Adaptive stochastic-deterministic chemical kinetic simulations

*Karan Vasudeva and Upinder S. Bhalla**

National Centre for Biological Sciences, TIFR, GKVK Campus, Bangalore 560065, India

ABSTRACT

Motivation: Biochemical signaling pathways and genetic circuits often involve very small numbers of key signaling molecules. Computationally expensive stochastic methods are necessary to simulate such chemical situations. Single-molecule chemical events often co-exist with much larger numbers of signaling molecules where mass-action kinetics is a reasonable approximation. Here, we describe an adaptive stochastic method that dynamically chooses between deterministic and stochastic calculations depending on molecular count and propensity of forward reactions. The method is fixed timestep and has first order accuracy. We compare the efficiency of this method with exact stochastic methods.

Results: We have implemented an adaptive stochastic-deterministic approximate simulation method for chemical kinetics. With an error margin of 5%, the method solves typical biologically constrained reaction schemes more rapidly than exact stochastic methods for reaction volumes >1–10 μm^3. We have developed a test suite of reaction cases to test the accuracy of mixed simulation methods.

Availability: Simulation software used in the paper is freely available from http://www.ncbs.res.in/kinetikit/download.html

Contact: bhalla@ncbs.res.in

Supplementary information: A GENESIS/Kinetikit implementation of models and the test suite used in this paper are available from http://www.ncbs.res.in/kinetikit/download.html and also from the DOQCS database, http://doqcs.ncbs.res.in

INTRODUCTION

Most current approaches to simulate chemical kinetics of signaling pathways involve mass-action approximations (Bhalla and Iyengar, 1999; Hartwell *et al*., 1999). Several studies have examined the role of stochasticity in reliability of biological switches (Bialek, 2001; Lisman and Goldring, 1988), in bacterial chemotaxis (Duke *et al*., 2001) and in probabilistic state switching of viral and gene expression events (Arkin *et al*., 1998; Kierzek *et al*., 2001; reviewed in Rao *et al*., 2002). Efficient exact stochastic methods have been devised that assume populations of indistinguishable molecules in each such state (Gillespie, 1977; Gibson and Bruck, 2000). These methods are used in several simulators, including the Systems Biology Workbench (Hucka *et al*., 2002, http://www.sbw-sbml.org/), BioSpice (http://biospice.lbl.gov/) and Stochastirator (http://opnsrcbio.molsci.org/stochastirator/). There are also more computationally intensive methods that explicitly represent each subunit and consider state transitions at the subunit level (Le Novere and Shimizu, 2001, http://www.zoo.cam.ac.uk/comp-cell/stochsim.html). Even the most efficient methods, however, take $O(N)$ time where N is the number of molecules in the system. Clearly, stochastic methods should scale to deterministic methods at large N, where mass-action assumptions hold and execution speed is independent of N. The idea of employing a mixed stochastic-deterministic method to approximate system dynamics is not new, and has been used by Holmes (2000). A method for performing mixed calculations has been described by Haseltine and Rawlings (2002), who in turn cite earlier work on elimination of fast relaxing variables by deriving reduced master equations in mesoscopic cases (Janssen, 1989; Vlad and Pop, 1988). Mixed approaches are theoretically based on the equivalence of stochastic and deterministic calculations at the thermodynamic limit, where N and volume become infinite, but concentration is finite (Kurtz, 1972). Recently two approximate methods for stochastic simulations have been reported. Resat *et al*. (2001) report a method for probability-weighted calculations. This gives a several fold speed-up but introduces differences in distributions, and is not independent of N. Gillespie (2001) has developed the tau-leap method that scales well with large N, but this method is still undergoing development.

Here, we describe a method that automatically and dynamically partitions reactions between deterministic and stochastic calculations. We focus on issues that arise during transitions between stochastic and deterministic methods when automatic switching is implemented, rather than on efficient numerical computations in either domain. We perform benchmarking and accuracy comparisons between our adaptive method with various deterministic and exact stochastic methods. For typical biologically constrained reaction parameters and a 5% accuracy criterion, our adaptive

---

*To whom correspondence should be addressed.

method is computationally faster than exact stochastic methods for cellular volumes > 1–$10\ \mu m^3$.

METHODS AND ALGORITHM

A simple scheme for stochastic transitions is attributed to Nakanishi (1976) as described by Gillespie (1977). It is outlined in brief here:

Assume a reaction

$$A \xrightarrow{kf} B \qquad (1)$$

where kf is the forward rate of the reaction, and A and B are molecular species.

