

Bounding the equilibrium distribution of Markov population models

Tuğrul Dayar¹, Holger Hermanns², David Spieler² and Verena Wolf^{2,*},[†]

¹*Bilkent University Department of Computer Engineering, TR-06800 Bilkent, Ankara, Turkey*

²*Computer Science, Saarland University, Saarbrücken, Germany*

SUMMARY

We propose a bounding technique for the equilibrium probability distribution of continuous-time Markov chains with population structure and infinite state space. We use Lyapunov functions to determine a finite set of states that contains most of the equilibrium probability mass. Then we apply a refinement scheme based on stochastic complementation to derive lower and upper bounds on the equilibrium probability for each state within that set. To show the usefulness of our approach, we present experimental results for several examples from biology. Copyright © 2011 John Wiley & Sons, Ltd.

Received 1 November 2010; Revised 10 May 2011; Accepted 4 June 2011

KEY WORDS: geometric bounds; stochastic complement; Lyapunov function; equilibrium probability distribution; continuous-time Markov chain

1. INTRODUCTION

Markov population models (MPMs) are continuous-time Markov chains that describe the dynamics and interactions of different populations. They have important applications in the life science domain, in particular in ecology, epidemics, and biochemistry. Depending on the system under study, a member of a population represents an individual that belongs to a certain biological species, an organism that suffers from an infectious disease, or a certain type of molecule in a living cell. Thus, if n is the number of distinct populations, the state space of the MPM is a subset of \mathbb{Z}_+^n , that is, a state is a vector $x = [x_1, \dots, x_n]$ of nonnegative integers, where the entry x_i is the size of the i th population. Typically, the transitions of an MPM are induced by a finite set of change vectors [1–3], here called transition classes [4, 5] such that each class specifies a possibly infinite number of edges in the underlying transition graph. For instance, we may have one class to represent the death of individuals. In biochemistry, each chemical reaction describes a transition class in the associated MPM. Often, the corresponding transition rates are state dependent, for example, the rate at which individuals of a certain population die depend on the population size.

The structural regularities of MPMs often enable accurate approximations of the model behavior. One such example is the widely used deterministic approximation of the dynamics of chemical reaction networks [6] representing the originally discrete state space as a continuum. If, however, one or more populations are small, a discrete representation of the population sizes becomes crucial because continuous approximations are often inaccurate [7]. This effect has also been observed experimentally in the context of chemical reactions [8–10]. The analysis of such MPMs is difficult because closed-form solutions are only possible in very particular and simplistic cases [11], whereas

*Correspondence to: Verena Wolf, Computer Science, Saarland University, Saarbrücken, Germany.

[†]E-mail: wolf@cs.uni-saarland.de

numerical solution techniques suffer from the problem that a very large or even infinite state space has to be explored. Therefore, Monte Carlo simulation is in widespread use [12] to estimate transient or equilibrium measures of MPMs.

Recently, progress has been made on numerically approximating the transient probability distribution of an MPM at particular time instants [13–15]. These approaches provide an approximate solution of the Kolmogorov differential equations and exploit the fact that only a relatively small subset of the state space is necessary for an accurate approximation. They concentrate on the most relevant states and perform a dynamical truncation of the state space that is found during the computation. For the long-run behavior of the system, however, a truncation of the state space is only possible if it is a priori known where the system stabilizes. Otherwise, attracting regions with a significant amount of probability might be neglected due to the truncation.

In this paper, we consider the problem of computing an accurate approximation of the equilibrium probability distribution of an ergodic MPM. We first derive geometric bounds for the equilibrium distribution, that is, we find those regions of the state space, where most of the probability mass is located in the limit. Then we perform a local refinement for these regions in order to bound the probabilities of individual states.

Let $\{X(t), t \geq 0\}$ be an ergodic MPM with state space $S \subseteq \mathbb{Z}_+^n$. The first step of our approach relies on the analysis of a drift function d that maps to each state x the expected change $d(x) = \frac{d}{dt}E[g(X(t)) | X(t) = x]$ with respect to a specific function g , assigning to each x a nonnegative real number. If the function g , called Lyapunov function, is a witness of ergodicity, then the drift is close to zero near attracting regions. We illustrate this by means of an example. Figure 1 shows the phase plot of the equilibrium distribution of an MPM that describes a gene regulatory network. The variables M and P represent the number of mRNA molecules and protein molecules, respectively. The higher the probability of a state, the darker its color in the plot. The system is monostable, that is, in the long-run the probability mass is mainly concentrated at one distinct region in S . In this region, the drift $d(x)$ is close to zero. For higher values of M and P , say $M > 800$ and $P > 500$, the drift is much smaller than zero. We determine geometric bounds (here depicted as a dashed line) by using a simple threshold on the drift, that is, we consider the set $W \subset S$ of all states where the drift is greater than a certain negative threshold. In other words, W contains those states, where most of the probability mass is located in equilibrium.

In the second step of our approach, we consider the stochastic complement of W [16], which allows us to derive bounds for the probabilities of individual states within W . The stochastic complement is a finite Markov chain with state space W , where each outgoing transition leading to a state not in W is redirected to W . This redirection is defined in such a way that the equilibrium probability distribution of the stochastic complement gives the conditional equilibrium probabilities




Figure 1. The equilibrium distribution of an MPM that describes a gene regulatory network and geometric bounds (dashed line) enclosing a subset W of the state space S with an equilibrium probability larger than 0.95.

for the states in W . Because the exact redirection probability can only be obtained from a complete solution of the infinite system, we compute upper and lower bounds on the conditional equilibrium probabilities based on the results developed by Courtois and Semal [17]. For this, we derive a number of linear equation systems that we solve using LU factorization. We exploit the structure of the MPM to reduce the number of equation systems to be solved. Together with the first step of the approach, this yields upper and lower bounds for the equilibrium probabilities of all states in W . To the best of our knowledge, this is the first technique that provides accurate bounds for the equilibrium probabilities of MPMS.


The paper is structured as follows: in the next section, we present the geometric bounding technique to compute a set W of states where most of the equilibrium probability mass is located. We then introduce the stochastic complement for MPMS in Section 3 and adapt the state-wise bounding technique in Section 4. We present a detailed case study from biology in Section 5 and give further exemplary numerical results in Section 6. We conclude the paper with Section 7.

Throughout this paper, we use capital letters for sets and boldface capital letters for matrices. We define matrices and vectors by square brackets. We indicate entries of matrices and vectors using subscripts. The lowercase letters x, y, z, w represent state vectors. Probability vectors are row vectors. Furthermore, \mathbf{I} is the identity matrix, $\mathbf{0}$ is the zero matrix, e is the column vector of 1s, e_i is the i th unit column vector, and T denotes matrix transposition.

2. GEOMETRIC BOUNDS

We describe the transitions of an MPM by transition classes. A transition class τ is a pair (α, v) , where $\alpha : S \rightarrow \mathbb{R}_{\geq 0}$ is a function that determines the (time-independent) transition rate and $v \in \mathbb{Z}^n$ is a change vector that determines the successor state of the transition. Thus, if $x \in S$ and $\alpha(x) > 0$, then there is a transition from state x to state $(x + v)$ with rate $\alpha(x)$. We assume that v has at least one nonzero entry and that $\alpha(x)$ is a polynomial in $x = [x_1, \dots, x_n]$. A MPM may be specified by a set $\{\tau_1, \dots, \tau_J\}$ of transition classes.

