Published in final edited form as:

Biotechnol Prog. 2008 ; 24(1): 56–61. doi:10.1021/bp070255h.

Algorithms and Software for Stochastic Simulation of Biochemical Reacting Systems *

Hong Li[†], Yang Cao[‡], Linda R. Petzold[§], and Daniel T. Gillespie[¶]

[†]*Department of Computer Science, University of California Santa Barbara. CA 93106. email: hongli@cs.ucsb.edu*

[‡]*Department of Computer Science, Virginia Tech. VA 24061. email: ycao@cs.vt.edu*

[§]*Department of Computer Science, University of California Santa Barbara. CA 93106. email: petzold@cs.ucsb.edu*

[¶]*Dan T. Gillespie Consulting, 30504 Cordoba Place. Castaic, CA 91384. email: GillespieDT@mailaps.org*

Abstract

Traditional deterministic approaches for simulation of chemically reacting systems fail to capture the randomness inherent in such systems at scales common in intracellular biochemical processes. In this article we briefly review the state of the art in discrete stochastic and multiscale algorithms for simulation of biochemical systems and we present the $S_{TOCHKIT}$ software toolkit.

1 Introduction

The time evolution of well stirred chemically reacting systems has traditionally been simulated by solving a set of coupled ordinary differential equations (ODEs). Although the deterministic formulation is sufficient in many cases, it fails to capture the inherent stochasticity in many biochemical systems formed by living cells,^{1–4} in which the small population of a few critical reactant species can result in discrete and stochastic behavior. The dynamics of those systems can be simulated accurately using the machinery of Markov process theory, specifically the stochastic simulation algorithm (SSA) of Gillespie.^{1, 2} SSA, an essentially exact procedure for generating realizations of the chemical master equation (CME), is in widespread use for the stochastic simulation of biochemical systems. But for many real biochemical systems, the computational cost of simulation by SSA can be prohibitively high. This is due to the fact that SSA must simulate every reaction event. When there are large populations of some chemical species, and/or fast reactions involved in the system, a great many reaction events must be simulated.

It is very important to have a fast SSA implementation. This algorithm is likely to be a part of any accelerated and multiscale discrete stochastic methodology. A number of authors have proposed successively more efficient formulations of SSA. The Next Reaction Method (NRM),⁵ Optimized Direct Method (ODM),⁶ Sorting Direct Method (SDM),⁷ fast Kinetic Monte Carlo Method (KMC),⁸ KMC with minimal searching⁹ and Logarithmic Direct Method (LDM)¹⁰ with sparse matrix state update are mathematically equivalent formulations which offer improved performance over the original SSA. Often SSA is used to generate ensembles

*This work was supported by the California Institute of Technology under DARPA Award No. F30602-01-2-0558, by the U. S. Department of Energy under DOE award No. DE-FG03-00ER25430, by the NIH under awards GM078993 and GM075297, and by the Institute for Collaborative Biotechnologies through grant DAAD19-03-D-0004 from the U. S. Army Research Office.

of stochastic realizations from which approximate probability density functions for species populations can be obtained. Parallel computation has been used to speed up the calculation of such ensembles.¹¹ Some of our current work and others¹² explores the use of novel computer architectures to speed up SSA simulation. However, there is a limit to how much SSA can be sped up, given that it must simulate every reaction event in the system and there are usually a great many reaction events.

Approximate accelerated stochastic methods^{13–17} have been developed to speed up the stochastic simulation. The first such method is *tau-leaping*.¹³ Tau-leaping advances the simulation by a pre-selected time τ which can often be chosen large enough to encompass more than one reaction event. Stiffness (the presence of both fast and slow reactions) is often an issue in approximate discrete stochastic simulation of chemically reacting systems, just as it is an important consideration in the deterministic simulation of chemically reacting systems. To improve the numerical stability for stiff problems, *implicit tau-leaping*¹⁸ and the *trapezoidal tau-leaping methods*¹⁹ have been proposed. Another, complementary means of accelerating discrete stochastic simulation is the use of hybrid methods,^{14, 20} which approximate the fast reactions involving species with large populations via ODEs, and the slow reactions involving species with small populations via SSA. Although hybrid methods can be quite effective in solving some problems, they cannot efficiently handle the situation of a system with fast reactions involving a chemical species for which the population is small. The slow-scale SSA (ssSSA)¹⁵ was developed to deal with this type of situation. The ssSSA makes use of a stochastic partial equilibrium approximation to advance the system at the scale of the slow reactions.