The probability of an individual molecule of A not changing state to B in a time dt is:

$$p\mathrm{Not} = \exp(-kf * dt) \qquad (2)$$

Suppose we have nA molecules capable of making the transition. Now the probability of zero transitions is:

$$p0 = p\mathrm{Not}^{nA} = \exp(-kf * dt * nA) \qquad (3)$$

To generalize to the reaction with $A1, A2, \ldots$ substrates, the number of possible combinations of reaction terms is

$$nA1 * nA2 * \cdots \qquad (4)$$

The probability of zero transitions becomes

$$\begin{aligned} p0 &= p\mathrm{Not}^{(nA1*nA2\cdots)} \\ &= \exp(-kf * dt * nA1 * nA2 \cdots) \end{aligned} \qquad (5)$$

Here, we apply the approximation

$$\exp(-x) \sim 1 - x \quad \text{when } x \ll 1 \qquad (6)$$

Note that this is true within 5% when $x < 0.2$, since the error goes as x^2 and higher order terms.

So if the propensity of the reaction is small, i.e.

$$kf * dt * nA1 * nA2 \cdots \ll 1 \qquad (7)$$

We can approximate

$$p0 \sim 1 - kf * dt * nA1 * nA2 \cdots \qquad (8)$$

Given that $p0$ is close to 1, the probability $p1$ of a single reaction is small, and the probability of more than one reaction is $O(p1^2)$ and higher order terms. Thus, we have a good approximation if

$$p1 \sim (1 - p0) \sim kf * dt * nA1 * nA2 \cdots \qquad (9)$$

and

$$p1 \ll 1 \qquad (10)$$

Note that condition (10) is identical to condition (6). The essence of the approximate stochastic method is to assume that the probability $p1$ of a forward reaction is given approximately by relation 9, and the error term goes as the square of $p1$.

The source for random numbers is a uniform deviate R in $(0,1)$. We use the SPRNG library (Mascagni and Srinivasan, 2000). The update rule is:

$$nA1(t + dt) = nA1(t) - (R < p1) \qquad (11)$$

where the inequality term has a value of either 0 or 1.

Combining stochastic with deterministic calculations

Using the formulation above, we check the probability term $p1$ for each reaction step [Equation (9)] to verify that $p1 \ll 1$. What happens when $p1$ violates this criterion? The simplest approach is to now require that the reaction be computed deterministically. This can readily be done if the calculations are performed in the same units as the stochastic computations, namely, in terms of number of molecules and time. This procedure neglects molecular fluctuations due to these reactions, as considered in the implementation issues below.

From mass action,

$$dA1/dt = -kf * nA1 * nA2 * \cdots \qquad (12)$$

and for a simple Euler integration scheme, this expands into an expression equivalent to Equation (9):

$$\Delta A1 = -kf * dt * nA1 * nA2 * \cdots = P \qquad (13)$$

where P is the propensity. Although (9) and (13) are similar in form, the interpretation of $p1$ as probability no longer applies in (13) as this is now a propensity relationship, and the number of molecules changing state can exceed 1.

Numerically, we can use the same propensity term for P both for the stochastic and the deterministic computation. The naïve approach is to compute P for each reaction transition for each timestep. If P is small we can use the stochastic calculation from Equation (11) to decide if a molecule changes sides, and if P is large we use the deterministic formulation in Equation (13). The threshold for P, $p\mathrm{Min}$, is set by our accuracy requirements, e.g. for 5% accuracy, $P < 0.2$ [Equation (6)].