When the MPM represents a network of chemical reactions, each reaction corresponds to one transition class. Chemical reactions are usually specified by stoichiometric equations of the form



where the components of $s = [s_1, \dots, s_n]$, $r = [r_1, \dots, r_n] \in \mathbb{Z}_+^n$ are stoichiometric coefficients, P_1, \dots, P_n represent types of chemical species (or molecules), and $\lambda > 0$ is the reaction rate constant. The coefficients s_i and r_i respectively denote the number of P_i species that is consumed and produced by the reaction for $i \in \{1, \dots, n\}$. If $x = [x_1, \dots, x_n]$ is the current state of the underlying MPM, then the rate $\alpha(x)$ at which this reaction occurs is proportional to the possible number of reactant combinations, that is, if $s_i \leq x_i$ for all $i \in \{1, \dots, n\}$, then

$$\alpha(x) = \lambda \prod_{i=1}^n \frac{x_i!}{s_i!(x_i - s_i)!}$$

and $\alpha(x) = 0$ otherwise. The corresponding transition class (α, v) has the change vector $v = r - s$.

Example 1

We present a network of chemical reactions where a gene is transcribed into mRNA (species M) and the mRNA is translated into a protein (species P) [8]. Both species can degrade. The corresponding stoichiometric equations and the transition classes are listed in Table I. The state space is given by $S = \mathbb{Z}_+^n$ with $n = 2$, where for a state $x = [x_1, x_2] \in S$ the entry x_1 encodes the number of molecules of type M and x_2 encodes the number of molecules of type P .

For a given set of transition classes $\{\tau_1, \dots, \tau_J\}$, we define the infinite infinitesimal generator matrix \mathbf{Q} underlying the MPM $\{X(t), t \geq 0\}$ such that the entry $\mathbf{Q}_{x, x+v}$ equals $\sum_{\{j | v^{(j)}=v\}} \alpha_j(x)$, where $\tau_j = (\alpha_j, v^{(j)})$ and $j \in \{1, \dots, J\}$. Moreover, we let the diagonal entries of \mathbf{Q} be equal to the

Table I. Transition classes of the gene expression model.

j	Stoichiometric equation	τ_j	$\alpha_j(x)$	$v^{(j)}$
1	$\emptyset \xrightarrow{\lambda} M$	$(\alpha_1, v^{(1)})$	λ	e_1^T
2	$M \xrightarrow{\mu} M + P$	$(\alpha_2, v^{(2)})$	μx_1	e_2^T
3	$M \xrightarrow{\delta_M} \emptyset$	$(\alpha_3, v^{(3)})$	$\delta_M x_1$	$-e_1^T$
4	$P \xrightarrow{\delta_P} \emptyset$	$(\alpha_4, v^{(4)})$	$\delta_P x_2$	$-e_2^T$

negative row sum of the off-diagonal entries. The matrix \mathbf{Q} has only finitely many nonzero entries in each row and in each column, that is, \mathbf{Q} is a row-/column-finite infinite matrix [18]. Note that $\sup_x |\mathbf{Q}_{x,x}|$ may be infinite and that the number of states reachable from a given initial state may be infinite.

The following well known theorem characterizes the ergodicity of X in terms of the drift with respect to the Lyapunov function g (see, for instance, [19]).

Theorem 1

The MPM X is ergodic iff there exists a function $g : S \rightarrow \mathbb{R}_{\geq 0}$, a finite subset $W \subset S$, and a constant $\gamma > 0$ such that

- (i) $\frac{d}{dt} E[g(X(t)) | X(t) = x] \leq -\gamma$ for all $x \in S \setminus W$,
- (ii) $\frac{d}{dt} E[g(X(t)) | X(t) = x] < \infty$ for all $x \in W$, and
- (iii) $\{x \in S | g(x) \leq l\}$ is finite for all $l < \infty$.

The value $d(x) = \frac{d}{dt} E[g(X(t)) | X(t) = x]$ is called the drift in state x and can be expressed directly as

$$d(x) = \sum_{j=1}^J \alpha_j(x)(g(x + v^{(j)}) - g(x)).$$

The drift can be used to derive bounds on the equilibrium probability of being in a certain subset of states.

Let us assume that g is as in Theorem 1, that is, g is a witness for the ergodicity of X . Because X is ergodic, it has a unique equilibrium probability distribution π , which satisfies $\pi \mathbf{Q} = 0$ subject to $\pi e = 1$. Moreover, there exists a positive real number c with $c \geq \max_{x \in S} d(x)$. Let the scaled Lyapunov function be

$$g^*(x) = \frac{g(x)}{c + \gamma}.$$

Then for the scaled drift function d^* we have

$$d^*(x) = \frac{d}{dt} E[g^*(X(t)) | X(t) = x] = \frac{d(x)}{c + \gamma}.$$

Thus, the first two conditions of Theorem 1 are now equivalent to

$$d^*(x) \leq \frac{c}{c + \gamma} - \bar{\chi}_W(x), \tag{1}$$

where $\bar{\chi}_W(x)$ is the characteristic function of the set $S \setminus W$, that is, $\bar{\chi}_W(x) = 1$ if $x \notin W$ and 0 otherwise. Now we multiply Equation (1) with π_x and sum over x . Then the left-hand side becomes zero and we arrive at

$$\sum_{x \notin W} \pi_x \leq \frac{c}{c + \gamma}.$$

Thus, equation (1) provides an upper bound on the probability mass outside W and, therefore, also a lower bound on the equilibrium probability mass inside W . Let us assume we seek a subset W , which satisfies $\sum_{x \in W} \pi_x > (1 - \epsilon)$ for $\epsilon \in (0, 1)$. Then, we need to choose an appropriate Lyapunov function $g(x)$, derive the drift $d(x)$ and let $\gamma > 0$ be such that $\epsilon = \frac{c}{c + \gamma}$. From this, we can compute the scaled drift $d^*(x) = \frac{\epsilon}{c}d(x)$ and finally choose

$$W = \{x \in S \mid d^*(x) > \epsilon - 1\}. \tag{2}$$

Of course, the choice of the original Lyapunov function g is crucial for the approach presented in the aforementioned equation. The function g must be a witness for ergodicity. Moreover, its choice may drastically influence the size of W and, therefore, also the tightness of the probability bounds. In this paper, we restrict the choice of g to multivariate polynomials on $\mathbb{R}_{\geq 0}$ that obey condition (iii) of Theorem 1. This implies that if in each state $x \in S$ the drift is lower than or equal to a constant c , then the set W of all states where the drift is at least $-\gamma$ must be finite. Thus, in that case, we can conclude that X is ergodic. Our experimental results indicate that the squared Euclidean norm often suffices to witness the ergodicity of X and to provide tight probability bounds (see Sections 5 and 6).

3. STOCHASTIC COMPLEMENTATION

In the following, we extend the theory of stochastic complementation, originally proposed for finite discrete time Markov chains (DTMCs) [16], to infinite continuous-time Markov chains (CTMCs). We assume that the CTMC $\{X(t), t \geq 0\}$, with state space S , is ergodic and has an irreducible row-/column-finite infinite infinitesimal generator matrix \mathbf{Q} . For a finite subset $W \subset S$, we partition \mathbf{Q} as

$$\mathbf{Q} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix},$$

where the square submatrix \mathbf{A} contains all transitions within the finite set W , and the possibly infinite matrices \mathbf{B} , \mathbf{C} , and \mathbf{D} contain all transitions from W to $\overline{W} = S \setminus W$, from \overline{W} to W , and within \overline{W} , respectively.

