
Theoretical advances in the modelling and
interrogation of biochemical reaction systems:
alternative formulations of the chemical Langevin
equation and optimal experiment design for model
discrimination

Bence Mélykúti

A thesis submitted in partial fulfilment
of the requirements for the degree of
Doctor of Philosophy at the University of Oxford

Keble College

University of Oxford

Michaelmas Term, 2010

Copyright © Bence Mélykúti, 2010

Abstract

This thesis is concerned with methodologies for the accurate quantitative modelling of molecular biological systems. The first part is devoted to the chemical Langevin equation (CLE), a stochastic differential equation driven by a multidimensional Wiener process. The CLE is an approximation to the standard discrete Markov jump process model of chemical reaction kinetics. It is valid in the regime where molecular populations are abundant enough to assume their concentrations change continuously, but stochastic fluctuations still play a major role. We observe that the CLE is not a single equation, but a family of equations with shared finite-dimensional distributions. On the theoretical side, we prove that as many Wiener processes are sufficient to formulate the CLE as there are independent variables in the equation, which is just the rank of the stoichiometric matrix. On the practical side, we show that in the case where there are m_1 pairs of reversible reactions and m_2 irreversible reactions, there is another, simple formulation of the CLE with only $m_1 + m_2$ Wiener processes, whereas the standard approach uses $2m_1 + m_2$. Considerable computational savings are achieved with this latter formulation. A flaw of the CLE model is identified: trajectories may leave the nonnegative orthant with positive probability.

The second part addresses the challenge when alternative, structurally different ordinary differential equation models of similar complexity fit the available experimental data equally well. We review optimal experiment design methods for choosing the initial state and structural changes on the biological system to maximally discriminate between the outputs of rival models in terms of L_2 -distance. We determine the optimal stimulus (input) profile for externally excitable systems. The numerical implementation relies on

sum of squares decompositions and is demonstrated on two rival models of signal processing in starving *Dictyostelium* amœbæ. Such experiments accelerate the perfection of our understanding of biochemical mechanisms.

Contents

Preface	7
Mathematical modelling in molecular systems biology	7
The structure of the thesis	8
Statement of the author's original results	10
 I Stochastic modelling for biochemical reaction kinetics	 13
 1 Stochastic phenomena in molecular systems biology	 15
 2 Stochastic models for biochemical reaction kinetics	 19
2.1 Intrinsic noise	21
2.2 Continuous time, discrete space Markov jump process formulation of chemical reaction kinetics	24
2.2.1 Preliminaries: notation and basic notions	24
2.2.2 Strong and weak characterisation of the process	26
2.2.3 Moment equations	28
2.3 The stochastic simulation algorithm and its approximative evolutions . .	29
 3 An Itô stochastic differential equation model for chemical reaction ki- netics: the chemical Langevin equation	 35

3.1	Gillespie's derivation of the chemical Langevin equation	37
3.2	Alternative motivation: comparison of moment equations for the discrete Markov process and an Itô diffusion	40
3.3	The CLE as a parametric family of SDEs: alternative, weakly equivalent formulations	42
3.3.1	Gillespie's original formulation	45
3.3.2	The minimal formulation	46
3.3.3	A general, state-independent reduction technique	50
3.3.4	State space reduction	53
3.4	Applications	55
3.4.1	A cyclical reaction system	56
3.4.2	A K^+ channel	58
3.4.3	The Goldbeter–Koshland switch	59
3.5	Simulations	61
3.5.1	A cyclical reaction system	61
3.5.2	A K^+ channel	63
3.5.3	The Goldbeter–Koshland switch	63
3.6	Criticism of the chemical Langevin equation	64
3.6.1	Species in the chemical Langevin equation can become negative	65
3.6.2	Problems to be addressed in any proof of existence and uniqueness	71
3.6.3	A strong existence and uniqueness proof for the chemical Langevin equation on the strictly positive orthant	74
3.6.4	Revisiting Proposition 3.3.1: the proof of the weak equivalence of our alternative formulations	76

4 Summary of the results of Part I 87

A	Appendix. Proofs deferred from Chapter 3	91
A.1	Proof of Proposition 3.2.1	91
A.2	Proof of Proposition 3.3.2	92
A.3	Proof of Proposition 3.3.3	93
A.4	Proof of Proposition 3.3.5	93
A.5	State space reduction for Construction 4	94
II	Optimal experiment design for discrimination between rival biochemical models	97
5	Motivation for the mathematical design of experiments	99
5.1	Previous work on optimal experiment design for biochemical reaction systems	100
5.1.1	Experiment design for system identification	100
5.1.2	Experiment design for model discrimination	101
5.1.3	Closing the cycle: model invalidation	103
5.2	Problem formulation	103
5.2.1	Overview	103
5.2.2	Mathematical formalisation of the problem	105
6	Initial condition design and its extension to the design of optimal structural changes for model discrimination	109
6.1	Initial condition design for model discrimination	109
6.1.1	Linear case	110
6.1.2	Nonlinear case	111
6.2	Design of structural changes for model discrimination	114

