, x_n) \in \mathbb{R}_{\geq 0}^n \mid x_1 + x_2 + \cdots + x_n = 1\}$$

$$\mathrm{ri}(\Delta^n) := \{(x_1, x_2, \ldots, x_n) \in \mathbb{R}_{> 0}^n \mid x_1 + x_2 + \cdots + x_n = 1\}.$$

*An $m \times n$ matrix $A = (a_{ij})_{m \times n}$ of integer entries is a **design matrix** iff all its column sums $\sum_i a_{ij}$ are equal. Let $a_j := (a_{1j}, a_{2j}, \ldots, a_{mj})^T$ be the j'th column of A. Define $\theta^{a_j} := \theta_1^{a_{1j}} \theta_2^{a_{2j}} \ldots \theta_m^{a_{mj}}$. Define the **parameter space** $\Theta := \{\theta \in \mathbb{R}_{> 0}^m \mid \theta^{a_1} + \theta^{a_2} + \cdots + \theta^{a_n} = 1\}$. The **toric model** of A is the map*

$$p_A = (p_1, p_2, \ldots, p_n) : \Theta \to \Delta^n \text{ given by } p_j(\theta) = \theta^{a_j} \text{ for } j = 1 \text{ to } n.$$

We could also have defined the parameter space Θ to be all of $\mathbb{R}_{> 0}^m$, in which case we would need to normalize the probabilities by the *partition function* $\theta^{a_1} + \theta^{a_2} + \cdots + \theta^{a_n}$ to make sure they add up to 1. For our present purposes, the current approach will prove technically more direct.

Note that here $p_j(\theta)$ specifies $\Pr[j \mid \theta]$, the conditional probability of obtaining outcome j given that the true state of the world is described by θ.

A central problem of statistical inference is the problem of **parameter estimation**. After performing several independent identical trials, suppose the **data vector** $u \in \mathbb{Z}_{\geq 0}^n$ is obtained as a record of how many times each outcome occurred. Let the norm $|u|_1 := u_1 + u_2 + \cdots + u_n$ denote the total number of trials performed. The **Maximum Likelihood** solution to the problem of parameter estimation finds that value of the parameter θ which maximizes the **likelihood function** $f_u(\theta) := \Pr[u \mid \theta]$, i.e.:

$$\hat{\theta}(u) := \arg\sup_{\theta \in \Theta} f_u(\theta) \tag{1}$$

is a **maximum likelihood estimator** or MLE for the data vector u. We will call the point $\hat{p}(u) := p_A(\hat{\theta}(u))$ a **maximum likelihood distribution**.

Definition 2. *Let A be an $m \times n$ design matrix, and u a data vector. Then the **sufficient polytope** is $P_A(u) := \{p \in \mathrm{ri}(\Delta^n) \mid Ap = A\frac{u}{|u|_1}\}$.*

The following theorem is a version of Birch's theorem from Algebraic Statistics. It provides a variational characterization of the maximum likelihood distribution as the unique maximum entropy distribution in the sufficient polytope. In particular the maximum likelihood distribution always belongs to the sufficient polytope, which justifies the name.

Theorem 1. *Fix a design matrix A of size $m \times n$.*

1. *If $u, v \in \mathbb{Z}_{\geq 0}^n$ are nonzero data vectors such that $Au/|u|_1 = Av/|v|_1$ then they have the same maximum likelihood estimator: $\hat{\theta}(u) = \hat{\theta}(v)$.*
2. *Further if $P_A(u)$ is nonempty then*
   (a) *There is a unique distribution $\tilde{p} \in P_A(u)$ which maximizes Shannon entropy $H(p) = -\sum_{i=1}^n p_i \log p_i$ viewed as a real-valued function from the closure $\overline{P_A(u)}$ of $P_A(u)$ with $0 \log 0$ defined as 0.*

(b) $\{\tilde{p}\} = P_A(u) \cap p_A(\Theta)$.

(c) $\tilde{p} = \hat{p}(u)$, the Maximum Likelihood Distribution for the data vector u.

Proof. 1. Fix a data vector u. Note that $f_u(\theta) = \frac{|u|_1!}{u_1!u_2!...u_n!}p_1(\theta)^{u_1}p_2(\theta)^{u_2}\ldots p_n(\theta)^{u_n} = \frac{|u|_1!}{u_1!u_2!...u_n!}\theta^{Au}$. Therefore the maximum likelihood estimator

$$\hat{\theta}(u) = \arg\sup_{\theta\in\Theta}\theta^{Au} = \arg\sup_{\theta\in\Theta}(\theta^{Au})^{1/|u|_1} = \arg\sup_{\theta\in\Theta}\theta^{Au/|u|_1}$$

where the second equality is true because the function $x \mapsto x^c$ is monotonically increasing whenever $c > 0$. It follows that if $v \in \mathbb{Z}_{\geq 0}^n$ is a data vector such that $Au/|u|_1 = Av/|v|_1$ then $\hat{\theta}(u) = \hat{\theta}(v)$.