From Equations (9) and (13), the adaptive method is equivalent to the simple Euler method when it switches to deterministic mode, and the accuracy is of first order. More sophisticated adaptive-time step integration methods such as Runge–Kutta and implicit methods are well known. However, these methods are more difficult to combine with a stochastic method using fixed time steps. We therefore use a more accurate variant of the simple Euler method, called the exponential Euler method, for our deterministic calculations (MacGregor, 1987). We also use the approach of predefined time step control to explicitly specify reduced time steps to improve accuracy when transients are expected, typically as
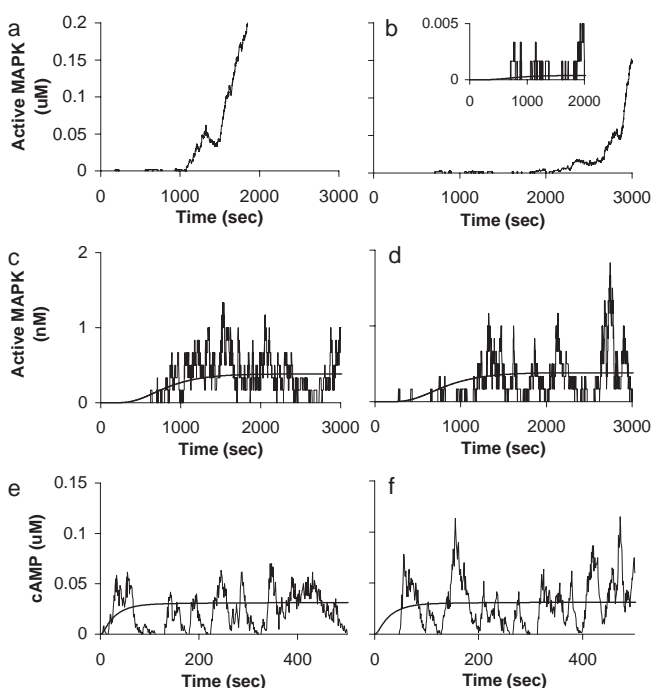
Fig. 1. Example stochastic simulations. Noisy lines are results from typical stochastic runs, smooth lines are deterministic solutions. Left column: results from GB method. Right column: results from adaptive method. (**a**, **b**) Macroscopic effects from microscopic fluctuations. Mitogen-activated protein kinase (MAPK) activity is shown stochastically crossing a threshold for positive feedback loop activation, leading to high (macroscopic) levels of MAPK activity in the system. The model is the large model specified in Table 2. (**c**, **d**) Baseline MAPK responses and fluctuations. (**e**, **f**) Baseline cyclic adenosine monophosphate (cAMP) responses, illustrating different temporal properties of stochastic responses. The model is the cAMP pathway model specified in Table 1.

the result of simulated inputs. Some example cases of results from our adaptive method are compared with results from the Gibson–Bruck (GB) method in Figure 1.
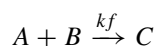
IMPLEMENTATION

Reaction situations

We now consider a set of interesting reaction situations that give rise to numerical problems with the naïve method. In each case, we identify elaborations to the naïve method that address these situations.

1. High propensity with low molecular count
   Consider a reaction

$$A + B \xrightarrow{kf} C$$

where kf is large, nA is large but nB is small. An example of this with typical quantities is when A is Calcium Calmodulin Kinase type II (50 µM), B is activated Calcium–Calmodulin

with four calcium ions bound ($2e - 5 \, \mu M$) and kf is $50 \, [1/(\mu M \cdot s)]$. We therefore have:

$$P = dt * 50 * 2e - 5 * 50 = 5e - 2 \, \mu M/s * dt \quad (14)$$

From our first-order Euler method criterion, the error term is the square of the probability of an individual molecule reacting. For an error term of 5%, we get:

$$(P/\text{total concentration})^2 \sim \text{Error} \sim 0.05 \quad (15)$$

To ensure that we are within this error bound,

$$P/2e - 5 < \sqrt{0.05} \quad (15a)$$

Substituting for P from 14 and reducing, we get a limit for dt:

$$dt < 8e - 5 \quad (16)$$

However, even if we take $dt = 5e - 5$ we get a propensity $P = 2.5e - 6 \, \mu M$ per timestep. For a volume of $1e - 15 \, m^3$, $P = 1.5$ molecules/timestep and $nB = 12$ molecules. It is clearly inappropriate to use deterministic calculations here, as small numbers of molecules are involved and we are very far from the thermodynamic limit. At the same time, our simple stochastic calculations fail because inequality (10) is not satisfied. Our solution is to revert to a pseudo-Euler formulation using stochasticity for the fractional parts:

$$nB(t + 1) = nB(t) + \text{floor}(P) + [(R < \text{frac}(P)] \quad (17)$$

where R is a linear deviate, as in Equation (11), and the inequality term has a value of either 0 or 1. Note that this is a slight elaboration on Equation (11). Equation (17) is similar in form to the chemical Langevin equation, of a deterministic value plus a random term. The random term here is designed to obtain the correct mean but is an approximation to higher moments.

This situation shows that we must set a cut-off number of molecules nMin below which all calculations involving that molecule should be done stochastically. Previous studies have suggested that an order of magnitude separation should be maintained between fast and slow reactions to avoid this situation (Haseltine and Rawlings, 2000). However, this separation is not always feasible when experimental constraints determine reactant numbers and reaction rates. Further, reaction partitioning is automatic and adaptive in our method, so changes during the course of a simulation may introduce this situation.

