Modelling and analysis of the sugar cataract development process using stochastic hybrid systems

D. Riley[1]   X. Koutsoukos[1]   K. Riley[2]

[1]ISIS/EECS, Vanderbilt University, Nashville, TN 37235, USA
[2]Yale University, HHMI, New Haven, CT 06520, USA
E-mail: derek.riley@vanderbilt.edu

Abstract: Modelling and analysis of biochemical systems such as sugar cataract development (SCD) are critical because they can provide new insights into systems, which cannot be easily tested with experiments; however, they are challenging problems due to the highly coupled chemical reactions that are involved. The authors present a stochastic hybrid system (SHS) framework for modelling biochemical systems and demonstrate the approach for the SCD process. A novel feature of the framework is that it allows modelling the effect of drug treatment on the system dynamics. The authors validate the three sugar cataract models by comparing trajectories computed by two simulation algorithms. Further, the authors present a probabilistic verification method for computing the probability of sugar cataract formation for different chemical concentrations using safety and reachability analysis methods for SHSs. The verification method employs dynamic programming based on a discretisation of the state space and therefore suffers from the curse of dimensionality. To analyse the SCD process, a parallel dynamic programming implementation that can handle large, realistic systems was developed. Although scalability is a limiting factor, this work demonstrates that the proposed method is feasible for realistic biochemical systems.

1   Introduction

Modelling and analysis of biochemical systems are important tasks because they can unlock insights into the complicated dynamics of systems, which are difficult, expensive or presently impossible to test experimentally. A variety of techniques have been used to model biochemical systems, but their effectiveness is often limited by trade-offs imposed by the modelling paradigms. Stochastic differential equations (SDEs) have been used to model biochemical reactions [1, 2]; however, analysis of these models has mainly been limited to simulation. Hybrid systems have also been used to model biochemical systems [3, 4]; however, verification methods based on deterministic hybrid systems fail to capture the probabilistic nature of some biochemical processes and therefore may not be able to correctly analyse certain systems. Stochastic hybrid systems (SHS) have been used to capture the stochastic nature of biochemical systems but have

previously only been used for simulations [5] or analysis of systems with simplified continuous dynamics [6].

In this paper, we present an SHS framework for modelling biochemical systems and demonstrate the approach for the sugar cataract development (SCD) process. An accumulation of sorbitol is thought to be a major factor in the development of a sugar cataract. An explanation of the theorised process and experimental data can be seen in [7]. Certain drugs have been developed to limit the efficiency of the enzymatic reactions, thus reducing the chance of cataract development. However, many such drugs have off-target or unpredicted effects, which perturb the system in other, often unexpected ways. The chemical reactions and kinetic coefficients for the model have been previously studied [8], but a model incorporating the effect of drug treatment has, to our knowledge, not been created. Understanding the exact conditions that lead to the development of sugar cataracts and

how drug treatment will potentially reverse the conditions will help better predict and prevent sugar cataracts from developing.

This work provides several new contributions. Our SHS framework can be used to model and simulate complex large-scale biochemical systems. We illustrate the approach using the SCD starting from an existing model to establish realistic parameters. We then present two new models of medication-controlled SCD that extend previous SHS models. The first model incorporates the medication administration policy and models the effect that the medication has on the chemical reactions. The second new model adds probabilistic delays to capture both the absorption of the drug into the system and the eventual drug metabolism.

We validate the SHS models by comparing trajectories computed by two simulation algorithms, the hybrid stochastic simulation algorithm (HSSA) and hybrid Euler–Maruyama (HEM) algorithm. These algorithms extend the stochastic simulation algorithm (SSA) and Euler–Maruyama (EM), respectively, by incorporating switching between discrete modes. Further, we show that the HEM algorithm is an order of magnitude more efficient than the HSSA algorithm and can achieve similar accuracy.

We use a dynamic programming verification method based on a discretisation of the state space to analyse the SCD problem [9]. We analyse both safety and reachability properties of the medicated and non-medicated models. In the context of SCD, we can use the safety verification to determine the probability that a cataract will develop from any state of the system. We can use reachability analysis to determine the probability that the patient will reach a safe condition without first developing a cataract. This is particularly useful when analysing the effectiveness of a drug. We apply the SHS verification methods described in [9] for analysing the SCD system. While the stochastic dynamics of biochemical processes in general can be accurately modelled by the chemical master equation (CME), the equation is impossible to solve for most practical systems [1]. The SSA is equivalent to solving the master equation, but if the number of molecules of any of the reactants is large, the SSA is not efficient [5]. It is computationally intractable to enumerate all possible states of the model employed by the SSA for formal verification because the reaction rates depend on the concentrations, and the SSA models individual molecules. Therefore our approach suggests starting with the continuous stochastic dynamics and generating discrete approximations with coarser (and variable) resolution.

The proposed verification method suffers from the curse of dimensionality, so we have developed a parallel dynamic programming implementation of the verification algorithm that can handle large systems. The parallel algorithm enhances performance by dividing the problem among many processors to distribute memory storage and computation. Although scalability is a limiting factor, this work demonstrates that the parallel technique is feasible for realistic biochemical systems. Further, we present a practical error analysis method which characterises the quality of the solution based on the resolution of the approximation.

The SHS modelling framework offers many advantages for modelling biochemical systems. It allows for continuous time and event-driven dynamics in multiple time scales in a stochastic context. The method has the ability to model multiple behavioural modes of a system as well as the switching dynamics between the modes in order to capture, for example, the effects of medication. The resulting framework can be used to model systems with complex nonlinear dynamics while still formally defining the execution and behaviour. The simulation algorithms offer efficiency and simplicity. SHS simulation is significantly more efficient than the SSA algorithm because it uses a coarser resolution without significantly affecting accuracy.

In addition to simulation, SHS offer a formal framework for probabilistic verification of reachability properties. In contrast to analysis based on Monte Carlo methods, formal verification of stochastic systems is an exhaustive technique that computes the probability that the system trajectory for arbitrary initial states will reach a target set while avoiding unsafe regions and it considers an infinite time horizon. Such analysis can generate insights that cannot be obtained by simple simulations. The parallel methods we have developed enable the analysis of systems with up to seven continuous dimensions that correspond to the number of reacting chemical species. The analysis method uses a discrete grid approximation and can currently handle the order of 10^9 discrete states. The large, fine grid ensures an accurate approximation and generates reachability probabilities for every possible initial condition of the system.

The organisation for the rest of the paper is as follows: Section 2 describes the related work, Section 3 describes modelling of biochemical systems and our three SHS models of SCD, Section 4 describes the probabilistic verification method, Section 5 presents our experimental results and Section 6 concludes the work.

2 Related work

SCD has been studied previously and modelled because it is an important problem, and the dynamics of the system are complex and difficult to test experimentally [8]. The structure of SDH has been used in a molecular modelling program (DOCK) to analyse nearly 250 000 compounds from the National Cancer Institute Database and predict those with high affinity for SDH [10]. Such compounds can be tested in diabetes models as potential therapeutics.

Experimental advances in the biological sciences can be incorporated into formal modelling and analysis techniques to benefit both fields mutually. Previous research has applied modelling and analysis techniques to biological

138

IET Syst. Biol., 2009, Vol. 3, Iss. 3, pp. 137–154

systems from species population evolution to molecular dynamics, but generally the models have been limited in size or accuracy because of computational restrictions. As computing power has increased, modelling and analysis approaches have evolved to take advantage of the increased power to improve accuracy and speed. Formal analysis of the models for these and other systems has provided new insights into complicated or difficult to test systems. Modelling the effect of drugs and specifically enzyme inhibitors is an important task because it can enhance the understanding prior to the execution of actual experiments, which have inherent disadvantages. Previous drug modelling has focused mainly on the physical interactions at the molecular level [11–13]. Modelling of enzyme inhibitors has not, to our knowledge been attempted in conjunction with chemical reactions at the scale that we present.

Stochastic π-calculus is a modelling framework that is able to express systems of concurrent components [14]. These components, or processes, each define a continuous time Markov chain (MC) and therefore they can be simulated using Gillespie's algorithm. Several tools have been created, which implement the semantics of Gillespie's algorithm for stochastic π-calculus models [15, 16]. The models are countably infinite, so verification can only be performed symbolically, and current symbolic verification techniques are unable to handle nonlinear dynamics. Nonlinear dynamics are common in biochemical systems, so verification of realistic systems with this technique is limited.

SDEs have been used for modelling cell signalling pathways and molecular motion [4, 2, 14]. Since only specialised cases of SDEs can be solved analytically, the vast majority of models are simulated using Monte Carlo techniques. Our models use SDEs to model the continuous dynamics; however, we also incorporate discrete dynamics to model phenomena such as medication application.

Nonstochastic hybrid systems have been used for modelling biological systems in order to capture the complicated dynamics using well-defined abstractions. Biomolecular network modelling is accomplished by using differential equations to model feedback mechanisms and discrete switches to model changes in the underlying dynamics [3]. Biological protein regulatory networks have been modelled with hybrid systems using linear differential equations to describe the changes in protein concentrations and discrete switches to activate or deactivate the continuous dynamics based on protein thresholds [4].

SHS further improve on the accuracy of hybrid systems by providing a more realistic probabilistic framework for modelling real-world biochemical systems. A modelling technique that uses SHS to construct models for chemical reactions involving a single reactant specie is presented in [6]. An algorithm for computing the reachability probability of SHS is presented in [17]. A genetic regulatory network was modelled with an SHS model and compared to a deterministic model in [18]. SHS models of biochemical systems have been developed and simulated using hybrid simulation algorithms in [5, 19].

Several modelling and simulation techniques have been developed that combine other modelling techniques to improve the overall accuracy of the methods. Methods that combine the use of ordinary differential equations with the SSA are described in [20–23]. A method that combines τ-leaping and the next reaction technique is described in [24]. The methods do not formally allow for the inclusion of external hybrid dynamics, nor do they include stochastic continuous dynamics. Tools for the stochastic simulation of chemical reactions have been developed in [25, 26].

The finite state projection (FSP) algorithm can be applied to the solution of the CME as an alternative to the SSA algorithm and Monte Carlo methods [27]. The approach has been extended to systems with multiple time scales [28]. FSP utilises a method of model reduction based on knowledge about a specific system. While this method is efficient, it is restricted to a finite time horizon. Further, small changes to the model (such as medication effects) may change the model reduction. The lattice of states generated by the FSP could be used by our verification method and incorporating the model reduction technique into our SHS framework is a promising research direction.

The analysis technique described in this paper employs a reachability analysis method based on discrete approximations. Discrete approximation methods based on finite differences have been studied extensively in [29]. On the basis of discrete approximations, the reachability problem can be solved using algorithms for discrete processes [30]. The approach has been applied for optimal control of SHS given a discounted cost criterion in [31]. For verification, the discount term cannot be used and convergence of the value function can be ensured only for appropriate initial conditions. A related grid-based method for safety analysis of stochastic systems with applications to air traffic management has been presented in [32]. Our approach is similar but using viscosity solutions we show the convergence of the discrete approximation methods [9].

Reachability analysis for SHS can be performed using Monte Carlo methods [33]. Multiple stochastic simulations are used to determine the reachability probability for an initial state of an SHS. Confidence intervals and accuracy probabilities can be selected by adjusting the number of simulations. Stochastic roadmap simulation extends the Monte Carlo technique by analysing multiple trajectories simultaneously. The analysis of these ensemble properties can significantly improve the understanding of the entire system [34]. Our technique extends this notion to analyse all possible starting values and trajectories simultaneously; however, scalability remains a limiting factor of our approach.