7	Input design for model discrimination	117
7.1	Designing an input profile using linearisation	119
7.2	Obtaining an upper bound on the L_2 -gain of the system to assess input performance	121
8	An application: signal sensing in <i>Dictyostelium discoideum</i>	123
8.1	Initial condition design for model discrimination	125
8.2	Design of structural changes for model discrimination	128
8.3	Input design for model discrimination	130
9	Summary of the results of Part II	137
B	Appendix. Methods deferred from Chapters 6 and 7	141
B.1	Sum of squares decompositions	141
B.2	SOS programmes for initial condition design	142
B.3	SOS programmes for optimal structural design	143
B.4	SOS programmes for optimal input design	144
10	Outlook and conclusion	147
10.1	Future directions in the modelling of chemical reaction kinetics with con- tinuous stochastic models	147
10.2	A different optimal experiment design problem for model discrimination based on the Hankel operator	150
10.3	Conclusion	151
	Acknowledgements	153
	List of abbreviations	155

Preface

Mathematical modelling in molecular systems biology

The past decade has brought many discoveries in molecular biological research. This has been facilitated by novel quantitative, often high-throughput experimental methods and research at the interface between the life and physical sciences that interprets these rich data sets. The interaction of these disciplines has immensely improved the efficiency with which we can interact with or modify intracellular biochemical processes. The motivation for such experiments may be to understand these biochemical processes better, and increasingly, to harness them for industrial purposes.

The mathematical modelling of biochemical reaction networks plays a central role in these developments (Murray, 2003; Szallasi, Stelling, and Periwal, 2006). Simple qualitative descriptions are proving increasingly insufficient for understanding the intricate dynamical complexity observed in biological processes. As a result, quantitative mathematical models are now routinely used in order to describe and analyse the complex dynamics generated by protein interactions (Cornish-Bowden, 2004), metabolic pathways (Heinrich and Schuster, 1996; Fell, 1997), the regulation of gene expression (Bower and Bolouri, 2004) and other biochemical processes. The study of the dynamics of

these biochemical systems using a combination of experimental work and mathematical–computational modelling is the remit of *molecular systems biology*.

The structure of the thesis

In this thesis we develop methodological results for two problems in the vast field of systems biology. The first question we examine is how best to model biochemical reaction kinetics by stochastic differential equations (SDEs). Part I of this thesis is devoted to this problem. The second question we investigate is how to set up an experiment if multiple dynamical models have been proposed for the same biochemical system and one wants to discriminate between them with the aim of invalidating incorrect models (Part II).

The introduction to Part I gives justification for our interest in stochastic models for intracellular reaction kinetics. We postulate that the inherently probabilistic nature of molecular collisions and interactions, especially in systems with low molecular populations (such as signal transduction pathways or gene regulatory networks, particularly in small prokaryotic cells), must be reflected in our mathematical models. Chapter 1 shows examples where stochasticity is demonstrably present in such intracellular chemical reaction networks. Chapter 2 discusses the most widely used stochastic models for chemical reaction kinetics. The most fundamental one — which is the starting point for all modelling efforts in this field — is a continuous time, discrete space Markov process. Several other processes have been developed; their development was mainly motivated by demand for faster numerical simulation, at a price of some loss of accuracy.

Chapter 3 contains the main results of Part I. We examine thoroughly a diffusion approximation, the so-called *chemical Langevin equation* (CLE), a multidimensional SDE,

to the previously mentioned discrete stochastic process. We formulate other SDEs which share the same finite-dimensional distributions and thereby achieve computational improvements and a better geometrical understanding of this continuous stochastic process. We observe that the CLE predicts negative molecular concentrations in many cases, even in simple ones. We do not attempt to define a new CLE which is free of this fault, but we discuss in great detail how this issue affects the validity of our preceding derivations. Chapter 4 summarises and interprets the results of Part I. Appendix A closes this part by giving a few proofs that were deferred from the main text.

Part II takes us from the study of the computational modelling ‘infrastructure’ for molecular biology to the model development stage, to the situation when one needs to pick the best from seemingly equally good models of similar complexity. We show how to design optimal experiments that provide the most information to help distinguish between the models. We will also make a transition in our modelling framework of choice from stochastic to deterministic dynamical systems. This problem has a rich structure even in the deterministic framework, and such a study is an informative first step towards any treatment of the problem in the stochastic setting.

Chapter 5 reviews earlier approaches to the design of optimal experiments and identifies our approach through our chosen mathematical definition of optimality. Chapter 6 is the immediate precursor to our results: it summarises how to choose the initial state and what systemic changes to introduce to the system of interest for a maximally discriminating experiment.

The main new theoretical developments of Part II are given in Chapter 7. We find that the optimal external stimulus function takes a sinusoidal profile. Practical considerations, such as how to modify the theoretically optimal input profile to be applicable in practice, are discussed in Chapter 8, which is a case study on the application of our

optimal experiment design techniques to a signal transduction pathway. We find that when a sinusoidal input cannot practically be applied, then a square wave input of the same frequency is still superior to a step input. The conclusion in Chapter 9 reiterates how this work fits into the wider landscape of the model design, analysis and redesign cycle. It also gives a reference to a study where our newly developed experiment design techniques were successfully implemented in practice. Appendix B introduces the concept of sum of squares (SOS) decompositions, a central technical tool for our solutions, and collects in one place its applications in the three discussed experiment design approaches.

Chapter 10 recapitulates the main findings of both parts and signposts directions for possible future research. The thesis is concluded with acknowledgements, a list of the abbreviations used in the text and the bibliography.

