Exact Stochastic Simulation of Chemical Reactions with Cycle Leaping*

Marc D. Riedel
University of Minnesota
200 Union St. S.E.
Minneapolis, MN 55455
mriedel@umn.edu

Jehoshua Bruck
California Institute of Technology
Mail Code 136-93
Pasadena, CA 91125
bruck@paradise.caltech.edu

Abstract— **The stochastic simulation algorithm (SSA), first proposed by Gillespie, has become the workhorse of computational biology. It tracks integer quantities of the molecular species, executing reactions at random based on propensity calculations. An estimate for the resulting quantities of the different species is obtained by averaging the results of repeated trials. Unfortunately, for models with many reaction channels and many species, the algorithm requires a prohibitive amount of computation time. Many trials must be performed, each forming a lengthy trajectory through the state space. With coupled or reversible reactions, the simulation often loops through the same sequence of states repeatedly, consuming computing time, but making no forward progress.**

We propose a algorithm that reduces the simulation time through *cycle leaping*: when cycles are encountered, the exit probabilities are calculated. Then, in a single bound, the simulation leaps directly to one of the exit states. The technique is exact, sampling the state space with the expected probability distribution. It is a component of a general framework that we have developed for stochastic simulation based on probabilistic analysis and caching.

I. Introduction

Randomness is inherent to all biochemical systems: at any given instant, the choice of which reaction fires next is a matter of chance. Certain biochemical systems appear to exploit this randomness for evolutionary advantage, choosing between different outcomes with a probability distribution – in effect, hedging their bets with a portfolio of responses that is carefully tuned to the environmental conditions. Examples include the lysis/lysogeny decision of the *lambda phage* [1] and the *pap pili* epigenetic response of bacteria [20].

Modeling and simulating the probabilistic behavior of such systems is a challenging problem. On the one hand, macroscopic-level modelling – say in terms of simple "on/off" activation levels – lacks sufficient detail to capture important aspects of the system behavior; on the other hand, microscopic-level simulations – in terms of the molecular dynamics – are computationally prohibitive. A successful approach has been to model systems at an intermediate level, sometimes called the *mesoscopic*. One assumes that non-reactive collisions occur far more frequently than reactive collisions, and so the medium is "well-stirred". The analysis tracks the quantities of the different molecular species, but not their spatial location.

The state of the system is modeled in terms of the whole (i.e., non-negative integer) quantities of the constituent molecules. As biochemical reactions fire, discrete state transitions occur. The behavior is that of a probabilistic, discrete-event system – or a *Markov chain*.

Gillespie proposed stochastic simulation (sometimes called *Monte Carlo*) to characterize such systems: beginning from an initial state, reactions are chosen at random, based on propensity calculations [13]. As reactions fire, the quantities of the different species change by integer amounts. An estimate for the resulting quantities of the different species is obtained by averaging the results of repeated trials.

The drawback of Gillespie's stochastic simulation algorithm (SSA), as it has become known, is the amount of computation required. Although the simulation does not track the spatial location of individual molecules, it executes each and every reaction that occurs, updating the quantities of species present. The simulation can be very lengthy since there are a multitude of reactions happening nearly in *parallel* and

*This work is supported by the "Alpha Project" at the Center for Genomic Experimentation and Computation, a National Institutes of Health Center of Excellence in Genomic Sciences (grant no. P50 HG02370).

these must all be executed *serially*. At each step, the choice of which reaction occurs next entails a probability calculation as well as generating a random number. Each trial consists of a long sequence of reactions; many such trials must be performed in order to obtain an accurate estimate. This adds up to significant computation time [23].

A. Cycles

Typical biochemical systems contain reversible reactions; also, the reactions are often highly coupled, that is, many of the molecular species appear both as reactants and products. Indeed, coupled systems produce the most interesting dynamics, including switch-like behavior. With coupled and reversible reactions, the simulation trajectories tend to be lengthy due to cycling. The simulation loops through the same sequence of states repeatedly, consuming computing time, but making no forward progress.