There is a great need for reliable and efficient software that makes these state of the art simulation tools available to the systems biology community. To this end, we have developed **STOCHKIT**, a software toolkit for discrete stochastic and multiscale simulation of chemically reacting systems. The aim of **STOCHKIT** is to make reliable and efficient stochastic and multiscale simulation accessible to practicing biologists and chemists, while remaining open to extension by accommodating new algorithms and implementations.

In this paper we review the recent work on algorithms for discrete stochastic and multiscale simulation of biochemical reaction networks. We introduce **STOCHKIT** and briefly describe some early success stories with **STOCHKIT**.

2 Discrete Stochastic Simulation Methods

Our concern here is with a system of molecules of N chemical species $\{S_1, \dots, S_N\}$ which interact through M chemical reactions channels $\{R_1, \dots, R_M\}$. We assume the system to be “well-stirred,” and confined to some constant volume Ω at a constant temperature. The state of the system at time t is specified by the vector $\mathbf{X}(t) \equiv (X_1(t), \dots, X_N(t))$, where $X_i(t)$ is the number of S_i molecules in the system at time t .

Each reaction channel R_j is assumed to be characterized mathematically by two quantities. The first is its *state-change vector* $\mathbf{v}_j \equiv (v_{1j}, \dots, v_{Nj})$, where v_{ij} is the change in the S_i molecular population caused by one R_j reaction; i.e., R_j causes the system to jump, essentially instantaneously, from its present state $\mathbf{x}$ to state $\mathbf{x} + \mathbf{v}_j$. The other characterizing quantity for R_j is its *propensity function* a_j . This is defined so that $a_j(\mathbf{x}) dt$ is the probability, given $\mathbf{X}(t) = \mathbf{x}$, that one R_j reaction will occur somewhere inside the system in the next infinitesimal time interval $[t, t + dt)$. For a reaction of the form $S_1 \rightarrow \text{product(s)}$, $a_j(\mathbf{x})$ typically has the form $c_j x_1$, where c_j is some constant. For a reaction of the form $S_1 + S_2 \rightarrow \text{product(s)}$, $a_j(\mathbf{x})$ typically has the form $c_j x_1 x_2$, or $c_j \frac{1}{2} x_1 (x_1 - 1)$ if $S_2 = S_1$, where again c_j is some constant.

2.1 The SSA and its Various Formulations

The basic procedure for generating trajectories or “realizations” of $\mathbf{X}(t)$ is called the *stochastic simulation algorithm* (SSA).^{1, 2} The theoretical basis for the SSA is the function $p(\tau, j | \mathbf{x}, t)$, which is defined so that $p(\tau, j | \mathbf{x}, t) d\tau$ is the probability, given $\mathbf{X}(t) = \mathbf{x}$, that the *next* reaction in the system will occur in the infinitesimal time interval $[t + \tau, t + \tau + d\tau)$, and will be an R_j reaction. This function is thus the joint probability density function of the two random variables “time to the next reaction” (τ) and “index of the next reaction” (j), given that the system is currently in state $\mathbf{x}$. It is not hard to show, by applying the laws of probability, that¹

$$p(\tau, j | \mathbf{x}, t) = a_j(\mathbf{x}) e^{-a_0(\mathbf{x})\tau} = a_0(\mathbf{x}) e^{-a_0(\mathbf{x})\tau} \times \frac{a_j(\mathbf{x})}{a_0(\mathbf{x})}, \quad (1)$$

where $a_0(\mathbf{x}) \equiv \sum_{k=1}^M a_k(\mathbf{x})$. The last form here implies that τ is an exponential random variable with mean $1/a_0(\mathbf{x})$, and j is a statistically independent integer random variable with point probabilities $a_j(\mathbf{x})/a_0(\mathbf{x})$.