3 Modelling biochemical reactions using SHS

In this section, we present techniques for modelling biochemical processes using SHS. First, we present the chemical reaction modelling method used and then we describe extensions to these techniques to model medications. We conclude by formally defining SHS, which we use to model the biochemical systems.

3.1 Dynamics of biochemical reactions

Chemical reactions are inherently probabilistic because of the unpredictability of molecular motion, so stochastic models are ideal for describing the dynamics accurately [35]. Slow reactions occur when reaction rates and concentrations are small enough and they can be modelled and simulated efficiently using discrete stochastic techniques. However, discrete simulations become inefficient when there are large concentrations of molecules and/or fast reaction rates. When discrete models become inefficient, these 'fast' reactions can be accurately modelled as continuous stochastic models [5].

Fast reactions occur at a rate that is fast enough to consider as occurring at a constant rate therefore eliminating the need to consider individual reactions. Such reactions can be modelled more efficiently as continuous stochastic models. The rate of change of each chemical species in a fast reaction is calculated using the chemical dynamics from the biochemical reactions using (1) [5, 36]

$$\mathrm{d}x_i = \sum_{j=1}^{M_{\mathrm{fast}}} v_{ji} a_j(x(t))\mathrm{d}t + \sum_{j=1}^{M_{\mathrm{fast}}} v_{ji} \sqrt{a_j(x(t))}\mathrm{d}w_j \qquad (1)$$

Biochemical systems can contain a mixture of both fast and slow reactions. When fast and slow dynamics must both be considered it is most efficient to use a combined, hybrid modelling approach to take advantage of the efficiency of continuous modelling for the fast reactions while still keeping the accuracy of discrete modelling for the slow reactions. Determining which reactions are fast or slow is based on analysis of the rates using the kinetic coefficients and chemical concentrations. To determine the slowest rate, the smallest possible concentrations for each chemical species are used. Similarly, the fastest rate can be determined by using the highest possible concentrations. Since the reaction rates depend on the concentrations, reactions may be classified as either fast or slow dynamically based on the system state.

3.2 Medication modelling

Understanding how a biochemical system will operate under normal conditions is important; however, in many systems, it is advantageous to understand how the system will act when it is perturbed by outside influences such as a medication. The interaction of a drug with a biochemical system is important to model and analyse because often the anticipated affect of the drug is altered by unforeseen influences, and theoretical modelling and testing can help to demonstrate the safety of a medication before it is tested on real subjects.

Drugs are administered to patients to improve their health by altering the equilibrium of the biochemical reactions responsible for their symptoms. There are several defining characteristics of drugs that are considered when modelling their behaviour. Drugs can generally be classified as either stimulants or inhibitors that increase or decrease reaction rates. The efficacy of a drug is the potential therapeutic response that a it might produce. Additionally, drugs are metabolised by the body at varying rates, so the decay of the drug must be understood to accurately model its behaviour.

The most direct drug modelling approach is to add the relevant chemical species and reactions. While this may appear to be a logical approach, it adds complexity to the system, and may not completely describe the effects of the drug if certain chemical reactions are not considered. A simpler technique is to model the behaviour of a drug as an inhibitor or stimulant and avoid increasing the number of chemical reactions or chemical species considered.

Because stimulants and inhibitors alter the reaction rates of certain reactions, modelling the effect of a drug on a given chemical reaction can be accomplished by altering the kinetic coefficients. The amount of change of the kinetic coefficients is determined by the efficacy and metabolism rate of the drug. For the SCD model, discrete modes describe the system under different drug influences, and discrete transitions model the application and metabolism of the drug.

SHSs are ideal for modelling the effect of medication in biochemical systems because they are able to model continuous and discrete dynamics in a stochastic framework which is both efficient and accurate. To describe our approach, we will next define a formal model of SHSs.

3.3 Stochastic hybrid systems

We adopt the model presented in [37]. To establish the notation, let Q be a set of discrete states. For each $q \in Q$, we consider the Euclidean space $\mathbb{R}^{d(q)}$ with dimension $d(q)$ and we define an invariant as an open set $X^q \subseteq \mathbb{R}^{d(q)}$. The hybrid state space is denoted as $S = \bigcup_{q \in Q} \{q\} \times X^q$. Let $\bar{S} = S \cup \partial S$ and $\partial S = \bigcup_{q \in Q} \{q\} \times \partial X^q$ denote the completion and the boundary of S, respectively. The Borel σ-field in S is denoted as $\mathcal{B}(S)$.

Definition 1: A SHS is defined as $H = ((Q, d, \mathcal{X}), \boldsymbol{b}, \boldsymbol{\sigma}, \mathrm{Init}, \lambda, R)$ where Q is a set of discrete states (modes), $d: Q \to \mathbb{N}$ is a map that defines the continuous state space

dimension for each $q \in Q$, $\mathcal{X}: Q \to \mathbb{R}^{d(\cdot)}$ is a map that describes the invariant for each $q \in Q$ as an open set $X^q \subseteq \mathbb{R}^{d(q)}$, $\boldsymbol{b}: Q \times X^q \to \mathbb{R}^{d(q)}$ and $\boldsymbol{\sigma}: Q \times X^q \to \mathbb{R}^{d(q) \times p}$ are drift vectors and dispersion matrices, respectively, Init: $\mathcal{B}(S) \to [0, 1]$ is an initial probability measure on S, $\lambda: \bar{S} \to \mathbb{R}_+$ is a non-negative transition rate function and $R: \bar{S} \times \mathcal{B}(\bar{S}) \to [0, 1]$ is a transition measure.

To define the execution of the system, we denote $(\Omega, \mathcal{F}, P)$ the underlying probability space, and consider an $\mathbb{R}^p$-valued Wiener process $w(t)$ and a sequence of stopping times $\{t_0 = 0, t_1, t_2, \ldots\}$. Let the state at time t_i be $s(t_i) = (q(t_i), x(t_i))$ with $x(t_i) \in X^{q(t_i)}$. When there is no confusion, we will use interchangeably the notation (q, x) and s for the hybrid state to simplify complex formulas and often we will use the notation $s_{t_i} = (q_{t_i}, x_{t_i})$ for brevity. While the continuous state stays in $X^{q(t_i)}$, $x(t)$ evolves according to the SDE

$$\mathrm{d}x = \boldsymbol{b}(q, x)\mathrm{d}t + \boldsymbol{\sigma}(q, x)\mathrm{d}w \qquad (2)$$

where the discrete state $q(t) = q(t_i)$ remains constant and the solution of (2) is understood using the Itô stochastic integral [38]. A sample path of the stochastic process is denoted by $x_t(\omega)$, $t > t_i$, $\omega \in \Omega$.

The next stopping time t_{i+1} represents the time when the system transitions to a new discrete state. The discrete transition occurs either because the continuous state x exits the invariant $X^{q(t_i)}$ of the discrete state $q(t_i)$ (guarded transition) or based on an exponential distribution with transition rate function λ (probabilistic transition). Therefore t_{i+1} can be defined as the minimum between two other stopping times: (i) The first hitting time of the boundary $\partial X^{q(t_i)}$ defined as $t_{i+1}^* = \inf\{t \geq t_i, x(t) \in \partial X^{q(t_i)}\}$ and (ii) a stopping time τ_{i+1} described by an exponential distribution with survivor function

$$M(t, \omega) = \exp\left(-\int_{t_i}^{t} \lambda(q(t_i), x_z(\omega))\mathrm{d}z,\right), \ \omega \in \Omega$$

Thus, the time of the next discrete transition t_{i+1} is a stopping time whose distribution is defined by the survivor function

$$F(t, \omega) = I_{(t < t_{i+1}^*)} \exp\left(-\int_{t_i}^{t} \lambda(q(t_i), x_z(\omega))\mathrm{d}z\right), \ \omega \in \Omega$$

where I denotes the indicator function. Given a set $A \in \mathcal{F}$ the indicator function is defined as $I_A(\omega) = 1$ if $\omega \in A$ and 0 if $\omega \notin A$.

At time t_{i+1} the system will transition to a new discrete state and the continuous state may jump according to the reset measure R. The trajectory of $x(t)$ is assumed to be left-continuous, so we denote $x(t_{i+1}^-)$ the solution of (2) at $t = t_{i+1}$ and $s(t_{i+1}^-) = (q(t_{i+1}^-), x(t_{i+1}^-))$ where $q(t_{i+1}^-) = q(t_i)$

the discrete state before the transition. If $t_{i+1} = \infty$, the system continues to evolve according to (2) with $q(t) = q(t_i)$. If $t_{i+1} < \infty$, the system jumps at t_{i+1} to a new state $s(t_{i+1}) = (q(t_{i+1}), x(t_{i+1}))$ according to the transition measure $R(s(t_{i+1}^-), A)$ with $A \in \mathcal{B}(S)$. The evolution of the system is then governed by the SDE (2) with $q(t) = q(t_{i+1})$ until the next stopping time.

Fig. 1 shows a generic SHS model with two states and two transitions (one probabilistic and one guarded). The continuous dynamics of each state are defined by the associated SDEs. The probabilistic transition fires at the firing rate λ, and the guarded transition fires when x hits the boundary $x \in \partial X^{q_2}$. The logical condition $x \in \partial X^{q_2}$ is often referred to as the guard of the transition. Upon firing of a transition, the state resets according to the map $R((q, x), A)$.

The following assumptions are imposed on the model. The functions $\boldsymbol{b}(q, x)$ and $\sigma(q, x)$ are bounded and Lipschitz continuous in x for every q, and thus the SDE (2) has a unique solution for every q. The transition rate function λ is a bounded and measurable function, which is assumed to be integrable for every $x_t(\omega)$. For the transition measure, it is assumed that $R(\cdot, A)$ is measurable for all $A \in \mathcal{B}(S)$, $R(s, \cdot)$ is a probability measure for all $s \in S$ and $R((q, x), \mathrm{d}z)$ is a stochastic continuous kernel. Also, the boundaries ∂X^q are assumed to be sufficiently smooth and the trajectories of the system satisfy a non-tangency condition with respect to the boundaries. A sufficient condition for the non-tangency assumption is that the diffusion term is non-degenerate, that is, $a(q, x) = \boldsymbol{\sigma}(q, x)\boldsymbol{\sigma}^{\mathrm{T}}(q, x)$ is positive definite. Furthermore, it is assumed that the set Q is finite and that X^q is bounded for every q. Let $N_t = \sum_i I_{\geq t_i}$ denote the number of jumps in the interval $[0, t]$. It is assumed that the expected number of jumps is finite for every initial state $s \in S$, that is $E_s[N_t] < \infty$. A sufficient condition for ensuring finitely many jumps can be formulated by imposing restrictions on the map $R(s, A)$ [9, 39].

3.4 Sugar cataract modelling

This section describes three SHS models of the biochemical process of SCD. The first model describes the biochemical process of SCD. The two subsequent models extend the first model to include the effect of medication on the system. The first medicated model assumes that the effect of the drug on the system is instantaneous, while the final
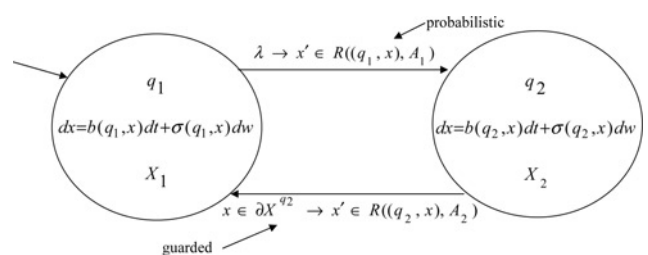

Figure 1 *Stochastic hybrid system*

model is designed to incorporate probabilistic delay to model absorption and metabolisation.