2. Stochastic lid
Assume a cut-off nMin for the number of molecules below which calculations are performed stochastically, and a reaction where the equilibrium number of molecules is just below nMin. Assume also that propensity is not limiting (i.e. propensity $> p$Min). In stochastic mode, the fluctuations in

molecule number will repeatedly encounter the cut-off nMin. Each time this happens, the calculations switch to deterministic mode. The system relaxes towards equilibrium, which is below nMin, and the process repeats. This process gives rise to a 'lid' to an otherwise stochastic response.

Consider the reaction

$$A \longleftrightarrow B$$
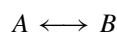
where $kf = 10$, $kb = 1$, $nA + nB = 1050$, $dt = 0.0002$; nMin $= 100$, pMin $= 0.1$.

The equilibrium solution is $nA \sim 95.5$; $nB \sim 954.5$. The expected fluctuations in nA and nB are approximately sqrt(nA) ~ 10. However, the algorithm will apply a 'lid' whenever the fluctuations take the values beyond $nA = 100$ (Fig. 2a). This distorts the expected distribution. One possible solution to this is to use hysteresis in the criteria for switching methods. When tested in complex reaction schemes, however, the hysteresis approach sometimes leads to semi-periodic cycling between the hysteresis bounds because of amplification and feedback in the system. Our solution is to specify a cut-off on molecular number nMin, which is sufficiently large so that the change in standard deviation (SD) is tolerable. In our simulations we use a default nMin of 100 and pMin of 0.1, and these values can be adjusted by the user. The tradeoff for using large nMin is that for large propensities, we use the simple Euler method for integration [Equation (17)] rather than the more accurate exponential Euler method.

3. Rounding

When concentration is converted into molecular number, the value is often fractional. Further, fractions may occur when molecular numbers are calculated at run-time using deterministic numerical integration methods. How should fractions be handled by a mixed-method simulation ? We considered:

— Rounding to the nearest integer (Haseltine and Rawlings, 2002). This may introduce a bias. For example, if the initial number of molecules is 1.2, it will always be rounded to 1. This bias may be amplified in adaptive methods performing run-time rounding. For example, the deterministic numerical method may introduce a small consistent shortfall in molecule concentration, say 0.1 molecule. A molecule cycling between deterministic and stochastic will be rounded up each cycle and the molecular concentration will rise.

— Retention of fractional part of molecule count, but permitting only integer transitions.

   Here, a molecule starting at $n = 1.2$ could take values 0.2 or 2.2 and so on. This is unphysiological. Further, it raises the question of whether the propensities should be calculated using the floating value or the integer value.

— Probabilistic rounding up or down. Here a molecule of count 1.2 is rounded up to 2 a fraction 20% of the time, and down to 1, 80% of the time. The average of
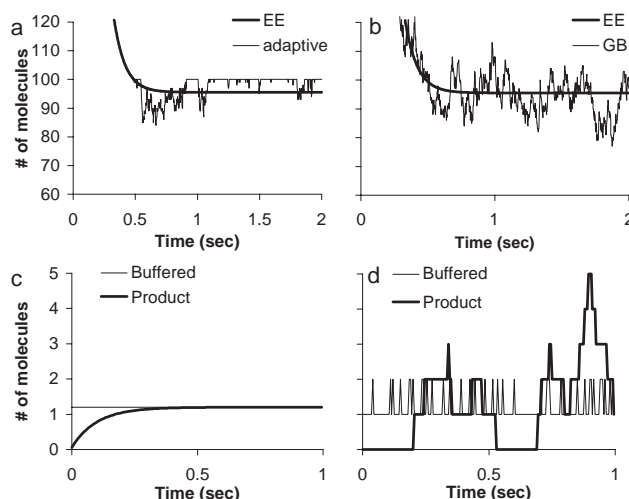

Fig. 2. Special cases at interface between stochastic and deterministic methods. (**a**) Stochastic lid when adaptive method switches between deterministic and stochastic calculations. The reaction is of the form $A \Longleftrightarrow B$ where $kf = 100$, $kb = 10$, $A0 = 1050$ molecules, $B0 = 0$. The thick line is the deterministic calculation using the Exponential Euler (EE) method. (**b**) Same reaction using the GB method, to show the correct distribution. (**c**) Handling of buffering in the deterministic case. Thin line: buffered molecule, thick line: product. The reaction is again of the form Buffered $\Longleftrightarrow$ Product, where $kf = 10$, $kb = 10$. (**d**) Same reaction using the adaptive method, showing the fluctuations in the buffered molecule.

the molecular counts is identical to that obtained using deterministic calculations. This option appeared most consistent, and was adopted.