Example 1 *Consider a system with three types of molecules X_1, X_2, and X_3. The* **state** *of the system is described by*

$$[x_1, x_2, x_3],$$

where x_1, x_2, and x_3 are integer *variables, assuming non-negative values corresponding to the number of molecules of types X_1, X_2, and X_3, respectively. For instance, the system might be in the state $[3, 3, 3]$ with three molecules of each type.*

Consider the three reactions:

$$\begin{aligned} R_1: &\quad 2X_1 + X_2 \rightarrow 3X_3 \\ R_2: &\quad X_1 + 2X_3 \rightarrow 3X_2 \\ R_3: &\quad X_2 + X_3 \rightarrow 2X_1 \end{aligned}$$

The types that are consumed are referred to as the reactants*, whereas those that are created are referred to as the* products*. Note that these reactions are* coupled*: the types appear both as reactants and products in different reactions.*

Suppose that the system is in the state $[5, 5, 5]$ and reaction R_3 fires. One molecule of type X_1 and one of type X_3 are consumed; two of type X_2 are produced. This results in the state transition:

$$[5, 5, 5] \xrightarrow{R_1} [4, 7, 4].$$

From this state, suppose reactions R_1, R_3 and R_2 fire, in this order. This results in the state transitions shown in Figure 1. Note that the system returns *to the state $[4, 7, 4]$ as a result.* □

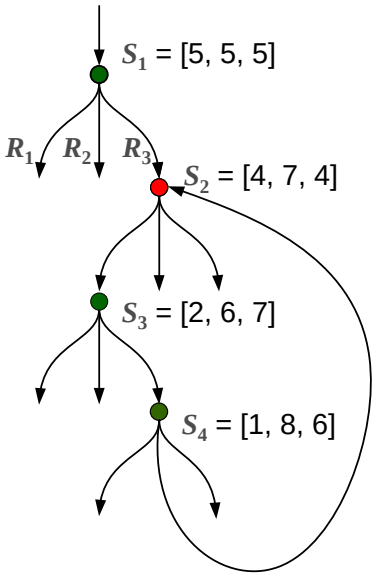

Fig. 1. A loop in stochastic simulation: the sequence of states S_2, S_3, S_4 is visited repeatedly.

Cycles in trajectories are not surprising, since the simulation is tracking the minutiae of the physical behavior: as molecular species are consumed and produced, once can expect the system to vacillate, returning to states that it has visited. And yet, in the stochastic simulation algorithm, computing time is frittered away as the calculations of the propensities and random numbers are replayed every time that the simulation cycles.

B. Cycle Leaping

We propose a algorithm that reduces the simulation time through *cycle leaping*: when cycles are encountered, the exit probabilities are calculated. Then, in a single bound, the simulation leaps directly to one of the exit states. This is illustrated in Figure 2. The technique is exact, sampling the state space with the expected probability distribution. It is a component of a general framework that we have developed for stochastic simulation based on probabilistic analysis and caching.

C. Related Work

Gibson and Bruck proposed algorithmic improvements to Gillespie's SSA [11][12]. Their method achieves significant speedups by structuring the computation through prioritized data structures and by using random numbers parsimoniously.

Also, Several methods have been proposed to expedite stochastic simulation through approximations. Gillespie proposed a technique called "tau-

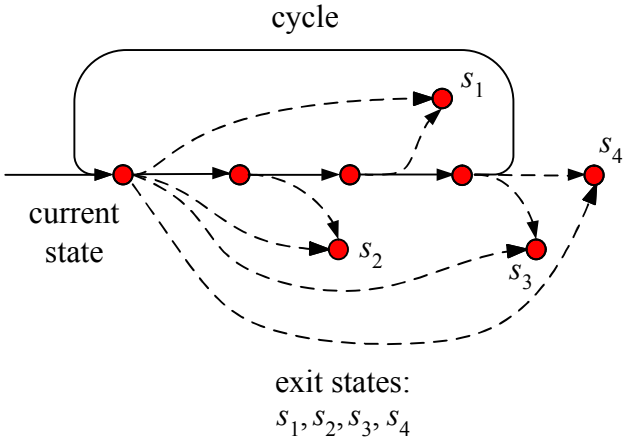

Fig. 2. Cycle Leaping

leaping" [16][17]; the technique was analyzed and refined in numerous follow-up papers [2][3][5][30][32]. Other approximate techniques include the *partial equilibrium* assumption [4] [27] and *quasi-steady state* analysis [27]. Unfortunately, such approximations are not always applicable, particularly for sensitive segments of the simulation where single-molecule events can affect the outcome; furthermore, with approximate methods, the resulting errors are generally difficult to quantify.

