Cell Cycle Modeling for Budding Yeast with Stochastic Simulation Algorithms

Tae-Hyuk Ahn[1], Layne T. Watson[1,2], Yang Cao[1]
Clifford A. Shaffer[1] and William T. Baumann[3]

Abstract: For biochemical systems, where some chemical species are represented by small numbers of molecules, discrete and stochastic approaches are more appropriate than continuous and deterministic approaches. The continuous deterministic approach using ordinary differential equations is adequate for understanding the average behavior of cells, while the discrete stochastic approach accurately captures noisy events in the growth-division cycle. Since the emergence of the stochastic simulation algorithm (SSA) by Gillespie, alternative algorithms have been developed whose goal is to improve the computational efficiency of the SSA. This paper explains and empirically compares the performance of some of these SSA alternatives on a realistic model. The budding yeast cell cycle provides an excellent example of the need for modeling stochastic effects in mathematical modeling of biochemical reactions. This paper presents a stochastic approximation of the cell cycle for budding yeast using Gillespie's stochastic simulation algorithm. To compare the stochastic results with the average behavior, the simulation must be run thousands of times. A load balancing algorithm improved overall performance on a parallel supercomputer.

Keywords: Stochastic simulation algorithm (SSA), cell cycle, budding yeast, parallel computing, load balancing.

1 Introduction

The cell-division cycle is the sequence of events that take place in a eukaryotic cell leading to its replication. A growing cell replicates all its components and di-

---

[1] Department of Computer Science, Virginia Polytechnic Institute and State University, Blacksburg, VA 24061-0106, USA.

[2] Department of Mathematics, Virginia Polytechnic Institute and State University, Blacksburg, VA 24061-0123, USA.

[3] Department of Electrical and Computer Engineering, Virginia Polytechnic Institute and State University, Blacksburg, VA 24061-0111, USA.

vides them into two daughter cells, so that each daughter has the information and machinery necessary to repeat the process [Murray and Hunt (1993)]. Mathematical modeling and computational methods are needed to understand complex yeast control systems. Deterministic mathematical modeling for the budding yeast cell cycle gives the average behavior of populations of dividing cells [Chen, Calzone, Csikasz-Nagy, Cross, Novak, Tyson (2004)]. However, some major regulatory proteins occur in small numbers such that minor changes in timing and reaction rates can have major inputs on outcomes. Thus, the stochastic approach provides more accurate results than does the deterministic one. In addition, when cell cycle controls are compromised by mutation, random fluctuations are important for modeling the effects of the mutants. Therefore, it is desirable to translate a deterministic cell cycle model into a stochastic model, and simulate the model with an appropriate stochastic method.

Gillespie's stochastic simulation algorithm (SSA) [Gillespie (1976); Gillespie (1977)] is a well-known algorithm using Monte Carlo methods to simulate the chemical reactions. The SSA is an asymptotically exact stochastic method to simulate chemical systems, but the SSA is often slow because it simulates every reaction. Since the SSA emerged, there have been many attempts to improve the computational efficiency [Gibson and Bruck (2000); Gillespie (2001); Cao, Gillespie, Petzold (2005)], however, the core principles remain the same. Other recent work in stochastic simulation includes Rajan and Raha (2008).

One notable attempt to improve the SSA is the tau-leaping method [Gillespie (2001)]. Tau-leaping attempts to achieve increased computational efficiency by leaping over many fast reactions. The implicit tau-leaping method compensates for difficulty with stiff systems [Rathinam, Petzold, Cao, Gillespie (2003)]. Stiff systems are characterized by well separated fast and slow time scales in a dynamic system, the fastest of which is stable. Some approaches try to reduce time consumption with different assumptions such as quasi steady-state approximation (QSSA) [Rao and Arkin (2003)] and total quasi steady-state approximation (tQSSA) [Cilibetro, Capuani, Tyson (2007)] for stiff systems. This paper compares computational efficiency and exactness between SSA, tau-leaping, implicit tau-leaping, QSSA, and tQSSA based on numerical experiments first with a model using simple chemical reactions and stiff systems, and then with the budding yeast model. StochKit [Li, Cao, Petzold, Gillespie (2008)] is used to do stochastic simulation of the budding yeast model. Because StochKit supports various approximate simulation methods based on the SSA such as explicit and implicit tau-leaping methods, the computational efficiency of the approximation methods can be compared easily by using StochKit.

Stochastic methods require that the model be cast in terms of population because

they consider reactions with individual molecules. The problem is, however, that Ordinary Differential Equation (ODE) models are usually based on concentration values. Therefore, the concentration-based model has to be changed into a population-based model to simulate using a stochastic method. Previous work [Wang, Randhawa, Shaffer, Cao, Baumann (2008)] explained the conversion process using JigCell [Virginia Tech (2009)] in detail. StochKit [Li, Cao, Petzold, Gillespie (2008)] is used to do stochastic simulation of the converted budding yeast model. StochKit supports various approximate simulation methods based on the SSA, but only the exact SSA is used to get precise results.

Because the SSA simulates every time step, the SSA is much slower than a deterministic simulation. Moreover, the simulation must be run thousands of times to generate enough data to determine the correct distribution of the behavior. Therefore, it is desirable to run many independent SSA simulations in parallel. StochKit supports MPI for parallel SSA runs, but the user assigns jobs for each processor. Sometimes, the processor times for individual runs are quite different, which cause a system resource inefficiency. This paper presents a dynamic load balancing algorithm that improves the parallel efficiency.

The SSA and approximation methods are explained in the next section. The budding yeast cell cycle model and the dynamic load balancing algorithm are presented next. Finally new biological results and numerical comparisons are given.

2   Stochastic simulation algorithms

2.1   SSA

Suppose a biochemical system or pathway involves N molecular species $\{S_1, ..., S_N\}$. $X_i(t)$ denotes the number of molecules of species S_i at time t. People would like to study the evolution of the state vector $X(t) = (X_1(t), ..., X_N(t))$ given that the system was initially in the state vector $X(t_0)$. Suppose the system is composed of M reaction channels $\{R_1, ..., R_M\}$. In a constant volume Ω, assume that the system is well-stirred and in thermal equilibrium at some constant temperature. There are two important quantities in reaction channels R_j: the state change vector $v_j = (v_{1j}, ..., v_{Nj})$, and propensity function a_j. v_{ij} is defined as the change in the S_i molecules' population caused by one R_j reaction, and $a_j(x)dt$ gives the probability that one R_j reaction will occur in the next infinitesimal time interval $[t, t+dt)$.