SCD model (SCD1): A sugar cataract distorts the light passing through the lens of an eye by attracting water to the lens when an excess of sorbitol is present. Often these cataracts are formed in the eyes of diabetic patients who have highly fluctuating blood sugar levels. Several factors affect the accumulation of sorbitol including the amount of the enzyme SDH. SDH catalyses the reversible oxidation of sorbitol and other polyalcohols to the corresponding keto-sugars [8]. There are eight chemical species involved in the reaction: NADH(x_1), E–NADH(x_2), NAD$^+$(x_3), E–NAD$^+$(x_4), SDH(x_5), Fructose(x_6), Sorbitol(x_7) and the inactive form of SDH (Z).

A SHS model for SCD (SCD1) has been previously presented in [5, 40]. The ranges are bounded and are estimated using realistic concentration values derived from experimental data and Michaelis–Menten constants (Km) defined as the rate of the reaction at half-maximal velocity [8]. Table 1 describes the seven reactions and rates involved in SCD. The rates are calculated based on the average concentrations for each chemical species and the kinetic coefficients presented in Table 1. The first six reactions are classified as fast and the last reaction is classified as slow because it is orders of magnitude slower than the other reactions. Further discussion on the classification of fast and slow reactions can be found in [5].

Each of the six fast reactions are modelled using the SDE (1). The inactive form of SDH (Z) is not a reactant in any of the chemical equations, so its concentration is not modelled. The equations describe the rates of change of the individual chemical species and are

$$dx_1 = (-k_1 x_1 x_5 + k_2 x_2)dt - \sqrt{k_1 x_1 x_5}dw_1 + \sqrt{k_2 x_2}dw_2$$

$$dx_2 = (k_1 x_1 x_5 - k_2 x_2 - k_3 x_2 x_6 + k_4 x_4 x_7)dt + \sqrt{k_1 x_1 x_5}dw_1$$
$$- \sqrt{k_2 x_2}dw_2 - \sqrt{k_3 x_2 x_6}dw_3 + \sqrt{k_4 x_4 x_7}dw_4$$

$$dx_3 = (k_5 x_4 - k_6 x_3 x_5)dt + \sqrt{k_5 x_4}dw_5 - \sqrt{k_6 x_3 x_5}dw_6$$

$$dx_4 = (k_3 x_2 x_6 - k_4 x_4 x_7 - k_5 x_4 + k_6 x_3 x_5)dt + \sqrt{k_3 x_2 x_6}dw_3$$
$$- \sqrt{k_4 x_4 x_7}dw_4 - \sqrt{k_5 x_4}dw_5 + \sqrt{k_6 x_3 x_5}dw_6$$

$$dx_5 = (-k_1 x_1 x_5 + k_2 x_2 + k_5 x_4 - k_6 x_3 x_5)dt - \sqrt{k_1 x_1 x_5}dw_1$$
$$+ \sqrt{k_2 x_2}dw_2 + \sqrt{k_5 x_4}dw_5 - \sqrt{k_6 x_3 x_5}dw_6$$

$$dx_6 = (-k_3 x_2 x_6 + k_4 x_4 x_7)dt - \sqrt{k_3 x_2 x_6}dw_3 + \sqrt{k_4 x_4 x_7}dw_4$$

$$dx_7 = (k_3 x_2 x_6 - k_4 x_4 x_7)dt + \sqrt{k_3 x_2 x_6}dw_3 - \sqrt{k_4 x_4 x_7}dw_4$$

The single slow reaction SDH $\rightarrow$ Z describes the conversion of the enzyme (SDH) into its inactive form at a rate of $k_7 x_5$. When the reaction occurs, the number of molecules of x_5 is decreased by one and the concentration is decreased by

Table 1 Sugar cataract reactions and kinetic coefficients

| Reaction	Kinetic coefficient	Rate
SDH + NADH $\rightarrow$ E $-$ NADH	$k_1 = 6.2$	31.1
E $-$ NADH $\rightarrow$ SDH + NADH	$k_2 = 33$	151
E $-$ NADH + F $\rightarrow$ E $-$ NAD$^+$ + S	$k_3 = 0.0022$	6
E $-$ NAD$^+$ + S $\rightarrow$ E $-$ NADH + F	$k_4 = 0.0079$	19.5
E $-$ NAD$^+$ $\rightarrow$ SDH + NAD$^+$	$k_5 = 227$	998
SDH + NAD$^+$ $\rightarrow$ E $-$ NAD$^+$	$k_6 = 0.61$	3.2
SDH $\rightarrow$ Z	$k_7 = 0.0019$	0.002

$d_1 = 10^{-21}$ μM. The SHS model can be seen in Fig. 2. The reset on the transition ($x_5- = d_1$) describes the effect of the single slow reaction on the concentration of x_5. For the SCD system, the classifications of the reactions do not change dynamically because the kinetic coefficients are significantly different and the chemical concentrations do not fluctuate widely.

SCD model with medication control (SCD2): Drugs can help patients who are at high risk of developing sugar cataracts. These drugs work by inhibiting the enzyme SDH thereby reducing the rate at which SDH reacts with other molecules in the system. This initially results in less sorbitol production; however, since the reversible reactions are tightly coupled, the results can have side effects such as increasing the fructose levels.

We have created a new SHS model (SCD2), shown in Fig. 3, of drug-modulated SCD to include the effect that drug has on the system. The application of the drug is represented as a new discrete mode that represents
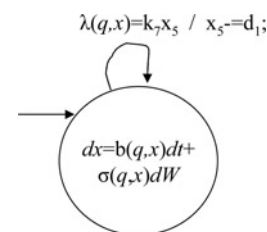
Figure 2 *SHS model of SCD1*
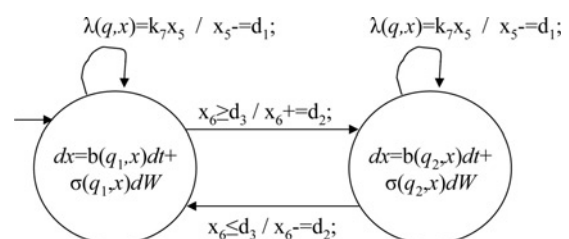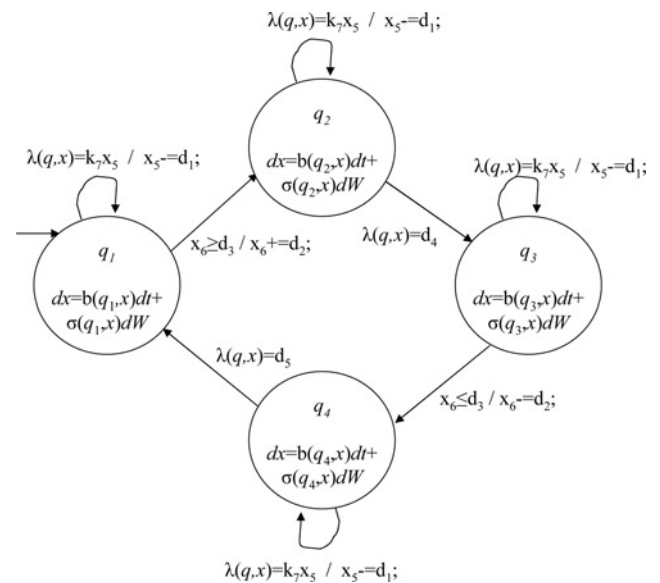

Figure 3 *SHS model of medication-controlled SCD2*

drug-influenced dynamics where the reaction rates k_1, k_6 and k_7 are reduced by 50% to model the inhibition of the enzyme. Since the drug is metabolised slowly and the amount that the rates are reduced is directly proportional to the concentration of the drug, modelling a constant concentration is a reasonable approximation.

We have modelled the drug administration based on an elevated level of fructose. It is assumed that patients self-monitor and self-administer the medication. When the amount of fructose in the blood rises above a threshold $d_3 = 250\,\mu M$, we use a guarded transition to drive the system to a new state, which introduces the effect of the drug. When the fructose level drops back below d_3, we use another guarded transition to transition to the original state effectively removing the effect of the drug. We also include resets on the mode transitions to avoid infinitely fast switching that arises due to the stochastic nature of the Wiener process. The reset increases or decreases the fructose concentration by $d_2 = 1\,\mu M$. Fast switching could also be avoided by using guards that do not overlap (a slightly larger value for increasing guards and a slightly smaller value for decreasing guards). We selected to use state resets in order to reduce the number of the states required by the discrete approximation of the SHS.

This type of guard eliminates the need for the reset on the transition, but we found that this was not appropriate for this model, and it adds complexity to the analysis of the system.

SCD model with probabilistically delayed medication effect (SCD3): The SCD2 model is effective for demonstrating the effect of medication on the reactions; however, realistically the effect of the drug will not be immediate because of variable drug metabolism rates. Drugs are generally administered in a form called a prodrug, which allows the transport of the actual drug to the appropriate cells. This prodrug is metabolised into an active form of the drug at different rates for different people. Furthermore, once a patient discontinues taking a drug, the body can metabolise the residual drug at variable rates depending on many factors.

We have developed a model (SCD3), seen in Fig. 4, which incorporates two new states to model the delay of the conversion from prodrug to drug (q_2) and metabolism after dosage is discontinued (q_4). We use guarded transitions to model exiting the medicated and non-medicated states and entering the respective delay states. We then use probabilistic transitions to model the exit from the delay states to model the stochastic nature of the conversion and metabolism rates. The value $d_4 = 0.05$ is the rate of an exponential distribution that models the delay incurred by the conversion of prodrug to drug, and $d_5 = 0.05$ is the value, which models the exponential distribution corresponding to the drug metabolism delay. These values were chosen so the average delay is on the order of one hour, which is reasonable for the SCD system, but the values could be easily changed to model other types of medications. SHS can also incorporate the

Figure 4 *SHS model of medication-controlled SCD3 with delays*

continuous state into the transition rate if such a model is necessary.

The continuous dynamics of the medicated (q_3) and non-medicated (q_1) states are consistent with SCD2. The dynamics of the delay state q_2 are the same as those in state q_1 to reflect the lack of change while the prodrug is being converted into the drug. The dynamics in the delay state q_4 model the metabolism of the drug after the administration is removed, so the kinetic coefficients are adjusted to reflect their half-life values. The coefficients can be adjusted to model various drugs.

3.5 Simulation results

To better understand and validate our models, we present simulation results using variants of the SSA and EM algorithms. The SSA simulates chemical reactions consuming reactants and creating products one reaction at a time. Individual reactions in a system are assigned probabilities of occurrence, and probability distributions are used to choose, which reaction fires at each iteration. Once a reaction fires, the quantities of reactants and products are updated [1]. The SSA is very accurate because of its level of precision, but it can be inefficient for large systems or fast reactions because many iterations must be completed before results can be observed. To efficiently handle practical systems, computational improvements such as τ-leaping or R-leaping have been devised for the SSA [41, 42]. R-leaping increases the number of reactants consumed and products produced in each step by a factor of R. This increases the efficiency of the approximation, but will degrade the accuracy for certain systems.