II. Stochasticity in Biochemical Reactions

Fixing environmental variables, such as temperature and external chemical gradients, we can assume that a cellular system behaves as a **Markov process**: The probability of future events depends only on the present state, not on the past sequence of events. Indeed, at each point in time, the probability of a given reaction occurring is a function of the current state only. It is proportional to the quantity of the reactants present as well as a rate constant.

A. Probabilistic Analysis

Consider a system consisting of M types of molecules $X_1, \ldots, X_M$, interacting through N reactions $R_1, \ldots, R_N$. For a reaction R_j let Q_j be the set of indices of the reactant types. With coefficients $q_1, \ldots, q_P$ for the reactant types, let

$$r_j = k_j \sum_{i \in Q_j} \binom{x_i}{q_i},$$

where x_i is the number of molecules of type X_i, k_j is the rate constant, and $\binom{x_i}{q_i}$ denotes the binomial coefficient. If any $x_i < q_i$, (i.e., there are insufficient molecules of a reactant for the reaction to proceed), then set $r_j = 0$. Now

$$p_j = \frac{r_j}{\sum_{k=1}^{N} r_k}$$

gives the probability that reaction R_j is the next one to fire, $j = 1, \ldots, N$.

Example 2 *For the reactions in Example 1, let the state be* $\mathbf{S} = [x_1, x_2, x_3]$. *The firing probabilities for* $R_1, R_2,$ *and* R_3 *are computed as follows:*

$$\begin{aligned}
p_1(x_1, x_2, x_3) &\equiv \frac{\frac{1}{2}x_1(x_1-1)x_2}{\frac{1}{2}x_1(x_1-1)x_2 + x_1x_3(x_3-1) + 3x_2x_3}, \\
p_2(x_1, x_2, x_3) &\equiv \frac{x_1x_3(x_3-1)}{\frac{1}{2}x_1(x_1-1)x_2 + x_1x_3(x_3-1) + 3x_2x_3}, \\
p_3(x_1, x_2, x_3) &\equiv \frac{3x_2x_3}{\frac{1}{2}x_1(x_1-1)x_2 + x_1x_3(x_3-1) + 3x_2x_3},
\end{aligned}$$

where x_1, x_2 *and* x_3 *denote the numbers of molecules of types* $X_1, X_2,$, *and* X_3, *respectively. Suppose that* $\mathbf{S} = [3, 3, 3]$. *Then the firing probabilities for* $R_1, R_2,$ *and* R_3 *are*

$$\begin{aligned}
p_1(3,3,3) &\equiv \frac{9}{9+18+27} = \frac{1}{6}, \\
p_2(3,3,3) &\equiv \frac{18}{9+18+27} = \frac{1}{3}, \\
p_3(3,3,3) &\equiv \frac{27}{9+18+27} = \frac{1}{2},
\end{aligned}$$

respectively. □

B. Biological Outcomes

Most existing methods for analysis focus on the change in the quantities of individual species as a function of time. In Gillespie's SSA, reactions are executed at random based on the probability calculations described in the previous section. The time between reactions is modelled as a Poisson process; time intervals in the simulation are obtained by sampling an exponential distribution [13]. Beginning from an initial state, the simulation is carried forward for a fixed time duration, say the average length of the cell cycle. Repeated trials are performed. The quantities of the individual species are estimated by averaging the results of repeated trials.