The SSA simulates every reaction event [Gillespie (1976); Gillespie (1977)]. With $X(t) = x$, $p(\tau, j|x, t)d\tau$ is defined as the probability 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. By letting $a_0(x) \equiv \sum_{j=1}^{M} a_j(x)$, the equation

$$p(\tau, j|x,t) = a_j(x) \exp(-a_0(x)\tau),$$ (1)

can be obtained. On each step of the SSA, two random numbers r_1 and r_2 are generated from the uniform $(0,1)$ distribution. From probability theory, the time for the next reaction to occur is given by $t + \tau$, where

$$\tau = \frac{1}{a_0(\mathbf{x})} \ln(\frac{1}{r_1}).$$ (2)

The next reaction index j is given by the smallest integer satisfying

$$\sum_{j'=1}^{j} a_{j'}(x) > r_2 a_0(x).$$ (3)

After τ and j are obtained, the system states are updated by $X(t+\tau) := x + v_j$, and the time is updated by $t := t + \tau$. This simulation iteration proceeds until the time t reaches the final time.

2.2   *Explicit tau-leaping*

The SSA is an exact stochastic method for chemical reactions, however, it is very slow for many practical systems because the SSA simulates one reaction at a time. One approximate simulation approach is *tau-leaping* [Gillespie (2001)]. The basic idea of the tau-leaping method is that many reactions can be simulated at each step with a preselected time τ. The tau-leaping method requires that the selected τ must be small enough to satisfy the "leap condition": The expected state change induced by the leap must be sufficiently small that propensity functions remain nearly constant during the time step τ.

$K_j(\tau;x,t)$ is defined as the number of times, given $X(t) = x$, that reaction channel R_j will fire in the time interval $[t, t+\tau)$ where $j = 1, \ldots, M$. If $X(t) = x$, then the state can be updated by

$$X(t+\tau) = x + \sum_{j=1}^{M} K_j(\tau;x,t)v_j.$$ (4)

$K_j(\tau;x,t)$ is modeled by a Poisson random variate. The explicit tau-leaping method assumes

$$K_j(\tau;x,t) = P_j(a_j(x)\tau),$$ (5)

where P_j is a Poisson random variate with mean and variance $a_j(x)\tau$.

In order to select the largest value of τ that satisfies the leap condition, the Jacobian matrices for the propensity functions are used [Gillespie (2001); Gillespie and Petzold (2003)]. One new approach is to select τ such that relative changes in the propensity functions are bounded [Cao, Gillespie, Petzold (2006)]. This new τ selection procedure is faster and more accurate than previous methods. Therefore, the explicit tau-leaping method proceeds as follows. Select a τ that satisfies the leap condition. Generate the Poisson random variables for each reaction and adjust the leap time by $t := t + \tau$ and the states by $X(t+\tau) := x + \sum_{j=1}^{M} K_j(\tau;x,t)v_j$. This simulation iteration also proceeds until the time t reaches the final time t_f.

2.3 *Implicit tau-leaping*

Implicit tau-leaping addresses the shortcomings of explicit tau-leaping when the systems are stiff [Rathinam, Petzold, Cao, Gillespie (2003)]. Stiff systems are characterized by well separated fast and slow time scales in a dynamic system, with the fast mode being stable. In a stiff system, solutions by explicit tau-leaping are unstable unless the time stepsize τ is kept smaller than the smallest (fastest) time scale in the system [Rathinam, Petzold, Cao, Gillespie (2003)].

Tau-leaping as described previously is an explicit method because the propensity functions a_j are evaluated at the current known state x. Therefore, the future state $X(t+\tau)$ is an explicit function of $X(t)$, and the states can be updated by the equation,

$$X^{et}(t+\tau) = x + \sum_{j=1}^{M} v_j P_j(a_j(x)\tau), \tag{6}$$

where the superscript "et" indicates that the explicit method is used. The implicit method is described by

$$X^{it}(t+\tau) = x + \sum_{j=1}^{M} v_j \Big[\tau a_j\big(X^{it}(t+\tau)\big) + P_j(a_j(x)\tau) - \tau a_j(x) \Big], \tag{7}$$

where the superscript "it" stands for the implicit method. The implicit equation is solved by Newton's method, and the floating point state $X^{it}(t+\tau)$ is rounded to the nearest integer values.

2.4 *QSSA*

The explicit and implicit tau-leaping methods achieve increased computational efficiency by attempting to leap over many fast reactions. The quasi-steady state approximation (QSSA) [Rao and Arkin (2003)] improves the efficiency by relying on this steady state assumption: the net rate of formation is approximately equal to

zero when the fast reacting species are in steady state. Consider common enzyme kinetic reactions using Michaelis-Menten kinetics. For substrate S, enzyme E, and product P, the Michaelis-Menten reaction is

$$E + S \underset{k_{-1}}{\overset{k_1}{\rightleftharpoons}} E : S \xrightarrow{k_2} P + E, \tag{8}$$

where k_1, k_{-1}, and k_2 are the rate constants. $E{:}S$ is the enzyme-substrate complex after the combination of substrate and enzyme. The rate equations corresponding to this reaction are

$$\frac{d[S]}{dt} = -k_1[S][E] + k_{-1}[E:S], \tag{9}$$

$$\frac{d[E:S]}{dt} = -\frac{d[E]}{dt} = -(k_{-1} + k_2)[E:S] + k_1[S][E], \tag{10}$$

$$\frac{d[P]}{dt} = k_2[E:S], \tag{11}$$

where $[X]$ denotes the concentration of the species X. Assume that total enzyme concentration $E_T = [E] + [E:S]$ and $[S] \gg [E:S]$. From the assumption, the characteristic time scale of $[S]$ is very slow in comparison to that of $[E:S]$, and $[E:S]$ reaches steady state quickly. From the quasi-steady state assumption, the rate equation for $[E:S]$ is approximated by

$$\frac{d[E:S]}{dt} = 0. \tag{12}$$

Mathematically, it is possible to obtain a single rate equation,

$$\frac{d[S]}{dt} = -k_2[E:S] = -\frac{k_2 E_T[S]}{K_m + [S]}, \tag{13}$$

where $K_m = (k_{-1} + k_2)/k_1$.

Similarly, the quasi-steady state assumption can be applied to the stochastic formulation. The QSSA in stochastic kinetics implies that the net rate of change for the conditional probability distribution of the fast reacting species is equal to zero. For the above Michaelis-Menten kinetics, the separate master equation for the enzyme-substrate complex is

$$\frac{dP([E:S] \mid [\hat{S}];t)}{dt} = 0, \tag{14}$$

where $[\hat{S}] = [S] + [E:S]$. From the above equation, one can easily derive the reduced system

$$\hat{S} \longrightarrow P, \tag{15}$$

with the propensity function

$$a(s) = \frac{k_2 E_T [\hat{S}]}{K_m + [\hat{S}]}. \tag{16}$$

Finally, the SSA is applied to this reduced system, and shows improved computational efficiency compared with using SSA on the original system.