For simulation of SCD2 and SCD3, we have created a new algorithm, the HSSA, which implements the SSA using R-leaping and discrete transitions between modes. The

standard *R*-leaping SSA is extended to incorporate the discrete dynamics, which are found in the SCD2 and SCD3 models. After each iteration of the SSA, the guards for all valid transitions are tested, and a transition which validates its guard conditions is fired if possible. Once the transition resets have been executed, the SSA algorithm resumes in the new state.

One iteration of the SSA describes the evolution of the chemical system over a very small unit of time and is considered to be equivalent to solving the CME for that state and timestep. Interrupting the multiple iterations of the SSA required to define a longer time span does not affect the convergence of the SSA, and is used in all fast/ slow chemical simulation techniques for decomposing the problem to make it more efficient. Therefore the HSSA can be considered as equivalent to solving the CME with hybrid switches.

The stochastic continuous dynamics of the SCD models presented in this work can be directly simulated using EM approximations [43]. EM approximations estimate the rate of change of the individual chemical species based on the current chemical concentrations and viable reactions. An appropriate step size must be chosen for the system to ensure an accurate approximation.

To accurately model the discrete transitions of the SCD models, we have developed a variant of the EM approximations called the HEM. In HEM, discrete transitions are incorporated into the EM approximations by analysing the state between steps of updating the continuous dynamics. If a guard condition on a transition is satisfied, then the transition is fired and resets are executed. Once the state is updated, the EM algorithm continues in the new state. Probabilistic transition firing is determined for both HEM and HSSA using the technique described in [39]. We draw a sample from a uniform distribution and test the exponential decay at various times to determine the jump time for each probabilistic transition. When the exponential decay is greater than or equal to the random value, the transition is fired.

Since the SSA algorithm accurately numerically simulates the stochastic time evolution of well-stirred coupled chemical reactions [42], we compare the results of our HSSA and HEM algorithms to demonstrate accuracy of the SHS models. In Fig. 5 we compare the sorbitol and fructose concentrations for the SCD1, SCD2 and SCD3, models, respectively. We chose sorbitol and fructose because they are the two chemicals that are most directly correlated with the development of cataracts. The initial conditions and parameters for all the experiments are shown in Table 2 from [44]. Figs. 5*a*–5*c* display the average concentration at each time step for sorbitol and fructose for 100 runs of the three models. The figures display the comparison between the HSSA and HEM approximation for sorbitol and fructose to demonstrate the

correctness of the SCD models. The 100 HSSA simulations completed in 98 h, and the 100 HEM simulations took 8 min on a 3 GHz desktop computer.

4 Probabilistic verification

In this section, we formulate the safety and reachability problems for the SCD system, we show that both can be characterised as viscosity solutions of a system of coupled HJB equations, and we present a numerical method for solving these equations.

4.1 Problem formulation

Biologists have determined that a ratio of sorbitol to fructose that is greater than one is correlated to the beginning stages of sugar cataract formation [45]. It has been shown that fructose and SDH play a significant role in the accumulation of sorbitol in the eye, which in turn begins the formation of sugar cataracts.

Simulations can determine whether or not a certain starting state will eventually lead to sugar cataract formation; however, it is much more useful to examine all possible starting states, which can be accomplished through verification. Since a ratio of sorbitol to fructose that is greater than one is correlated to the beginning stages of sugar cataract formation, we have identified those states, which meet that criteria as the set of unsafe states. States with a sorbitol to fructose ratio greater than 0.5 but less than 1 correspond to eyes, which are possibly at risk, but not at high risk of cataract formation. States with a ratio of less than 0.5 are at low risk and therefore are desirable, or target states. The unsafe and target sets are depicted in Fig. 6.

Medication can be administered to inhibit the enzyme SDH, which is intended to help keep sugar cataracts from forming in high-risk patients. The safety probability for the medicated model describes the probability that an initial state will transition to an unsafe state given the administration policy for the drug. The reachability probability for a state describes the probability that the patient will transition from the current state to a safer state without first reaching the unsafe states under the given drug administration policy.

4.2 Reachability analysis

Safety is a special case of the reachability problem, so we will formulate the reachability problem first and then the safety problem. The target and unsafe sets for SHS can be described as unions of target and unsafe sets, respectively, for multiple modes. Let $T = \cup_{q \in Q_T} \{q\} \times T^q$ and $U = \cup_{q \in Q_U} \{q\} \times U^q$ be subsets of S representing the set of target and unsafe states, respectively. We assume that T^q and U^q are proper open subsets of X^q for each q, that is,
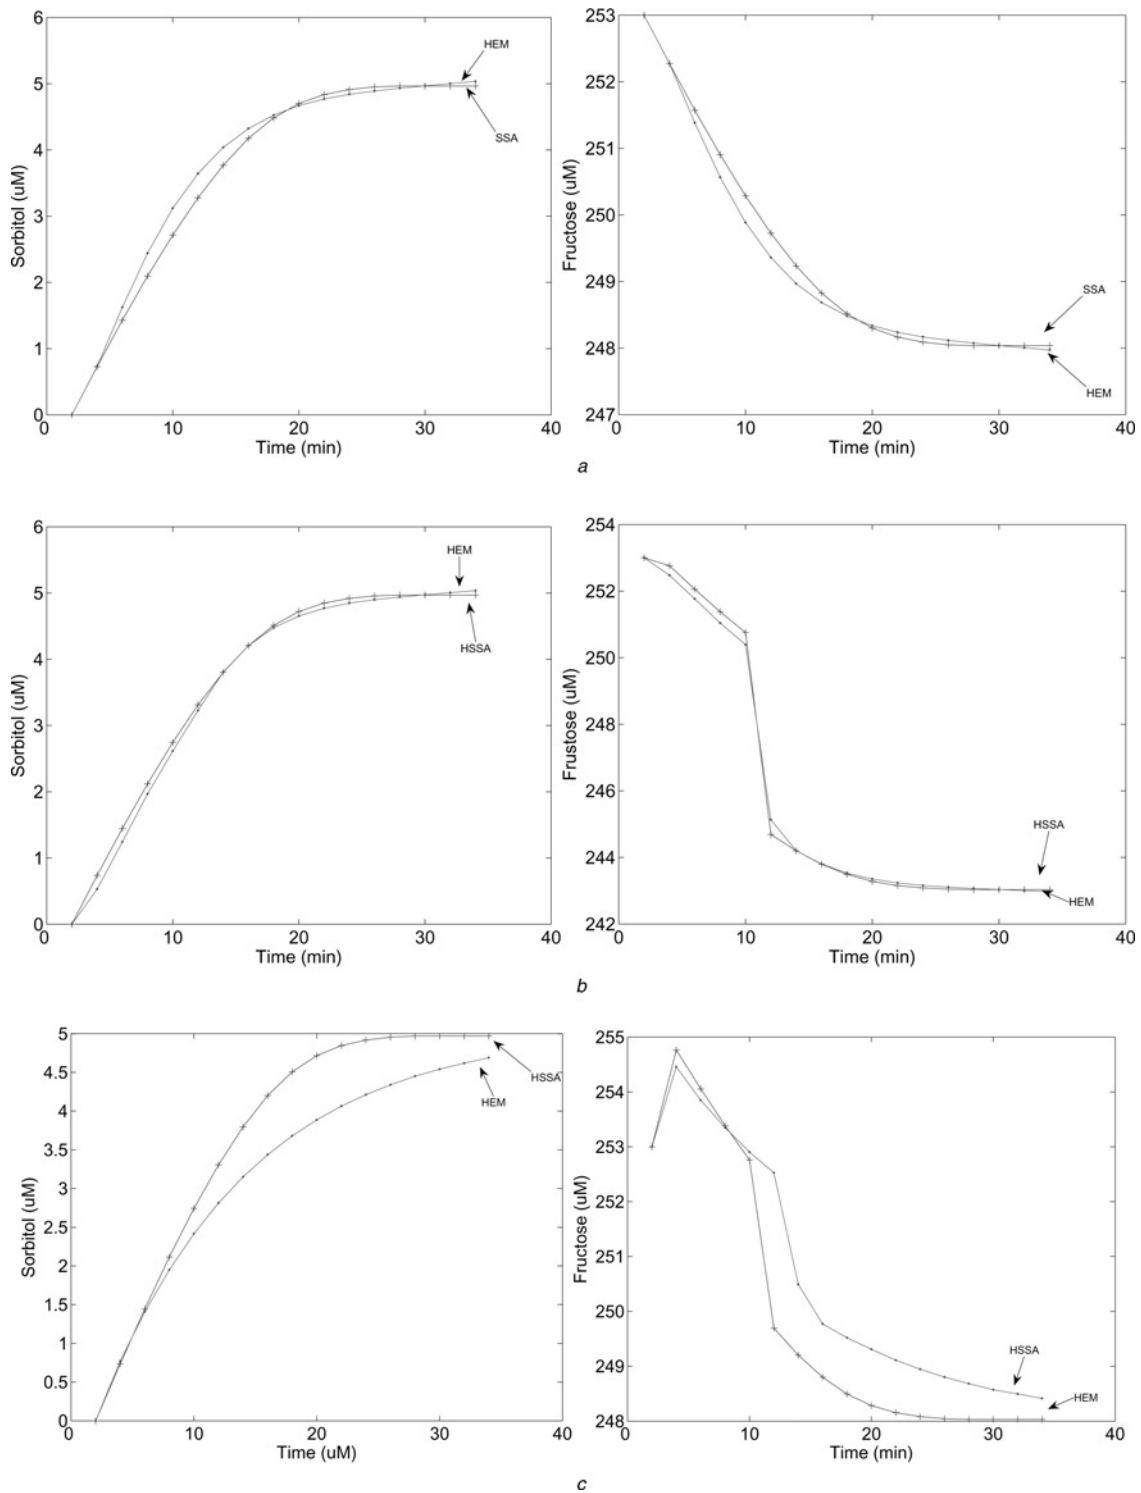
Figure 5 *SCD simulation results*
*a* SCD1
b SCD2
c SCD3

$\partial T^q \cap \partial X^q = \partial U^q \cap \partial X^q = \emptyset$ and the boundaries ∂T^q and ∂U^q are sufficiently smooth. We define $\Gamma^q = X^q \setminus (\bar{T}^q \cup \bar{U}^q)$ and $\Gamma = \cup_{q \in Q}\{q\} \times \Gamma^q$. The initial state (which, in general, can be a probability distribution) must lie outside the sets T and U. The transition measure $R(s, A)$ is

assumed to be defined so that the system cannot jump directly to U or T.