There are several statistically equivalent formulations of SSA for generating samples of τ and j according to these distributions. Perhaps the simplest is the so-called *direct method* (DM):^{1, 2} We draw two random numbers r_1 and r_2 from the uniform distribution in the unit-interval and take

$$\tau = \frac{1}{a_0(\mathbf{x})} \ln \left(\frac{1}{r_1} \right), \quad (2a)$$

$$j = \text{the smallest integer satisfying } \sum_{k=1}^j a_k(\mathbf{x}) > r_2 a_0(\mathbf{x}). \quad (2b)$$

Therefore, for the system in state $\mathbf{x}$ at time t , we can advance the system to the next reaction by computing τ and j according to these formulas and then replacing $t \leftarrow t + \tau$ and $\mathbf{x} \leftarrow \mathbf{x} + \mathbf{v}_j$.

There are several other exact algorithms for moving the system forward in time to the next reaction according to the joint density function (1). The *First reaction method* (FRM)^{1, 2} generates a putative time $\tau_k = a_k^{-1}(\mathbf{x}) \ln(1/r_k)$ to the next R_k event for $k = 1, \dots, M$, and then chooses τ to be the smallest of the τ_k 's, and j the index of the smallest. This method is usually slower than the DM because it requires M unit-interval uniform random numbers r_k . Execution time is in fact the main limitation of the SSA, because generating every reaction event one at a time can be very time-consuming if, as is usually the case, an enormous number of reaction events occurs during the times of interest. But since most accelerated “multiscale” or “hybrid” methods use the SSA as a core component, it is important to make the SSA as efficient as possible.

The *next reaction method* (NRM)⁵ is a heavily modified version of the FRM that in some cases is more efficient than the DM. It essentially saves the putative next firing times of all reaction channels in an indexed binary tree priority queue, and it uses a dependency graph to determine which propensities need to be updated after each reaction. However, the maintenance overhead for these data structures limits the benefits of the FRM to very large, very loosely coupled

systems. The *optimized direct method* (ODM)⁶ increases the efficiency of the reaction-selection step (2b), a key bottleneck in the DM, by pre-ordering the reactions so that those with larger propensity functions have smaller index values. The *sorting direct method* (SDM)⁷ carries this strategy one step further by dynamically re-ordering the reactions via a bubble-up sorting method, in which each time a reaction fires it swaps indices with the next lower-indexed reaction. Formally, the computation times of both the ODM and the SDM are $O(M)$.

SSA may be viewed as a special kind of Kinetic Monte Carlo algorithm which is applied to well-mixed chemically reacting systems. Since SSA is always applied to well-mixed chemically reacting systems, it has been possible to use the structure of these problems to put it on a solid theoretical foundation. A fast algorithm for KMC has been proposed⁹ which fixes the search depth to be $O(\log M)$ by reusing the intermediate data during the calculation of a_0 . This algorithm is independent of the order of reactions. Further efficiency can be achieved in the implementation by using sparse matrix techniques in the system state update stage.¹⁰

2.2 Tau-leaping and Beyond

Significant speedups in the SSA will inevitably involve *approximations* of one kind or another. One prominent approximate acceleration procedure is the *tau-leaping* simulation method,¹³ which also provides the theoretical connection between the SSA and the deterministic ODEs of traditional chemical kinetics.²¹ The basic idea of tau-leaping is to advance the system by a *pre-selected* time τ during which many reactions occur. To do this accurately, we must choose τ small enough that no propensity function changes “appreciably” during τ . To the extent that this *leap condition* is satisfied, then given the system in state $\mathbf{x}$ at time t , the number of times that reaction R_j will fire in $[t, t + \tau)$ will be approximately $\mathcal{P}_j(a_j(\mathbf{x})\tau)$, the Poisson random variable with mean (and variance) $a_j(\mathbf{x})\tau$. This leads to the basic (approximate) update formula for tau-leaping:

$$\mathbf{X}(t+\tau) \doteq \mathbf{x} + \sum_{j=1}^M \mathcal{P}_j(a_j(\mathbf{x})\tau) \nu_j. \quad (3)$$