In addition to initial-value rounding, the same situation may occur continuously in the simulation when molecules are numerically buffered to a set concentration to approximate an external time-varying input or chemical buffering agent. We adopt the same run-time rounding solution (Fig. 2c and d). More accurate distributions around the buffered value could be computed, but were felt to be unnecessarily complex for this simple method. Rounding issues are also important for exact stochastic methods if they are fed with a time-series of numbers scaled from continuous values of concentration terms. The Gillespie methods can be updated regularly with these stochastically rounded values, but the GB method may need additional changes since it retains transition times in a tree structure, and would not normally update the propensities if the stochastic rounding process occurred on an infrequently visited branch.

Test suite

In the course of testing the implementation of the mixed stochastic and exact stochastic methods, we devised a set of reaction situations (Table 1 in supplementary material). These include simple reaction cases as well as numerically difficult

cases, and also the benchmark cases below. This test suite may be useful for validating other approximate stochastic methods.

Benchmarks

We carried out benchmarks on simulations of different sizes to compare the speed and accuracy of different stochastic methods. All benchmarks were done using the GENESIS simulator. The tests were run on an Athlon 2000+ (1.66 GHz) PC with 1 GB RAM running Red Hat Linux 7.3. Although the absolute values of these benchmarks are hardware-dependent, it is likely that the relative speeds of different methods will scale similarly on other systems. Further, these run-times may be valuable in computational resource planning for different scales of biological models.

The simplest simulation was of end-product inhibition. The medium model was the cyclic adenosine monophosphate (cAMP) signaling pathway (Bhalla, 2002b). The large model was of synaptic signaling (Bhalla, 2002a). The latter two models are based on biochemical parameters from the literature and are available in the Database of Quantitative Cellular Signaling (Sivakumaran *et al.*, 2003) (Table 2 in supplementary material). Each model was run at a range of volumes with fixed initial concentrations. The smaller models were run for 0.1, 1, 10, 100 and 1000 μm^3. Due to computational limitations the large model was only run for smaller volumes. Computational times are reported for a single run of at least 10 s. If the computational time was <10 s then we recorded the time for 100 runs in order to estimate the time for a single run. As run-time was sensitive to the number of enzymes in an active state, we performed each series of tests at two different levels of input activation.

We report the speed of each method as clocked run-times divided by simulated time and by the number of chemical reactions in the model, to estimate the time taken to simulate each reaction for 1 s. Four points stand out from these benchmark results.

First, the exact stochastic methods have an almost linear dependence of run time on the volume, and hence on the number of molecules (N) in the model. This is expected since the number of molecular transitions also scales with N, and the exact stochastic methods calculate at least one new random number for each molecular transition. For very small N, setup times dominate the benchmarks, so this linear relationship flattens out. The deterministic solutions do not use random numbers and are therefore somewhat faster.

Second, the run-time for each of the exact stochastic methods increases sharply at greater activation levels within the physiological range. This is because molecular transitions occur more often at high levels of activity.

Third, the exact stochastic methods scale differently with model complexity. The GB method incurs greater housekeeping overhead for small models, but its more efficient use of random numbers gives it 5-fold greater speed than Gillespie's direct method for the large synaptic model (Fig. 3e and f).
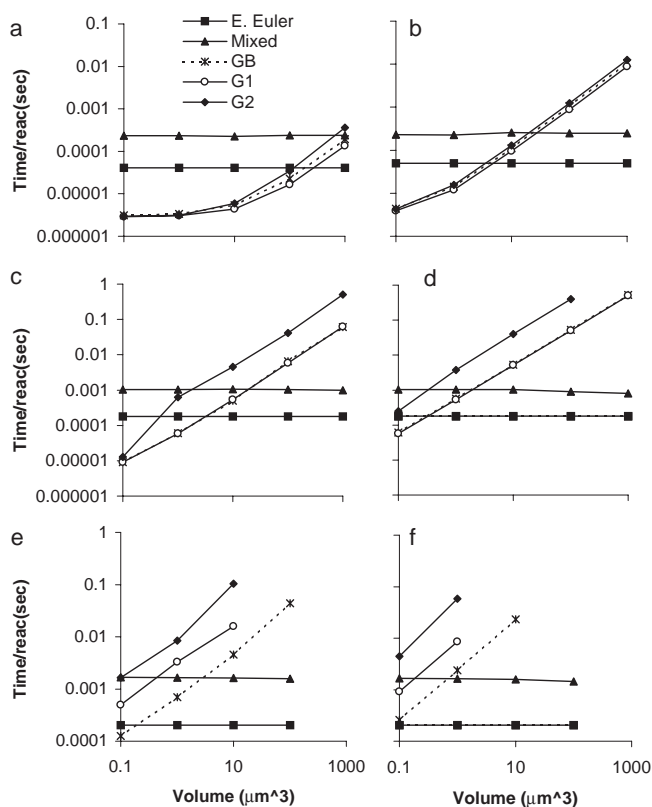