2.(a) Suppose $P_A(u)$ is nonempty. A local maximum of the restriction $H|_{\overline{P_A(u)}}$ of H to the polytope $\overline{P_A(u)}$ can not be on the boundary $\partial\overline{P_A(u)}$ because for $p \in \partial\overline{P_A(u)}$, moving in the direction of arbitrary $q \in P_A(u)$ increases H, as can be shown by a simple calculation:

$$\lim_{\lambda\to 0}\frac{d}{d\lambda}H((1-\lambda)p + \lambda q) \to +\infty.$$

Since H is a continuous function and the closure $\overline{P_A(u)}$ is a compact set, H must attain its maximum value in $P_A(u)$. Further H is a strictly concave function since its Hessian is diagonal with entries $-1/p_i$ and hence negative definite. It follows that $H|_{\overline{P_A(u)}}$ is also strictly concave, and has a unique local maximum at $\tilde{p} \in P_A(u)$, which is also the global maximum.

(b) By concavity of H, the maximum $\tilde{p}$ is the unique point in $P_A(u)$ such that $\nabla H(\tilde{p})$ is perpendicular to $P_A(u)$. We claim that $q \in P_A(u) \cap p_A(\Theta)$ iff $\nabla H(q) = (-1 - \log q_1, -1 - \log q_2, \ldots, -1 - \log q_n)$ is perpendicular to $P_A(u)$. Since all column sums are equal, this is equivalent to requiring that $\log q$ be in the span of the rows of A, which is true iff $q \in p_A(\Theta)$. Hence $P_A(u) \cap p_A(\Theta) = \{\tilde{p}\}$.

(c) To compute the Maximum Likelihood Distribution $\hat{p}(u)$, we proceed as follows:

$$\hat{p}(u) = p_A(\hat{\theta}(u)) = p_A(\arg\sup_{\theta\in\Theta}\theta^{Au}) = p_A(\arg\sup_{\theta\in\Theta}\theta^{Au/|u|_1})$$

$$= p_A(\arg\sup_{\theta\in\Theta}\theta^{A\tilde{p}}) = \arg\sup_{p\in p_A(\Theta)}p^{\tilde{p}} = \arg\sup_{p\in p_A(\Theta)}\sum_{i=1}^{n}\tilde{p}_i\log p_i = \tilde{p}$$

where the fourth equality uses $A\tilde{p} = Au/|u|_1$ and the last equality follows because $\sum_{i=1}^{n}\tilde{p}_i\log p_i$ viewed as a function of p attains its maximum in all of Δ^n, and hence in $p_A(\Theta)$, at $p = \tilde{p}$.

This theorem already exposes the core of our idea. We will design reaction systems that maximize entropy subject to the "correct" constraints capturing the polytope $P_A(u)$. Then because the reactions also proceed to maximize entropy, the equilibrium point of our dynamics will correspond to the maximum likelihood distribution. Most of the technical work will go in proving convergence of trajectories to these equilibrium points.

3 Reaction Networks

According to [20], "In building a design theory for chemistry, chemical reaction networks are usually the most natural intermediate representation - the middle of the hourglass [11]. Many different high

level languages and formalisms have been and can likely be compiled to chemical reactions, and chemical reactions themselves (as an abstract specification) can be implemented with a variety of low level molecular mechanisms."

In Subsection 3.1, we recall the definitions and results for reaction networks which we will need for our main results. For a comprehensive presentation of these ideas, see [13]. In Subsection 3.2, we prove a new result in reaction network theory. We extend a previously known global convergence result to the case of perturbations.

3.1 Brief review of Reaction Network Theory

For vectors $a = (a_i)_{i \in S}$ and $b = (b_i)_{i \in S}$, the notation a^b will be shorthand for the formal monomial $\prod_{i \in S} a_i^{b_i}$. We introduce some standard definitions.

Definition 3 (Reaction Network).
  *Fix a finite set S of **species**.*

1. *A **reaction** over S is a pair (y, y') such that $y, y' \in \mathbb{Z}_{\geq 0}^S$. It is usually written $y \to y'$, with **reactant** y and **product** y'.*
2. *A **reaction network** consists of a finite set S of species, and a finite set $\mathcal{R}$ of reactions.*
3. *A reaction network is **reversible** iff for every reaction $y \to y' \in \mathcal{R}$, the reaction $y' \to y \in \mathcal{R}$.*
4. *A reaction network is **weakly reversible** iff for every reaction $y \to y' \in \mathcal{R}$ there exists a positive integer $n \in \mathbb{Z}_{>0}$ and n reactions $y_1 \to y_2, y_2 \to y_3, \dots, y_{n-1} \to y_n \in \mathcal{R}$ with $y_1 = y'$ and $y_n = y$.*
5. *The **stoichiometric subspace** $H \subseteq \mathbb{R}^S$ is the subspace spanned by $\{y' - y \mid y \to y' \in \mathcal{R}\}$, and $H^\perp$ is the orthogonal complement of H.*
6. *A **siphon** is a set $T \subseteq S$ of species such that for all $y \to y' \in \mathcal{R}$, if there exists $i \in T$ such that $y_i' > 0$ then there exists $j \in T$ such that $y_j > 0$.*
7. *A siphon $T \subseteq S$ is **critical** iff $v \in H^\perp \cap \mathbb{R}_{\geq 0}^S$ with $v_i = 0$ for all $i \notin T$ implies $v = 0$.*

Definition 4. *Fix a weakly reversible reaction network $(S, \mathcal{R})$. The **associated ideal** $I_{(S,\mathcal{R})} \subseteq \mathbb{C}[x]$ where $x = (x_i)_{i \in S}$ is the ideal generated by the binomials $\{x^y - x^{y'} \mid y \to y' \in \mathcal{R}\}$. A reaction network is **prime** iff its associated ideal is a prime ideal.*

The following theorem follows from [13, Theorem 4.1, Theorem 5.2].