2.5 tQSSA

In the previous section, the quasi-steady state approximation (QSSA) eliminates the fastest reacting variable under some assumptions. In Michaelis-Menten kinetics, the necessary condition for the QSSA is $S_0 \gg E_T$, where S_0 is the initial substrate concentration, and E_T is the total enzyme concentration. In a protein interaction network (PIN), however, the enzymes and substrates often swap their roles [Blüthgen and Herzel (2003)]. Therefore, the QSSA condition will not be true for such a PIN. Borghans et al. [Borghans, de Bore, Segel (1996)] proposed that the proper slow timescale variable is $[\hat{S}] = [S] + [E:S]$ instead of $[S]$. In terms of this variable, the deterministic equations are

$$E_T = [E] + [E:S] = \text{constant}, \tag{17}$$

$$[E:S]^2 - \left(E_T + K_M + [\hat{S}]\right)[E:S] + E_T[\hat{S}] = 0, \tag{18}$$

$$\frac{d[\hat{S}]}{dt} = -k_2[E:S]. \tag{19}$$

This is called the total quasi-steady state approximation (tQSSA). To derive the equations for the stochastic simulation under the tQSSA, reduce the system to

$$\hat{S} \longrightarrow P \tag{20}$$

with the propensity function

$$a(s) = k_2[E:S], \tag{21}$$

where

$$[E:S] = \frac{\left(E_T + K_m + [\hat{S}]\right) - \sqrt{\left(E_T + K_m + [\hat{S}]\right)^2 - 4E_T[\hat{S}]}}{2}. \tag{22}$$

Finally, the SSA is applied to this reduced system, and shows improved computational efficiency compared with using SSA on the original system. Moreover, the tQSSA overcomes the modeling shortcomings of the QSSA.
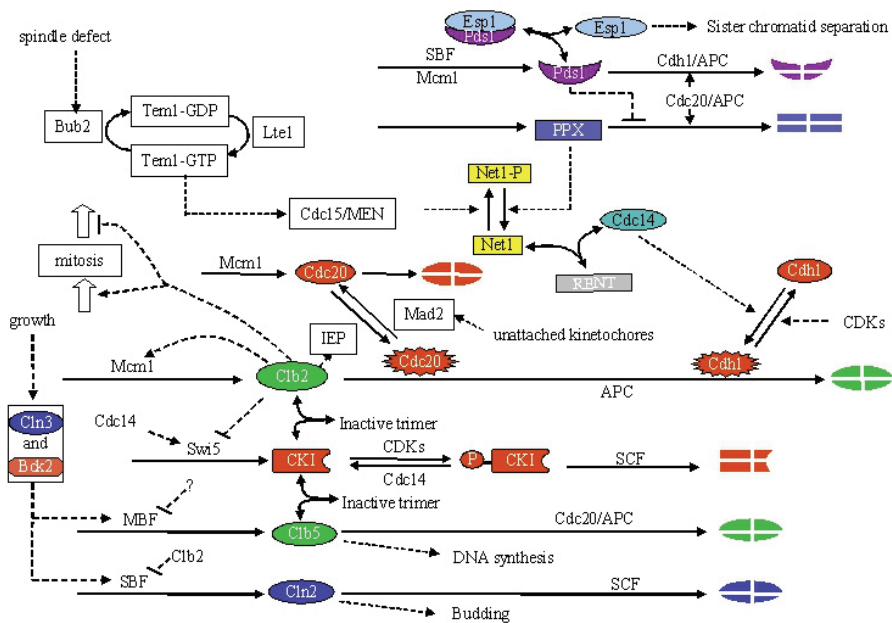
Figure 1: Wiring diagram of budding yeast.

3   A stochastic cell cycle model

The molecular machinery of eukaryotic cell cycle control is known in more detail for budding yeast, *Saccharomyces cerevisiae*, than for any other organism. Therefore, the unicellular budding yeast is an excellent organism for which to study cell cycle regulation. Molecular biologists have dissected and characterized individual cell cycle components and their interactions to derive a consensus picture of the regulatory network of the budding yeast. Figure 1 shows the wiring diagram for the budding yeast model [Virginia Tech (2009)]. The diagram should be read from the bottom-left toward the top-right. Solid arrows represent biochemical reactions, and dashed lines represent how components may influence one another. Empirical results of simulation runs for a stochastic formulation of the budding yeast model are presented in the following section.

Each run of the stochastic simulation provides potentially different results, and the real goal is to deduce the true population distribution for the potential outcomes. This requires potentially many thousands of simulation runs. Each simulation run requires different amounts of processor time. When hundreds of simulation runs are assigned statically to a processor, the total simulation times can be quite different

across the processors, causes a serious load imbalance. A dynamic load balancing algorithm was developed that more evenly distributes work to the processors.

The basic idea of the load balancing algorithm is a master/slave paradigm that dynamically adjusts the task chunk size. Pseudocode (with constants tuned for the problem at hand) for the load balancing follows.

> *Algorithm GetTask(n, p, task)*
1. ▷ **Input:** n = number of remaining tasks,
2. p = number of processors
3. ▷ **Output:** *chunk* = number of tasks for a processor
4. **begin**
5. setPoint ← floor (n/p)
6. **if** *setPoint* > 99 **then**
7. *chunk* ← floor $(setPoint \times 0.8)$
8. **else if** *setPoint* > 5 **then**
9. *chunk* ← floor $(setPoint \times 0.5)$
10. **else then**
11. *chunk* ← 1
12. **end if**
13. **end**

The idea of varying the task chunk size as the size of the task queue decreases is well known in parallel computing, where it is called guided self-scheduling. Figures 2 and 3 clearly show the advantage of the load balancing algorithm. Using Virginia Tech's 2200 processor (2.3 GHz PowerPC 970FX) System X, 100 worker processors were used for 10,000 stochastic simulations for a budding yeast prototype double mutant. Figure 2 shows the time that each processor required to complete its assigned tasks. The variance between processors is high. Figure 3 shows processor times using the dynamic workload distribution. Here, the variance in processor times is negligible, even though individual runs of the simulation have large fluctuations in their cell cycles.

Table 1: Execution times for static distribution.

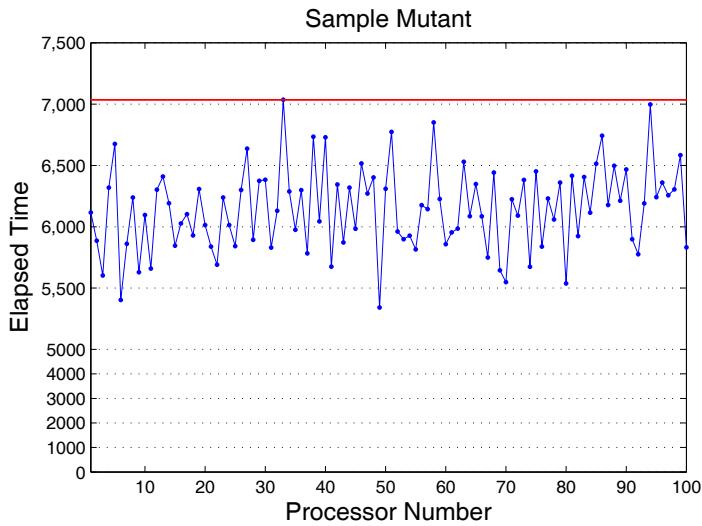| Parallel Runtime	Time (seconds)
T_p(average)	6144.93
Wall clock time	**7040.72**