Practical considerations cause the actual implementation of tau-leaping to be more complicated than simply substituting into formula (3) the current state $\mathbf{x}$ and generating M Poisson random numbers with the indicated means. First is the problem of choosing the leap time τ . The latest recipe for doing this²² efficiently estimates the largest τ for which the expected fractional change in every propensity function will be bounded by a user-specified amount ε (typically $\varepsilon = 0.04$). This is tricky to do since the change induced in any propensity function during a leap by τ will be a random variable, so the estimation procedure needs to take account of not only the mean of that change but also the fluctuations in that change about its mean. Second is the problem of ensuring that no reactant population be driven negative in a leap. This problem was originally attributed to the unbounded character of the Poisson random variable; however it was subsequently found to be more often a consequence of two or more reaction channels independently consuming the same reactant. One resolution of this problem²³ is to classify as “critical” all reactions that are within a user-specified number n_c firings of exhausting some reactant, and then to allow *no more than one* firing of a critical reaction during a leap (typically $n_c = 10$). This procedure²³ has the additional advantage that, as all reactions become critical, or equivalently as $n_c \rightarrow \infty$, it segues smoothly to the exact SSA. This is useful since (3) as it stands does not segue to the SSA in a computationally efficient way.

The foregoing “explicit” tau-leaping procedure has been shown to work well provided the time-scales of all the reactions are not too different. But it will seem very slow when applied to systems that evolve on widely separated timescales, because the leap condition generally

restricts τ to the smallest (fastest) timescale. Much work has been devoted recently to finding approximate workarounds for such systems. For overtly “stiff” systems – systems for which the fastest evolving modes are stable – a variety of methods has been proposed. The *implicit tau-leaping* method¹⁸ is a generalization of the implicit Euler method for handling stiff ordinary differential equations. It basically replaces (3) with

$$\mathbf{X}(t+\tau) \doteq \mathbf{x} + \sum_{j=1}^M [\mathcal{P}_j(a_j(\mathbf{x})\tau) - a_j(\mathbf{x})\tau + a_j(\mathbf{X}(t+\tau))\tau] \nu_j. \quad (4)$$

This formula must be solved implicitly for the state at time $t + \tau$; usually this is done numerically using Newton's method. A slight modification of (4) will keep the state of the system on the integer lattice. Implicit tau-leaping avoids the stability limitations of explicit tau-leaping, and thus allows much larger values of τ to be used. Its main drawback is that it overdamps the natural fluctuations of the *fast* species, and requires that a strategy called “down-shifting” be applied to recover those fast species fluctuations whenever required. The *trapezoidal tau-leaping method*¹⁹ puts a factor of 1/2 in front of the last two terms in brackets in (4). It has been shown to improve simulation accuracy for some systems. The most recent improvement in tau-leaping is the *adaptive explicit-implicit tau-leaping* procedure,²⁴ which automatically identifies when stiffness is present and dynamically chooses between the explicit and implicit tau-leaping schemes.

Another approach to handling stiff systems involves a stochastic generalization of the quasi-steady state and partial equilibrium methods of deterministic chemical kinetics.^{14–16, 20, 25–31} The theoretical basis for these methods is captured most comprehensively by the *slow-scale SSA*^{15,27,31} and its practical implementation as the *multiscale SSA*.¹⁶ The slow-scale SSA is a systematic procedure for partitioning the system into fast and slow reactions and species, and then simulating *only the slow reactions*, using specially modified propensity functions.

Implementing this procedure requires one to estimate, either theoretically or numerically, certain very specific features of the process consisting of the fast species undergoing *only* the fast reactions – the so-called the “virtual fast process”. When this can be done with sufficient accuracy, enormous speedups can be achieved for very stiff systems without appreciable loss of accuracy. The important Michaelis-Menten enzyme-substrate reaction typically falls within the provenance of the slow-scale SSA.

3 StochKit– Stochastic Simulation Toolkit

3.1 What is StochKit

STOCHKIT is an efficient, extensible stochastic simulation toolkit developed in C++ that aims to make state of the art stochastic simulation algorithms accessible to biologists and chemists, while remaining open to extension via new stochastic and multiscale algorithms. STOCHKIT consists of a suite of software for stochastic simulation. The STOCHKIT core implements the simulation algorithms. Additional tools are provided for the convenience of the simulation and analysis. A typical simulation process of STOCHKIT is shown in Figure 3.1

The intended audience for STOCHKIT can be divided into two distinct groups: those doing research on and development of stochastic simulation methods, and those seeking to employ such methods to further their biological or chemical research. Thus the package is designed to be both simple to use and easy to extend. STOCHKIT is freely available for download at www.engr.ucsb.edu/~cse.