Definition 1 (Embedded Markov Chain)

The transition probability matrix \mathbf{E} of the *embedded DTMC* associated with X is given by

$$\mathbf{E} = \mathbf{I} + \text{dinv}(\mathbf{Q})\mathbf{Q} = \begin{bmatrix} \tilde{\mathbf{A}} & \tilde{\mathbf{B}} \\ \tilde{\mathbf{C}} & \tilde{\mathbf{D}} \end{bmatrix},$$

where $\text{dinv}(\mathbf{Q})$ is a diagonal matrix with

$$\text{dinv}(\mathbf{Q})_{x,y} = \begin{cases} \frac{1}{|\mathbf{Q}_{x,x}|} & \text{if } x = y, \\ 0 & \text{otherwise.} \end{cases}$$

Note that the product $\text{dinv}(\mathbf{Q})\mathbf{Q}$ of the two infinite matrices is well-defined because both factors are row-/column-finite infinite matrices. By $\text{dinv}(\mathbf{D})$ we denote the block of $\text{dinv}(\mathbf{Q})$ that corresponds to block \mathbf{D} in \mathbf{Q} .

Definition 2 (Stochastic Complement)

The stochastic complement $\overline{\mathbf{Q}}$ of \mathbf{A} is defined as

$$\overline{\mathbf{Q}} = \mathbf{A} + \mathbf{B} \sum_{i=0}^{\infty} \tilde{\mathbf{D}}^i \tilde{\mathbf{C}}.$$

Note that the infinite series $\sum_{i=0}^{\infty} \tilde{\mathbf{D}}^i$ in Definition 2 is meant to be element-wise convergent. The stochastic complement for a finite block of an infinite generator matrix has already been defined in [20] (cf. page 22). However, in this definition, the matrix inverse of the infinite block $-\mathbf{D}$ is used,

but no proof of well-definedness is given. Our definition of the stochastic complement circumvents this problem by using embedding. This allows us to prove well-definedness of $\bar{\mathbf{Q}}$ as presented in the following.

Theorem 2

The stochastic complement $\bar{\mathbf{Q}}$ of \mathbf{A} is a well-defined infinitesimal generator matrix and it is irreducible.

Proof

For $\bar{\mathbf{Q}}$ to be well-defined, we need to ensure that $\sum_{i=0}^{\infty} \tilde{\mathbf{D}}^i$ converges element-wise. Now, because \mathbf{Q} is ergodic and therefore irreducible, \mathbf{E} is irreducible as well. Consequently, we can follow [21] to reason that $\tilde{\mathbf{D}}$ is strictly substochastic, and with Corollary 2 of Lemma 5.4^D of [22], we have that $\sum_{i=0}^{\infty} \tilde{\mathbf{D}}^i$ is element-wise finite.

To prove that $\bar{\mathbf{Q}}$ is an infinitesimal generator matrix, we will show that

- (i) $\bar{\mathbf{Q}}e = 0$ and
- (ii) $\bar{\mathbf{Q}}_{x,y} \geq 0$ for $x \neq y$.

Now, observe that

$$\tilde{\mathbf{D}}e + \tilde{\mathbf{C}}e = e$$

can be rewritten as

$$\tilde{\mathbf{C}}e = (\mathbf{I} - \tilde{\mathbf{D}})e.$$

We use this equality to show that $\sum_{i=0}^{\infty} \tilde{\mathbf{D}}^i \tilde{\mathbf{C}}$ is a stochastic matrix. Obviously it is nonnegative and we are left to show that the row sums converge element-wise and equal one. We begin with considering the finite sequence

$$\sum_{i=0}^N \tilde{\mathbf{D}}^i \tilde{\mathbf{C}}e = (\mathbf{I} + \tilde{\mathbf{D}} + \tilde{\mathbf{D}}^2 + \dots + \tilde{\mathbf{D}}^N)(\mathbf{I} - \tilde{\mathbf{D}})e = (\mathbf{I} - \tilde{\mathbf{D}}^{N+1})e = e - \tilde{\mathbf{D}}^{N+1}e$$

and have to show $\lim_{N \rightarrow \infty} \tilde{\mathbf{D}}^{N+1}e = 0$, that is, the sequence $\tilde{\mathbf{D}}^{N+1}e$ converges element-wise to the zero vector. For that, we will reason about the k th row sum of $\tilde{\mathbf{D}}^N$, that is, we have

$$e_k^T \tilde{\mathbf{D}}^N e = (e_k^T \tilde{\mathbf{D}}^N e)^T = e^T (\tilde{\mathbf{D}}^T)^N e_k = \|(\tilde{\mathbf{D}}^T)^N e_k\|_1, \tag{3}$$

which holds because $\tilde{\mathbf{D}}$ and, therefore, also $\tilde{\mathbf{D}}^N$ are row/column-finite for all N . Now, we observe that $\lim_{N \rightarrow \infty} (\tilde{\mathbf{D}}^T)^N e_k = 0$ element-wise because all entries of $\tilde{\mathbf{D}}^N$ tend to zero. By Schur's Theorem [23] (cf. also [24], page 137) according to which in the sequence space

$$\ell^1 := \{x = (\xi_i)_{i=1}^{\infty} \mid \|x\|_1 := \sum_{i=1}^{\infty} |\xi_i| < \infty\},$$

convergence in the 1-norm coincides with element-wise convergence, we have that $\lim_{N \rightarrow \infty} \|(\tilde{\mathbf{D}}^T)^N e_k\|_1 = 0$ element-wise for all k . Therefore, in combination with Equation (3), we also have $\lim_{N \rightarrow \infty} \tilde{\mathbf{D}}^{N+1}e = 0$ which finally results in $\sum_{i=0}^{\infty} \tilde{\mathbf{D}}^i \tilde{\mathbf{C}}e = e$. Consequently, for the row sums of $\bar{\mathbf{Q}}$ we have

$$\bar{\mathbf{Q}}e = \mathbf{A}e + \underbrace{\mathbf{B} \sum_{i=0}^{\infty} \tilde{\mathbf{D}}^i \tilde{\mathbf{C}}e}_e = (\mathbf{A} + \mathbf{B})e = 0.$$

The nonnegativity of the off-diagonal entries of $\bar{\mathbf{Q}}$ follows from $\mathbf{B} \geq 0$, $\sum_{i=0}^{\infty} \tilde{\mathbf{D}}^i \tilde{\mathbf{C}} \geq 0$, and $\mathbf{A}_{x,y} \geq 0$ for $x \neq y$.

To show the irreducibility of $\bar{\mathbf{Q}}$, we notice that if \mathbf{A} is irreducible, $\bar{\mathbf{Q}}$ is trivially irreducible as well. Otherwise, \mathbf{A} can be symmetrically permuted into a block lower-triangular form in which the diagonal blocks are irreducible [25]. Without loss of generality, consider partitionings of \mathbf{A} and $\bar{\mathbf{Q}}$ with two diagonal blocks as in

$$\mathbf{Q} = \left[\begin{array}{cc|c} \mathbf{A}^{(11)} & \mathbf{0} & \mathbf{B}^{(1)} \\ \mathbf{A}^{(21)} & \mathbf{A}^{(22)} & \mathbf{B}^{(2)} \\ \hline \mathbf{C}^{(1)} & \mathbf{C}^{(2)} & \mathbf{D} \end{array} \right] \quad \bar{\mathbf{Q}} = \left[\begin{array}{cc} \bar{\mathbf{Q}}^{(11)} & \bar{\mathbf{Q}}^{(12)} \\ \bar{\mathbf{Q}}^{(21)} & \bar{\mathbf{Q}}^{(22)} \end{array} \right].$$