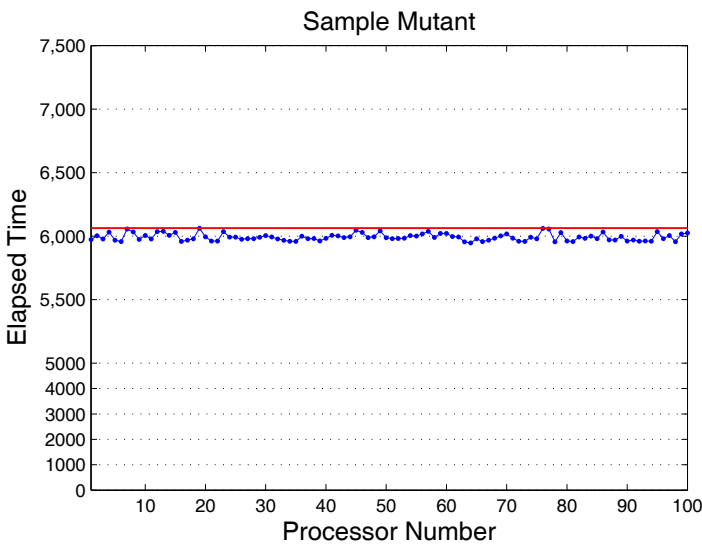Figure 2: Static workload distribution.

Figure 3: Dynamic workload distribution.

Tables 1 and 2 show the average and total times for the static and dynamic workload distributions. It is observed that dynamic workload distribution reduces system

Table 2: Execution times for load balanced distribution.

Parallel Runtime	Time (seconds)
T_p(average)	5992.43
Wall clock time	**6066.84**

resource use by approximately 14%.

4 Numerical experiments for simple reactions

The irreversible isomerization system and the Goldbeter-Koshland switch are used here to compare various stochastic algorithms. The irreversible isomerization shows the computational efficiency of the explicit tau-leaping method compared with the SSA on a nonstiff system. The GK switch is a suitable stiff model to compare various stochastic algorithms (see, e.g., [Barik, Paul, Baumann, Cao, Tyson (2008)]).


Figure 4: The SSA simulation (solid lines) and the explicit tau-leaping simulation (dotted lines) for the irreversible isomerization. The error control parameter ε is 0.03 (left) and 0.15 (right).

4.1 Irreversible isomerization

The first application is the simplest chemical reaction, the irreversible isomerization