And yet, one-dimensional averages as a function of time are not always informative. For a variety of systems, the result is bimodal. Consider the *lambda*

bacteriophage, a virus that infects the *E. coli* bacteria. It chooses one of two survival strategies: either it integrates its genetic material with that of its host and then replicates when the bacterium divides (*lysogeny*); or else it manipulates the molecular machinery of its host to make many copies of itself, killing the bacterium in the process, and thereby releasing its progeny into the environment. The choice of which strategy to pursue, while based on environmental inputs, is probabilistic: in some cases, the virus chooses the first strategy, say with probability 0.33, and the second with probability 0.67, while in other cases the probabilities are reversed [25]. Clearly the virus is hedging its bets, an approach that provides significant advantages in an evolutionary context. Other examples include the *pap pili* epigenetic response of bacteria [20] and the lentiviral positive-feedback loop in the HIV virus [35].

We advocate a framework for analysis that focuses the probability distribution of biological *outcomes*, as summarized in Figure 3. Such outcomes are indicated by thresholds in certain molecular quantities. For instance, the decision of the lambda virus is indicated by thresholds on two of its constituent types, Cro and cII [1]: the decision to go into lysogeny is indicated by $Cro > 55$, while the decision to go lysis by $cII > 145$. These two conditions are mutually exclusive; however, this need not be the case in general. We note that the outcomes need not be simple threshold conditions; they can be arbitrarily complex *logical* functions defined on the state space.

Example 3 *For the set of reactions in Example 1, define the following outcomes:*

- C_1*: states* $\mathbf{S} = [x_1, x_2, x_3]$ *with* $x_1 \geq 12, x_2 < 12, x_3 < 12$,
- C_2*: states* $\mathbf{S} = [x_1, x_2, x_3]$ *with* $x_2 \geq 12, x_1 < 12, x_3 < 12$,
- C_3*: states* $\mathbf{S} = [x_1, x_2, x_3]$ *with* $x_3 \geq 12, x_1 < 12, x_2 < 12$.

Beginning from the state $\mathbf{S} = [3, 3, 3]$*, we perform 500 trials of stochastic simulation, Of these,*

- *36 ended in a state from which no further reactions were possible,*
- *18 ended in a state satisfying* C_1*,*
- *357 ended in a state satisfying* C_2*,*
- *89 ended in a state satisfying* C_3*.*

We conclude:

$$\Pr(C_1) = \frac{18}{500} = 0.04$$

$$\Pr(C_2) = \frac{357}{500} = 0.71,$$

$$\Pr(C_3) = \frac{89}{500} = 0.18.$$

□

III. Cycle Leaping

With a focus on outcomes, we can apply probabilistic analysis to expedite the simulation. Suppose that a trajectory enters a cycle, that is, a sequence of states through which it loops repeatedly. If there is at least one transition with non-zero probability that exits the cycle, then we can assert that the trajectory will eventually exit. If there is no such exit transition, then the simulation will remain trapped in this cycle indefinitely. We call such a cycle terminal; it corresponds to a *quasi*-equilibrium terminal condition [16].

In our algorithm, whenever a non-terminal cycle is encountered, the exit probabilities are computed and the simulation leaps directly to one of the exit states. We illustrate with an example. Consider the reactions in Figure 4. Note that the rates of the reactions on the left-hand side are much larger (by a factor of a thousand) than those on the right-hand side. Consider an initial state with a single molecule of the species A and none of the other species.

By inspection, it is apparent that reactions R_1, R_2 and R_3 are likely to fire in sequence many times – on the order of a thousand times – before either reaction R_4 or R_5 fires. Once either R_4 or R_5 fires, producing a molecule of X or Y, respectively, then the trajectory terminates since no further reactions are possible.

$$\begin{array}{lcl} R_1\text{:} & A & \stackrel{1}{\rightarrow} B \\ R_2\text{:} & B & \stackrel{2}{\rightarrow} C \\ R_3\text{:} & C & \stackrel{3}{\rightarrow} A \end{array} \qquad \begin{array}{lcl} R_4\text{:} & B & \stackrel{0.001}{\rightarrow} X \\ R_5\text{:} & C & \stackrel{0.002}{\rightarrow} Y \end{array}$$

Fig. 4. A coupled set of biochemical reactions.