Theorem 2. *A weakly reversible prime reaction network $(S, \mathcal{R})$ has no critical siphons.*

We now recall the mass-action equations which are widely employed for modeling cellular processes [29,26,28,30] in Biology.

Definition 5 (Mass Action System). *A **reaction system** consists of a reaction network $(S, \mathcal{R})$ and a **rate function** $k : \mathcal{R} \to \mathbb{R}_{>0}$. The **mass-action equations** for a reaction system are the system of ordinary differential equations in* concentration *variables $\{x_i(t) \mid i \in S\}$:*

$$\dot{x}(t) = \sum_{y \to y' \in \mathbb{R}} k_{y \to y'} \, x(t)^y \, (y' - y) \tag{2}$$

where $x(t)$ represents the vector $(x_i(t))_{i \in S}$ of concentrations at time t.

Note that $\dot{x}(t) \in H$, so affine translations of H are invariant under the dynamics of Equation 2. We recall the well known notions of detailed balanced and complex balanced reaction system.

Definition 6. *A reaction system $(S, \mathcal{R}, k)$ is*

1. ***Detailed balanced** iff it is reversible and there exists a point $\alpha \in \mathbb{R}_{>0}^S$ such that for every $y \to y' \in \mathcal{R}$:*
$$k_{y \to y'}\, \alpha^y\, (y' - y) = k_{y' \to y}\, \alpha^{y'}\, (y - y')$$
   *A point $\alpha \in \mathbb{R}_{>0}^S$ that satisfies the above condition is called a **point of detailed balance**.*
2. ***Complex balanced** iff there exists a point $\alpha \in \mathbb{R}_{>0}^S$ such that for every $y \in \mathbb{Z}_{\geq 0}^S$:*

$$\sum_{y \to y' \in \mathcal{R}} k_{y \to y'}\, \alpha^y\, (y' - y) = \sum_{y'' \to y \in \mathcal{R}} k_{y'' \to y}\, \alpha^{y''}\, (y - y'')$$

   *A point $\alpha \in \mathbb{R}_{>0}^S$ that satisfies the above condition is called a **point of complex balance**.*

The following observations are well known and easy to verify.

– A complex balanced reaction system is always weakly reversible.
– If all rates $k_{y \to y'} = 1$ and the network is weakly reversible then the reaction system is complex balanced with point of complex balance $(1, 1, \ldots, 1) \in \mathbb{R}^S$; if the network is reversible then the reaction system is also detailed balanced with point of detailed balance $(1, 1, \ldots, 1) \in \mathbb{R}^S$.
– Every detailed balance point is also a complex balance point, but there are complex balanced reversible networks that are not detailed balanced.

It is straightforward to check that every point of complex balance (respectively, detailed balance) is a fixed point for Equation 2. The next theorem, which follows from [2, Theorem 2] and [15], states that a converse also exists: if a reaction system is complex balanced (respectively, detailed balanced) then every fixed point is a point of complex balance (detailed balance). Further there is a unique fixed point in each affine translation of H, and if there are no critical siphons then the basin of attraction for this fixed point is as large as possible, namely the intersection of the affine translation of H with the nonnegative orthant.

Theorem 3 (Global Attractor Theorem for Complex Balanced Reaction Systems with no critical siphons). *Let $(S, \mathcal{R}, k)$ be a weakly reversible complex balanced reaction system with no critical siphons and point of complex balance α. Fix a point $u \in \mathbb{R}_{>0}^S$. Then there exists a point of complex balance β in $(u + H) \cap \mathbb{R}_{>0}^S$ such that for every trajectory $x(t)$ with initial conditions $x(0) \in (u + H) \cap \mathbb{R}_{>0}^S$, the limit $\lim_{t \to \infty} x(t)$ exists and equals β. Further the function $g(x) := \sum_{i=1}^n x_i \log x_i - x_i - x_i \log \alpha_i$ is strictly decreasing along non-stationary trajectories and attains its unique minimum value in $(u + H) \cap \mathbb{R}_{\geq 0}^S$ at β.*

It is not completely trivial to show, but nevertheless true, that this theorem holds with weakly reversible replaced by "reversible" and "complex balance" replaced by "detailed balance." What is to be shown is that the point of complex balance obtained in $(u + H) \cap \mathbb{R}_{\geq 0}^S$ by minimizing $g(x)$ is actually a point of detailed balance, and this follows from an examination of the form of the derivative $\frac{d}{dt} g(x(t))$ along trajectories $x(t)$ to Equation 2.