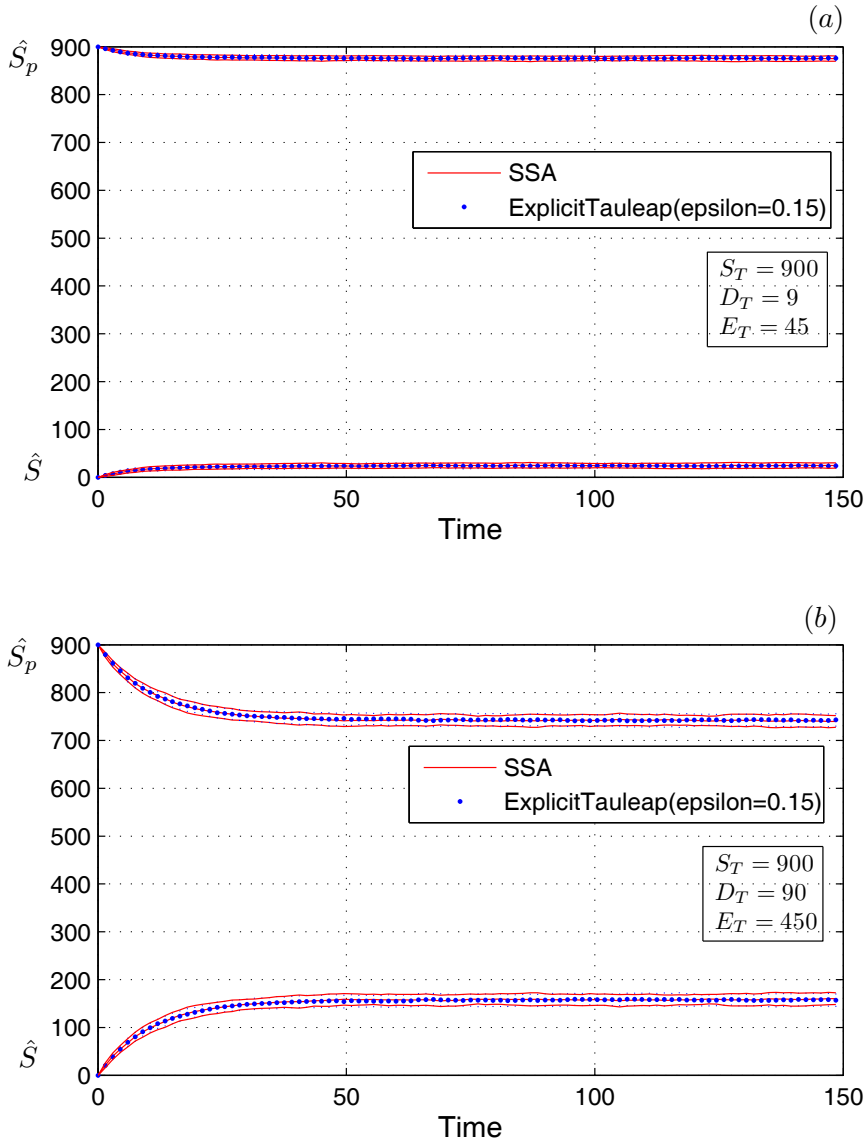$$S_1 \xrightarrow{c_1} 0. \tag{23}$$

Figure 5: The SSA simulation (solid lines) and the explicit tau-leaping simulation (dotted lines) for the GK switch. State space (S_T, D_T, E_T) is (a):(900, 9, 45) and (b):(900, 90, 450).
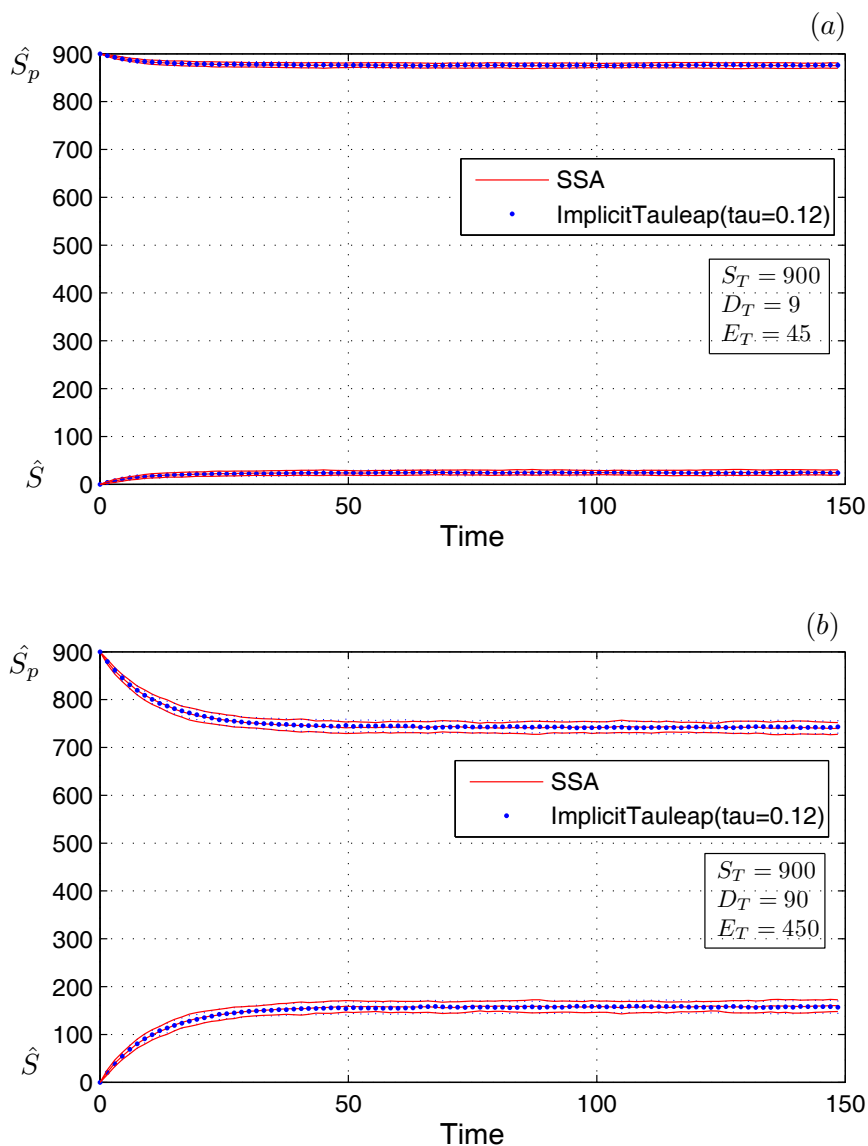
Figure 6: The SSA simulation (solid lines) and the implicit tau-leaping simulation (dotted lines) for the GK switch. State space (S_T, D_T, E_T) is (a):(900, 9, 45) and (b):(900, 90, 450).
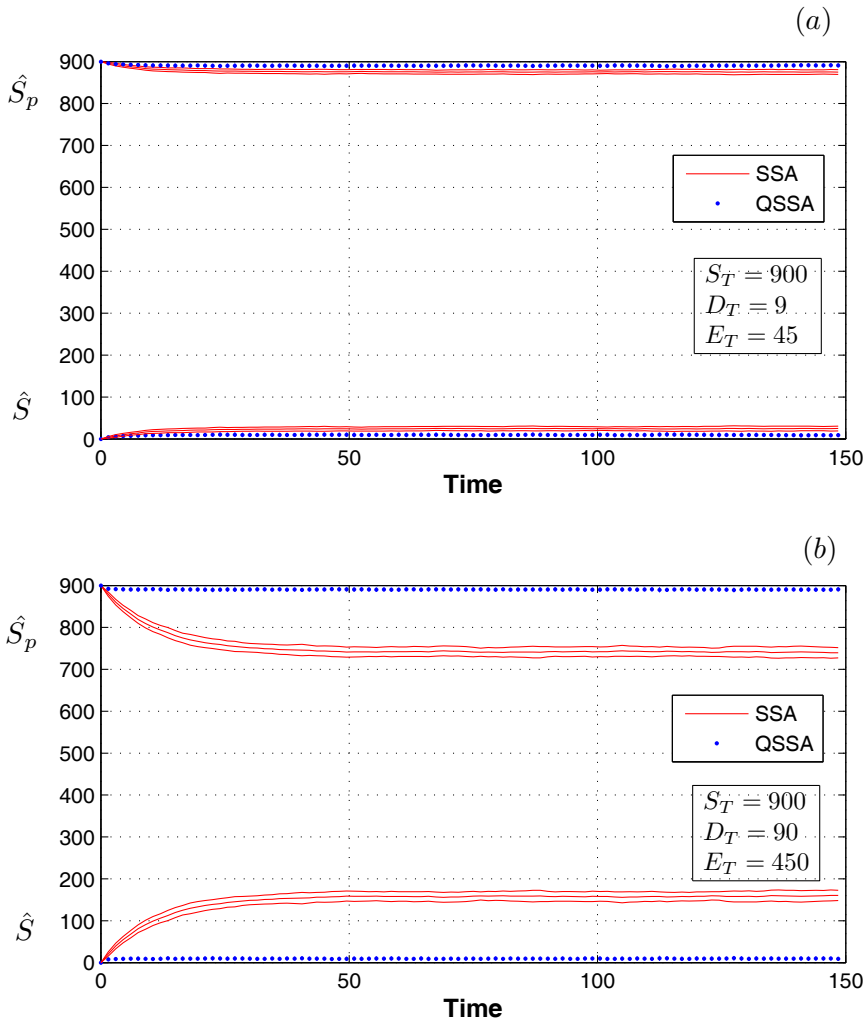
Figure 7: The SSA simulation (solid lines) and the QSSA simulation (dotted lines) for the GK switch. State space (S_T, D_T, E_T) is (a):(900, 9, 45) and (b):(900, 90, 450).
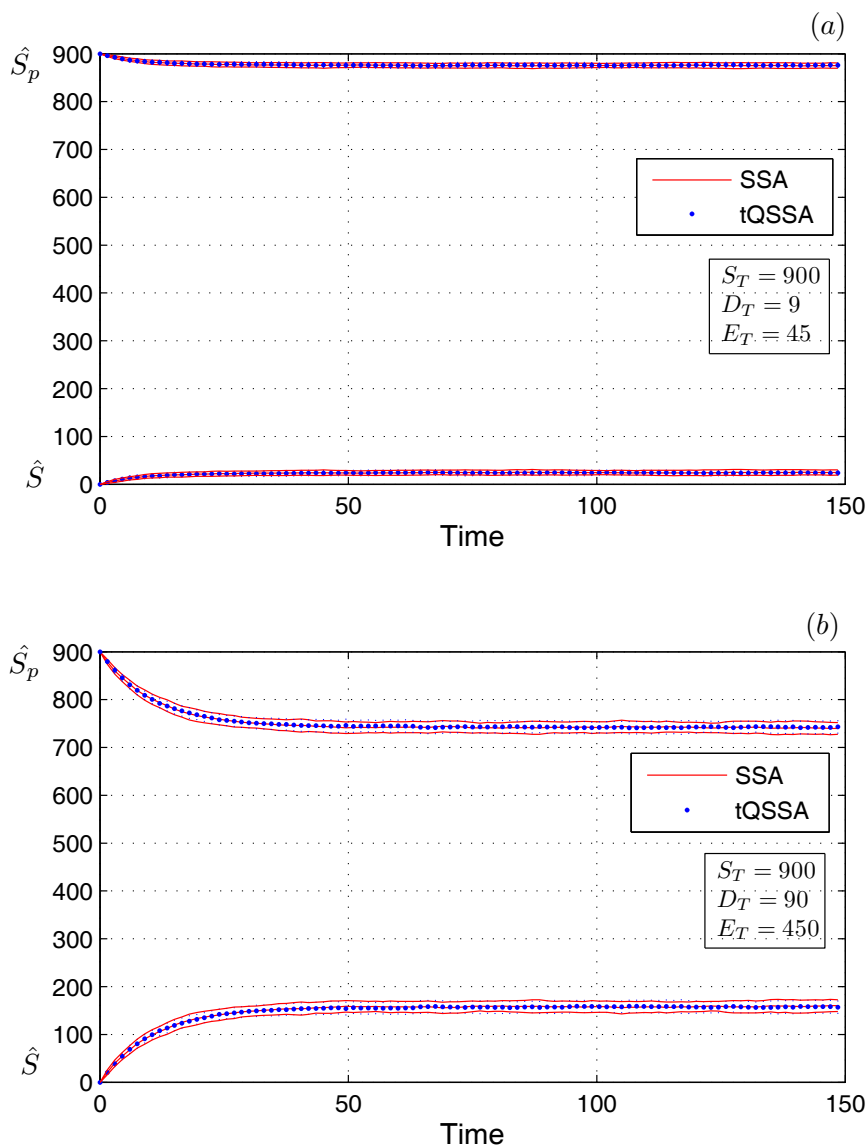
Figure 8: The SSA simulation (solid lines) and the tQSSA simulation (dotted lines) for the GK switch. State space (S_T, D_T, E_T) is (a):(900, 9, 45) and (b):(900, 90, 450).