In this partitioning, $\mathbf{A}^{(11)}$ and $\mathbf{A}^{(22)}$ are irreducible, and therefore, $\bar{\mathbf{Q}}^{(11)}$ and $\bar{\mathbf{Q}}^{(22)}$ are irreducible as well. Let $Pr[x \rightsquigarrow y]$ denote the probability of reaching state y from state x independent of time. Obviously, there exist states $x, w \in W$ and $y, z \in \bar{W}$ such that

$$\mathbf{B}_{x,y}^{(1)} > 0, \tilde{\mathbf{C}}_{z,w}^{(2)} > 0, \text{ and } Pr[y \rightsquigarrow z] > 0,$$

otherwise \mathbf{Q} must have been reducible. Here, $\tilde{\mathbf{C}}^{(2)}$ corresponds to the block $\mathbf{C}^{(2)}$ of \mathbf{Q} in \mathbf{E} assuming the same partitioning. Moreover, there exists an $i \geq 0$ such that

$$\tilde{\mathbf{D}}^i(y, z) > 0$$

because $Pr[y \rightsquigarrow z] > 0$. Consequently, we have

$$\bar{\mathbf{Q}}_{x,w}^{(12)} = \left(\mathbf{B} \sum_{i=0}^{\infty} \tilde{\mathbf{D}}^i \tilde{\mathbf{C}} \right)_{x,w} > 0.$$

If $\mathbf{A}^{(21)} = \mathbf{0}$ or there are more than two irreducible diagonal blocks, similar arguments can be used to prove that appropriate off-diagonal blocks in the block upper-triangular part of \mathbf{Q} are nonzero. \square

Theorem 3

Let $\bar{\pi}^W$ be the unique equilibrium probability distribution that corresponds to the generator matrix $\bar{\mathbf{Q}}$. Then for $x \in W$, we have that

$$\bar{\pi}_x^W = \frac{\pi_x}{\sum_{y \in W} \pi_y}.$$

Proof

First, let us prove the identity

$$\left[\begin{array}{cc} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{array} \right] \left[\begin{array}{cc} \mathbf{I} & \mathbf{0} \\ \sum_{i=0}^{\infty} (\tilde{\mathbf{D}})^i \tilde{\mathbf{C}} & \mathbf{I} \end{array} \right] = \left[\begin{array}{cc} \mathbf{A} + \mathbf{B} \sum_{i=0}^{\infty} (\tilde{\mathbf{D}})^i \tilde{\mathbf{C}} & \mathbf{B} \\ \mathbf{0} & \mathbf{D} \end{array} \right]. \tag{4}$$

We observe that

$$(\mathbf{I} - \mathbf{T}) \sum_{i=0}^{\infty} \mathbf{T}^i = \mathbf{I}$$

for any strictly substochastic row-/column-finite infinite matrix \mathbf{T} where the infinite series $\sum_{i=0}^{\infty} \mathbf{T}^i$ is meant to converge element-wise. Now, we choose $\mathbf{T} = \tilde{\mathbf{D}} = \mathbf{I} + \text{dinv}(\mathbf{D})\mathbf{D}$ for which strict substochasticity has been shown in the proof of Theorem 2, and we obtain

$$(\mathbf{I} - (\mathbf{I} + \text{dinv}(\mathbf{D})\mathbf{D})) \sum_{i=0}^{\infty} (\mathbf{I} + \text{dinv}(\mathbf{D})\mathbf{D})^i = \mathbf{I},$$

which reduces to

$$\text{dinv}(\mathbf{D})\mathbf{D} \sum_{i=0}^{\infty} (\mathbf{I} + \text{dinv}(\mathbf{D})\mathbf{D})^i = -\mathbf{I}.$$

Now, let the infinite matrix $\text{diag}(\mathbf{D})$ be defined as

$$\text{diag}(\mathbf{D})_{x,y} = \begin{cases} |\mathbf{D}_{x,y}| & \text{if } x = y, \\ 0 & \text{otherwise.} \end{cases}$$

Then, multiplying the previous equality on both sides by $\text{diag}(\mathbf{D})$ from the left gives

$$\mathbf{D} \sum_{i=0}^{\infty} (\mathbf{I} + \text{dinv}(\mathbf{D})\mathbf{D})^i = -\text{diag}(\mathbf{D})$$

because $\text{diag}(\mathbf{D})\text{dinv}(\mathbf{D}) = \mathbf{I}$. Exploiting this equality finally yields

$$\begin{aligned} \mathbf{C} + \mathbf{D} \sum_{i=0}^{\infty} (\tilde{\mathbf{D}})^i \tilde{\mathbf{C}} &= \mathbf{C} + \mathbf{D} \sum_{i=0}^{\infty} (\mathbf{I} + \text{dinv}(\mathbf{D})\mathbf{D})^i (\text{dinv}(\mathbf{D})\mathbf{C}) \\ &= \mathbf{C} + \underbrace{(-\text{dinv}(\mathbf{D}))\text{dinv}(\mathbf{D})}_{-\mathbf{I}} \mathbf{C} = \mathbf{0}. \end{aligned}$$

Note that the multiplications involved in taking the powers of $\tilde{\mathbf{D}}$ and in the product of $\text{dinv}(\mathbf{D})$ and $\text{diag}(\mathbf{D})$ are well-defined because these matrices are row-/column-finite. Now, letting $\pi = [\pi^W \ \pi^{\bar{W}}]$ we have

$$\pi^W \bar{\mathbf{Q}} + \pi^{\bar{W}} \mathbf{0} = 0$$

if we multiply Equation (4) with π from the left. With Theorem 2, $\bar{\mathbf{Q}}$ is irreducible and, therefore, π^W is the unique solution of $\pi^W \bar{\mathbf{Q}} = 0$ up to unit 1-norm. Thus, $\bar{\pi}^W = \pi^W \left(\sum_{y \in W} \pi_y\right)^{-1}$. \square

4. STATE-WISE BOUNDS

To derive upper and lower bounds for the equilibrium probability of each state in the finite subset W , we adapt the results in [17] and [26] to our setting. As in Section 2 and 3, we assume that \mathbf{Q} is the generator matrix of an ergodic MPM, and that $W \subset S$ is a finite subset of the state space. Let the substochastic matrix \mathbf{C} be defined as

$$\mathbf{C} = \mathbf{I} + \frac{1}{u} \mathbf{A},$$

where $u > \max_x |\mathbf{A}_{x,x}|$. We define the transition probability matrices \mathbf{C}_y for $y \in W$ as $\mathbf{C}_y = \mathbf{C} + (e - \mathbf{C}e)e_y^T$, that is, we increase the elements of the column corresponding to state y in \mathbf{C} to obtain a stochastic matrix.

We remark that \mathbf{A} , and hence, \mathbf{C} may be reducible. In that case, the quality of the bounds computed using \mathbf{C}_y deteriorate [27, Section 6]. However, in all problems we considered, \mathbf{A} turned out to be irreducible, implying the irreducibility of \mathbf{C}_y . Therefore, in the following derivations we assume that \mathbf{A} is irreducible.

Theorem 4

Let the unique equilibrium probability distribution of \mathbf{Q} be partitioned according to the state space partition as in $\pi = [\pi^W \ \pi^{\bar{W}}]$, and let \mathbf{C}_y be irreducible. Then for all $x \in W$

$$\min_{y \in W} \pi_x^{\mathbf{C}_y} \leq \frac{\pi_x^W}{\sum_{y \in W} \pi_y} \leq \max_{y \in W} \pi_x^{\mathbf{C}_y},$$

where $\pi^{\mathbf{C}_y}$ is the unique equilibrium probability distribution of the DTMC associated with \mathbf{C}_y .