We can ask: beginning with a single molecule of A, what is the probability that we get a molecule of X versus the probability that we get a molecule of Y? We could answer this question – approximately – through stochastic simulation. We would perform N trials, and count the number of times that we get an X (call this C_X) versus the number of times that we get a Y (call this C_Y). We would then estimate the probability of each event as

$$\begin{aligned} P(S_X) &\approx \frac{C_X}{N} \\ P(S_Y) &\approx \frac{C_Y}{N} \approx 1 - P(S_X). \end{aligned}$$

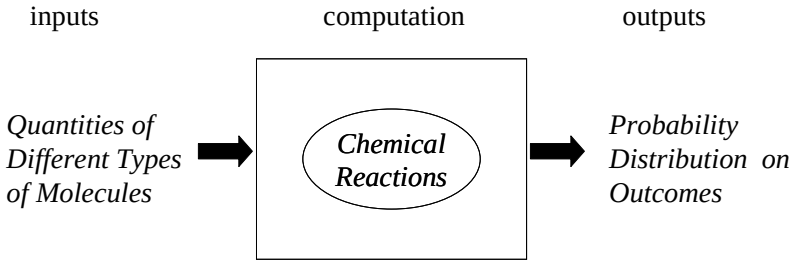

Fig. 3. Focusing on biological outcomes.

And yet, such a crude application of computing power is unnecessary in this case. We can *calculate* the probabilities exactly as follows. Denote the state with a single molecule of A as S_A and so on. As shown in Figure 5, beginning from S_A, there is only one possible transition, to S_B, so this transition has probability 1. From S_B, the probability of transitioning to S_C is

$$p = \frac{2b}{2b + (0.001)b} \approx 0.9995$$

where b denotes the number of molecules of B (here $b = 1$). The probability of transitioning to S_X is

$$1 - p \approx 0.0005.$$

From S_C the probability of transitioning to S_A is

$$q = \frac{3c}{3c + (0.002)c} \approx 0.9993$$

where c denotes the number of molecules of C (here $c = 1$). The probability of transitioning to S_Y is

$$1 - q \approx 0.0007.$$

For a trajectory that returns to S_A, we can reason as follows: in a second pass, the probability of exiting to S_X and that of exiting to S_Y will be exactly the same. Indeed, this is true for any number of passes. Accordingly, the probability that we will *eventually* exit to S_X versus exiting to S_Y is simply the relative probabilities of each occurring in single pass:

$$Pr(S_X) = \frac{Pr(S_X|\text{1st pass})}{Pr(S_X|\text{1st pass}) + Pr(S_Y|\text{1st pass})}.$$

Accordingly,

$$Pr(S_X) = \frac{1-p}{(1-p) + p(1-q)} = \frac{1-p}{1-pq} \approx 0.4287,$$

$$Pr(S_Y) = \frac{p(1-q)}{(1-p) + p(1-q)} = \frac{p(1-q)}{1-pq} \approx 0.5713.$$

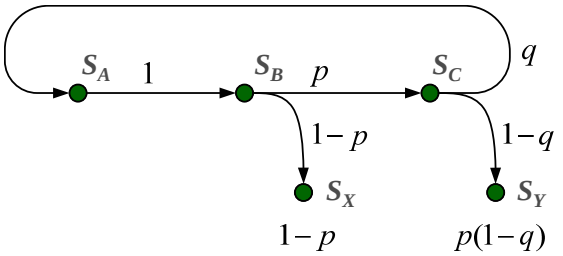

Fig. 5. Cycle in the state space for the reactions in Figure 4.

Conceptually, upon entering the cycle at S_A, we can break the transition to S_B and instead introduce two new transitions to S_X and S_Y with probabilities $P(S_X)$ and $P(S_Y)$ as shown in Figure 6.