3.2 The Core Package of StochKit

The `STOCHKIT` core implements state of the art stochastic simulation algorithms through a unified interface. These algorithms include Gillespie's SSA algorithm,^{1,2} the optimized SSA algorithm,⁶ tau-leaping methods (explicit tau-leaping,¹³ implicit tau-leaping¹⁸ and trapezoidal tau-leaping¹⁹ with fixed stepsize, and adaptive stepsize, nonnegativity-preserving explicit tau-leaping method^{22,23}), slow scale SSA¹⁵ and multiscale SSA¹⁶ methods. These algorithms were implemented as modules. Users who are not experts in stochastic simulation can simply use the default option (1 for SSA and 0 for adaptive explicit tau-leaping method) to let `STOCHKIT` automatically select the simulation algorithm and the corresponding stepsize. Advanced users can select different options for better performance or accuracy. Those developing new algorithms (or simply refining existing ones) need only supply a new module that captures their particular innovation (i.e. stepsize selection, single step execution, data management, etc.). Basic matrix and vector operations are provided in the Math library of `STOCHKIT`, which facilitates the development of new modules. Examples are provided to explain how to use and extend these algorithms (see 11). The capacity to run an ensemble of stochastic simulations is provided.

Besides the convenience of usage and extension, special attention has been paid to the accuracy of simulation results by `STOCHKIT`. An essential concern is the pseudorandom number generator. High quality pseudorandom number generation is the cornerstone of any stochastic simulation system. In fact, statistical results can only be relied upon if the independence of the samples can be guaranteed. Thus, a high quality pseudorandom number generator is crucial for `STOCHKIT`. The standard library routines `rand()` from C could in principle be used for a uniform random number generator in our simulations. However, favoring speed over quality, `rand()` usually produces a short random sequence period which leads to a realistic possibility of random sequence repetition for simulations that require huge amounts of random numbers. Thus `STOCHKIT` uses the Scalable Parallel Random Number Generators Library (SPRNG),^{32, 33} which provides multiple high-quality pseudorandom uniform number generators. In addition, SPRNG provides a facility for generating uncorrelated random numbers in parallel. Non-uniform distributions are generated by *ranlib.c*,³⁴ which provides generators for a wide variety of distributions. Nevertheless, we adapted *ranlib.c* to use SPRNG's linear congruential generator as its uniform generator, to further minimize the probability of sequence repetition. A simple C++ wrapper was provided for the relevant *ranlib.c* routines to improve user-friendliness.

3.3 Useful Tools

In addition to the core simulation package, `STOCHKIT` provides three useful tools to support stochastic simulation and analysis. These tools include a simple converter to translate an SBML model file to the input files required by `STOCHKIT`, a data analyzer to calculate and compare the statistical information from simulation results, and a convenient MPI interface which enables the Monte Carlo simulation ensemble to run on a parallel cluster.

SBML2StochKit Converter—SBML (System Biology Markup Language)³⁵ is a computer-readable format for representing models of biochemical reaction networks. Many biochemical models have been represented with SBML files. For the convenience of SBML users, `STOCHKIT` provides this tool to convert an SBML³⁵ file to the input files required by `STOCHKIT`. With this converter, users can conveniently construct their problem files using a separate SBML model builder, make the conversion and run the simulation using `STOCHKIT`. The current version of `STOCHKIT` supplies only a command-line implementation for the translation. A more user-friendly GUI is under development.