Table 3: The number of runs and elapsed CPU time (sec) with the SSA and explicit tau-leaping method, where $t_f = 5$, $c_1 = 1$, $X_1 = 10^4$, and $\varepsilon = 0.03$.

Number of runs	1000	5000	10000	50000
SSA	48.74	240.40	487.76	2433.28
Explicit Tau-leaping	2.73	14.21	28.37	143.29

The initial parameters for this system are reaction rate constant $c_1 = 1$, and 10^4 S_1 molecules at time 0. Figure 4 shows the SSA and explicit tau-leaping results. The final time is $t_f = 5$ and the error control parameters are $\varepsilon = 0.03$ (left) and $\varepsilon = 0.15$ (right). In order to compare the exactness, 5000 runs are used for both methods. Figure 4 shows mean and mean $\pm$ one standard deviation. These three lines in the figure are visually identical. Figure 4 shows that increasing ε, while making tau-leaping faster, affects the accuracy. The exact SSA method requires 9925 steps to reach the final time ($t_f = 5$). The tau-leaping method requires 167 (34) leap steps for $\varepsilon = 0.03$ ($\varepsilon = 0.15$). Table 3 compares the computational efficiencies of the explicit tau-leaping method and the SSA. Explicit tau-leaping is about 17 times faster than the SSA with an appropriate approximation for this model.

Table 4: The number of runs and elapsed CPU time (sec) for the SSA, explicit tau-leaping, implicit tau-leaping, QSSA, and tQSSA algorithms.

Number of runs	1000	5000	10000	50000
SSA	44.57	216.77	434.10	2175.45
Explicit Tau-leaping	4.75	18.23	46.57	181.31
Implicit Tau-leaping	18.43	89.13	181.29	887.44
QSSA	2.63	13.23	26.33	132.47
tQSSA	2.67	13.33	26.79	133.68

4.2 Goldbeter-Koshland switch

The Goldbeter-Koshland switch (GK switch) consists of a substrate-product pair (S and S_p) that is interconverted by two enzymes (E and D):

$$D + S_p \underset{k_{-1d}}{\overset{k_{1d}}{\rightleftarrows}} D : S_p \xrightarrow{k_{2d}} S + D, \tag{24}$$

$$E + S \underset{k_{-1e}}{\overset{k_{1e}}{\rightleftarrows}} E : S \xrightarrow{k_{2e}} S_p + E. \tag{25}$$

The parameter set for the GK switch came from [Barik, Paul, Baumann, Cao, Tyson (2008)]. The parameter values are $S_T = 900$, $k_{1d} = 0.05555\text{min}^{-1}$, $k_{-1d} = 0.83\text{min}^{-1}$, $k_{2d} = 0.17\text{min}^{-1}$, $k_{1e} = 0.05\text{min}^{-1}$, $k_{-1e} = 0.8\text{min}^{-1}$, and $k_{2e} = 0.1\text{min}^{-1}$. In order to observe how the relationship between S_T and E_T affects the results, the two cases $(D_T, E_T) = (9, 45)$ and $(D_T, E_T) = (90, 450)$ are compared. $\hat{S}_p$ and $\hat{S}$ are defined by $\hat{S}_p = S_p + D : S_p$ and $\hat{S} = S + E : S$.

Figures 5 and 6 show the mean and mean $\pm$ one standard deviation trajectories of $\hat{S}_p$ and $\hat{S}$ for the explicit and implicit tau-leaping approximation algorithms and the SSA. When $\hat{S}_p$ reaches steady state, the standard deviations of the SSA are approximately 4.3 (a) and 13.2 (b). In the explicit method, the error parameter value ($\varepsilon = 0.15$) was chosen so that the standard deviation matched that of the SSA. The fixed step value ($\tau = 0.12$) was chosen similarly for the implicit method. Table 4 shows that the explicit method is faster than the implicit method for the same accuracy, due to the cost of the Newton iteration in the implicit method.

Figures 7 and 8 show the mean and mean $\pm$ one standard deviation trajectories of $\hat{S}_p$ and $\hat{S}$ for the QSSA and tQSSA algorithms. Figure 7(b) shows that the results with the QSSA are different from those for the SSA. The explanation in the QSSA section implies that the QSSA requires $S_T \gg E_T$. If $S_T \approx E_T$ or $S_T \ll E_T$, then the results from the QSSA algorithm are not reliable. Figure 7(a) shows that if $S_T \gg E_T$, then the results of the QSSA algorithm are similar to those from the SSA. In contrast with Fig. 7, Fig. 8 shows that the tQSSA algorithm works when $S_T \approx E_T$ or $S_T \ll E_T$.

In terms of CPU time, Tab. 4 shows that the QSSA and tQSSA algorithms are the fastest approximate algorithms. The QSSA and tQSSA algorithms are almost 20 times faster than the SSA. The explicit and implicit tau-leaping approximations also have improved computational time over the SSA.