3.2 A Perturbatively-Stable Global Attractor Theorem

Global attractor results usually assume that the reaction network is weakly reversible. We are going to describe our scheme in the next section. Our scheme will employ reaction networks that are not

weakly reversible, yet we will prove global attractor results for them. The key idea we use is that our reaction network can be broke into a reversible part, and an irreversible part. The reversible part acts on, but evolves independent of, the irreversible part. So we get to use the global attractor results "as is" on the reversible part. Further, as the reversible part approaches equilibrium, our irreversible part behaves as a perturbation of a reversible detailed-balanced network. The closer the reversible part gets to equilibrium, the smaller the perturbation of the irreversible part from the dynamics of a certain reversible detailed-balanced network.

To make this proof idea work out, we will need a perturbative version of Theorem 3. The next lemma shows that if the rates are perturbed slightly then, outside a small neighborhood of the detailed balance point, the strict Lyapunov function $g(x)$ from Theorem 3 continues to decrease along non-stationary trajectories.

Lemma 1. *Let $(S, \mathcal{R}, k)$ be a weakly reversible complex balanced reaction system with no critical siphons and point of complex balance α. For every sufficiently small $\epsilon > 0$ there exists $\delta > 0$ such that for all x' outside the ϵ-neighborhood of α in $(\alpha + H) \cap \mathbb{R}_{\geq 0}^S$, the derivative $\frac{d}{dt} g(x(t))|_{t=0} < -\delta$, where $x(t)$ is a solution to the Mass-Action Equations 2 with $x(0) = x'$.*

Proof. Let B_ϵ be the open ϵ ball around α in $(\alpha + H) \cap \mathbb{R}_{\geq 0}^S$, with ϵ small enough so that B_ϵ does not meet the boundary $\partial \mathbb{R}_{\geq 0}^S$. Consider the closed set $\bar{S} := (\alpha + H) \cap \mathbb{R}_{\geq 0}^S \setminus B_\epsilon$. Define the orbital derivative of g at x' as $\mathcal{O}_k g(x') := \frac{d}{dt} g(x(t))|_{t=0}$, where $x(t)$ is a solution to the mass-action equations 2 with $x(0) = x'$. Define $\delta := \inf_{x' \in \bar{S}} (-\mathcal{O}_k g(x'))$. If $\delta \leq 0$ then since $\bar{S}$ is a closed set, and $\mathcal{O}_k g$ is a continuous function, there exists a point x' such that $\mathcal{O}_k g(x') \geq 0$, which contradicts Theorem 3. ∎

We formalize the notion of perturbation using **differential inclusions**. Recall that differential inclusions model uncertainty in dynamics in a nondeterministic way by generalizing the notion of vector field. A differential inclusion maps every point to a subset of the tangent space at that point.

Definition 7. *Let $(S, \mathcal{R}, k)$ be a reaction system and let $\delta > 0$. The δ-**perturbation** of $(S, \mathcal{R}, k)$ is the differential inclusion $V : \mathbb{R}_{\geq 0}^S \to 2^{\mathbb{R}^S}$ that at point $x \in \mathbb{R}_{\geq 0}^S$ takes the value*

$$V(x) := \left\{ \sum_{y \to y' \in \mathcal{R}} k'_{y \to y'} x^y (y' - y) \;\middle|\; k'_{y \to y'} \in (k_{y \to y'} - \delta, k_{y \to y'} + \delta) \text{ for all } y \to y' \in \mathcal{R} \right\}.$$

*A **trajectory** of V is a tuple (I, x) where $I \subseteq \mathbb{R}$ is an interval and $x : I \to \mathbb{R}_{\geq 0}^S$ is a differentiable function with $\dot{x}(t) \in V(x(t))$.*

Theorem 4 (Perturbatively-Stable Global Attractor Theorem for Complex Balanced Reaction Systems with no critical siphons). *Let $(S, \mathcal{R}, k)$ be a weakly reversible complex balanced reaction system with no critical siphons. Fix a point $u \in \mathbb{R}_{>0}^S$. Then there exists a point of complex balance β in $(u + H) \cap \mathbb{R}_{>0}^S$ such that:*

1. *For every sufficiently small $\varepsilon > 0$, there exists $\delta > 0$ such that every trajectory of the form $(\mathbb{R}_{\geq 0}, x)$ to the δ-perturbation of $(S, \mathcal{R}, k)$ with initial conditions $x(0) \in (u + H) \cap \mathbb{R}_{\geq 0}^S$ eventually enters an ε-neighborhood of β and never leaves.*
2. *Consider a sequence $\delta_1 > \delta_2 > \cdots > 0$ and a sequence $0 < t_1 < t_2 < \dots$ such that $\lim_{i \to \infty} \delta_i = 0$ and $\lim_{i \to \infty} t_i = +\infty$, and a trajectory $(\mathbb{R}_{\geq 0}, x)$ with $x(0) \in (u + H) \cap \mathbb{R}_{\geq 0}^S$ such that $((t_i, \infty), x)$ is a trajectory of the δ_i-perturbation of $(S, \mathcal{R}, k)$. Then the limit $\lim_{t \to \infty} x(t) = \beta$.*

Proof (Proof sketch). 1. Fix $\varepsilon > 0$ such that the ε-ball B_ε around β does not meet the boundary $\partial \mathbb{R}^S_{\geq 0}$. By Lemma 1, outside B_ε, there exists $\delta_\varepsilon > 0$ such that the function $\mathcal{O}_k g < -\delta_\varepsilon$. Since $\mathcal{O}_k g$ is a continuous function of the specific rates k, a sufficiently small perturbation $\delta > 0$ in the rates will not change the sign of $\mathcal{O}_k g$. Hence, outside B_ϵ, the function g is strictly decreasing along trajectories $x(t)$ to Equation 2. It follows that eventually every trajectory must enter B_ϵ.