Consider the stopping time $\tau = \inf\{t \geq 0 : s(t) \in \partial T \cup \partial U\}$ corresponding to the first hitting time of the boundary

Table 2 Initial conditions and constants for the SCD models

Initial Cond.	Value, µM	Constant	Value
x_1	5.0	d_1	10^{-21} µM
x_2	0.0	d_2	5 µM
x_3	5.0	d_3	250 µM
x_4	0.0	d_4	0.05
x_5	1.0	d_5	0.05
x_6	253.0	HEM step	0.0001
x_7	0.0	R	0.10
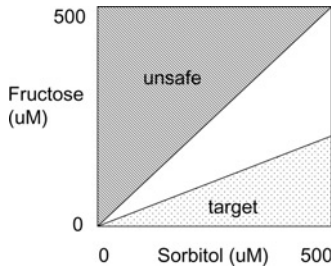
Figure 6 *Graphical depiction of the unsafe and target sets*

of the target or unsafe set. Let s be an initial state in Γ, then we define the function $V : \bar{\Gamma} \to \mathbb{R}_+$ by

$$V(s) = \begin{cases} E_s[I_{(s(\tau-)\in\partial T)}], & s \in \Gamma \\ 1, & s \in \partial T \\ 0, & s \in \partial U \end{cases}$$

where E_s denotes the expectation of functionals given the initial condition s and I denotes the indicator function. The function $V(s)$ can be interpreted as the probability that a trajectory starting at s will reach the set T while avoiding the set U. If the state hits the boundary of the unsafe or target set, then the value function will take the value 0 and 1, respectively, and it is assumed that the execution of the SHS terminates.

Given the assumptions on the sets T and U and their boundaries, we can construct a bounded function $c : \bar{S} \to \mathbb{R}_+$ continuous in x such that

$$c(q, x) = \begin{cases} 1, & \text{if } x \in \partial T^q \\ 0, & \text{if } x \in \partial U^q \cup \partial X^q \end{cases}$$

We define a counting process p^* by

$$p^*(t) = \sum_{i=1}^{\infty} I_{(t \geq t_i)} I_{(s(t_i-)\in\partial S)}$$

The process $p^*(t)$ counts the number of times the trajectory hits the boundary ∂S and jumps up to time t [46]. Then,

the value function V can be written as

$$V(s) = E_s\left[\int_0^{\infty} c(q_{t-}, x_{t-}) dp^*(t)\right] \quad (3)$$

The formulation of the reachability problem described above can be modified to describe safety. In a safety problem, we are given a set of safe states and we want to compute the probability that the system execution from an arbitrary (safe) initial state will go outside the safe set. Let $B = \cup_{q\in Q_B} \{q\} \times B^q$ be a subset of S representing the set of safe states. We assume that the set of unsafe states $X^q \backslash B^q$ for each q is a proper subset of X^q, that is, $\partial X^q \cap \partial B^q = \emptyset$. The initial state must lie inside the safe set B and the transition measure $R(s, A)$ is defined so that the system cannot jump out of the safe set directly to the unsafe set. We can transform the safety problem to a reachability problem by defining the target set as $T^q = X^q \backslash B^q$ and the unsafe set as $U^q = \emptyset$. Note that in this case, the definition of Γ^q becomes $\Gamma^q = X^q \backslash (T^q \cap U^q) = B^q$. Clearly with this transformation, the probability that the system is unsafe can be computed as the value function described by (3) similarly to the reachability problem.

Using a dynamic programming argument, it can be shown that the value function V for the reachability problem of SHSs is similar to the value function for the exit problem of a standard stochastic diffusion, but the running and terminal costs depend on the value function V itself. A detailed proof of the derivation can be found in [9]. We define $L^V(q, x) = \lambda(q, x) \int_{\Gamma} V(y) R((q, x), dy)$, $\psi^V(q, x) = c(q, x) + \int_{\Gamma} V(y) R((q, x), dy)$ and $\Lambda(t) = \exp\{-\int_0^t \lambda(q_0, x_z) dz\}$. Then for $s \in \Gamma$

$$V(s) = E_s\left[\int_0^{t_1^*} \Lambda(t) L^V(q_{t-}, x_{t-}) dt + \Lambda(t_1^*)\psi^V(q_{t_1^*}, x_{t_1^*})\right]$$
$$(4)$$

Equation (4) is similar to the discounted cost criterion with a target set of a standard stochastic diffusion [29]. The main difference is that the running cost $L^V(q, x)$ and the terminal cost $\psi^V(q, x)$ depend on the value function. It should be noted that the SHS satisfies the strong Markov property, and the same procedure can be repeated every time a jump occurs. Further, it can be shown under the non-degeneracy assumption that V is bounded and continuous [9, 44]. Then, based on the results of [29] V can be characterised as the viscosity solution of a system of HJB equations. In particular, V is the unique viscosity solution of the system of equations

$$\mathcal{H}_V((q, x), V, D_x V, D_x^2 V) = 0 \text{ in } \Gamma^q, \quad q \in Q \quad (5)$$

with boundary conditions

$$V(q, x) = \psi^V(q, x) \text{ on } \partial\Gamma^q, \quad q \in Q \quad (6)$$

146

IET Syst. Biol., 2009, Vol. 3, Iss. 3, pp. 137–154

where

$$\mathcal{H}_V\big((q, x), V, D_x V, D_x^2 V\big) = \boldsymbol{b}(q, x) D_x V$$
$$+ \frac{1}{2} \mathrm{tr}(a(q, x) D_x^2 V) + \lambda(q, x) V + L^V(q, x)$$

Equation (5) describes a set of coupled second-order partial differential equations (one for each discrete state), with boundary conditions given by (6), which can be viewed as a set of HJB equations associated with the reachability problem for the SHS. The coupling between the equations arises because the value function in a particular mode depends on the value function in the adjacent modes and is formally captured by the dependency of the running and terminal costs $L^V(q, x)$ and $\psi^V(q, x)$ on the value function V.

4.3 Numerical methods for reachability analysis

One of the advantages of characterising reachability as a viscosity solution is that for computational purposes we can use well-known numerical algorithms. In this paper, we employ the finite difference method presented in [29] to compute locally consistent MCs that approximate the original stochastic process while preserving local mean and variance. We consider a discretisation of the state space denoted by $\bar{S}^h = \cup_{q \in Q}\{q\} \times \bar{S}_q^h$ where $\bar{S}_q^h$ is a set of discrete points approximating X^q and $h > 0$ is an approximation parameter characterising the distance between neighbouring points. By abuse of notation, we denote the sets of boundary and interior points of $\bar{S}_q^h$ by ∂S_q^h and S_q^h, respectively. By the boundness assumption, the approximating MC will have finitely many states, which are denoted by $s_n^h = (q_n^h, \xi_n^h)$, $n = 1, 2, \ldots, N$.

First, we consider the continuous evolution of the SHS between jumps and assume that the state is (q, x). The local mean and variance given by the SDE (1) on the interval $[0, \delta]$ are

$$E[x(\delta) - x] = \boldsymbol{b}(q, x)\delta + o(\delta)$$
$$E[(x(\delta) - x)(x(\delta) - x)^{\mathrm{T}}] = a(q, x)\delta + o(\delta)$$

Let $\{q_n^h = q, \xi_n^h\}$ describe the MC on $S_q^h \subset X^q$ with transition probabilities denoted by $p_D^h((q, x), (q', x'))$. A locally consistent MC must satisfy

$$E[\Delta \xi_n^h] = \boldsymbol{b}(q, x)\Delta t^h(q, x) + o(\Delta t^h(q, x))$$

and

$$E[(\Delta \xi_n^h - E[\Delta \xi_n^h])(\Delta \xi_n^h - E[\Delta \xi_n^h])^{\mathrm{T}}]$$
$$= a(q, x)\Delta t^h(q, x) + o(\Delta t^h(q, x))$$

where $\Delta \xi_n^h = \xi_{n+1}^h - \xi_n^h$, $\xi_n^h = x$ and $\Delta t^h(q, x)$ are appropriate interpolation intervals (or the 'holding times') for the MC.

The diffusion transition probabilities $p_D^h((q, x), (q', x'))$ and the interpolation intervals can be computed systematically from the parameters of the SDE (details can be found in [29]). For a uniform grid with e_i denoting the unit vector in the ith direction, the transition probabilities are

$$p_D^h((q, x), (q, x \pm he_i)) = \frac{a_{ii}(q, x)/2 + hb_i^{\pm}(q, x)}{Q(q, x)}$$

$$p_D^h((q, x), (q, x + he_i + he_j)) = p_D^h((q, x), (q, x - he_i - he_j))$$
$$= \frac{a_{ij}^+(q, x)}{2Q(q, x)}$$

$$p_D^h((q, x), (q, x - he_i + he_j)) = p_D^h((q, x), (q, x + he_i - he_j))$$
$$= \frac{a_{ij}^-(q, x)}{2Q(q, x)}$$

and the interpolation intervals are $\Delta t(q, x) = h^2/Q(q, x)$ where $Q(q, x) = \sum_i a_{ii}(q, x) - \sum_{i, j: i \neq j} |(a_{ij}(x)|/2) + \sum_i h|b_i(q, x)|$, and $a^+ = \max\{a, 0\}$ and $a^- = \max\{-a, 0\}$ denote the positive and negative parts of a real number.

Next, we consider the jumps with transition rate $\lambda(q, x)$ and transition measure $R((q, x), A)$. Suppose that at time t the state is $\{q_n^h = q, \xi_n^h = x\}$. The probability that a jump will occur on $[t, t + \delta)$ conditioned on the past data can be approximated by

$$P[(q, x) \text{ jumps on } [t, t + \delta)|q(s), x(s), w(s), s \leq t]$$
$$= \lambda(q, x)\delta + o(\delta)$$

The ith jump of the approximating process is denoted by $\zeta((q, x), \rho_i)$ where ρ_i are independent random variables with distribution $\bar{R} = \{\rho: \zeta((q, x), \rho_i) \in A\} = R((q, x), A)$ with compact support Π. Let ζ^h be a bounded measurable function such that $|\zeta^h((q, x), \rho) - \zeta((q, x), \rho)| \to 0$ as $h \to 0$ uniformly in x for each ρ and which satisfies $\zeta^h((q, x), \rho) \in \bar{S}^h$.

If $x \in S_q^h$, then with probability $p_{\text{jump}}^h(q, x) = \lambda(q, x) \Delta t^h(q, x) + o(\Delta t^h(q, x))$ there is a jump and the next state is $(q_{n+1}^h, \xi_{n+1}^h) = \zeta^h((q, x), \rho_i)$ and with probability $1 - p_{\text{jump}}^h(q, x)$ the next state is determined by the diffusion probabilities p_D^h, thus the transition probabilities are given by

$$p^h((q, x), (q', x')) = (1 - p_{\text{jump}}^h(q, x))p_D^h((q, x), (q', x'))$$
$$+ p_{\text{jump}}^h(q, x)\bar{R}\{\rho: \zeta^h((q, x), \rho) = (q', x' - x)\} \quad (7)$$

For the points $x \in \partial S_q^h$ in the boundary, the next state is determined by $\zeta^h((q, x), \rho_i)$ with probability 1 and the transition probabilities are given by

$$p^h((q, x), (q', x')) = \bar{R}\{\rho: \zeta^h((q, x), \rho) = (q', x' - x)\} \quad (8)$$

Let $\bar{T}^h = \bar{S}^h \cap \bar{T}$ and $\bar{U}^h = \bar{S}^h \cap \bar{U}$ denote the discretised target and unsafe sets respectively. We denote by n_i the times of the jumps between modes and ν_h the stopping time

© The Institution of Engineering and Technology 2009

representing that $(q_n^h, \xi_n^h) \in \bar{T}^h \cup \bar{U}^h$, then the value function V can be approximated by

$$V^h(s) = E_s \left[\sum_{n=0}^{v_b} c(q_n^h, \xi_n^h) I_{(n=n_i)} \right]$$

The function V^h can be computed using a value iteration algorithm using value iteration assuming appropriate initial conditions, and it converges to the value function V of the SHS as $h \to 0$. The proof of the convergence can be found in [9].