Proof

From Theorem 2, we know that $\bar{\mathbf{Q}}$ is well-defined and irreducible. Furthermore, Theorem 3 ensures that π conditioned on W is identical to the equilibrium distribution of $\bar{\mathbf{Q}}$, which also coincides with the equilibrium probability distribution of

$$\bar{\mathbf{C}} = \mathbf{I} + \frac{1}{u} \bar{\mathbf{Q}}.$$

Note that

$$u > \max_{x \in W} |\mathbf{A}(x, x)| \geq \max_{x \in W} |\bar{\mathbf{Q}}(x, x)|$$

implies that $\bar{\mathbf{C}}$ is a stochastic matrix. The same arguments as in [17] can be used to conclude that the equilibrium distribution of $\bar{\mathbf{C}}$ is contained in the convex hull of the unit 1-norm Perron vectors of the irreducible matrices \mathbf{C}_y (see also [26, 28]). \square

From Equation (1), we know that W can be determined such that

$$1 - \epsilon < \sum_{x \in W} \pi_x \leq 1.$$

Consequently, Theorem 4 allows us to bound the (unconditional) individual equilibrium probabilities for states $x \in W$ as

$$(1 - \epsilon) \min_{y \in W} \pi_x^{\mathbf{C}_y} < \pi_x \leq \max_{y \in W} \pi_x^{\mathbf{C}_y}. \tag{5}$$

Equation (5) above implies an algorithmic procedure for the computation of lower and upper bounds on the equilibrium probabilities. We determine the set W as suggested in Section 2. Then, for each $y \in W$, we construct \mathbf{C}_y and, provided it is irreducible, compute the distribution $\pi^{\mathbf{C}_y}$. Finally, for every $x \in W$, we derive bounds according to Equation ((5)). For the states $x \in \bar{W}$, we have the trivial bound $0 < \pi_x \leq \epsilon$. The computational effort of this approach can be decreased by exploiting the structure of \mathbf{Q} and by ignoring matrices \mathbf{C}_y for which y is a state that has no incoming transitions from \bar{W} , that is, that have zero columns in \mathbf{C} . For a given set of transition classes, it is easy to compute the set $H \subseteq W$ of states that have incoming transitions from \bar{W} , because

$$H = \{x \in W \mid \exists j \in \{1, \dots, J\} : \alpha_j(x - v^{(j)}) > 0, (x - v^{(j)}) \in \bar{W}\}.$$

The use of similar arguments as in [17] and again, provided that \mathbf{C}_y is irreducible for each $y \in H$, we have

$$(1 - \epsilon) \min_{y \in H} \pi_x^{\mathbf{C}_y} < \pi_x \leq \max_{y \in H} \pi_x^{\mathbf{C}_y}.$$

To compute the equilibrium probability distributions $\pi^{\mathbf{C}_y}$ of \mathbf{C}_y for $y \in H$ efficiently, we consider $\pi^{\mathbf{C}_y} \mathbf{C}_y = \pi^{\mathbf{C}_y}$ subject to $\pi^{\mathbf{C}_y} e = 1$. Because this equation needs to be solved $|H|$ times for \mathbf{C}_y that differ from each other in two columns, we uncouple the slack probability mass added to column y and, assuming that \mathbf{A} is irreducible, obtain the irreducible stochastic matrix

$$\mathbf{P}_y = \begin{bmatrix} \mathbf{C} & e - \mathbf{C}e \\ e_y^T & 0 \end{bmatrix}$$

of order $|W| + 1$. Now, let $p_y = [p_y^{\mathbf{C}} \ p'_y]$ denote the equilibrium probability distribution of \mathbf{P}_y , that is, $p_y \mathbf{P}_y = p_y$ subject to $p_y e = 1$. Multiplying $(\mathbf{P}_y - I)$ on the left-hand side with p_y and equating to 0, we obtain $p_y^{\mathbf{C}} (\mathbf{C} - \mathbf{I}) = -p'_y e_y^T$, which can be solved subject to $p_y^{\mathbf{C}} e = 1$ for $y \in H$. We remark that the coefficient matrix does not depend on y and, therefore, can be LU factorized once. Consequently, the time complexity of the whole procedure of computing conditional bounds reduces to that of a sparse LU factorization of a matrix of order $|W|$ plus $|H|$ forward and backward substitutions.

5. CASE STUDY: EXCLUSIVE SWITCH

We developed a prototypical tool called *Geobound* [29] written in Java 6. It implements both the computation of geometric bounds via Lyapunov functions and the computation of the conditional probabilities using the approach presented in this section.

Geobound takes as input a transition class description of an MPM and tries to prove ergodicity using a user-specified Lyapunov function. If a multivariate polynomial $g(x)$ is chosen as Lyapunov function, then the drift is a multivariate polynomial as well. If the latter tends to $-\infty$ with increasing distance to the origin, then only finitely many states have a nonnegative drift and the first two conditions of Theorem 1 hold.

To determine the maximum drift, we restrict the domain of the search for the extrema to $\mathbb{R}_{\geq 0}^n$. We compute all extrema by equating the gradient $\nabla d(x)$ with the zero vector. This alone would not yield all local extrema. Therefore, we additionally solve $\nabla d(x) = 0$ for every projection of $d(x)$ onto each subspace of \mathbb{R}^n by setting all combinations of variables x_i with $i \in \{1, \dots, n\}$ to zero, to obtain all the local extrema located at the borders of the domain. Finally, we filter out all extrema outside of $\mathbb{R}_{\geq 0}^n$. The resulting nonlinear equation systems are solved using the *HOM4PS-2.0* package [30], an implementation of the polyhedral homotopy continuation method.

The global maximum c as the maximum over all local extrema found in the previous step is used to retrieve the scaled drift polynomial $d^*(x)$ (see Section 2). Note that traditional global optimization techniques such as gradient descent or simulated annealing cannot be used because we need to guarantee that the found maximum is indeed the global maximum. Moreover, we also need the local maxima as initial points when we determine W in the next step. More precisely, we start with the local maxima and recursively add all neighboring points of the integer grid that satisfy the condition in Equation (2) for the scaled drift.

To compute H , we calculate for each $x \in W$ and each transition class $(\alpha_j, v^{(j)})$ the predecessor $(x - v^{(j)})$. If $(x - v^{(j)}) \notin W$ and $\alpha_j(x - v^{(j)}) > 0$, then we add x to H . Next, we compute the nonzero entries of the finite matrix \mathbf{A} and compute the conditional bounds of π_x for $x \in W$ as explained in the previous text. For this, the tool employs the *SuperLU* [31] package for LU factorization. All computations are performed on a dual-core 2.66 GHz machine with 3 gigabytes of memory under Ubuntu 10.04.

In the following, we consider a case study from biology called exclusive switch. In previous work, the equilibrium distribution of this gene regulatory network has been approximated by Monte Carlo simulations [32], but no execution times are reported. Here, we present a direct method to truncate the state space using the Lyapunov approach presented in Section 2, and we approximate the equilibrium by computing state-wise bounds as explained in Section 4.

5.1. Model

The exclusive switch is a gene regulatory network consisting of two genes that share a promoter region (cf. Figure 2). The state of the promoter region determines whether the genes are used to produce proteins or not. Let P_1 be the product of the first gene, and let P_2 be the product of the second gene. Each of the proteins can bind to the promoter region and thereby inhibit the production of the other protein. More precisely, if the promoter region is free, both P_1 and P_2 are produced.




Figure 2. Illustration of the exclusive switch in section 5. The picture has been adapted from [32].