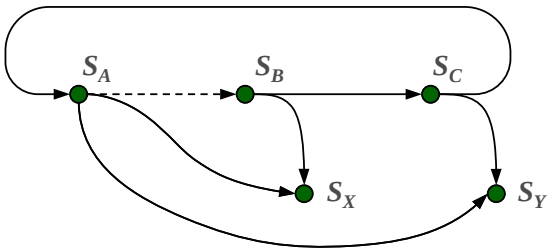

Fig. 6. Breaking the cycle in Figure 5.

Motivated by this simple example, we can consider the technique of **cycle leaping** in stochastic simulation. As trajectories are formed, a history of the states is recorded. When a cycle is encountered, the exit probabilities are computed. Then, based upon a single random number, the simulation leaps directly to one of the exit states, as illustrated in Figure 2.

Denote the transition probability from a state S_X to a state S_Y as $Pr[S_X, S_Y]$.

Algorithm 1: Cycle Leaping

For a cycle $S_1, \ldots, S_n$, compute the transition probabilities to the *exit* states:

let $p := 1$;
for i from 1 to n do

with exit states $R_1, \ldots, R_m$ from S_i,
for j from 1 to m do
let $e_{i,j} := p \times Pr[S_i, R_j]$;
end
let $p := p \times Pr[S_i, S_{i+1}]$;
end
remove the transition $S_1 \rightarrow S_2$;
let $E = \sum_{i,j} e_{i,j}$;
for i from 1 to n do
with exit states $R_1, \ldots, R_m$ from S_i,
for j from 1 to m do
add a transition from S_1 to R_j
with probability $\frac{e_{i,j}}{E}$;
end
end

□

It should be noted that there is overhead in applying cycle leaping:

- A history of the states must be maintained. For each state that is visited, a check must be performed to see if it is in the history.
- When a cycle is detected, the calculations in Algorithm 1 must be performed.

In practice, this overhead is minimal. One only targets small cycles, say 100 states in length. Accordingly, the history is a small sliding window. The calculations in Algorithm 1 are linear in the length of the cycle and linear in the number of transitions per state (to the exit states). The trajectories in many models spend 99% or more of their time in loops. With cycle leaping, the trajectories are shortened to 1% or less of their original length, and so the overhead is easily justified.

IV. Results

We discuss the application of cycle leaping in the simulation three model systems:

- The pheromone-response pathway in Baker's yeast (*Saccharomyces cerevisiae*) [1].
- Developmental Pathway Bifurcation in Phage *Lambda*-Infected *E. Coli* Cells [34].
- Self-perpetuating Epigenetic Pili Switches in Bacteria [20].

`insert table`

V. Discussion and Further Directions

With cycle leaping, multiple reactions are, in effect, executed in a single step, shortening the trajectories and resulting in more efficient utilization of random numbers. The concept of exploring reaction sequences before committing to a random choice can be applied in a more general context. Note that for sequences of several reactions, the probabilities are *multiplicative*. For instance, to compute the probability of a sequence of reactions R_1, R_2 through states S_1, S_2, we would simply multiply the probability of R_1 occurring from S_1 by the probability of R_2 occurring from S_2. If two different sequences merge to the same state, then the probabilities are *additive*. This is illustrated in Figure 7.

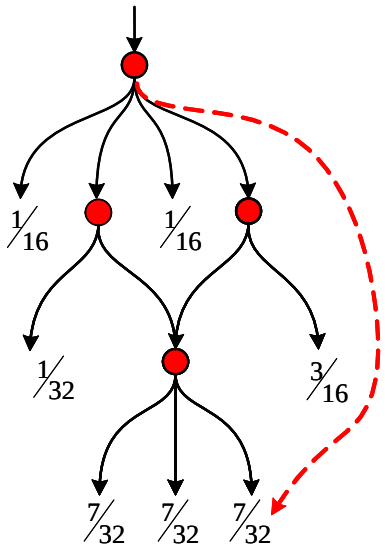

Fig. 7. Exploring sequences of reactions and leaping forward.

It is apparent that even short reaction sequences can sprawl over larger regions of the state space. If the simulation is likely to revisit the same portion of the state space, then it is judicious to record the probability calculations and later retrieve them, if needed.