Analysis of the computational complexity of value iteration algorithms is usually based on the contraction property of the iteration operator. The iteration operator used for verification of SHS corresponds to an undiscounted criterion and showing that it is a contraction mapping is more involved. We have proved that the iteration operator restricted to an appropriate set is a contraction mapping with respect to some weighted infinity norm and the polynomial-time complexity of the algorithm [9]. Reachability analysis of SHSs is polynomial on the number of states of the approximating Markov process, however, this number grows exponentially with the dimension of the continuous state space. Therefore application of the approach is limited to low-dimensional systems. Although scalability is a limiting factor, using parallel methods the approach is feasible for realistic systems as shown for the SCD process in the next section.

5   Experimental results

In this section, we present the implementation details and the results of the verification of the sugar cataract models. We also describe performance characteristics of the system at various resolutions of the state space.

5.1   Parallel implementation

The SCD1, SCD2 and SCD3 models are implemented using the constants presented in Table 2. The resolutions are presented in Table 3. In order to apply the approach described in this paper, we under-approximate each discrete

Table 3 Chemical species resolution scaling for experiments (μM)

Reactant	Resolution scaling c_i micromolar
NADH	1.0
$E - $ NADH	1.0
NAD$^+$	1.0
$E - $ NAD$^+$	1.0
SDH (E)	0.1
fructose (F)	25.0
sorbitol (S)	25.0

region X^q by $\bar{X}^q$ by considering a smooth boundary $\partial \bar{X}^q$. The discrete approximations must be created using the finest resolution possible to ensure an accurate result, but increasing the fineness of the resolution causes a significant increase in the number of states, so a balance of accuracy and efficiency must be found. Using scaling parameters similar to the resolution of measurement equipment resulted in reasonable results for the SCD system. Table 3 displays the resolution scaling parameters for the SCD system, and these resolutions result in an MC with $\sim$550 million states.

Storing the values for the value iteration algorithm requires several gigabytes of memory, so we have developed a parallel value iteration implementation to improve the scalability of the algorithm. Parallel algorithms traditionally cannot take full advantage of the increased computing capabilities because slow intra-computer communication delays computation. Some algorithms require more communication than others, and increased communication further decreases efficiency of the algorithm. Therefore algorithms that minimise communication maximise parallel algorithm performance.

Dynamic programming algorithms are very natural to parallelise because of the repetitive nature of the algorithms and the minimal communication required; however, care must be taken to ensure that the algorithm will converge to the correct solution in a parallel implementation. The value iteration algorithm is guaranteed to converge in a parallel implementation as long as communication between the partitions happens periodically [47]. The hybrid switches in our system affect the structure of the state space, but do not change the overall convergence results. The way the state space is partitioned directly affects the efficiency of the parallel method, so the partitioning must be chosen carefully to minimise the communication required.

To partition the problem for multiple processors, we divide the state space into 32 equally sized partitions and assign each partition to a processor. Each partition is then executed independently and the values at the boundaries that are shared with other partitions are periodically updated to guarantee convergence to the solution. The processors routinely calculate the collective amount of change to determine when value iteration can be completed. We use a parallel communication formalism message passing interface (MPI) to execute the communication between processors to ensure efficient communication.

To divide the state space in an effort to minimise communication required, we choose five dimensions and split each into two parts. The first split defines two partitions and one communication boundary, the second defines a total of four partitions and four communication boundaries, the third defines a total of 8 partitions and 12 communication planes, the fourth defines a total of 16 partitions and 32 communication hyperplanes and the fifth defines all 32 partitions and 80 communication hyperplanes. The 32 processors are assigned to the 32 partitions to

minimise the required communication to minimise the overhead required by the technique. More processors could be used to further enhance the scalability of this technique, but we found that 32 was an adequate number for this size of a problem.

The Advanced Computing Center for Research and Education (ACCRE) at Vanderbilt University provides the parallel computing resources for our experiments (www.accre.vanderbilt.edu). The computers form a cluster of 348 JS20 IBM PowerPC nodes running at 2.2 GHz with 1.4 GB of RAM per machine. We use C++ as the implementation language because ACCRE supports MPI compilers for C++. We use the MPI standard for communication between processors because it provides an efficient protocol for message passing middleware for distributed memory parallel computers. Currently, the bottlenecks of this approach are the memory size and speed, but as computing hardware improves, the size of systems that our method can efficiently handle will increase as well.

To visualise our results, we plot projections of the data for different concentrations of the chemicals involved. Specifically, these projections show the safety probability for the entire range of sorbitol and fructose levels for certain values of the five other variables. Multiple selections of the five other variables can be chosen to show a more comprehensive view of the data.

5.2 Safety and reachability analysis

Figs. 7 and 8 show projections of the value function for the safety and reachability results where $x_1 = 1.0$, $x_2 = 1.0$, $x_3 = 1.0$, $x_4 = 1.0$ and $x_5 = 0.1$. These figures show the safety or reachability analysis of the non-medicated SCD1 model (Fig. 7a), medicated SCD2 model (Fig. 7b) and medicated with delay SCD3 (Fig. 7c) model.

The differences between the safety verification results are shown in Fig. 9 to highlight the differences between the analysis of the three models. Fig. 9a shows the difference between the SCD1 and SCD2 models. The difference between the value functions for these models is negligible for fructose values under 250 uM corresponding to the fact that the drug is not administered below 250 uM. Fig. 9b displays the difference that including the drug absorption and metabolisation creates. The difference between SCD2 and SCD3 is especially large for situations where the concentrations of fructose and sorbitol are low.
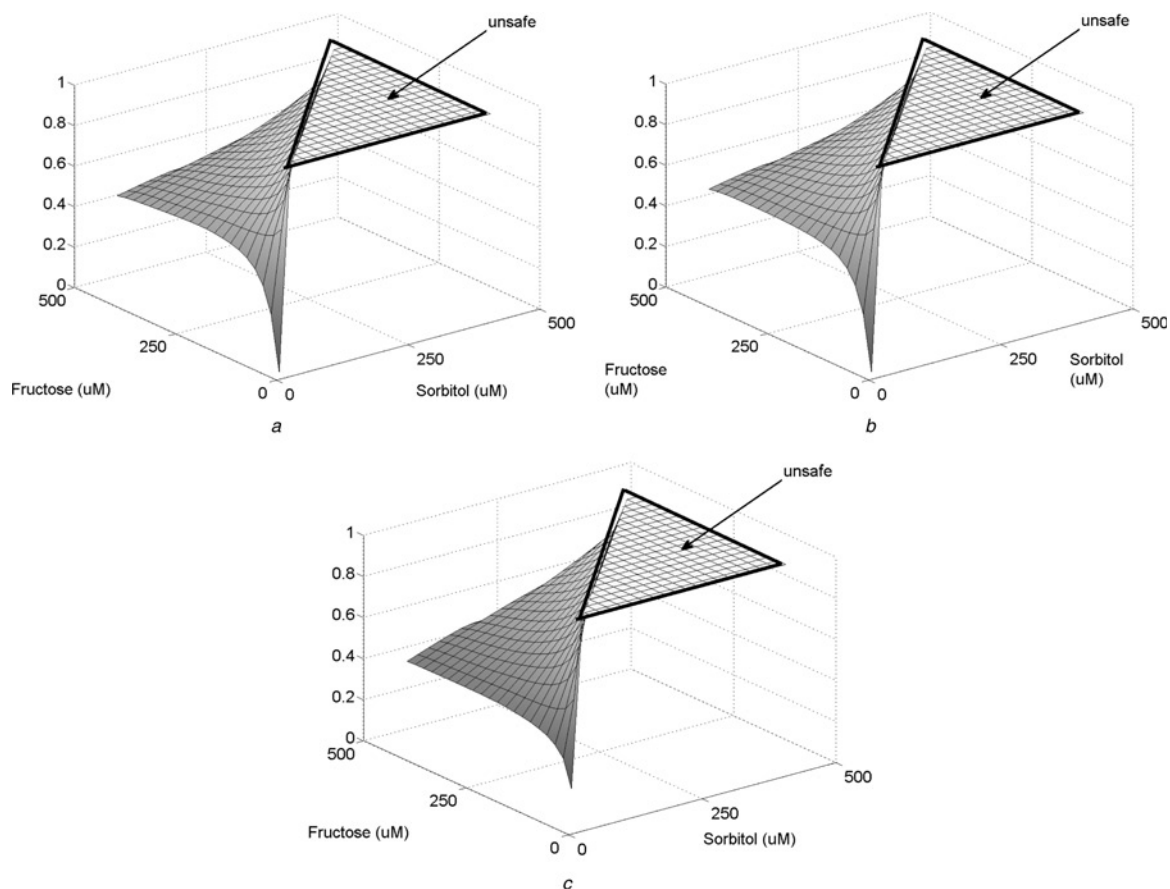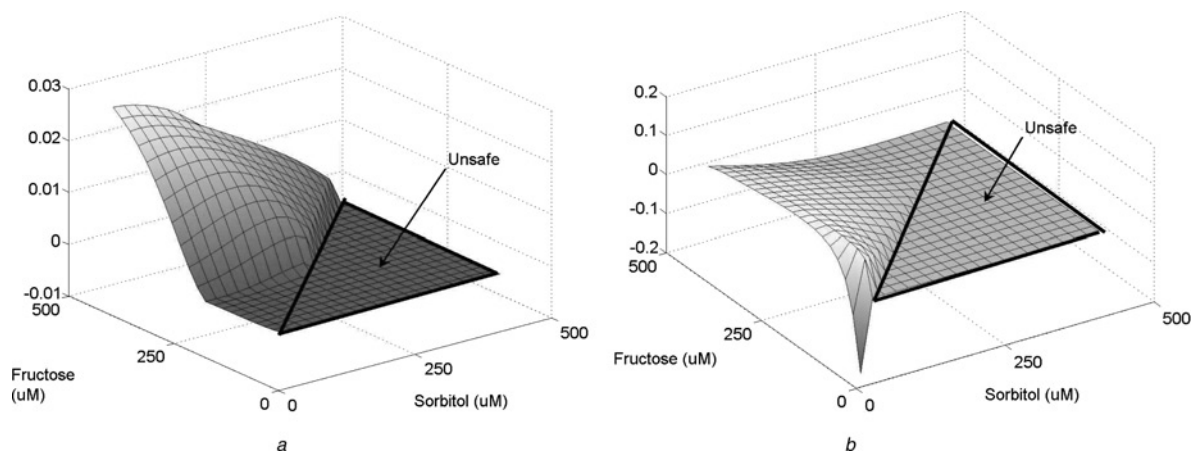