Fig. 3. Benchmarks. Graphs are plotted as time taken to simulate the progress of one chemical reaction for a period of 1 s, against the volume of the reaction mixture. For reference, at 1 μm^3 a concentration of 1 μM is equal to approx. 600 molecules. Left column: simulations with low levels of activation in model. Right column: simulations with high levels of activation in model. (**a, b**) small model (five chemical reactions). All the exact stochastic methods are equivalent. (**c, d**) Medium model (54 reaction steps). Gillespie's First Reaction method (G2) is now slower than Gillespie's Direct method (G1) or the Gibson–Bruck method (GB). (**e, f**) Large model (344 reaction steps). The GB method is clearly the fastest exact method here. It runs considerably slower under high-activation conditions (f). The adaptive stochastic method is the fastest over 1–10 μm^3 (5% error criterion).

Fourth, the benchmarks show that our implementation of the adaptive method is more computationally efficient than even the GB method at volumes greater than about 10 μm^3, when using biologically-constrained molecular concentrations and rate constants and an error limit of 5%. This occurs despite our using first-order calculations. In practical terms, anything larger than an *Escherichia Coli* or a dendritic spine of a typical mammalian neuron would be a candidate for an adaptive method.

Accuracy

A key question in going from an exact stochastic method to an approximate method is to determine the accuracy of the
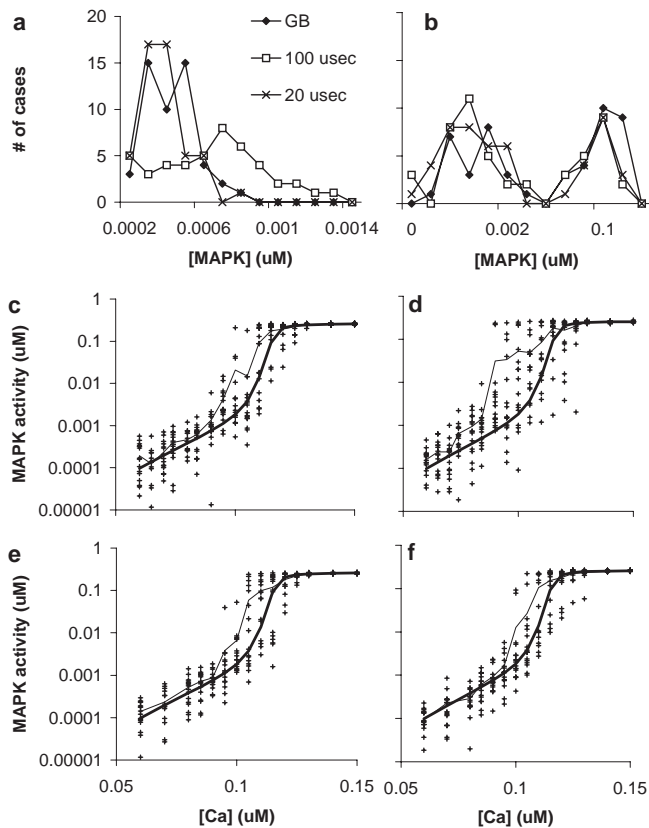
Fig. 4. Accuracy and scatter for MAPK activity in large model, after equilibrating for 3000 s. (**a**) Distribution of MAPK activity (averaged over 1000 s) for 10 μm^3 at basal calcium. (**b**) Distribution of MAPK activity for 1 μm^3. Distributions are based on the results from 50 runs. (**c–f**) MAPK activity as a function of fixed calcium levels. Solid lines are the deterministic solution. Plus symbols are 100 s means from 16 individual stochastic runs. Thin line is average of stochastic runs. (**c**) GB method. The mean of the stochastic runs is larger than the deterministic plot because of the asymmetry in the distribution. (**d–f**) Scatter using the adaptive stochastic method at timesteps of 100, 200 and 20 μs, respectively. The 100 μs case in (d) has greater scatter.

method under a range of conditions. We compared the mean and SD (first and second moments of the distributions) of the results from the GB method with those from the adaptive method, utilizing our large model and 50 runs in order to generate the distributions. The comparisons considered 17 molecules in the model, following 3000 s. equilibration time, for runs at a volume of 10 μm^3 (data not shown). When the adaptive method was run at a timestep of 100 μs, we observed differences only in the distribution for mitogen-activated protein kinase (MAPK) (*T*-test, $p < 1e - 4$). The MAPK distributions were not significantly different for a timestep of 20 μs (Fig. 4a).