5 Cell cycle results

5.1 Event implementation

Stochastic methods require the model to be in terms of population because they consider reactions with individual molecules. Because the original budding yeast model is based on normalized concentration values, a conversion process from an ODE model into a model in terms of number of molecules is needed. The conversion process is done using JigCell [Wang, Randhawa, Shaffer, Cao, Baumann (2008)], and consists of two phases, *unit checking* and *model conversion*. Unit checking verifies physical unit consistency inside the model. Model conversion converts the model by changing values of species and parameters based on the unit information.

After creating the population-based budding yeast model, there is technical issue that must be addressed. In Systems Biology Markup Language (SBML), an *event* is triggered when some condition is met. There are events defined to divide the cell or mark checkpoints within the cell cycle stages. A typical deterministic event has the form:

```
if (X > threshold)
    then (Event is triggered)
```

Because of the random nature of stochastic simulation, as illustrated in Fig. 9, unwanted events can be triggered when a deterministic SBML event is used for the stochastic model. In Fig. 9, an unwanted event (B) can be triggered with a wanted event (A) by using a deterministic event handling equation.


Figure 9: Event handling.

To prevent unwanted events, the event logic has to be rewritten to tolerate the situation where the value of X oscillates around the threshold. A second threshold value can be defined from the threshold and the direction of the test (greater than or less than). For budding yeast, this second threshold equals .5*threshold (for a greater-than test) or 1.5*threshold (for a less-than test). For instance, the event code above would be changed to:

```
if (X < second threshold)
    then (EventFlag ← TRUE)
if (X > threshold AND EventFlag = TRUE)
```

```
then (event is triggered;
      EventFlag ← FALSE)
```

StochKit [Li, Cao, Petzold, Gillespie (2008)] was used to do stochastic simulation of the converted budding yeast model, using the SSA option for the most precise results. JigCell can generate the StochKit model file by using the population based budding yeast model file.


Figure 10: Deterministic cell cycles.

Figure 11: Stochastic cell cycles.

5.2   Wild type simulation results with the SSA

The deterministc budding yeast model from Chen's paper [Chen, Calzone, Csikasz-Nagy, Cross, Novak, Tyson (2004)] was converted to a stochasitc model. To compare the stochastic results with the deterministic cell cycle simulation, mass and several representative species' trajectories are shown in Figs. 10 and 11. The deterministic result is from the XPP ODE simulator using JigCell. For comparison, the stochastic simulation results are converted back to normalized concentrations.

Table 5: Execution time comparisons with different methods for wild type budding yeast.

Method	Time(seconds)	The Number of Steps
Direct SSA	25.03	2,495,000
Explicit Tau-leaping(1) with $\varepsilon = 0.01$	82.893	2,444,000
Explicit Tau-leaping(1) with $\varepsilon = 0.05$	312.07	1,232,000
Explicit Tau-leaping(1) with $\varepsilon = 0.15$	618.41	290,000
Explicit Tau-leaping(2) with $\varepsilon = 0.01$	28.90	2,442,000
Explicit Tau-leaping(2) with $\varepsilon = 0.05$	29.44	2,319,000
Explicit Tau-leaping(2) with $\varepsilon = 0.15$	29.16	1,835,000
Implicit Tau-leaping with $\tau = 0.001$	7965.18	500,000
Implicit Tau-leaping with $\tau = 0.005$	1920.34	100,000
Implicit Tau-leaping with $\tau = 0.010$	1012.87	50,000

5.3   Wild type results comparisons

Table 5 shows the results of the stochastic budding yeast simulation with direct SSA, explicit tau-leaping, and implicit tau-leaping methods. The final time is $t_f = 500$. In order to select the largest value of τ that satisfies the leap condition, the Jacobian matrices for the propensity functions are used for the explicit tau-leaping(1) method [Gillespie and Petzold (2003)]. Another approach, explicit tau-leaping(2), is to select τ such that relative changes in the species are bounded [Cao, Gillespie, Petzold (2006)]. For the implicit method, the package StochKit does not provide an adaptive stepsize selection strategy, thus fixed step values were used in the experiment.

For direct SSA, the average time to run the simulation is approximately 25 seconds with 2,495,000 steps. From the results with simple chemical reactions, the explicit and implicit tau-leaping methods improved computational time over the SSA while obtaining equally accurate simulation results. However, for the budding yeast

model, the tau-leaping methods took more simulation time. The Jacobian matrices for the propensity functions are used for explicit tau-leaping(1), and the time was increased for a reduced number of steps. Thus the computational time for handling Jacobian matrices is larger than the computational time for steps. Even for the new tau selection method tau-leaping(2), the computational time did not improve over that of the SSA. This unexpected result is caused by the stiffness present in the cell cycle model. Ideally, for a stiff system the implicit tau-leaping or QSSA methods should be applied. However, here the implicit tau-leaping method took even more time than the explicit tau-leaping method, even with many fewer steps. One reason is because the implicit tau-leaping method has to solve a nonlinear system of equations in each step, which is quite expensive when the system size is large. Another reason is that currently there is no general stepsize selection algorithm for the implicit tau-leaping method, thus here only the fixed time step option in StochKit can be used, which severely limits the efficiency of the implicit tau-leaping method. The QSSA and tQSSA methods have great performance when a system has clear separation between fast and slow scales. However, for an oscillating system such as the cell cycle model, the fast and slow scales do not separate so well. Moreover, the scales of many reactions in the system switch in the cell cycle oscillation. Currently there is no known good way to dynamically partition this type of complex system and apply QSSA or tQSSA. Therefore, at present, direct SSA is the most appropriate method for the budding yeast model, and mutants in the next section were tested with the direct SSA method.

5.4 Mutants

To further test their understanding of the cell cycle process, modelers go beyond comparing simulation outputs from the model with the behavior of the wild type yeast cell. They also study the behavior for genetic variants, created by gene knockouts form the wild type cell. These variations are expressed in the model by changing some initial conditions, parameters values, or equations from the wild type budding yeast model. Deterministic models can only show that all cells die, reproduce endlessly, or die after some number of divisions. In the stochastic simulation, different runs of the simulation can yield different outcomes. The number of cycles until stopping division is different for different trajectories. To generate enough data to perform a statistical analysis, 10,000 independent simulations are executed from the same initial point using the load balancing parallel algorithm.

Let the random variable X denote the number of cell divisions before the cell stops dividing, and assume that the probability p of not dividing is constant and independent of the cell's previous history. Then X has a modified geometric distribution

given by

$$P(X = n) = p(1-p)^n \quad \text{where} \quad n = 0, 1, \ldots \tag{26}$$

Further, the probability of having *n* or more cycles is given by

$$P(X \geq n) = (1-p)^n \quad \text{where} \quad n = 0, 1, \ldots \tag{27}$$

If either probability is plotted on a log scale against *n*, it will be a straight line with a slope of $\log(1-p)$. In the analysis of the mutant simulations, a best least squares fit straight line is used to extract the slope and estimate the value of *p*.


Figure 12: Probability of getting a colony.