2. Fix a sequence $\varepsilon_1 > \varepsilon_2 > \cdots > 0$ with ε_1 small enough so that the ε_1-ball around β does not meet the boundary $\partial \mathbb{R}^S_{\geq 0}$ and $\lim_{i \to \infty} \varepsilon_i \to 0$. For each ε_i, there exists j such that δ_j is small enough as per part (1) of the theorem. So every trajectory will eventually enter the ϵ_i neighborhood of β, and never leave. Since this is true for every i and $\lim_{i \to \infty} \varepsilon_i \to 0$, the result follows.

4   Main Result

The next definition makes precise our scheme, which takes a design matrix A to a reaction system $\mathcal{S}_{MLE}$ depending on A. The choice of this reaction system is not unique, but depends on two choices of basis. We proceed in two stages. In the first stage, we construct the reaction system $\mathcal{S}_{MLD}$ which solves the problem of finding the maximum likelihood distribution. In the second stage, we add reactions to solve for θ from the algebraic relations between the θ and X variables, obtaining $\mathcal{S}_{MLE}$.

Definition 8. *Fix a design matrix $A = (a_{ij})_{m \times n}$, a basis B for the free group $\mathbb{Z}^n \cap \ker A$, and a maximal linearly-independent subset B' of the columns of A.*

1. *The reaction network $\mathcal{R}_{MLD}(A, B)$ consists of n species $X_1, X_2, \ldots, X_n$ and for each $b \in B$, the reversible reaction:*
$$\sum_{j:b_j > 0} b_j X_j \rightleftharpoons \sum_{j:b_j < 0} -b_j X_j$$
2. *The reaction system $\mathcal{S}_{MLD}(A, B)$ consists of the reaction network $\mathcal{R}_{MLD}(A, B)$ with an assignment of rate 1 to each reaction.*
3. *The reaction network $\mathcal{R}_{MLE}(A, B, B')$ consists of $m + n$ species $\theta_1, \theta_2, \ldots, \theta_m, X_1, X_2, \ldots, X_n$, and in addition to the reactions in $\mathcal{R}_{MLD}$, the following reactions:*
   – *For each column $j \in B'$ of A, a reaction $\sum_{i=1}^m a_{ij}\theta_i \to 0$.*
   – *For each column $j \in B'$ of A, a reaction $X_j \to X_j + \sum_{i=1}^m a_{ij}\theta_i$.*
4. *The reaction system $\mathcal{S}_{MLE}(A, B, B')$ consists of the reaction network $\mathcal{R}_{MLE}(A, B, B')$ with an assignment of rate 1 to each reaction.*

Note that by the rank-nullity theorem of linear algebra, the dimension of the kernel plus the rank of the matrix equals the number of columns of the matrix. Hence counting the reversible reactions as two irreversible reactions, our scheme yields a reaction system whose number of reactions is twice the number of columns of A.

It is clear from the definition of $\mathcal{S}_{MLE}$ that the reactions that come from $\mathcal{R}_{MLD}$ are reversible and evolve without being affected by the other reactions. Hence we first prove global convergence of the reaction system $\mathcal{S}_{MLD}$ to the maximum likelihood distribution. This part is fairly straightforward. The key point is to verify that the reaction network $\mathcal{R}_{MLD}$ has no critical siphons. In fact, we show in the next lemma that $\mathcal{R}_{MLD}$ is prime, which will imply "no critical siphons" by Theorem 2.

Lemma 2. *Fix a design matrix $A = (a_{ij})_{m \times n}$ and a basis B for the free group $\mathbb{Z}^n \cap \ker A$. Then the reaction network $\mathcal{R}_{MLD}(A, B)$ is prime and $\mathcal{S}_{MLD}(A, B)$ is detailed balanced. Consequently, the reaction system $\mathcal{S}_{MLD}(A, B)$ is globally asymptotically stable.*

9

Proof. $\mathcal{R}_{MLD}(A, B)$ is prime by [17, Corollary 2.15]. The idea is to look at the toric model p_A as a ring homomorphism $\mathbb{C}[x_1, x_2, \ldots, x_n] \to \mathbb{C}[\mathbb{N}A]$ with $x_j \mapsto \theta^{a_j}$. (Here $\mathbb{N}A$ is the affine semigroup generated by the columns of A.) The kernel of this ring homomorphism is the associated ideal of $\mathcal{R}_{MLD}(A, B)$ by [17, Proposition 2.14], and the codomain is an integral domain, so the kernel must be prime.

To verify that $\mathcal{S}_{MLD}(A, B)$ is detailed balanced, note that the point $(1, 1, \ldots, 1) \in \mathbb{R}^n$ is a point of detailed balance since all rates are 1. Global asymptotic stability now follows from Theorem 2 and Theorem 3.