Otherwise, if P_1 (P_2) is bound to the region, only the first (second) gene is active and thus only P_1 (P_2) is produced, respectively. The overall system consists of five species, hence, the state space $S \subset \mathbb{Z}_+^5$. In a state $x = [x_1, \dots, x_5] \in S$, x_1 (x_2) represents the number of P_1 (P_2) molecules. The entry x_3 is one if the promotor region is free and zero otherwise. Similarly, x_4 (x_5) is one if a molecule of type P_1 (P_2) is bound to the promotor region, respectively, and zero otherwise. We describe the exclusive switch by the set of stoichiometric equations in Table II. Additionally, we list the corresponding transition classes. Note that species G represents both genes.

5.2. Geometric bounds

For the sake of simplicity, we choose the squared Euclidean norm as the Lyapunov function for our case study. Note that this choice influences the accuracy of the probability bounds. Obviously, it is possible to take as Lyapunov function a multivariate polynomial and optimize over the coefficients. Our experimental results, however, indicate that for the models that we consider the squared Euclidean norm performs well. Thus, the Lyapunov function $g(x)$ is defined as

$$g(x) = x_1^2 + \dots + x_5^2.$$

The drift function $d(x)$ then becomes

$$\begin{aligned} d(x) &= \sum_{j=1}^{10} \alpha_j(x)(g(x + v_j) - g(x)) \\ &= \lambda x_3(2x_1 + 1) + \lambda x_3(2x_2 + 1) + \delta x_1(-2x_1 + 1) + \delta x_2(-2x_2 + 1) \\ &\quad + \beta x_1 x_3(-2x_1 - 2x_3 + 2x_4 + 3) + \beta x_2 x_3(-2x_2 - 2x_3 + 2x_5 + 3) \\ &\quad + \nu x_4(-2x_4 + 2x_1 + 2x_3 + 3) + \nu x_5(-2x_5 + 2x_2 + 2x_3 + 3) \\ &\quad + \lambda x_4(2x_1 + 1) + \lambda x_5(2x_2 + 1). \end{aligned}$$

With the initial condition $x = [0, 0, 1, 0, 0]$, it is easy to see that $x_k \in \{0, 1\}$ for $k \in \{3, 4, 5\}$ and $x_3 + x_4 + x_5 = 1$, that is, at any time the promotor region is either free, a molecule of type P_1 is bound, or P_2 is bound. For each state of the promotor, we consider one drift function:

$$\begin{aligned} d_3(x_1, x_2) &= d(x_1, x_2, 1, 0, 0) = -2(\delta + \beta)(x_1^2 + x_2^2) + (2\lambda + \delta + \beta)(x_1 + x_2) + 2\lambda \\ d_4(x_1, x_2) &= d(x_1, x_2, 0, 1, 0) = -2\delta(x_1^2 + x_2^2) + \delta(x_1 + x_2) + 2(\nu + \lambda)x_1 + \lambda + \nu \\ d_5(x_1, x_2) &= d(x_1, x_2, 0, 0, 1) = -2\delta(x_1^2 + x_2^2) + \delta(x_1 + x_2) + 2(\nu + \lambda)x_2 + \lambda + \nu \end{aligned}$$

Table II. Chemical reactions and transition classes of the exclusive switch.

j	Stoichiometric equation	τ_j	$\alpha_j(x)$	$v^{(j)}$
1	$G \xrightarrow{\lambda} G + P_1$	$(\alpha_1, v^{(1)})$	λx_3	e_1^T
2	$G \xrightarrow{\lambda} G + P_2$	$(\alpha_2, v^{(2)})$	λx_3	e_2^T
3	$P_1 \xrightarrow{\delta} \emptyset$	$(\alpha_3, v^{(3)})$	δx_1	$-e_1^T$
4	$P_2 \xrightarrow{\delta} \emptyset$	$(\alpha_4, v^{(4)})$	δx_2	$-e_2^T$
5	$G + P_1 \xrightarrow{\beta} G.P_1$	$(\alpha_5, v^{(5)})$	$\beta x_1 x_3$	$(-e_1 - e_3 + e_4)^T$
6	$G + P_2 \xrightarrow{\beta} G.P_2$	$(\alpha_6, v^{(6)})$	$\beta x_2 x_3$	$(-e_2 - e_3 + e_5)^T$
7	$G.P_1 \xrightarrow{\nu} G + P_1$	$(\alpha_7, v^{(7)})$	νx_4	$(-e_4 + e_1 + e_3)^T$
8	$G.P_2 \xrightarrow{\nu} G + P_2$	$(\alpha_8, v^{(8)})$	νx_5	$(-e_5 + e_2 + e_3)^T$
9	$G.P_1 \xrightarrow{\lambda} G.P_1 + P_1$	$(\alpha_9, v^{(9)})$	λx_4	e_1^T
10	$G.P_2 \xrightarrow{\lambda} G.P_2 + P_2$	$(\alpha_{10}, v^{(10)})$	λx_5	e_2^T

After fixing the model parameters to $\lambda = 0.05$, $\delta = 0.005$, $\beta = 0.01$, and $\nu = 0.008$, it turns out that the global maximum drift $c = \max_{x \in S} d(x) = 0.42465$ is in states

$$x^{(m1)} = [6.05, 0.25, 0, 1, 0] \text{ and } x^{(m2)} = [0.25, 6.05, 0, 0, 1].$$

Note that the coefficients of the highest order terms are negative, which implies that $W = \{x \in S \mid d(x) \geq -\gamma\}$ is finite. We choose $\epsilon = 0.1$, which yields a set of states that contains at least 90% of the equilibrium probability mass. We derive the scaled drift functions

$$d_k^*(x) = \frac{\epsilon}{c} d_k(x) \text{ with } k \in \{3, 4, 5\}$$

and finally compute

$$\begin{aligned} W &= \{x \in S \mid d^*(x) > -0.9\} \\ &= \{(x_1, x_2, 1, 0, 0) \in S \mid d_3^*(x_1, x_2) > -0.9\} \\ &\cup \{(x_1, x_2, 0, 1, 0) \in S \mid d_4^*(x_1, x_2) > -0.9\} \\ &\cup \{(x_1, x_2, 0, 0, 1) \in S \mid d_5^*(x_1, x_2) > -0.9\}. \end{aligned}$$

We illustrate the set W in Figure 3, left, with respect to the state space projected on P_1 and P_2 . The set W contains 1,680 states altogether.

5.3. Conditional probabilities

We apply Theorem 4 to bound the individual equilibrium probabilities conditioned on being in W . For an illustration of the conditional lower bounds, we refer to Figure 3, left. To compute unconditional bounds, the lower bounds should be multiplied by $(1 - \epsilon) = 0.9$. We list further results for the exclusive switch in Table III where we vary ϵ (cf. first column). The columns ‘ $|W|$ ’ and ‘ $|H|$ ’ list the size of the sets W and H . In the two columns with heading ‘Time (s)’, we give the execution time in seconds that was needed to determine the sets W and H , as well as the execution time for the




Figure 3. State-wise lower bounds of the conditional equilibrium distribution of the exclusive switch and a set that contains at least 90% of the probability mass (left). The dashed lines confine the three sets whose union is W . The right plot shows the difference between state-wise lower and upper bounds.

Table III. Numerical results for the exclusive switch.

ϵ	$ W $	$ H $	Time (s)			ρ
			Computing W and H	Computing bounds	Memory (MB)	
1.0e-1	1,680	108	1.0	5.0	43.5	2.9e-3
5.0e-2	2,934	144	1.4	5.7	64.2	1.4e-3
1.0e-2	11,994	297	1.4	83.2	304.6	2.8e-4

LU factorization that was used to compute the bounds on the conditional probabilities. Moreover, we list our memory requirements (in MegaBytes) in the column ‘Memory (MB)’. Finally, in the last column we give the maximum absolute difference of the bounds of the unconditional equilibrium probabilities, ρ , that is,

$$\rho = \max_{x \in W} (\max_{y \in H} \pi_x^{C_y} - (1 - \epsilon) \min_{y \in H} \pi_x^{C_y}).$$

In Figure 3, right, we illustrate the difference between upper and lower bounds of the conditional probabilities. Note that the maximum absolute difference of the conditional probabilities is smaller than ρ . To compute bounds on the unconditional probabilities, each lower bound of the conditional probability is multiplied by $(1 - \epsilon)$, whereas the upper bound remains the same.