We perform this caching not only for each trajectory, but also across successive trajectories. As larger and larger swathes of the state space become known, longer and longer leaps are made. The overhead of caching structured information about the state space can be considerable; accordingly, it must be carefully managed. If too much information is cached, the burden of indexing it and retrieving it can outweigh the cost of recalculating it. Nevertheless, we have found that this approach, integrated with cycle and event leaping, provides very significant improvements in the running time.

References

[1] A. Arkin, J. Ross, and H. McAdams, "Stochastic Kinetic Analysis of Developmental Pathway Bifurcation in Phage Lambda-Infected E. Coli Cells," *Genetics*, Vol. 149, No. 1633, 1998.

[2] Y. Cao, H. Li, and L. Petzold, "Efficient Formulation of the Stochastic Simulation Algorithm for Chemically Reacting Systems," *Journal of Chemical Physics*, Vol. 121, No. 9, pp. 4059–4067, 2004.

[3] Y. Cao, L. Petzold, M. Rathinam, and D. Gillespie, "The Numerical Stability of Leaping Methods for Stochastic Simulation of Chemically Reacting Systems," *Journal of Chemical Physics*, Vol. 121, No. 24, pp. 12169–12178, 2004.

[4] Y. Cao, D. Gillespie, and L. Petzold, "Multiscale Stochastic Simulation Algorithm with Stochastic Partial Equilibrium Assumption for Chemically Reacting Systems," *Journal of Computational Physics*, Vol. 206, No. 2, pp. 395–411, 2005.

[5] Y. Cao, D. Gillespie, and L. Petzold, "Avoiding Negative Populations in Explicit Poisson Tau-Leaping," *Journal of Chemical Physics*, Vol. 123, No. 5, 2005.

[6] Y. Cao and L. Petzold, "Accuracy Limitations and the Measurement of Errors in the Stochastic Simulation of Chemically Reacting Systems," *Journal of Computational Physics*, Vol. 212, No. 1, pp. 6–24, 2006.

[7] I. Chang, E. Gilbert, N. Eliashberg, and J. Keasling. "A Three-Dimensional, Stochastic Simulation of Biofilm Growth and Transport-related Factors that Affect Structure," *Microbiology*, Vol .149, 2003

[8] K. Chen, L. Calzone, A. Csikasz-Nagy, F. Cross, B. Novak, and J. Tyson, "Integrative Analysis of Cell Cycle Control in Budding Yeast," *Molecular Biology of the Cell,* Vol. 15, pp. 3841–3862, 2004.

[9] M. Elowitz, A. Levine, E. Siggia, and P. Swain, "Stochastic Gene Expression in a Single Cell," *Science*, Vol. 297, No. 5584, pp. 1183–1186, 2002.

[10] D. Endy and R. Brent, "Modelling Cellular Behaviour," *Nature*, Vol. 409, No. 6818, pp. 391–395, 2001.

[11] M. Gibson and J. Bruck, "Efficient Exact Stochastic Simulation of Chemical Systems with Many Species and Many Channels," *Journal of Physical Chemistry A*, No. 104, pp. 1876–1889, 2000.

[12] M. Gibson, "Computational Methods for Stochastic Biological Systems," *Ph.D. Disssertation*, Adviser J. Bruck, California Institute of Technology, 2001.

[13] D. Gillespie, "Exact Stochastic Simulation of Coupled Chemical Reactions," *Journal of Physical Chemistry*, No. 81, pp. 2340–2361, 1977.

[14] D. Gillespie, "A Rigorous Derivation of the Chemical Master Equation," *Physica A*, Vol. 188, No. 1-3, pp. 404–425, 1992.

[15] D. Gillespie, "The Chemical Langevin Equation," *Journal of Chemical Physics*, Vol. 113, No. 1, pp.297–306, 2000.

[16] D. Gillespie "Approximate Accelerated Stochastic Simulation of Chemically Reacting Systems," *Journal of Chemical Physics*, Vol. 115, No. 4, pp. 1716–1733, 2001.