In this model, the MAPK pathway participates in a positive-feedback loop, which is bistable under the conditions of the

model (Bhalla *et al*., 2002). We asked if positive feedback could amplify differences in distributions between the GB and the adaptive method. We performed a series of runs at increasing calcium levels for the model using the deterministic, the GB and the adaptive method at various timesteps. Each of the stochastic runs was repeated 16 times to generate the distributions (Fig. 4c–f). As expected, reducing the timestep improves the accuracy with which the distribution is calculated (Fig. 4f). From Equations (7), (9) and (10) it is straightforward to estimate the timestep that gives a desired accuracy for any given model using the adaptive method, and we have implemented this in our simulator. In practice we find this is a highly conservative estimate.

DISCUSSION

We have implemented and extensively tested a method that adaptively selects stochastic or deterministic calculations to numerically simulate the progress of chemical reactions. Our method is simple, approximate and adaptive. Although we are aware of the limitations of fixed time step methods, we feel there is value in having a well-grounded and tested simple method as a step towards more efficient methods that address the complexities of adaptive stochastic calculations. The adaptive functioning of the method has two advantages: First, it automatically partitions complex reaction schemes between stochastic and deterministic cases, thus minimizing the need for user input. Second, it sets an upper bound to computation time in the large-volume limit as the system automatically switches to deterministic calculations. We have analyzed issues of switching between methods, performed efficiency benchmarks, and examined the accuracy of the method.

Optimization

With a 5% error allowance, our implementation of this adaptive method is faster than sophisticated exact stochastic methods such as the Gibson–Bruck method (Gibson and Bruck, 2000) for typical biological signaling systems in volumes over 1–10 μm^3. For comparison, *E.coli* has a volume of about 3 μm^3 and a typical mammalian cell is 1000 μm^3 or larger. Our dynamic stochastic/deterministic framework could be applied to improved stochastic and deterministic methods. For example, Gillespie (2001) has described an approximate version of his next step method, called the tau-leap method. It would be very interesting to combine such a method with the Runge–Kutta method. Based on the speed-up of the tau-leap method and the greater efficiency of Runge–Kutta integration for the deterministic domain, one could expect an improvement of 10–1000-fold.

Further detail

The current approach lacks microscopic and spatial detail. At the microscopic level, our method treats each pool of

molecules of a given type as indistinguishable for the purposes of computing transitions. The simulator STOCHSIM handles cases where molecular complexes are distinguishable (Le Novere and Shimizu, 2001). A more complete implementation of hybrid methods could be a three-tier hybrid scheme, utilizing distinguishable molecular stochastic transitions at the finest level, then stochastic calculations based on indistinguishable molecules, and finally deterministic calculations.

Spatial detail is also omitted from the current algorithm. A fully stochastic spatial approach is used by the simulator MCell (Stiles and Bartol, 2000). The opposite extreme is to use deterministic calculations in a reaction-diffusion system obeying Fick's law, as is done by VCell (Schaff *et al.*, 1997). Our current analysis suggests that there is considerable scope for devising adaptive methods that perform automatic scaling between these extremes.

ACKNOWLEDGEMENTS

REFERENCES

Arkin,A., Ross,J. and McAdams,H.H. (1998) Stochastic kinetic analysis of developmental pathway bifurcation in phage lambda-infected *Escherichia coli* cells. *Genetics*, **149**, 1633–1648.

Bhalla,U.S. (2002a) Mechanisms for temporal tuning and filtering by postsynaptic signaling pathways. *Biophys. J.*, **83**, 740–752.

Bhalla,U.S. (2002b) Use of Kinetikit and GENESIS for modeling signaling pathways. In Hildebrandt,J.D. and Iyengar,R. (eds), *Methods Enzymol*. Academic Press, pp. 3–23.

Bhalla,U.S. and Iyengar,R. (1999) Emergent properties of networks of biological signaling pathways. *Science*, **283**, 381–387.