Figure 7 *Safety results for SCD1, SCD2 and SCD3*
a SCD1
b SCD2
c SCD3
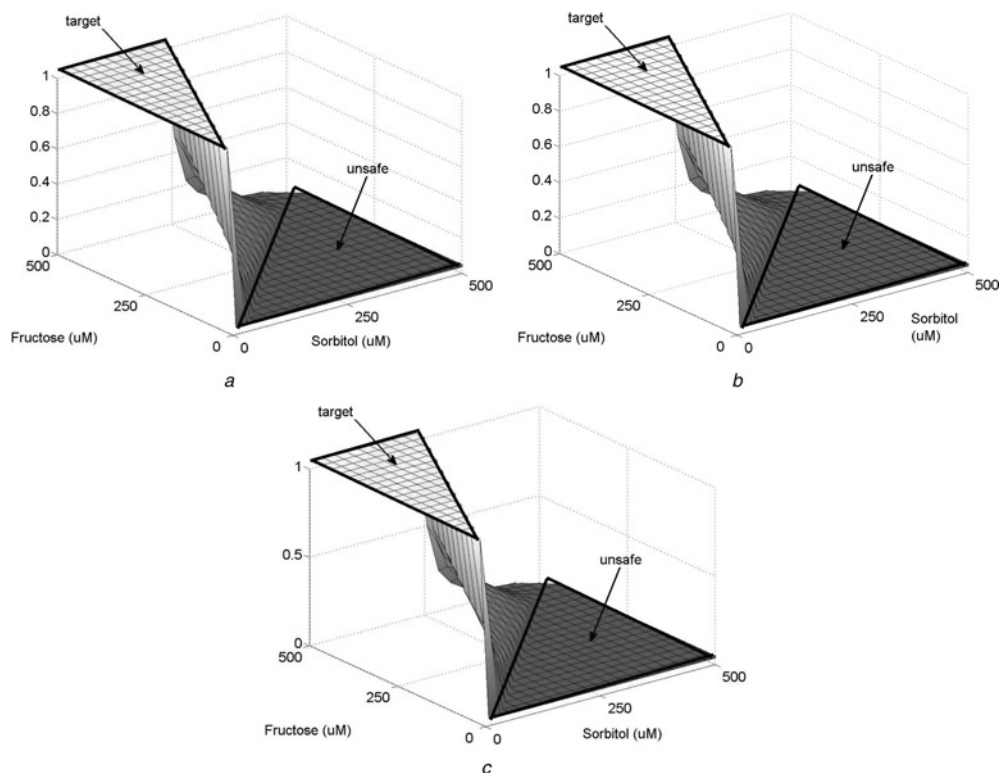
Figure 8 *Reachability results for SCD1, SCD2 and SCD3*
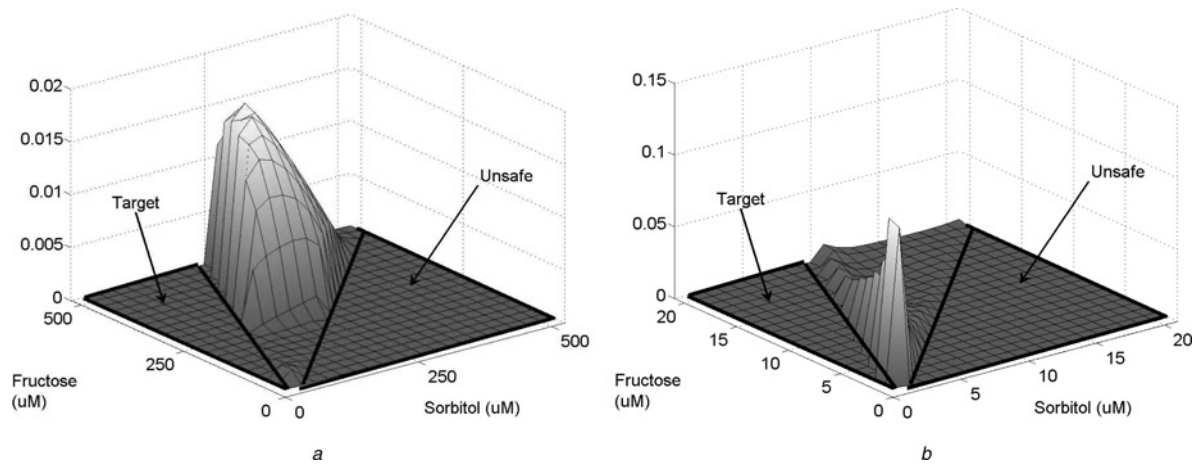
a SCD1
b SCD2
c SCD3

Figure 9 *Differences between the safety results for the SCD models*

a SCD1 and SCD2
b SCD2 and SCD3

Fig. 10 displays the differences between the calculated reachability values for SCD1, SCD2 (Fig. 10*a*) and SCD2, SCD3 (Fig. 10*b*). Figs. 10*a* and 10*b* show that the differences between the models are greater closer to the target set than the unsafe set. This implies that the effects of the drug are most influential to patients who are less likely to develop cataracts. This is most likely caused by the side effects of the drug that can adversely affect the patients fructose levels.

Analysing the data generated by these experiments could possibly help predict sugar cataracts by demonstrating where the safest and most unsafe concentrations exist. Examining the differences between the various medication models could also possibly help further the understanding of how drugs are converted from prodrugs and metabolised. It could also give guidance for choosing the most effective or economical treatment to avoid cataract development. Furthermore, analysis

Figure 10 *Differences between the reachability results for the SCD models*
a SCD1 and SCD2
b SCD2 and SCD3

can possibly guide doctors to better understand and predict why some drugs are more effective than others in different situations and better predict side effects for individual patients.

If the original dynamics of the model are changed by the user, the MDP must be regenerated with the new dynamics and the value iteration must be rerun. The computed value function from the original model can be used as the initial value function for the altered model to significantly improve the efficiency of the value iteration convergence. The amount of the improvement will depend on the dynamics of the system and the changes that are made. The difference between the original model and the new model can be easily quantified using statistical methods by comparing the original value function and the newly computed result.

5.3 Scalability analysis

In this section, we present the effect of the resolution choice on the results and performance of the algorithm. The base resolution for each variable is presented in Table 3 and is the finest resolution that we consider.

Our results show that the efficiency of the parallel implementation increases as the number of states increases demonstrating the polynomial complexity of the algorithm with respect to the number of states. We divided the states as evenly as possible among the processors in a regular configuration to minimise the required communication between processors. As the number of states grows, the ratio of communication time to computation time decreases. This can be seen in Fig. 11*c* where smaller the resolution (h) parameter, the lower the ratio.

The bottleneck of our approach is memory, not processing speed. Each of the 32 processors that we used had 1.4 GB of ram on them and we have been able to analyse an MDP of up to 1.3 billion states. For every GB of memory used, approximately 28 million states can be analysed. Therefore larger systems could easily be analysed by increasing the amount of memory on the machines performing the analysis.

We have been able to analyse systems of up to eight continuous dimensions using 32 processors. To analyse systems with more continuous dimensions, more processors
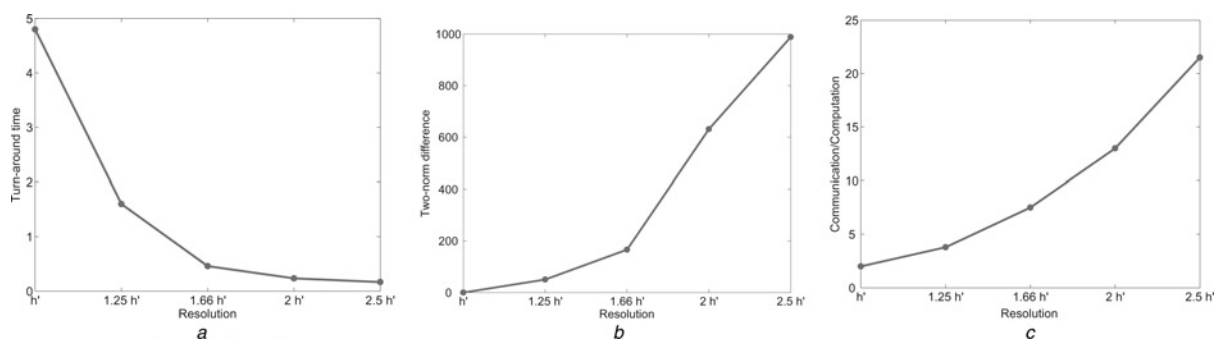
Figure 11 *Performance results*
a Turn around times
b Two-norm
c Communication against computation

must be used. For every unit increase in the number of dimensions, the number of processors used must double. While this is an exponential requirement, trends in high performance computing have been tending towards lower cost, more powerful machines.

We have run tests using different resolutions to demonstrate the scalability of our approach. The resolution parameter h describes the scaling in comparison to the base resolution in Table 3, and it is multiplied by c_i to give the final resolution in each dimension. The turn-around-time of the verification algorithm at various resolutions can be seen in Fig. 11a.

5.4 Error analysis

Characterisation of the error in finite difference methods is useful for indicating the accuracy of the approximation. Local error is introduced by inaccuracies of the individual grid approximation and decreases as the order of the approximation method increases. Our analysis method uses a second-order method with a uniform grid for accuracy and paralleliseability, and the local truncation error is $O(h^2)$. Exact characterisation of the global error is very hard because of the boundary conditions as well as the specific problem constraints of the system. However, practical error analysis can be used to give an indication of the error of an approximation, which is useful for determining an appropriate resolution.

In practical error analysis, error estimates can be obtained by comparing fine grid approximations with coarser grid approximations to give an indication of the error of the approximation. An estimate for the order of accuracy for the approximation can be obtained by comparing the results of three different resolutions. Because a fine grid is used instead of the exact solution, the computed error will decrease faster than the actual error as the compared step sizes decrease; however, the characterisation of the error is still a good indicator of the actual error.

We assume that a solution is a pth order approximation, that is, the error is decreasing as $O(h^p)$. We can obtain an estimate of p using the solutions at three different resolutions. If we assume that we have solutions on grids with resolution parameters h, $h/2$ and $h/4$, then we can compare the coarser solutions to determine an estimate of the order. Let $h' = h/4$, the approximate error for $4h'$ is $\overline{E}(h) = E(h) - E(h') \simeq (4^p - 1)Ch'^p$. Similarly for $h/2 = 2h'$, $\overline{E}(h/2) \simeq (2^p - 1) Ch'^p$. The ratio of error is $\overline{R}(h) = (\overline{E}(h)/\overline{E}(h/2)) \simeq 2^p + 1$, therefore an estimate for the order p is $p = \log_2 (\overline{R}(h) - 1)$.

Coarse grid approximations do not always generate overlapping grids, so interpolation methods must be used to facilitate error analysis for these systems. We extend the notion of linear interpolation to higher dimensions to compare the error in the value functions at different resolutions. Fig. 12 shows a two-dimensional example of bilinear interpolation. The coarse value function in two
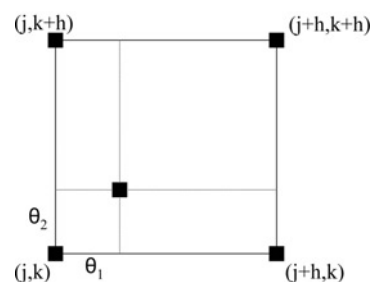

Figure 12 *Two-dimensional value function interpolation*

dimensions at $x_1 = j$, $x_2 = k$ is $V^{j,k}$. The interpolated value is given by

$$V_i^{j+\theta_1,k+\theta_2} = (1 - \theta_1)(1 - \theta_2)V^{j,k} + (\theta_1)(1 - \theta_2)V^{j+h,k} + (1 - \theta_1)(\theta_2)V^{j,k+h} + (\theta_1)(\theta_2)V^{j+h,k+h}$$

This method of interpolation can be extended to any dimension by multiplying θ_l values in each dimension l to the corresponding value function $V^{j,\dots,l,\dots,n}$ and adding all possible combinations thereof.

We calculate the error for our system using the two-norm difference (D_V) between the interpolated value function V_{int} and the finer approximation V_f, $D_V = (1/N)\sqrt{\sum_i (V_f - V_{int})^2}$ where N is the number of states. The results of this analysis are shown in Fig. 11b. Using the ratios of the three finest resolutions, we have computed $p = 2.4$ as an estimate for our order of accuracy. This is a reasonable estimate for our system, but is an overestimate of the order because the finest grid is not an exact solution.