Statement of the author’s original results

Part I of this thesis is concerned with the justification, the form and shortcomings of the CLE. Its core is the peer-reviewed article Mélykúti, Burrage, and Zygalakis (2010). The author of this thesis developed the theory, did the calculations for the examples, and wrote the paper with the exception of the literature review. The motivation for the work came from Prof. Kevin Burrage’s observation, that the first two moment equations derived from the chemical master equation, Eqs. (2.2.3) and (2.2.4), should impose constraints on the CLE. It was Dr. Konstantinos C. Zygalakis’ observation that the original CLE meets these constraints (Section 3.3.1) and he conducted the numerical simulations for this paper. Section 3.6 of this thesis contains new insights and arguments that were not part of the article.

In Part II, we discuss three approaches to the optimal design of experiments aimed at model discrimination. This part is based on the peer-reviewed article Mélykúti, August, Papachristodoulou, and El-Samad (2010). The main novelty, the *Input design for model discrimination*, was developed by the author of this thesis. The first approach, the *Initial condition design for model discrimination*, originally appeared in conference proceedings (Papachristodoulou and El-Samad, 2007). It was subsequently generalised by Dr. Elias August, a member of Dr. Papachristodoulou’s group, to include optimised modifications of the biological system (*Design of structural changes for model discrimination*).

Part I

Stochastic modelling for biochemical reaction kinetics

Chapter 1

Stochastic phenomena in molecular systems biology

Activities and the observed behaviour of cells most often reflect ongoing intracellular biochemical processes. Therefore it is impossible to understand the behaviour of microbes or the subcellular basis of physiological processes of complex multicellular organisms without understanding the underlying biochemistry.

Many cell colonies, in a somewhat puzzling manner, exhibit substantial phenotypic variation even when the cells share an identical genetic background and the same environment. Mounting evidence points to the role of stochasticity in biochemical reaction outcomes as the source of such diversity. There is uncertainty in the frequency and timing of biochemical reactions involved in gene expression and its control, and inevitably there are minuscule differences between individual cells in the concentrations of cellular components and in their local environments. The cumulative effect of these differences is sufficient to result in different cell fates.

The first experimental proof that stochastic fluctuations are a prevalent phenomenon

in gene expression was published by Elowitz *et al.* (2002). The authors expressed cyan and yellow alleles of green fluorescent protein (gfp) controlled by identical promoters in *Escherichia coli* cells. The cell colonies were observed by microscopy imaging with a magnification that allowed the measurement of fluorescence in individual cells. The authors compared the relative fluorescence levels of these two proteins in various experiments. They showed that at strong constitutive expression of both proteins, both the relative difference between the two fluorescence intensities and the overall cell-cell variation were low. On the other hand, in wild-type (*lacI*⁺) *E. coli* strains, where the artificial lac-repressible promoters are repressed, the gfp expression fell to 3–6%, and noise rose approximately fivefold. (The measure of noise was defined as standard deviation divided by the mean.) This effect was reversible: with the addition of saturating amounts of isopropyl β -D-thiogalactopyranoside (IPTG), which binds and inactivates the *lac* repressor, both the levels of fluorescent proteins and the noise therein returned to the levels observed in the first experiment. This and further experiments with different levels of added IPTG, and the *lac* repressor expressed from plasmids or by a synthetic oscillatory network (the Repressilator; Elowitz and Leibler, 2000), prove that noise increase in gene expression is directly correlated to the repression of transcription.

Whether a bacteriophage λ infection of an *E. coli* cell results in lysis or lysogeny has long been suspected to be decided by stochasticity in the regulation of viral gene expression. *Lysis* is the outcome when the nucleic acid and the structural proteins of the virus are replicated and assembled into complete viruses within the *E. coli* cell, and the infected cell eventually bursts (*lyses*) and releases these new, infectious viruses. The other possibility is *lysogeny*, when the viral genome is integrated into the bacterial chromosome. At cell division it is also replicated and passed on to the daughter cell. Certain events can trigger the release of the viral genome, initiating its proliferation via

the lytic cycle. Arkin, Ross, and McAdams (1998) developed a stochastic model of this genetic circuit, relying on many studies that had built up a good understanding of the underlying biology. Their simulation results support the view that, broadly, competition between two mutually repressing regulatory proteins, Cro and CI, determine the outcome of the phage λ infection. If CI concentration reaches a high enough level soon after infection, then the cell will most probably continue to lysogeny. However, if the Cro feedback loop is established, then CI production will be repressed and the cell will follow the lytic path. It is ultimately chance influenced by the circumstances that determines the outcome. But this paradigm has been challenged. St-Pierre and Endy (2008) showed in experiments that preexisting variation, in particular the cell volume at infection, has a great impact on the ratio between lytic and lysogenic outcomes. They speculate that variability between host cells may be sufficient to provide variability in the outcome even under a deterministic model. However, there is still room for stochasticity to have a decisive role in determining infection outcome: first, the finding about the importance of cell volume is not inconsistent with the quantitative stochastic kinetics model which can incorporate volume (and thus concentration) information, and second, as of yet, the infection outcome cannot be completely controlled in experiments. Quantifying the contribution of stochasticity of reaction kinetics to the lysis–lysogeny decision circuitry remains an open question.