We can now obtain global convergence for $\mathcal{S}_{MLD}$.

Theorem 5 (The reaction system $\mathcal{S}_{MLD}(A, B)$ computes the Maximum Likelihood Distribution). *Fix a design matrix $A = (a_{ij})_{m \times n}$, a basis B for the free group $\mathbb{Z}^n \cap \ker A$, and a nonzero data vector $u \in \mathbb{Z}_{\geq 0}^n$. Let $x(t) = (x_1(t), x_2(t), \ldots, x_n(t))$ be a solution to the mass-action differential equations for the reaction system $\mathcal{S}_{MLD}(A, B)$ with initial conditions $x(0) = u/|u|_1$. Then $x(\infty) := \lim_{t \to \infty} x(t)$ exists and equals the maximum likelihood distribution $\hat{p}(u)$.*

Proof. For the system $\mathcal{S}_{MLD}(A, B)$, note that $(x(0) + H) \cap \mathbb{R}_{>0}^n = P_A(u/|u|_1)$. By Theorem 3, $x(\infty)$ exists, and the function $\sum_{i=1}^n x_i \log x_i - x_i - x_i \log 1$ attains its unique minimum in $P_A(u/|u|_1)$ at $x(\infty)$. Since the system is mass-conserving, $\sum_{i=1}^n x_i$ is constant on $P_A(u/|u|_1)$, so this is equivalent to the fact that Shannon entropy $H(x) = -\sum_{i=1}^n x_i \log x_i$ is increasing, and attains its unique maximum value in $P_A(u/|u|_1)$ at $x(\infty)$. By Theorem 1, the point $x(\infty)$ must be the maximum likelihood distribution $\hat{p}(u)$.

As the reversible reactions in $\mathcal{S}_{MLE}$ approach closer and closer to equilibrium, we wish to absorb the values of the X variables into reaction rates and pretend that the irreversible reactions are reactions only in the θ variables. This has the advantage that we can treat this pretend reaction system in the θ variables as a perturbation of a reversible, detailed balanced system. We can then hope to employ Theorem 4 and conclude global convergence for these irreversible reactions, and hence for $\mathcal{S}_{MLE}$.

One small technical point deserves mention. The pretend reaction system in the θ variables is not a reaction system since the rates are not real numbers but functions of time. This will not trouble us. We have already provisioned for this in Definition 7 by allowing perturbations of reaction systems to be differential inclusions.

Theorem 6 (The reaction system $\mathcal{S}_{MLE}(A, B, B')$ computes the Maximum Likelihood Estimator). *Fix a design matrix $A = (a_{ij})_{m \times n}$, a basis B for the free group $\mathbb{Z}^n \cap \ker A$, and a nonzero data vector $u \in \mathbb{Z}_{\geq 0}^n$. Let $x(t) = (x_1(t), x_2(t), \ldots, x_n(t), \theta_1(t), \theta_2(t), \ldots, \theta_m(t))$ be a solution to the mass-action differential equations for the reaction system $\mathcal{S}_{MLE}(A, B, B')$ with initial conditions $x(0) = u/|u|_1$ and $\theta(0) = 0$. Then $x(\infty) := \lim_{t \to \infty} x(t)$ exists and equals the maximum likelihood distribution $\hat{p}(u)$, and $\theta(\infty) := \lim_{t \to \infty} \theta(t)$ exists and equals the maximum likelihood estimator $\hat{\theta}(u)$.*

Proof (Proof sketch). Fix u and let $\hat{p} = \hat{p}(u)$ and $\hat{\theta} = \hat{\theta}(u)$. Note that for the species $X_1, X_2, \ldots, X_n$, the differential equations for $\mathcal{S}_{MLE}(A, B)$ and $\mathcal{S}_{MLD}(A, B, B')$ are identical, since these species appear purely catalytically in the reactions that belong to $\mathcal{R}_{MLE}(A, B, B') \setminus \mathcal{R}_{MLD}(A, B)$. Hence $x(\infty) = \hat{p}(u)$ follows from Theorem 5.

To see that $\theta(\infty) = \hat{\theta}$, let us first allow the X species to reach equilibrium, then treat the θ system with replacing the X species by rate constants representing their values at equilibrium. The

system $\Theta_{MLE}(A, B, B', x(\infty))$ obtained in this way in only the θ species is a reaction system with the reactions

- For each column $j \in B'$ of A, a reaction $\sum_{i=1}^{m} a_{ij}\theta_i \to 0$ of rate 1
- For each column $j \in B'$ of A, a reaction $0 \to \sum_{i=1}^{m} a_{ij}\theta_i$ of rate $x_j(\infty)$.

This is a reversible reaction system, and the maximum likelihood estimators $\hat{\theta}$ are precisely the points of detailed balance for this system, where we are using the fact that B' was a maximal linearly-independent set of the columns of A. In addition, this system has no siphons since if species θ_i is absent, and $a_{ij} > 0$ then θ_i will immediately be produced by the reaction $0 \to \sum_{i'=1}^{m} a_{i'j}\theta_{i'}$. (We are assuming A has no 0 row. If A has a 0 row, we can ignore it anyway.) It follows from Theorem 3 that this system is globally asymptotically stable, and every trajectory approaches a maximum likelihood estimator $\hat{\theta}$.