In wet lab experiments, the viability of mutants is assessed by determining whether single mutant cells could grow into a colony. To determine the relationship between viability and the probability of ceasing to divide, simulation is used to determine the probability of one cell producing a colony of size greater than 10^6 in 32 division cycles. 32 division cycles correspond roughly to incubating the cells on a plate for two days. A 1 mm^3 colony has about 20×10^6 cells. Therefore, it is assumed that 0.05 mm^3 can be visible. The results of these simulations are summarized in Fig. 12. From this figure, the ability to observe colonies has a roughly switch-like characteristic with the switching point near $p = 0.25$.

Figure 13: $P(X \geq n)$.

Figure 14: $\log(P(X \geq n))$.

It is interesting that some mutants have different fates in different growth media. For example, some mutants of budding yeast are inviable in "rich" medium (fast growth rate) but partially viable in "poor" medium (slow growth rate). "Partially viable" means that cells growing in poor medium have some probability of not dividing ($0 < p < 0.5$, p as defined above). In rich medium, $p = 1$. Partial viability is not an option in a deterministic model, for which all cells behave identically (either $p = 0$ or $p = 1$).

Simulation results for a prototype model of a partially viable, budding yeast mutant are reported in Figs. 13 and 14. For fast growth rates, every simulated cell arrests in the cell cycle ($p = 1$). If the growth rate is sufficiently slow, then simulated cells begin to divide with a certain probability $1 - p$. For the case in Fig. 14, $p = 0.22$. From Fig. 12, this implies that the probability of a single cell forming a colony is 65%, which could be compared to experimental observations for a particular mutant. The point is that stochastic simulation accounts for the property of partial viability that a deterministic simulation cannot.

6   Conclusions and future work

The simulation results reported here, while limited, show important characteristics of each approximation algorithm. The explicit tau-leaping method improves computational efficiency, compared to the SSA, for nonstiff systems, but can be unstable on stiff systems. The implicit tau-leaping method is stable, but much slower than the explicit method. The tQSSA algorithm produces excellent agreement with the SSA and is more efficient by an order of magnitude. Some of these approaches are simulated with the budding yeast model, and were shown to be impractical for this realistic cell cycle model.

The budding yeast stochastic simulation results reported here, while limited, show important characteristic aspects of cell cycle empirical data, such as mixed mutant viability. Because random fluctuations are important to accurately simulate mutants, some major regulatory proteins occur in small numbers, the stochastic approach is more realistic and accurate than the deterministic approach for modeling the budding yeast cell cycle. The guided self-scheduling load balancing algorithm is effective for managing the large numbers of SSA trajectories required by the stochastic approach.

The data collected here are only the first step toward calculating population doubling times, and probabilities for successful colony formation of the various mutants. Those probabilities are one of the tests for verifying that the simulations match experimental results. The results are also perhaps skewed by the fact that all simulation runs begin with a static initial condition that might not be represen-

tative of the true population. In future work, the authors will calculate population statistics and use them to generate appropriate initial conditions.

References

Barik, D.; Paul, M.R.; Baumann, W.T.; Cao, Y.; Tyson, J.J. (2008): Stochastic simulation of enzyme-catalyzed reactions with disparate timescales. *Biophys. J.*, vol. 95, no. 8, pp. 3563–74.

Borghans, J.A.M.; deBore, R.J.; Segel, L.A. (1996): Extending the quasi-steady state approximation by changing variables. *Bull. Mathematical Biology*, vol. 58, pp. 43–63.

Blüthgen, N.; Herzel, H. (2003): How robust are switches in intracellular signaling cascades?. *Journal of Theoretical Biology*, vol. 225, pp. 293–300.

Cao, Y.; Gillespie, D.T.; Petzold, L.R. (2005): The slow-scale stochastic simulation algorithm. *Journal of Chemical Physics*, vol. 122, pp. 014116+.

Cao, Y.; Gillespie, D.T.; Petzold, L.R. (2006): Efficient stepsize selection for the tau-leaping simulation. *Journal of Chemical Physics*, vol. 124, pp. 144109+.

Chen, K.C.; Calzone, L.; Csikasz-Nagy, A.; Cross, F.R.; Novak, B.; Tyson, J.J. (2004): Integrative analysis of cell cycle control in budding yeast. *Mol. Biol. Cell*, vol. 15, pp. 3841–3862.

Cilibetro, A.C.; Capuani, F.; Tyson, J.J. (2007): Modeling networks of coupled enzymatic reactions using the total quasi-steady state approximation. *PLoS Computational Biology*, vol. 3, pp. 463–472.

Cross, F.R. (2003): Two redundant oscillatory mechanisms in the yeast cell cycle. *Dev. Cell*, vol. 4, pp. 741–752.

Gibson, M.A.; Bruck, J. (2000): Efficient exact stochastic simulation of chemical systems with many species and many channels. *Journal of Physical Chemistry*, vol. 104, pp. 1876–1889.

Gillespie, D.T. (1976): A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics*, vol. 22, pp. 403–434.

Gillespie, D.T. (1977): Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry*, vol. 81, pp. 2340–2361.

Gillespie, D.T. (2001): Approximate accelerated stochastic simulation of chemically reacting systems. *Journal of Chemical Physics*, vol. 115, pp. 1716–1733.

Gillespie, D.T.; Petzold, L.R. (2003): Improved leap-size selection for accelerated stochastic simulation. *Journal of Chemical Physics*, vol. 119, pp. 8229–8234.

Li, H.; Cao, Y.; Petzold, L.R.; Gillespie, D.T. (2008): Algorithms and software for stochastic simulation of biochemical reacting systems. *Biotechnology Prog.*, vol. 24, pp. 56–61.

McAdams, H.H.; Arkin, A. (1997): Stochastic mechanisms in gene expression. *Proc. Natl. Acad. Sci.*, vol. 94, pp. 814–819.

Murray, A.; Hunt, T. (1993): *The Cell Cycle. An Introduction.* Oxford University Press, USA.

Rajan, N.; Raha, S. (2008): The stochastic alpha method: a numerical method for simulation of noisy second order dynamical systems. *CMES: Comput. Modeling Engrg. Sci.*, vol. 23, pp. 91–116.

Rao, C.V.; Arkin, A.P. (2003): Stochastic chemical kinetics and the quasi-steady-state assumption: application to the Gillespie algorithm. *Journal of Chemical Physics*, vol. 118, pp. 4999–5010.

Rathinam, M.; Petzold, L.R.; Cao, Y.; Gillespie, D.T. (2003): Stiffness in stochastic chemically reacting systems: the implicit tau-leaping method. *Journal of Chemical Physics*, vol. 119, pp. 12784–12794.

Virginia Tech (2009): The JigCell website, http://jigcell.biol.vt.edu .

Voit, E.O. (2000): *Computational Analysis of Biochemical Systems.* Cambridge University Press, USA.

Wang, P.; Randhawa, R.; Shaffer, C.A.; Cao, Y.; Baumann, W.T (2008): Converting macromolecular regulatory models from deterministic to stochastic formulation. In *Proceedings of the 2008 Spring Simulation Multiconference*, SpringSim '08. ACM, New York, pp. 385–392.