Di Talia *et al.* (2007) studied the variability in cell cycle length in budding yeast, and how much of it is attributable to stochasticity in gene expression. They found that compared to haploid cells, diploids and tetraploids exhibit ever lower noise in the time from division to budding (G_1 phase). If cell cycle timing is dependent on some protein concentration, then decreasing fluctuation in this concentration will decrease relative fluctuation in cell cycle length. Therefore their experimental finding is consistent with

the hypothesis emerging from stochastic modelling, that protein noise should decrease with the square root of the copy number of genes (*ploidy*). The authors decomposed the G_1 phase into two stages, which are separated by the exit of the transcriptional repressor Whi5 from the nucleus. The length of the first stage is primarily size dependent: for smaller birth cell volume it lasts for longer. But there is further unexplained variability in the first stage, and also in the second stage, which together decreases with increasing ploidy as we have just seen. The authors conclude that their analysis demonstrates that stochasticity in chemical reactions has a role in generating variability in the timing of cellular transitions.

But it is not only gene expression, and the related areas of development and differentiation, where the importance of stochasticity in reaction outcomes has been recognised; ion channel dynamics is also heavily influenced by it. For particulars, see Brennan, Fink, and Rodriguez (2009) and references therein.

The next chapter discusses in detail how to interpret and capture the stochasticity at the molecular level of which we have seen examples. The conceptual tool in this effort will be, no doubt, mathematical modelling, that will give an unambiguous language to define and analyse noise in chemical reaction networks.

Chapter 2

Stochastic models for biochemical reaction kinetics

Mathematical modellers of intracellular reactions justifiably view cells as closed containers in which reactions occur in the cytoplasm in the same way reactions occur in other solvents. This stance is usually extended with three assumptions, whose role is made clear in Gillespie (1992). The first is that the cytoplasm is a *well-stirred* solution, that is, it is spatially homogeneous with random fluctuations; the position of a randomly selected molecule is a uniformly distributed random variable. The second assumption states that the biological system is in *thermal equilibrium*, that is, at a constant absolute temperature; the velocity of a randomly selected molecule follows the Maxwell–Boltzmann distribution with a fixed temperature. The third assumption is that cell volume is constant. Over short time periods this is a reasonable assumption, although over their life cycle from birth until division, cells grow to roughly double their initial size. However, if needed, this change in volume and the corresponding dilution of macromolecules can be easily accounted for in a model. This will be discussed in Section 2.2.1. Under these

assumptions, mathematical models that are widely used in chemistry and chemical engineering directly apply to biochemistry. In the entirety of this work, these assumptions will be made.

One shortcoming of this framework is that by assuming well-stirredness, the spatial structure of a cell is ignored. In reality, there are molecules confined to certain locations (e.g. transmembrane proteins), and in eukaryotic cells, there is a complex internal membrane structure which hinders the free diffusion of molecules (Alberts *et al.*, 2008). There is also evidence that a previously unbound enzyme that has not had the time to diffuse away from a substrate molecule can rebind with a higher probability, and that this effect can give a markedly different dynamic behaviour to that predicted by the reaction rate equations (Takahashi, Tănase-Nicola, and ten Wolde, 2010). Even the cell membrane has considerable lateral structure and is compartmentalised on multiple scales. Lipid rafts and cytoskeletal structures impede the diffusion of membrane proteins; and the heterogeneous collection of these compartments, of the microdomains, is thought to selectively concentrate proteins, for the assembly of signalling complexes (Burrage *et al.*, 2007).

Since it is not a focus of this work, we mention it here that the modelling of spatially structured biochemical processes is a vibrant area of computational biology. *Reaction-diffusion equations* have been the traditional tool for this; these are partial differential equations that describe the spatially varying concentration of reacting molecular species (Murray, 2003). Models that are more similar in spirit to the ones that we will study are stochastic in nature. Some of them partition the reaction volume into subvolumes; each such subvolume is assumed to be well-stirred and molecules can diffuse between neighbouring subvolumes (Hattne, Fange, and Elf, 2005; Drawert, Lawson, Petzold, and Khammash, 2010). A different strategy is to keep track of the location of each individual

molecule as it diffuses in the cytoplasm. In terms of computational cost, probably the most applicable implementation is the so-called *Green’s function reaction dynamics* (van Zon and ten Wolde, 2005; Takahashi *et al.*, 2010). Transforming the standard kinetic rates for use in these spatial models is a delicate task (Erban and Chapman, 2009; Fange *et al.*, 2010). While keeping the importance of spatial structure in mind, we note that an enormous body of research over several decades has shown that for reactions that involve freely diffusing molecules in the cytoplasm, well-stirredness is a reasonable modelling assumption that gives vastly useful models.

2.1 Intrinsic noise

The standard reaction kinetics models used in chemistry are multidimensional ordinary differential equations (ODEs) that describe the time evolution of the concentrations of molecular species as they interact and react with each other. However, living cells are much smaller than the reactors for which traditional chemistry developed its models, hence the population sizes of interacting molecules are significantly lower than was the case in traditional applications. Certain proteins are present in a cell in numbers in the order of magnitude of hundreds or only dozens. Particular genes are usually present in even lower numbers. There has been a report (Brenner and Tomizawa, 1991) of an average level of only 3–11 unbound RNA II regulatory molecules of the ColE1 system present in a single *Escherichia coli* cell (7 nM concentration; cell volume can be estimated as $0.6\text{--}2.7 \times 10^{-15}$ ℓ , hence 0.6–2.8 nM concentration corresponds to one molecule per cell). For such chemical species a model in which their concentrations change continuously is clearly inaccurate. Instead, a discrete model where integer molecular counts are the variables is preferable.