6. NUMERICAL RESULTS

In this section, we list further numerical results for various models from biology. All models have previously been analyzed using the Monte Carlo simulations, and execution times are not reported [8, 32, 33]. Here, we approximate their equilibrium distribution using the techniques presented previously. Unless stated otherwise, we again use the squared Euclidean norm

$$g(x) = \sum_{i=1}^n x_i^2$$

as a Lyapunov function.

6.1. Toggle switch

The toggle switch [34] is a gene regulatory network that is similar to the exclusive switch. Instead of a single promotor, it consists of two promotor regions. If the first promotor \bar{R}_B is free (variable x_1), the production of protein A (variable x_3) is enabled. The other promotor \bar{R}_A (variable x_2) controls the production of protein B (variable x_4). Both proteins can repress the expression of the other protein by binding to the promotor of the corresponding gene. In the repressed state, the promotors are represented by the species R_A and R_B . Proteins bound to the promotor can unbind and protein molecules can degrade over time. We list the stoichiometric equations in Table IV (cf. [33]) as well as the transition classes. We choose parameters $\lambda_A = \lambda_B = 6.0$, $\delta_A = \delta_B = 0.4$, $\beta_0 = 1.0$, and $\beta_1 = 0.5$. The results are presented in Table V. The global maximum drift $c = 124.225$ was found at $x_3 = x_4 = 8.375$, where both promotor regions are free.

For this system we could verify tristability, that is, in addition to the two regions with a high number of protein molecules of one of the two species, we could identify another peak in the equilibrium

Table IV. Transition classes for the toggle switch model.

j	Stoichiometric equation	τ_j	$\alpha_j(x)$	$v^{(j)}$
1	$\bar{R}_B \xrightarrow{\lambda_A} \bar{R}_B + A$	$(\alpha_1, v^{(1)})$	$\lambda_A(1 - x_1)$	e_3^T
2	$A \xrightarrow{\delta_A} \emptyset$	$(\alpha_2, v^{(2)})$	$\delta_A x_3$	$-e_3^T$
3	$\bar{R}_A + A \xrightarrow{\beta_0} R_A$	$(\alpha_3, v^{(3)})$	$\beta_0(1 - x_2)x_3$	$e_2^T - e_3^T$
4	$R_A \xrightarrow{\beta_1} \bar{R}_A + A$	$(\alpha_4, v^{(4)})$	$\beta_1 x_2$	$e_3^T - e_2^T$
5	$\bar{R}_A \xrightarrow{\lambda_B} \bar{R}_A + B$	$(\alpha_5, v^{(5)})$	$\lambda_B(1 - x_2)$	e_4^T
6	$B \xrightarrow{\delta_B} \emptyset$	$(\alpha_6, v^{(6)})$	$\delta_B x_4$	$-e_4^T$
7	$\bar{R}_B + B \xrightarrow{\beta_0} R_B$	$(\alpha_7, v^{(7)})$	$\beta_0(1 - x_1)x_4$	$e_1^T - e_4^T$
8	$R_B \xrightarrow{\beta_1} \bar{R}_B + B$	$(\alpha_8, v^{(8)})$	$\beta_1 x_1$	$e_4^T - e_1^T$

Table V. Numerical results for the toggle switch model.

ϵ	$ W $	$ H $	Time (s)		Memory (MB)	ρ
			Computing W and H	Computing bounds		
1.0e-1	7,964	292	1.5	29.0	227.3	1.9e-3
5.0e-2	14,016	384	1.9	65.8	433.2	9.7e-4

distribution in a region with a low number of protein molecules. In [33], the authors explained that peak by a deadlock situation were caused by both promotor regions being bound simultaneously.

6.2. Gene expression

Next, we consider the gene expression in Example 1 with the parameters $\lambda = 100.0$, $\mu = 0.01$, $\delta_M = 0.2$, and $\delta_P = 0.02$. These parameters are chosen such that the attracting regions are located far from the origin. This implies that the state space must be truncated from above and additionally from below because otherwise the number of states in the set W becomes intractably large. In the examples presented so far, the set W always included the origin. The reason is that all states ‘below’ an attractor have positive drift if the squared Euclidean norm is used. If we take as a Lyapunov function the distance to attracting regions, it is possible to derive a set W that does not include the origin. The reason is that a state that is located far ‘below’ an attractor may have a negative drift. In general, it is difficult to guess attracting regions. One possibility is to first consider the multi-dimensional drift function

$$\tilde{d}(x) = \frac{d}{dt} E[X(t) | X(0) = x]$$

and find regions where it is maximal. Then, another Lyapunov function is used to bound the equilibrium probability, namely, the function that measures the distance to these regions. Thus, we set $\tilde{d}(x)$ to zero and obtain

$$\lambda - \delta_M x_1 = 0 \quad \mu x_1 - \delta_P x_2 = 0,$$

which has the unique solution $(\tilde{x}_1, \tilde{x}_2) = \left(\frac{\lambda}{\delta_M}, \frac{\lambda\mu}{\delta_M\delta_P}\right)$. Next we consider the Lyapunov function

$$g(x) = a(x_1 - \tilde{x}_1)^2 + b(x_2 - \tilde{x}_2)^2,$$

that is, the weighted squared distance to the possible attractor $(\tilde{x}_1, \tilde{x}_2)$ computed before. We tested different weights and found that in the choice $a = 1.0$, $b = 20.0$ provides good numerical results, which we list in Table VI. Note that for $a = b = 1.0$ and $\epsilon = 0.1$, the set W contains 52,316 states (vs 23,770 for $a = 1.0$, $b = 20.0$) and the results are slightly more accurate (not shown). This indicates that the choice of the Lyapunov function is particularly important if the attracting regions must be bounded from below as well.

We have applied our technique to several other models from systems biology, for example, protein synthesis [35], and could achieve bounding precisions up to a magnitude of 10^{-6} at moderate computation times of less than 8 s depending on the model structure.

Table VI. Numerical results for the gene expression model.

ϵ	$ W $	$ H $	Time (s)		Memory (MB)	ρ
			Computing W and H	Computing bounds		
1.0e-1	23,770	515	0.7	19.0	605.2	1.6e-4
5.0e-2	47,517	728	1.3	90.1	866.3	4.6e-8

7. CONCLUSION

This paper has set the algorithmic ground for the computation of state-wise equilibrium probability bounds for Markov population models, a general class of structured infinite-state continuous-time Markov chains. This is achieved by combining Lyapunov theory with numerical approximation and bounding techniques. Special attention has been paid to the rigorous derivation of the results in the context of the infinite state space setting especially when extending the theory of stochastic complementation to Markov population models. The presented results have been successfully implemented in a software tool, which is particularly suited to problems originating from the systems biology. It is the first tool of its kind, and empirical results show its computational strengths.

The Lyapunov function used as a witness for ergodicity plays an important role in the approach we follow. Thus far the function is constructed manually. As we have discussed, often the Euclidean norm works well, but more elaborate multivariate polynomials give better results with respect to both tightness of the bounds and computational effort. For future work, we will investigate the automated synthesis and coefficient optimization of the polynomial based on the transition class description. This seems a very promising direction to further improve the runtime and the quality of the computed bounds.