DataAnalyzer—For stochastic modeling, it is important to collect the statistical information from an ensemble of many independent simulations. The DataAnalyzer is a simple `MATLAB` package to generate and plot statistical information from an ensemble. Moreover, the DataAnalyzer provides functions to evaluate the distribution differences³⁶ between multiple ensembles. For example, the accuracy of different algorithms with different parameters can be measured by calculating the distribution distance³⁶ between the probability density functions (PDFs) generated via the ensemble from the simulation and the corresponding PDFs from the experimental data, or from an ensemble of an “exact” simulation such as SSA. Figure 2 shows a plot generated by DataAnalyzer, which gives the histogram distance between ensembles of the state variable generated from the SSA and the explicit tau-leaping method with fixed tau value $\tau = 0.4$ for the Schlögl model.^{37,38}

MPI Parallel Toolbox—In many applications, one has to run a large number of stochastic realizations to collect the ensemble and study the statistics. This task is naturally suited to parallel computation. `STOCHKIT` provides an MPI toolbox for running many simulations on multiple processors using MPI protocol. We recommend using SPRNG to generate the random numbers since it provides better performance and accuracy, which is particularly important in parallel random number generation on a parallel machine.

4 Success Stories

Although the development of `STOCHKIT` is still in its early stages, there are already around 100 `STOCHKIT` users worldwide, including model developers who use `STOCHKIT` to run stochastic simulation and algorithm developers who use and modify `STOCHKIT` to conduct research on stochastic simulation algorithms. Success stories come from both sides.

Most of the initial success stories come from the side of algorithm development. `STOCHKIT` provides a simulation framework and a unified interface to help stochastic simulation algorithm research. The modular feature of this package makes the algorithm extension very convenient. For example, if a new stepsize selection formula or even a new simulation formula is to be developed, one only needs to replace the corresponding module in `STOCHKIT` with the newly developed one and make full use of other modules that are already in the package to generate the simulation and make the comparison. In this way, we have successfully developed the adaptive tau-leaping^{22,23} algorithm using `STOCHKIT`. The convergence³⁹ and stability⁴⁰ of tau-leaping methods have also been studied with the help from `STOCHKIT`. The convenience of extension and the ability to use parallel computation to generate a large number of independent simulations makes the research cycle dramatically shorter. Besides our own algorithm development, success stories also come from other research groups. Gunawan et al⁴¹ has used `STOCHKIT` to generate ensembles with different parameters and conducted the parametric sensitivity analysis. Kim et al⁴² has used `STOCHKIT` to speed up their research on the spectral method for sensitivity analysis. Munsky and Khammash⁴³ have developed a numerical algorithm to approximately solve the chemical master equation by comparing ensemble data generated from `STOCHKIT`.

On the other side, model developers have made use of this powerful simulation and analysis package to speed up their model development. The most successful application is in the model development of the gene regulatory networks in the heat-shock response (HSR) of *E. Coli*.^{44,45} This model exhibits a multiscale and stochastic nature, which makes the system very stiff. By using multiscale SSA¹⁶ in `STOCHKIT`, the numerical simulation of the HSR model is 100 times faster than the direct SSA. The time savings in the numerical simulation helped

¹The Schlögl model is a famous chemical reaction system that for certain parameters exhibits bistable behavior in the state variable. Details are available in the references.^{37, 38}

shorten the total time of model development. Another ongoing example is the stochastic cell cycle model of budding yeast⁴⁶ by Tyson's research group. They have successfully combined the simulation power of *StochKit* with the convenient GUI from their model development tool *JigCell*.⁴⁷ With the help of *StochKit* they have been able to model, simulate and compare the statistics given by the model and the experimental data.⁴⁸

Acknowledgement

We would like to thank Andrew Hall, Sotiria Lampoudi and the rest of the *StochKit* team for their contributions to the development of the *StochKit* package.