The well-stirredness assumption implicitly entails that the model does not track molecules individually; therefore their collisions and reactions are random events under the model. The model is not meant to describe and predict when and what kind of reaction occurs, only their average intensities. The numbers of firings of separate reaction channels and the population sizes may broadly, on average, follow trajectories of deterministic ODE models, but only with random fluctuations. Importantly, there is the possibility of major divergences from the deterministic trajectory. These considerations motivate the use of stochastic models for biochemical reaction networks. Wolkenhauer *et al.* (2004) discuss the justification for choosing either the continuous deterministic or the discrete stochastic model, and clarify common misconceptions about the ranges of their applicability.

This fluctuation in molecular counts around average values that is inherent to the system due to the above phenomena is called *intrinsic* (or *internal*) *noise*. It is well understood, and it is known how to model it. Also, it is distinct from fluctuations caused by environmental changes and interference with other unmodelled intracellular processes, collectively called *extrinsic* (or *external*) *noise* (Gillespie, 2000; Paulsson and Ehrenberg, 2001; van Kampen, 2007). Extrinsic noise is more elusive than intrinsic, perhaps because its sources are nondescript and diverse, and its modelling is in its infancy.

The review paper by Wilkinson (2009) gives an alternative introduction to these issues. It provides further examples to motivate the need for stochastic modelling of biochemical networks: stochasticity leads to phenotypic diversity even in a uniform genetic and external environment, and this improves the overall survival chance of a bacterial population; stochastic events lie at the heart of cellular ageing. The paper discusses different modelling approaches and the statistical challenges in fitting a model to exper-

imental data. An introductory textbook by the same author is also available (Wilkinson, 2006).

We think that to capture and model extrinsic noise accurately, its different sources need to be identified, and these different types need to be treated with due separation. It is naive to search for a universal form of extrinsic noise for mathematical models. The nomenclature also needs to be enriched to reflect separation between the sources of extrinsic noise.


Some types of extrinsic noise, such as thermal or pH fluctuations, or fluctuations in the availability of unmodelled enzymes, impact on the propensities directly. We speculate that some other types of extrinsic noise should be seen as duals to intrinsic noise, following the duality between molecular population sizes and numbers of reaction firings. The second duality means that the system state can be identified not only by the molecular population sizes, but an equivalent description is to use the initial state and how many reactions occurred in each reaction channel separately — these two uniquely determine the population sizes. The latter formalism is called *degree of advancement* (or *extent of reaction*) coordinates (Goutsias, 2005). The duality we propose between intrinsic and specific types of extrinsic noise is the following. Intrinsic noise is the uncertainty in the mapping from the population sizes to how many reactions of different kinds actually occur in a given time interval. Then these other types of extrinsic noise could be dually *defined* as the uncertainty in the mapping from the number of reactions dictated by the core model to how population sizes actually change. This proposed definition covers changes in molecular populations due to the molecules’ participation in unmodelled reactions (‘cross-talk’).

Throughout this work from this point onwards, our focus will be exclusively on intrinsic noise.

2.2 Continuous time, discrete space Markov jump process formulation of chemical reaction kinetics

2.2.1 Preliminaries: notation and basic notions

Let us suppose that there are n chemical species $\mathcal{S}_1, \dots, \mathcal{S}_n$, reacting through m chemical reaction channels $\mathcal{R}_1, \dots, \mathcal{R}_m$. These chemical reactions are represented by formulæ like



where $\alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n \in \mathbb{N}$ (nonnegative integers), and $k \in]0, \infty[$ is a *reaction rate constant*, whose role will become clear shortly. All state vectors will be written as column vectors. Let $X_i(t)$ ($i \in \{1, \dots, n\}$) denote the number of molecules of species $\mathcal{S}_i$ at time t , and let $X(t) = (X_1(t), \dots, X_n(t))^T$. Equivalently, the state can be given by the vector of concentrations $c(t) = (c_1(t), \dots, c_n(t))^T$, a formalism usually used in the continuous ODE models.

Any set of chemical reactions is uniquely characterised by two sets of quantities. The first is the *stoichiometric matrix* $\nu \in \mathbb{Z}^{n \times m}$, which encodes the combinatorial aspect of reactions: a single firing of the j th reaction channel changes the count of the i th species by ν_{ij} . If reaction (2.2.1) is $\mathcal{R}_j$, then $\nu_{ij} = \beta_i - \alpha_i$. The columns $\nu_{\cdot 1}, \dots, \nu_{\cdot m}$ of ν are called the *stoichiometric* (or *update*) *vectors*. The second set of quantities required are the *propensity* (or *intensity*) *functions* $a(x) = (a_1(x), \dots, a_m(x))^T$ that reflect the probabilities of the reactions to occur: if the chemical system is in state x , then the probability of a single firing of reaction channel $\mathcal{R}_j$ in an infinitesimal time interval of length h is $a_j(x)h + o(h)$. (Here the standard o notation is used for an unspecified one-variable

real-valued function for which $\lim_{h \rightarrow 0} o(h)/h = 0$.) The propensity functions decide the model dynamics, and there are different ways of endowing them with a numerical value. A very commonly followed principle is to demand that the propensity is proportional to the number of distinct subsets of the available molecules that can form the left-hand side (input) of the reaction: for reaction (2.2.1),

$$a_j(x) = k \prod_{\substack{1 \leq i \leq n \\ \alpha_i > 0}} \binom{x_i}{\alpha_i}.$$