[17] D. Gillespie and L. Petzold, "Improved Leap-Size Selection for Accelerated Stochastic Simulation," *Journal of Chemical Physics*, Vol. 119, No. 16, pp. 8229–8234, 2003.

[18] R. Heinrich and S. Shuster, "The Regulation of Cellular Systems," Chapman and Hall, 1996.

[19] A. Hernday, M. Krabbe, B. Braaten, and D. Low, "Self-perpetuating Epigenetic Pili Switches in Bacteria," *Proceedings National Academy of Sciences*, Vol. 99, No. 4, 2002.

[20] A. Hernday, B. Braaten, and D. Low, "The Intricate Workings of a Bacterial Epigenetic Switch," *Advances in Experimental Medicine and Biology*, Vol. 547, No. 83-9, 2004.

[21] D. Knuth, "The Art of Computer Programming, Vol. 2: Seminumerical Algorithms," Addison-Wesley, 1997.

[22] H. Kuwahara, C. Myers, N. Barker, M. Samoilov, and A. Arkin, "Asynchronous Abstraction Methodology for Genetic Regulatory Networks," *Workshop on Computational Methods in Systems Biology*, 2005.

[23] L. Lok, "The Need For Speed in Stochastic Simulation," *Nature Biotechnology*, Vol. 22, No. 8, 2004.

[24] L. Lok and R. Brent, "Automatic Generation of Cellular Reaction Networks with Moleculizer 1.0," *Nature Biotechnology*, Vol. 23, 2005.

[25] H. McAdams and A. Arkin, "Stochastic Mechanisms in Gene Expression," *Proceedings of the National Academy of Sciences*, Vol. 94, No. 3, pp. 814–819, 1997.

[26] B. Munsky and M. Khammash, "The Finite State Projection Algorithm for the Solution of the Chemical Master Equation," *Journal of Chemical Physics*, Vol. 124, No. 4, 2006.

[27] C. Rao and A. Arkin, "Stochastic Chemical Kinetics and the Quasi-Steady-State Assumption: Application to the Gillespie Algorithm," *Journal of Chemical Physics*, Vol. 118, No. 11, pp. 4999–5010, 2003.

[28] C. Rao, D. Wolf, and A. Arkin, "Control, Exploitation and Tolerance of Intracellular Noise," *Nature*, Vol. 420, No. 6912, pp. 231–237, 2002.

[29] J. Raser and E. O'Shea, "Noise in Gene Expression: Origins, Consequences, and Control," *Science*, Vol. 309, No. 5743, pp. 2010–2013, 2005.

[30] M. Rathinam, L. Petzold, Y. Cao, and D. Gillespie, "Stiffness in Stochastic Chemically Reacting Systems: The Implicit Tau-leaping Method," *Journal of Chemical Physics*, Vol. 119, No. 24, pp. 12784–12794, 2003.

[31] L. Salwinski and D. Eisenberg, "In Silico Simulation of Biological Network Dynamics," *Nature Biotechnology*, Vol. 22, No. 8, 2004.

[32] T. Tian and K. Burrage, "Binomial Leap Methods for Simulating Stochastic Chemical Kinetics," *Journal of Chemical Physics*, Vol. 121, No. 21, pp. 10356–10364, 2004.

[33] M. Tomita, "Whole-Cell Simulation: A Grand Challenge of the 21st Century," *Trends in Biotechnology*, Vol. 19, No. 6, 2001.

[34] T. Thomson, `http://yeastpheromonemodel.org/wiki/Main_Page`

[35] L. Weinberger, J. Burnett, J. Toettcher, A. Arkin, and D. Schaffer, "Stochastic Gene Expression in a Lentiviral Positive-Feedback Loop: HIV-1 Tat Fluctuations Drive Phenotypic Diversity," *Cell*, Vol. 122, 2005.

[36] E. Voit, "Computational Analysis of Biochemical Systems: A Practical Guide for Biochemists and Molecular Biologists," Cambridge Univ. Press, 2000.