ACKNOWLEDGEMENTS

We would like to thank Michele Benzi for his valuable comments on the convergence proofs.

This research has been partially funded by the German Research Council (DFG) as part of the Cluster of Excellence on Multimodal Computing and Interaction at Saarland University and the Transregional Collaborative Research Center ‘Automatic Verification and Analysis of Complex Systems’ (SFB/TR 14 AVACS).

REFERENCES

1. Ball K, Kurtz TG, Popovic L, Rempala G. Asymptotic analysis of multiscale approximations to reaction networks. *The Annals of Applied Probability* 2006; **16**(4):1925–1961.
2. Engblom S. Numerical solution methods in stochastic chemical kinetics. *PhD Thesis*, Uppsala University, Division of Scientific Computing, Numerical Analysis, 2008.
3. Grassmann WK. Finding transient solutions in Markovian event systems through randomization. In *Numerical Solution of Markov Chains*, Stewart WJ (ed.). Marcel Dekker: New York, 1991; 357–371.
4. Heegaard PE, Sandmann W. Ant-based approach for determining the change of measure in importance sampling. In *Proceedings of the Winter Simulation Conference, WSC 2007*. IEEE Press: Piscataway, New Jersey, 2007; 412–420.
5. Henzinger TA, Jobstmann B, Wolf V. Formalisms for specifying Markovian population models. In *Proceedings of the 3rd International Workshop on Reachability Problems*, Vol. 5797, LNCS. Springer-Verlag: Berlin, Heidelberg, 2009.
6. Kurtz TG. The relationship between stochastic and deterministic models for chemical reactions. *Journal of Chemical Physics* 1972; **57**(7):2976–2978.
7. Henzinger TA, Mateescu M, Mikeev L, Wolf V. Hybrid numerical solution of the chemical master equation. In *Proceedings of the 8th International Conference on Computational Methods in Systems Biology, CMSB '10*. ACM Digital Library: New York, 2010.
8. Thattai M, van Oudenaarden A. Intrinsic noise in gene regulatory networks. *Proceedings of the National Academy of Sciences of the United States of America* 2001; **98**(15):8614–8619.
9. Swain PS, Elowitz MB, Siggia ED. Intrinsic and extrinsic contributions to stochasticity in gene expression. *Proceedings of the National Academy of Sciences of the United States of America* 2002; **99**(20):12795–12800.
10. Paulsson J. Summing up the noise in gene networks. *Nature* 2004; **427**(6973):415–418.
11. Jahnke T, Huisinga W. Solving the chemical master equation for monomolecular reaction systems analytically. *Journal of Mathematical Biology* 2007; **54**(1):1–26.
12. Ramsey S, Orrell D, Bolouri H. Dizzy: stochastic simulation of large-scale genetic regulatory networks. *Journal of Bioinformatics and Computational Biology* 2005; **3**(2):415–436.
13. Didier F, Henzinger TA, Mateescu M, Wolf V. Fast adaptive uniformization of the chemical master equation. *IET Systems Biology Journal* 2010. To appear.
14. Hegland M, Burden C, Santoso L, MacNamara S, Booth H. A solver for the stochastic master equation applied to gene regulatory networks. *Journal of Computational and Applied Mathematics* 2007; **205**(2):708–724.
15. Sidje RB, Burrage K, MacNamara S. Inexact uniformization method for computing transient distributions of Markov chains. *SIAM Journal on Scientific Computing* 2007; **29**(6):2562–2580.
16. Meyer CD. Stochastic complementation, uncoupling Markov chains, and the theory of nearly reducible systems. *SIAM Review* 1989; **31**(2):240–272. DOI: 10.1137/1031050.

17. Courtois P-J, Semal P. Bounds for the positive eigenvectors of nonnegative matrices and for their approximations by decomposition. *Journal of the ACM* 1984; **31**(4):804–825. DOI: 10.1145/1634.1637.
18. Halmos PR. *A Hilbert Space Problem Book*, 2nd ed., Graduate Texts in Mathematics. Springer: New York, Berlin, Heidelberg, 1982.
19. Tweedie RL. Sufficient conditions for regularity, recurrence and ergodicity of Markov processes. *Mathematical Proceedings of the Cambridge Philosophical Society* 1975; **78**(01):125–136. DOI: 10.1017/S0305004100051562.
20. Riska A. Aggregate matrix-analytic techniques and their applications. *PhD Thesis*, The College of William and Mary, 2002. AAI3081205.
21. Ramaswami V, Taylor PG. Some properties of the rate operators in level dependent quasi-birth-and-death processes with countable number of phases. *Communications in Statistics. Stochastic Models* 1996; **12**(1):143–164.
22. Seneta E. *Non-negative matrices; an introduction to theory and applications*. Allen & Unwin: London, 1973.
23. Schur J. Über lineare Transformationen in der Theorie der unendlichen Reihen. *Journal für die reine und angewandte Mathematik (Crelle's Journal)* 1921; **1921**(151):79–111. <http://dx.doi.org/10.1515/crll.1921.151.79>.
24. Banach S. *Théorie des opérations linéaires*. Chelsea Publishing Company: New York, 1955.
25. Duff IS, Reid JK. An implementation of Tarjan's algorithm for the block triangularization of a matrix. *ACM Transactions on Mathematical Software* 1978; **4**(2):137–147. DOI: 10.1145/355780.355785.
26. Courtois P-J. Analysis of large Markovian models by parts. Applications to queueing network models. *Messung, Modellierung und Bewertung von Rechensystemen, 3. GI/NTG-Fachtagung*, 1985; 1–10.
27. Dayar T, Stewart WJ. Quasi lumpability, lower-bounding coupling matrices, and nearly completely decomposable Markov chains. *SIAM Journal on Matrix Analysis and Applications* 1997; **18**(2):482–498. DOI: 10.1137/S0895479895294277.
28. Muntz RR, de Souza e Silva E, Goyal A. Bounding availability of repairable computer systems. In *Sigmetrics '89: Proceedings of the 1989 ACM Sigmetrics International Conference on Measurement and Modeling of Computer Systems*. ACM: New York, NY, USA, 1989; 29–38. DOI: 10.1145/75108.75376.
29. Geobound. <http://alma.cs.uni-sb.de>, Last site access on 9 March 2011.
30. Lee TL, Li TY, Tsai CH. HOM4PS-2.0: A software package for solving polynomial systems by the polyhedral homotopy continuation method. *Computing* 2008; **83**(2–3):109–133. DOI: 10.1007/s00607-008-0015-6.
31. Demmel JW, Eisenstat SC, Gilbert JR, Li XS, Liu JWH. A supernodal approach to sparse partial pivoting. *SIAM Journal on Matrix Analysis and Applications* 1999; **20**(3):720–755.
32. Loinger A, Lipshtat A, Balaban NQ, Biham O. Stochastic simulations of genetic switch systems. *Physical Review E* 2007; **75**(2):021904.
33. Lipshtat A, Loinger A, Balaban NQ, Biham O. Genetic toggle switch without cooperative binding. *Physical Review Letters* 2006; **96**(18):188101. DOI: 10.1103/PhysRevLett.96.188101.
34. Gardner TS, Cantor CR, Collins JJ. Construction of a genetic toggle switch in *Escherichia coli*. *Nature* 2000; **403**(6767):339–342. DOI: 10.1038/35002131.
35. Goss PJE, Peccoud J. Quantitative modeling of stochastic systems in molecular biology by using stochastic Petri nets. *Proceedings of the National Academy of Sciences of the United States of America* 1998; **95**(12):6750–6755.