A slight modification — basically an adaptation to the continuous setting or a simplification for large molecular population sizes — is the *law of mass action*. This definition of the propensity is

$$a_j(x) = k \prod_{i=1}^n x_i^{\alpha_i}$$

for reaction (2.2.1). The conversion of the reaction rate constants between these two cases or between the concentrations and the counts formalisms is straightforward and can be found in various sources, for example in Wolkenhauer *et al.* (2004) or in Section 6.6 of Wilkinson (2006). These references also give guidance as to how to compute propensities in a solution with changing volume, where the amount of solvent varies but reactant counts are unaffected by this. It is often argued that in any reaction at most two molecules interact at a given time (that is, the reactions are at most bimolecular or second order, so $\sum_{i=1}^n \alpha_i \leq 2$), or can be replaced by successive second-order reactions, because the probability of three or more molecules colliding almost simultaneously is very small (Wilkinson, 2006). Under these assumptions, the propensity functions are polynomials of degree at most 2, and specifically are of the form $k_r, k_s x_i, k_u x_i x_k$ and $k_v x_i (x_i - 1)$ or, in the mass action kinetics case, $k_v x_i^2$ (with $k_r, k_s, k_u, k_v > 0$ constants, $i, k \in \{1, \dots, n\}$).

2.2.2 Strong and weak characterisation of the process

These preparations have foreshadowed how the Markov jump process is constructed. Suppose that at time t_0 the chemical system is in state $X(t_0) = X_0 \in \mathbb{N}^n$. For simplicity, we will always use deterministic initial states; the generalisation to random initial states is straightforward. Each reaction channel has some probability of firing, specifically, a reaction $\mathcal{R}_j$ will occur after an exponentially distributed random waiting time with parameter $a_j(X(t_0))$. More accurately, only the earliest of these m reactions will occur (the one which has the smallest of the realised waiting times), say $\mathcal{R}_\ell$ with waiting time τ . Then the state will be defined as $X(t_0)$ for times in $[t_0, t_0 + \tau[$, and $X(t_0) + \nu_\ell$ at $t_0 + \tau$. This update step is then repeated over and over, first with parameters $a_j(X(t_0) + \nu_\ell)$ for the m exponential waiting times, and then with the propensity function evaluated at the subsequent states.

To this pathwise construction corresponds the following description of the probability distribution of the state. For $P(X, t \mid X_0, t_0)$, the probability of the state being in X at time t , given that it was in X_0 at time t_0 , a well-known argument from the theory of Poisson processes yields

$$\begin{aligned} P(X, t + h \mid X_0, t_0) &= P(X, t \mid X_0, t_0) \left(1 - \sum_{j=1}^m a_j(X)h + o(h) \right) \\ &\quad + \sum_{j=1}^m P(X - \nu_j, t \mid X_0, t_0) (a_j(X - \nu_j)h + o(h)) + o(h) \end{aligned}$$

(Feller, 1957; Karlin, 1966). This equation reflects that the event that the system is in X some infinitesimal time h after time t is the disjoint union of three kinds of events. One is that the system was in X at t and then no reaction occurred in the interval up to $t + h$. The second kind of event is that the system at t was in a state one update vector away from X , in $X - \nu_j$ for some j , and jumped into state X by a single firing

of reaction channel $\mathcal{R}_j$. The third kind of event is one with a small probability, that the state jumped to X via multiple jumps. Simple algebraic rearrangement and passage to the limit $h \rightarrow 0$ gives the *chemical master equation* (CME):

$$\begin{aligned} \frac{\partial}{\partial t} P(X, t \mid X_0, t_0) = & - \sum_{j=1}^m a_j(X) P(X, t \mid X_0, t_0) \\ & + \sum_{j=1}^m a_j(X - \nu_{\cdot j}) P(X - \nu_{\cdot j}, t \mid X_0, t_0). \end{aligned} \quad (2.2.2)$$

The CME is a *forward equation* that describes the distribution of the continuous time, discrete space Markov process. It is a system of ODEs where variables are transition probabilities. There is one equation for each configuration X of the state space. This system of ODEs is linear, but is typically very large or even infinite, in the case when no upper bound can be established for some variable. Even if all variables are bounded above, due to the large number of equations, the numerical solution of the CME is computationally very challenging. Munsky and Khammash (2006) introduced the *finite state projection* (FSP) method for the approximate solution of the CME when direct solution is not possible. The FSP method finds a state space truncation such that the Markov process that is confined to this reduced state space (and stopped in a cemetery state upon leaving the reduced state space) approximates the original process. For a given error tolerance and a fixed final time, the approximation by the FSP method has a probability distribution function at the final time that is pointwise within the required error range from that of the original process.