It can be seen that the two-norm difference decreases significantly as smaller step sizes are used, but the reduction in error becomes less and less as the stepsize shrinks. Since smaller stepsizes significantly increase computational cost, trade-offs in accuracy and efficiency must be carefully considered. As computational resources increase in speed and capacity, larger SHS models will be able to be analysed with these methods with a higher degree of accuracy and speed.

6 Conclusions

Biochemical system modelling and analysis are important but challenging tasks, which hold promise to unlock secrets of complicated biochemical systems. SHS are an ideal modelling paradigm for biochemical systems because they are scalable and they incorporate probabilistic dynamics into hybrid systems to capture the inherent stochastic nature of the biochemical systems. The SCD problem is an excellent example of a system that is modelled effectively using the presented modelling methods. Our simulation algorithms also demonstrate the accuracy of the SHS models.

While the computational complexity of the proposed verification method results in a somewhat expensive analysis,

the results of the analysis are exhaustive, and the approach is virtually automated. Improvements in parallel approaches and computational hardware are steadily increasing the ability of this technique to handle larger and larger systems. While the curse of dimensionality limits the ultimate scalability of this method, it is useful for many realistic systems, and it provides a comprehensive analysis.

7 Acknowledgments

8 References

[1] GILLESPIE D.: 'A general method for numerically simulating the stochastic time evolution of coupled chemical reactions', *J. Comput. Phys.*, 1976, **22**, pp. 403–434

[2] BARBUTI R., MAGGIOLO-SCHETTINI A., MILAZZO P., TRIONA A.: 'An alternative to gillespie's algorithm for simulating chemical reactions', *Comput. Methods Syst. Biol.*, 2005

[3] ALUR R., BELTA C., IVANICIC F., *ET AL*.: 'Hybrid modeling and simulation of biomolecular networks'. Hybrid Systems Computation and Control, 2001, (*LNCS*, **2034**), pp. 19–33

[4] GHOSH R., TOMLIN C.: 'Symbolic reachable set computation of piecewise affine hybrid automata and its application to biological modeling: delta-notch protein signalling', *Syst. Biol.*, 2004, **1**, pp. 170–183

[5] SALIS H., KAZNESSIS Y.: 'Accurate hybrid stochastic simulation of a system of coupled chemical or biochemical reactions', *J. Chem. Phys.*, 2005, **122**, pp. 54–103

[6] HESPANHA J., SINGH A.: 'Stochastic models for chemically reacting systems using polynomial stochastic hybrid systems', *Int. J. Robust Control*, 2005, **15**, pp. 669–689 (Special Issue on Control at Small Scales)

[7] PATTERSON J., PATTERSON M., KINSEY V., REDDY D.: 'Lens assays on diabetic and galactosemic rats receiving diets that modify cataract development', *Invest. Ophthalm. Vis. Sci.*, 1965, **4**, pp. 98–103

[8] MARINI I., BUCCHIONI L., BORELLA P., CORSO A.D., MURA U.: 'Sorbitol dehydrogenase from bovine lens: purification and properties', *Arch. Biochem. Biophys.*, 1997, **340**, pp. 383–391

[9] KOUTSOUKOS X., RILEY D.: 'Computational methods for verification of stochastic hybrid systems', *IEEE Trans. Syst. Man Cybern. A*, 2008, **38**, (2), pp. 385–396

[10] DARMANIN C., IWATA T., CARPER D., EL-KABBANI O.: 'Discovery of potential sorbitol dehydrogenase inhibitors from virtual screening', *Med. Chem.*, 2006, **2**, pp. 239–242

[11] RAMOS M., MELO A., HENRIQUES E., *ET AL*.: 'Modeling enzyme-inhibitor interactions in serine proteases', *Int. J. Quantum. Chem.*, 1999, **74**, (3), pp. 299–314

[12] AFZELIUS L., ZAMORA I., RIDDERSTRM M., ANDERSSON T., KARLN A., MASIMIREMBWA C.: 'Competitive cyp2c9 inhibitors: enzyme inhibition studies, protein homology modeling, and three-dimensional quantitative structure-activity relationship analysis', *Mol. Pharmacol.*, 2001, **59**, (4), pp. 909–919

[13] FISH W., MADIHALLY S.: 'Modeling the inhibitor activity and relative binding affinities in enzyme-inhibitor-protein systems: application to developmental regulation in a pg-pgip system', *Biotechnol. Prog.*, 2004, **20**, (3), pp. 721–727

[14] KWIATKOWSKA M., NORMAN G., PARKER D., TYMCHYSHYN O., HEATH J., GAFFNEY E.: 'Simulation and verification for computational modeling of signaling pathways'. Proc. Winter Simulation Conf., 2006, pp. 1666–1674

[15] PHILLIPS A., CARDELLI L.: 'A correct abstract machine for the stochastic pi-calculus', TCSB VII, 2006, (*LNCS*, **4230**)

[16] PRIAMI C., REGEV A., SHAPIRO E., SILVERMAN W.: 'Application of a stochastic name-passing calculus to representation and simulation of molecular processes', *Inf. Process. Lett.*, 2001, **80**, pp. 25–31

[17] ABATE A., AMIN S., PRANDINI M., LYGEROS J., SASTRY S.: 'Computational approaches to reachability analysis of stochastic hybrid systems'. Hybrid Systems Computation and Control, 2007, (*LNCS*, **4416**), pp. 4–17

[18] HU J., WU W., SASTRY S.: 'Modeling subtilin production in bacillus subtilis using stochastic hybrid systems'. Hybrid Systems Computation and Control, 2004, (*LNCS*, **2993**), pp. 417–431

[19] HASELTINE E., RAWLINGS J.: 'Approximate simulation of coupled fast and slow reactions for stochastic chemical kinetics', *J. Chem. Phys.*, 2002, **117**, pp. 6959–6969

[20] ALFONSI A., CANCES E., TURINICI G., VENTURA B.D., HUISINGA W.: 'Adaptive simulation of hybrid stochastic and deterministic models for biochemical systems'. ESAIM: Proc. 2005, vol. 14, pp. 1–13

[21] BENTELE M., EILS R.: 'General stochastic hybrid method for the simulation of chemical reaction processes in cells', *Comput. Methods Syst. Biol.*, 2005, **3082**, pp. 248–251

[22] KIEHL T., MATTHEYSES R., SIMMONS M.: 'Hybrid simulation of cellular behavior', *Bioinformatics*, 2004, **20**, (3), pp. 316–322

[23] GRIFFITH M., COURTNEY T., PECCOUD J., SANDERS W.: 'Dynamic partitioning for hybrid simulation of the bistable hiv-1 transactivation network', *Bioinf. Syst. Biol.*, 2006, **22**, (22), pp. 2782–2789

[24] HARRIS L., CLANCY P.: 'A partitioned leaping approach for multiscale modeling of chemical reaction dynamics', *J. Chem. Phys.*, **125**, (14)

[25] HOOPS S., SAHLE S., GAUGES R., ET AL.: 'Copasi-a complex pathway simulator', *Bioinf. Syst. Biol.*, 2006, **22**, (24), pp. 3067–3074

[26] ADALSTEINSSON D., MCMILLEN D., ELSTON T.: 'Biochemical network stochastic simulator (bionets): software for stochastic modeling of biochemical networks', *BMC Bioinformatics*, 2004, **5**, (24)

[27] MUNSKY B., KHAMMASH M.: 'The finite state projection algorithm for the solution of the chemical master equation', *J. Chem. Phys.*, 2006, **124**

[28] MUNSKY B., KHAMMASH M.: 'A multiple time-step finite state projection algorithm for the solution to the chemical master equation'. Technical Report CCDC-06-1130, Center for Control, Dynamical Systems and Computation, University of California at Santa Barbara, 2006

[29] KUSHNER H., DUPUIS P.: 'Numerical methods for stochastic control problems in continuous time' (Springer-Verlag, 2001)

[30] PUTERMAN M.: 'Markov decision processes-discrete stochastic dynamic programming' (Wiley, 2005)

[31] KOUTSOUKOS X.: 'Optimal control of stochastic hybrid systems based on locally consistent markov decision processes', *Int. J. Hybrid Syst.*, 2004, **4**, pp. 301–318

[32] HU J., PRANDINI M., SASTRY S.: 'Probabilistic safety analysis in three dimensional aircraft flight'. Proc. 42nd IEEE Conf. Decision and Control, 2003, pp. 5335–5340

[33] BLOM H., LYGEROS J., EVERDIJ M., LOIZOU S., KYRIAKOPOULOS K.: 'Stochastic hybrid systems: theory and safety critical applications', *Lect. Notes Control Inf. Sci.*, 2006, **337**

[34] APAYDIN M., BRUTLAG D., GUESTRIN C., HSU D., LATOMBE J., VARMA C.: 'Stochastic roadmap simulation: an efficient representation and algorithm for analyzing molecular motion', *J. Comput. Biol.*, 2003, **10**, (3–4), pp. 257–281

[35] ELOWITZ M., LEVINE A., SIGGIA E., SWAIN P.: 'Stochastic gene expression in a single cell', *Science*, 2002, **297**, p. 1183

[36] DE JONG H.: 'Modeling and simulation of genetic regulatory systems: a literature review', *J. Comput. Biol.*, 2002, **9**, (1), pp. 67–103

[37] BUJORIANU M., LYGEROS J.: 'Theoretical foundations of general stochastic hybrid systems: modeling and optimal control'. IEEE Conf. Decision and Control, 2004

[38] JAZWINSKI A.: 'Stochastic processes and filtering theory' (Academic Press, 1970)

[39] BERNADSKIY M., SHARYKIN R., ALUR R.: 'Structured modeling of concurrent stochastic hybrid systems'. FORMATS, 2004, (*LNCS*, **3253**), pp. 309–324

[40] RILEY D., KOUTSOUKOS X., RILEY K.: 'Safety analysis of sugar cataract development using stochastic hybrid systems'. Hybrid Systems Computation and Control, 2007, (*LNCS*, **4416**), pp. 758–761

[41] GILLESPIE D.: 'Approximate accelerated stochastic simulation of chemically reacting systems', *J. Chem. Phys.*, 2001, **115**, pp. 1716–1733

[42] AUGER A., CHATELAIN P., KOUMOUTSAKOS P.: 'R-leaping: accelerating the stochastic simulation algorithm by reaction leaps', *J. Chem. Phys.*, 2006, **125**, pp. 84–103

[43] KLOEDEN P., PLATEN E.: 'Numerical solution of stochastic differential equations' (Springer-Verlag, 1992)

[44] RILEY D., KOUTSOUKOS X., RILEY K.: 'Verification of biochemical processes using stochastic hybrid systems'. Intelligent Control, 2007, pp. 100–105

[45] BARBUTI R., CATAUDELLA S., MAGGIOLO-SCHETTINI A., MILAZZO P., TRIONA A.: 'A probabilistic model for molecular systems', *Fundam. Inform.*, 2005, **XX**, pp. 1–15

[46] DAVIS M.: 'Markov models and optimization' (Chapman and Hall, 1993)

[47] BERTSEKAS D., TSITSIKLIS J.: 'Parallel and distributed computation: numerical methods' (Prentice-Hall, 1989)