Our actual system may be viewed as a perturbation of the system $\Theta_{MLE}(A, B, B', x(\infty))$. Consider any trajectory $(x(t), \theta(t))$ to $\mathcal{S}_{MLE}(A, B, B')$ starting at $(u/|u|_1, 0)$. We are going to consider the projected trajectory $(\mathbb{R}_{\geq}, \theta)$. We now show that it is possible to choose appropriate t_i and δ_i so that $((t_i, \infty), \theta(t))$ is a trajectory of a δ_i-perturbation of $\Theta_{MLE}(A, B, B', x(\infty))$, for $i = 1, 2, \ldots$.

Wait for a sufficiently large time t_1 till $x(t)$ is in a sufficiently small δ_1 neighborhood of $x(\infty)$ which it will never leave. After this time, we obtain a differential inclusion in the θ species with the mass-action equations 2 for the reactions

- For each column j of A, a reaction $\sum_{i=1}^{m} a_{ij}\theta_i \to 0$ of rate 1
- For each column j of A, a reaction $0 \to \sum_{i=1}^{m} a_{ij}\theta_i$ with time-varying rate lying in the interval $(x_j(\infty) - \delta_1, x_j(\infty) + \delta_1)$.

Continuing in this way, we choose a decreasing sequence $\delta_1 > \delta_2 > \cdots > 0$ with $\lim_{i \to \infty} \delta_i \to 0$, and corresponding times $t_1 < t_2 < t_3 \ldots$ with $\lim_{i \to \infty} t_i \to \infty$ such that after time t_i, $x(t)$ is in a δ_i neighborhood of $x(\infty)$ which it will never leave. Then $((t_i, \infty), \theta(t))$ is a trajectory of the δ_i-perturbation of $\Theta_{MLE}(A, B, B', x(\infty))$. Hence $\theta(t)$ satisfies the conditions of Theorem 4. Hence $\lim_{t \to \infty} \theta(t) = \hat{\theta}$.

5   Related Work and Conclusions

The mathematical similarities of both log-linear statistics and reaction networks to toric geometry have been pointed out before [9,17]. Craciun et al. [9] refer to the steady states of complex-balanced reaction networks as *Birch points* "to highlight the parallels" with algebraic statistics. This paper develops on these observations, and serves to flesh out this mathematical parallel into a scheme for molecular computation.

Various building blocks for molecular computation that assume mass-action kinetics have been proposed before. We briefly review some of these proposals.

In [19], Napp and Adams model molecular computation with mass-action kinetics, as we do here. They propose a molecular scheme to implement message passing schemes in probabilistic graphical models. The goal of their scheme is to convert a factor graph into a reaction network that encodes the single-variable marginals of the joint distribution as steady state concentrations. In comparison, the goal of our scheme is to do statistical inference and compute maximum likelihood estimators for log-linear models. Napp and Adams focus on the "forward model" task of how a given data-generating process (a factor graph) can lead to observed data, whereas our focus is on the "backward model"

task of inference, going from the observed data to the data-generating process. Further our scheme couples the deep role that MaxEnt algorithms play in Machine Learning with MaxEnt's roots in the Second Law of Thermodynamics whereas Napp and Adams are drawing their inspiration from variable elimination implemented via message passing which has its roots in Boolean constraint satisfaction problems.

Qian and Winfree [23,24] have proposed a DNA gate motif that can be composed to build large circuits, and have experimentally demonstrated molecular computation of a Boolean circuit with around 30 gates. In comparison, our scheme natively employs a continuous-time dynamical system to do the computation, without a Boolean abstraction.

Taking a control theory point of view, Oishi and Klavins [20] have proposed a scheme for implementing linear input/output systems with reaction networks. Note that for a given matrix A, the set of maximum likelihood distributions is usually not linear, but log-linear.

Daniel et al.[10] have demonstrated an in vivo implementation of feedback loops, exploiting analogies with electronic circuits. It is possible that the success of their schemes is also related to the toric nature of mass-action kinetics.

Buisman et al. [5] have proposed a reaction network scheme for computation of algebraic functions. The part of our scheme which reads out the maximum likelihood estimator from the maximum likelihood distribution bears some similarity to their work.

One limitation of our present work is that the number of columns of the matrix A can become very large, for example $2^{|V|}$ for a graphical model with V nodes. Since the number of species and number of reactions both depend on the number of columns of A, this can require an exponentially large reaction network which may become impractical. One direction for future work is to extend our scheme by specifying a reaction network that computes maximum likelihood for graphical models.

We have some freedom in our scheme in the choice of basis sets B and B'. In any chemical implementation of this work, there might be opportunity for optimization in choice of basis.