The above stochastic process is the gold standard in stochastic chemical reaction network modelling. Its study was pioneered in the 1960s (McQuarrie, 1967) and 1970s (Gillespie, 1976, 1977), and it was rigorously justified in the early 1990s by a physical model that uses colliding spheres to represent interacting molecules (Gillespie, 1992). Strictly speaking, the physical model only applies to gas-phase reaction systems, but it

is customary to accept without proof that the resulting mathematical model also applies to liquid-phase systems. This stochastic process can also be described as a stochastic equation driven by m independent unit rate Poisson processes undergoing random time changes (Kurtz, 1978; Ball, Kurtz, Popovic, and Rempala, 2006).

2.2.3 Moment equations

Given the CME (2.2.2), ODEs can be derived that describe the evolution of, for example, the mean and the covariance matrix for fixed propensities. In particular, by multiplying Eq. (2.2.2) by X_i and summing over all its possible values, one can show

$$\frac{d}{dt}E(X(t)) = \sum_{j=1}^m \nu_j E(a_j(X(t))) = \nu E(a(X(t))), \quad (2.2.3)$$

where the last term is the product of an $n \times m$ matrix and an m -dimensional column vector. Similarly, the time evolution of the second moment is given by

$$\begin{aligned} \frac{d}{dt}E(X(t)X(t)^T) &= E(\nu a(X(t))X(t)^T) + E(X(t)a(X(t))^T \nu^T) \\ &\quad + E(B(X(t))), \end{aligned} \quad (2.2.4)$$

where the *diffusion matrix* $B(x)$ is defined by

$$B(x) = \nu \operatorname{diag}(a(x)) \nu^T$$

(van Kampen, 2007; Tomioka *et al.*, 2004). Here $\operatorname{diag}(a(x)) \in \mathbb{R}^{m \times m}$ is the diagonal matrix with the entries of $a(x)$ in its diagonal. The equation for the covariance matrix follows easily.

Under the law of mass action kinetics, both Eqs. (2.2.3) and (2.2.4) are closed only if all reactions are at most first order (that is, all propensity functions a_j have degree at most one). If there is at least one bimolecular reaction, then the corresponding propensity function is a quadratic polynomial which will render the right-hand side of Eq. (2.2.3)

dependent on the appropriate second moment. At the same time, Eq. (2.2.4) will become dependent on third moments through the $E(\nu a(X(t))X(t)^T)$ and $E(X(t)a(X(t))^T\nu^T)$ terms. Similarly, the equation for any higher moment will be dependent on even higher moments. In such a case, we say that the moment equations are not closed.

There are a few proposals on how to relax this problem through approximation to get a finite, closed set of ODEs (see e.g. Singh and Hespanha, 2006; Gómez-Uribe and Verghese, 2007; Ullah and Wolkenhauer, 2009). These methods may give spectacularly good results. In a genetic regulatory system which oscillates in the stochastic modelling setup (for a certain choice of rate parameters), the method by Gómez-Uribe and Verghese (2007) predicts oscillation, whereas the standard ODE model, the reaction rate equation, fails to do so and settles to a fixed level. The ability to use approximate information about the second moments may make such a difference. However, it does not seem possible to know *a priori* about any of these closure methods whether they will be accurate, only *a posteriori*, by comparing their predictions to the full stochastic model. Then it is questionable how useful they are; whether they can save the expensive computations needed to work with the stochastic models. They might have utility in accelerating the exploration of the parameter space: when interesting behaviour is found, those parameter values can be checked by the full stochastic model. But other interesting parameter combinations may be overlooked due to the inaccuracy of the moment closure method.

2.3 The stochastic simulation algorithm and its approximative evolutions

Gillespie (1976, 1977) considered first how to algorithmically generate realisations of the stochastic process given in Section 2.2.2. Gillespie's *direct method*, also known as

the *stochastic simulation algorithm* (SSA), is easy to implement, and has become a very popular computational tool.

Given $X(t)$ at time t , the SSA draws a random waiting time τ to the next reaction from an exponential distribution with parameter $a_0(X(t)) = \sum_{j=1}^m a_j(X(t))$. This can be done efficiently using a uniformly distributed random variable on $[0, 1]$, U , and the inverse cumulative distribution function of the exponential distribution, to get $\tau = -\ln U/a_0(X(t))$. Then the algorithm randomly selects which reaction to occur based on the relative sizes of $a_1(X(t)), \dots, a_m(X(t))$: reaction $\mathcal{R}_j$ is chosen with probability $a_j(X(t))/a_0(X(t))$. If reaction $\mathcal{R}_j$ is selected indeed, then the state vector is updated as

$$X(t + \tau) = X(t) + \nu_{\cdot j},$$

and the algorithm repeats. It is easy to see that this algorithm is equivalent to the strong characterisation in Section 2.2.2, but it only needs two independent uniformly distributed random variables in each step. The algorithmic implementation of that characterisation, which in every step generates an exponential waiting time for each reaction channel and realises the reaction with the shortest waiting time, is called the *first reaction method* (Gillespie, 1976).