References

1. Gillespie DT. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J. Comp. Phys* 1976;22:403–434.
2. Gillespie DT. Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem* 1977;81:2340–2361.
3. McAdams HH, Arkin A. Stochastic mechanisms in gene expression. *Proc. Natl. Acad. Sci. USA* 1997;94:814–819. [PubMed: 9023339]
4. Arkin A, Ross J, McAdams HH. Stochastic kinetic analysis of developmental pathway bifurcation in phage λ -infected *E. Coli* cells. *Genetics* 1998 Aug;149:1633–1648. [PubMed: 9691025]
5. Bruck J, Gibson M. Efficient exact stochastic simulation of chemical systems with many species and many channels. *J. Phys. Chem* 2000;105:1876–1889.
6. Li H, Cao Y, Petzold L. Efficient formulation of the stochastic simulation algorithm for chemically reacting systems. *J. Phys. Chem* 2004;121(9):4059–4067.
7. McColluma JM, Peterson GD, Coxc CD, Simpson ML, Samatova NF. The sorting direct method for stochastic simulation of biochemical systems with varying reaction execution behavior. *J. Comput. Biol. Chem* 2005 Feb.;30:39–49.
8. Beichl I, Blue J, Sullivan F. Faster Monte Carlo simulations. *Physical Rev. E* 1995;51:867–868.
9. Schulze TP. Kinetic Monte Carlo simulations with minimal searching. *Physical Review E* 2002;65(3):036704.
10. Li, H.; Petzold, L. Technical report. Santa Barbara: Department of Computer Science, University of California; 2006. Logarithmic direct method for discrete stochastic simulation of chemically reacting systems. <http://www.engr.ucsb.edu/~cse>
11. StochKit Team. User's guide for stochkit. <http://www.engr.ucsb.edu/~cse>
12. Yoshimi M, Osana Y, Iwaoka Y, Funahashi A, N-Hiroi, Shibata Y, Iwanaga N, Kitano H, Amano H. The design of scalable stochastic biochemical simulator on FPGA. *Proc. of I. C. on Field Programmable Technologies (FPT2005)* 2005:139–140.
13. Gillespie DT. Approximate accelerated stochastic simulation of chemically reacting systems. *J. Chem. Phys* 2001;115(4):1716–1733.
14. Rao C, Arkin A. Stochastic chemical kinetics and the quasi steady-state assumption: application to the Gillespie algorithm. *J. Chem. Phys* 2003;118:4999–5010.
15. Gillespie D, Cao Y, Petzold L. The slow-scale stochastic simulation algorithm. *J. Chem. Phys* 2005;122:014116.
16. Gillespie D, Cao Y, Petzold L. Multiscale stochastic simulation algorithm with stochastic partial equilibrium assumption for chemically reacting systems. *J. Comp. Phys* 2005;206(2):395–411.
17. Frankowicz M, Moreau M, Szczesny PP, Tóth J, Vicente L. Fast variables elimination in stochastic kinetics. *J. Phys. Chem* 1993;97:1891–1895.
18. Rathinam M, Cao Y, Petzold LR, Gillespie DT. Stiffness in stochastic chemically reacting systems: The implicit tau-leaping method. *Journal of Chemical Physics*. 2003
19. Cao Y, Petzold L. Trapezoidal tau-leaping formula for the stochastic simulation of biochemical systems. *Proceedings of Foundations of Systems Biology in Engineering* 2005:149–152.FOSBE