References

1. Alan Agresti. *Categorical Data Analysis*. John Wiley & Sons, 2013.
2. David Angeli, Patrick De Leenheer, and Eduardo D. Sontag. A Petri net approach to the study of persistence in chemical reaction networks. *Math. Biosci.*, 210(2):598–618, 2007.
3. Yaakov Benenson, Binyamin Gil, Uri Ben-Dor, Rivka Adar, and Ehud Shapiro. An autonomous molecular computer for logical control of gene expression. *Nature*, 429(6990):423–429, 2004.
4. YMM Bishop, Stephen Feinberg, and Paul Holland. Discrete multivariant analysis. 1975.
5. HJ Buisman, Huub MM ten Eikelder, Peter AJ Hilbers, and Anthony ML Liekens. Computing algebraic functions with biochemical reaction networks. *Artificial life*, 15(1):5–19, 2009.
6. Luca Cardelli. Strand Algebras for DNA Computing. *Natural Computing*, 10:407–428, 2011.
7. Luca Cardelli. Two-domain DNA strand displacement. *Mathematical Structures in Computer Science*, 23(02):247–271, 2013.
8. Ronald Christensen and R Christensen. *Log-linear models and logistic regression*, volume 168. Springer New York, 1997.
9. Gheorghe Craciun, Alicia Dickenstein, Anne Shiu, and Bernd Sturmfels. Toric dynamical systems. *Journal of Symbolic Computation*, 44(11):1551–1565, 2009. In Memoriam Karin Gatermann.

10. Ramiz Daniel, Jacob R Rubens, Rahul Sarpeshkar, and Timothy K Lu. Synthetic analog computation in living cells. *Nature*, 497(7451):619–623, 2013.

11. John Doyle and Marie Csete. Rules of engagement. *Nature*, 446(7138):860–860, 2007.

12. Stephen E Fienberg, Alessandro Rinaldo, et al. Maximum likelihood estimation in log-linear models. *The Annals of Statistics*, 40(2):996–1023, 2012.

13. Manoj Gopalkrishnan. Catalysis in reaction networks. *Bulletin of Mathematical Biology*, 73:2962–2982, 2011. 10.1007/s11538-011-9655-3.

14. Manoj Gopalkrishnan, Ezra Miller, and Anne Shiu. A geometric approach to the global attractor conjecture. In preparation.

15. Friedrich J. M. Horn. The dynamics of open reaction systems. In *Mathematical aspects of chemical and biochemical problems and quantum chemistry*, volume VIII of *Proc. SIAM-AMS Sympos. Appl. Math.*, New York, 1974.

16. Steffen L Lauritzen. *Graphical models*. Oxford University Press, 1996.

17. Ezra Miller. Theory and applications of lattice point methods for binomial ideals. In *Combinatorial Aspects of Commutative Algebra and Algebraic Geometry*, pages 99–154. Springer, 2011.

18. Ron Milo, Shai Shen-Orr, Shalev Itzkovitz, Nadav Kashtan, Dmitri Chklovskii, and Uri Alon. Network motifs: simple building blocks of complex networks. *Science*, 298(5594):824–827, 2002.

19. Nils E Napp and Ryan P Adams. Message passing inference with chemical reaction networks. In *Advances in Neural Information Processing Systems*, pages 2247–2255, 2013.

20. Kevin Oishi and Eric Klavins. Biomolecular implementation of linear I/O systems. *Systems Biology, IET*, 5(4):252–260, 2011.

21. L. Pachter and B. Sturmfels. *Algebraic Statistics for Computational Biology*. Number v. 13 in Algebraic Statistics for Computational Biology. Cambridge University Press, 2005.

22. Lulu Qian, David Soloveichik, and Erik Winfree. Efficient turing-universal computation with DNA polymers. In *DNA computing and molecular programming*, pages 123–140. Springer, 2011.

23. Lulu Qian and Erik Winfree. A simple DNA gate motif for synthesizing large-scale circuits. *J. R. Soc. Interface*, 2011.

24. Lulu Qian and Erik Winfree. Scaling up digital circuit computation with DNA strand displacement cascades. *Science*, 332(6034):1196–1201, 2011.

25. Ehud Shapiro and Yaakov Benenson. Bringing DNA computers to life. *Scientific American*, 294(5):44–51, 2006.

26. Guy Shinar and Martin Feinberg. Structural sources of robustness in biochemical reaction networks. *Science*, 327(5971):1389–1391, 2010.

27. David Soloveichik, Georg Seelig, and Erik Winfree. DNA as a universal substrate for chemical kinetics. *Proceedings of the National Academy of Sciences*, 107(12):5393–5398, 2010.

28. Eduardo D. Sontag. Structure and stability of certain chemical networks and applications to the kinetic proofreading model of T-cell receptor signal transduction. *IEEE Trans. Autom. Control*, 46:1028–1047, 2001.

29. Matthew Thomson and Jeremy Gunawardena. Unlimited multistability in multisite phosphorylation systems. *Nature*, 460(7252):274–277, 2009.

30. John J Tyson, Katherine C Chen, and Bela Novak. Sniffers, buzzers, toggles and blinkers: dynamics of regulatory and signaling pathways in the cell. *Current opinion in cell biology*, 15(2):221–231, 2003.

31. Boyan Yordanov, Jongmin Kim, Rasmus L Petersen, Angelina Shudy, Vishwesh V Kulkarni, and Andrew Phillips. Computational design of nucleic acid feedback control circuits. *ACS synthetic biology*, 3(8):600–616, 2014.