Bhalla,U.S., Ram,P.T. and Iyengar,R. (2002) MAP kinase phosphatase as a locus of flexibility in a mitogen-activated protein kinase signaling network. *Science*, **297**, 1018–1023.

Bialek,W. (2001) Stability and noise in biochemical switches. *Adv. Neural Information Processing Systems*, **13**, 103–109.

Duke,T.A.J., Le Novere,N. and Bray,D. (2001) Conformational spread in a ring of proteins: a stochastic approach to allostery. *J. Mol. Biol.*, **308**, 541–553.

Gibson,M.A. and Bruck,J. (2000) Efficient exact stochastic simulation of chemical systems with many species and many channels. *J. Phys. Chem. A*, **104**, 1876–1889.

Gillespie,D.T. (1977) Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.*, **81**, 2340–2361.

Gillespie,D.T. (2001) Approximate accelerated stochastic simulation of chemically reacting systems. *J. Chem. Phys.*, **115**, 1716–1733.

Hartwell,L.H., Hopfield,J.J., Leibler,S. and Murray,A.W. (1999) From molecular to modular cell biology. *Nature*, **402**, C47–C50.

Haseltine,E.L. and Rawlings,J.B. (2002) Approximate simulation of coupled fast and slow reactions for stochastic chemical kinetics. *J. Chem. Phys.*, **117**, 6959–6969.

Holmes,W.R. (2000) Models of calmodulin trapping and CaM kinase II activation in a dendritic spine. *J. Comput. Neurosci.*, **8**, 65–85.

Hucka,M., Finney, A., Sauro,H.M., Bolouri,H., Doyle,J. and Kitano,H. (2002) The ERATO Systems Biology Workbench: enabling interaction and exchange between software tools for computational biology. *Pac. Symp. Biocomput.*, **2002**, 450–461.

Janssen,J.A.M. (1989) The elimination of fast variables in complex chemical reactions II. Mesoscopic level (reducible case). *J. Stat. Phys.*, **57**, 171-198.

Kierzek,A.M., Zaim,J. and Zielenkiewicz,P. (2001) The effect of transcription and translation initiation frequencies on the stochastic fluctuations in prokaryotic gene expression. *J. Biol. Chem.*, **276**, 8165–8172.

Kurtz,T.G. (1972) The relationship between stochastic and deterministic models for chemical reactions. *J. Chem. Phys.*, **57**, 2976–2978.

Le Novere,N. and Shimizu,T.S. (2001) STOCHSIM: modelling of stochastic biomolecular processes. *Bioinformatics*, **17**, 575–576.

Lisman,J.E. and Goldring,M.A. (1988) Feasibility of long-term storage of graded information by the Ca^{2+}/calmodulin-dependent protein kinase molecules of the postsynaptic density. *Proc. Natl Acad. Sci. USA*, **85**, 5320–5324.

MacGregor,R.J. (1987) *Neural and Brain Modeling*. Academic Press, San Diego.

Mascagni,M. and Srinivasan,A. (2000) Algorithm **806**: SPRNG: a scalable library for pseudorandom number generation. *ACM Trans. Math. Soft.*, **26**, 436–461.

Nakanishi,T. (1976) *J. Phys. Soc. Jpn*, **40**, 1232.

Rao,C.V., Wolf,D.M. and Arkin,A.P. (2002) Control, exploitation and tolerance of intracellular noise. *Nature*, **420**, 231–237.

Resat,H., Wiley,H.S. and Dixon,D.A. (2001) Probability-weighted dynamic Monte Carlo method for reaction kinetics simulations. *J. Phys. Chem. B*, **105**, 11026–11034.

Schaff,J., Fink,C.C., Slepchenko,B., Carson,J.H. and Loew,L.M. (1997) A general computational framework for modeling cellular structure and function. *Biophys. J.*, **73**, 1135–1146.

Sivakumaran,S., Hariharaputran,S., Mishra,J. and Bhalla,U.S. (2003) The database of quantitative cellular signaling: repository and analysis tools for chemical kinetic models of signaling networks. *Bioinformatics*, **19**, 408–415.

Stiles,J.R. and Bartol,T.M. (2000) Monte Carlo methods for simulating realistic synaptic microphysiology using MCell. In De Schutter,E. (ed.), *Computational Neuroscience: Realistic Modeling for Experimentalists*. CRC Press, Boca Raton, pp. 87–128.

Vlad,M.O. and Pop,A. (1989) A physical interpretation of age-dependent master equations. *Physica A*, **155**, 276–310.