20. Haseltine EL, Rawlings JB. Approximate simulation of coupled fast and slow reactions for stochastic chemical kinetics. *J. Chem. Phys* 2002;117(15):6959–6969.
21. Gillespie DT. The chemical Langevin equation. *J. Chem. Phys* 2000;113:297–306.
22. Gillespie D, Cao Y, Petzold L. Efficient stepsize selection for the tau-leaping method. *J. Chem. Phys* 2006;124:044109. [PubMed: 16460151]
23. Gillespie D, Cao Y, Petzold L. Avoiding negative populations in explicit tau leaping. *J. Chem. Phys* 2005;123:054104–054112. [PubMed: 16108628]
24. Gillespie D, Cao Y, Petzold L. The adaptive explicit-implicit tau-leaping method with automatic tau selection. *J. Chem. Phys.* 2006submitted to
25. Goutsias J. Quasiequilibrium approximation of fast reaction kinetics in stochastic biochemical systems. *J. Chem. Phys* 2005;122(18):184102. [PubMed: 15918689]
26. Samant A, Vlachos DG. Overcoming stiffness in stochastic simulation stemming from partial equilibrium: A multiscale Monte Carlo algorithm. *J. Chem. Phys* 2005;123
27. Gillespie DT, Cao Y, Petzold LR. Accelerated stochastic simulation of the stiff-enzyme-substrate reaction. *J. Chem. Phys* 2005;123
28. Haseltine EL, Rawlings JB. On the origins of approximations for stochastic chemical kinetics. *J. Chem. Phys* 2005;123
29. Liu D, E W, Vanden-Eijnden E. Nested stochastic simulation algorithm for chemical kinetic systems with disparate rates. *J. Chem. Phys* 2005;123
30. Kaznessis Y, Salis H. An equation-free probabilistic steady-state approximation: dynamic application to the stochastic simulation of biochemical reaction networks. *J. Chem. Phys* 2005;123(21):214106. [PubMed: 16356038]
31. Cao Y, Gillespie D, Petzold L. Comment on nested stochastic simulation algorithm for chemical kinetic systems with disparate rates [*J. Chem. Phys.* 123, 194107 (2005)]. *J. Chem. Phys* 2007;126:137101. [PubMed: 17430072]
32. Mascagni, M. Proceedings of the Ninth SIAM Conference on Parallel Processing for Scientific Computing. San Antonio: Texas; 1999. SPRNG: A scalable library for pseudorandom number generation.
33. Mascagni M, Srinivasan A. SPRNG: A scalable library for pseudorandom number generation. *ACM Transactions on Mathematical Software* 2000;volume 26:436–461.
34. Brown, BW.; Lovato, J.; Russell, K. RANLIB.C Library of C Routines for Random Number Generation. M.D.: Anderson Cancer Center, The University of Texas; 1991.
35. Hucka M, Finney A, Sauro HM, Bolouri H, Doyle JC, Kitano H. The systems biology markup Language (SBML): a medium for representation and exchange of biochemical network models. *Bioinformatics* 2003;19(4):524–531. [PubMed: 12611808]
36. Cao Y, Petzold L. Accuracy limitations and the measurement of errors in the stochastic simulation of chemically reacting systems. *J. Comput. Phys* 2006;212:6–24.
37. Schlöogl F. On thermodynamics near a steady state. *Zeitschrift für Physik* 1971;248:446–458.
38. Gillespie, D. Markov Processes: An Introduction for Physical Scientists. Academic Press; 1992.
39. Cao Y, Rathinam M, Petzold L, Gillespie D. Consistency and stability of tau leaping schemes for chemical reaction systems. *SIAM Multiscale Modeling* 2005;volume 4:867–895.
40. Rathinam M, Cao Y, Petzold L, Gillespie D. The numerical stability of leaping methods for stochastic simulation of chemically reacting systems. *J. Chem. Phys* 2004;121(24):12169–12178. [PubMed: 15606235]
41. Petzold L, Gunawan R, Cao Y, Doyle FJ III. Sensitivity analysis of discrete stochastic systems. *J. Biophys* 2005;88:2530–2540.*
42. Debusschere BJ, Kim D, Najm HN. Spectral methods for parametric sensitivity in stochastic dynamical systems. *J. Biophys* 2007;92:379–393.
43. Khammash M, Minsky B. The finite state projection algorithm for the solution of the chemical master equation. *J. Chem. Phys* 2006;124(4):044104. [PubMed: 16460146]
44. Kurata, H.; El-Samad, H.; Yi, T.; Khammash, M.; Doyle, J. Feedback regulation of the heat shock response in *E. Coli*. Proceedings of the 40th IEEE conference on Decision and Control; 2001. p. 837–842.

45. Khammash, M.; Kurata, H.; Doyle, J. Stochastic analysis of the heat shock response in E. Coli. 3rd International Conference on Systems Biology; 2002.
46. Tyson JJ, Novak B. Regulation of the eukaryotic cell cycle: Molecular antagonism, hysteresis, and irreversible transitions. J. Theoretical Biology. 2001
47. Vass M, Allen N, Shaffer CA, Ramakrishnan N, Watson LT, Tyson JJ. The jigcell model builder and run manager. Bioinformatics 2004;18:3680–3681. [PubMed: 15273159]
48. Tyson, JJ. , . Personal communication.

Figure 1.
Simulation Process of **StochKit**

Figure 2.

A histogram distance plot for the Schlögl model.^{37,38} This plot is based on 10,000 samples of the state variable generated from the SSA and the explicit tau-leaping method with $\tau = 0.4$.