The *next reaction method* by Gibson and Bruck (2000) cuts down computation time compared to the SSA through careful algorithm design. For each reaction channel, it stores a tentative waiting time. Similarly to the first reaction method, whichever waiting time is the shortest defines the time to the next reaction and the type of reaction. After this reaction occurred, the algorithm calculates the new propensities. For the one reaction channel that fired, a new random waiting time is drawn. For the other $m - 1$ channels, the waiting times are updated in a deterministic fashion. This strategy reduces the random number draws per reaction from two in the SSA to just one (in addition to

the m required for the initialisation of the next reaction method), while it still maintains the independence of the m waiting times. Further saving in computational time can be achieved by, first, storing the tentative waiting times in a binary tree structure (called an indexed priority queue) that is sorted such that the shortest waiting time is always at the root. Thus the root defines the time and type of the next reaction. The second trick is to calculate only those propensities that have changed.

Anderson (2007) used the stochastic equations formalism introduced by Kurtz (1978, which we have already mentioned at the end of Section 2.2.2) to analyse the next reaction method. This approach provides a clearer explanation of why the waiting times can be updated the way proposed in the next reaction method. It also lends itself more naturally to generalisations to time-dependent propensity functions or to chemical systems with delayed reactions.

These three methods (the direct, the first reaction and the next reaction methods) are all exact, that is, they generate samples of the Markov jump process of Section 2.2.2. Their weakness is the computational cost: since each method simulates every single reaction, if at least some of the reacting species is abundant, that species will generate many reactions. This happens even when those reactions do not contribute much to the dynamics, like the reversible enzyme-substrate complex formation in an enzymatic reaction or the dimerisation and dissociation of a protein. This is not a wise allocation of computational resources, which is the limiting factor when one needs to sample a large batch of simulation runs to explore the distribution of possible system dynamics.

Much effort has been concentrated on developing approximative but faster algorithms. We review some of the most commonly known ones. Gillespie (2001) introduced the τ -leap method, which takes time steps (*leaps*) of length $\tau > 0$. In one such leap multiple reactions are allowed to occur. τ must be chosen short enough so that neither of the

propensities change considerably in time intervals of length τ . Then the number of reaction firings in the time interval for the different reaction channels can be approximated with independent Poisson random variables. Thus the algorithm generates m Poisson random variables in each leap and updates the state accordingly. This assumption will resurface in Section 3.1 as Condition 1. The choice of τ is a challenging problem and has been addressed in subsequent papers: the longer the time intervals are, the faster the algorithm will be, but the less accurate. Also, the Poisson distribution has an infinitely long tail. The larger τ is, the more likely it is that in a given time interval the proposed reaction number is greater than there are reactant molecules. The execution of such a state update would lead to unrealistic negative molecular counts.

Tian and Burrage (2004) proposed to replace Poisson random variables with binomial ones whose maximal possible values are the smallest respective reactant population sizes so that no variable will become negative in any reaction. The case when a species appears as a reactant in two or more reactions is also dealt with. Then the binomial variables are such that their sum is not greater than the population size of the shared reactant species. Independently, Chatterjee, Vlachos, and Katsoulakis (2005) also proposed binomial reaction numbers. Their solution in the case of a species reacting in multiple channels is to generate firing numbers in a predetermined (or random) sequence of all reaction channels at each leap, such that the upper bound on firings takes account of how many molecules were used in reactions that were earlier in this sequence. This approach introduces a bias into sampling, but the authors report the bias to be small in simulations compared to the error introduced by leaping instead of simulating each individual reaction.

The *R-leaping* method by Auger, Chatelain, and Koumoutsakos (2006) (somewhat similarly to the *k_a-leap method* by Gillespie, 2001) generates a fixed number L of reac-

tions in each leap and randomly generates the leap length. Deciding how many of these L reactions will be $\mathcal{R}_1, \dots, \mathcal{R}_{m-1}$ or $\mathcal{R}_m$ is done by sequential draws from binomial distributions, which is an unbiased method in this setting.

The next chapter studies how one can approximate the discrete stochastic process of Section 2.2.2 with a continuous one that is given by an SDE. Any numerical integration scheme applied to such an SDE is essentially a further alternative approximation of the SSA. Although they might have arisen from different mathematical concepts, in practice they can be used for the same applications in computational modelling.

Chapter 3

An Itô stochastic differential equation model for chemical reaction kinetics: the chemical Langevin equation

In this chapter we introduce and study the standard Itô *stochastic differential equation* (SDE) analogue of the reaction kinetics models we discussed in the previous chapter. Due to space limitations, we assume that the reader is familiar with Itô SDEs. We recall the notion of a *Wiener process* or *Brownian motion* only. Detailed elaboration of these topics can be found in standard stochastic analysis textbooks, such as Karatzas and Shreve (1998) or Øksendal (2007).

Definition 3.0.1 A (*standard, one-dimensional*) *Wiener process* (or *Brownian motion*) is a continuous stochastic process $(W(t))_{t \geq 0}$ adapted to some filtration $(\mathcal{F}_t)_{t \geq 0}$ on a probability space $(\Omega, \mathcal{F}, P)$ with the following properties:

- $W(0) = 0$ almost surely,
- the increment $W(t+h) - W(t)$ ($h > 0$) is independent of $\mathcal{F}_t$, and
- this increment is normally distributed with mean 0 and variance h : $W(t+h) - W(t) \sim \mathcal{N}(0, h)$. ♣

The chemical Langevin equation (CLE), introduced by Gillespie (2000), is the standard SDE model for chemical reaction kinetics. The drift (deterministic) term of the CLE is just the right-hand side of the standard ODE model (of the reaction rate equation), while the diffusion (stochastic) term takes a very special form. Gillespie captured this stochastic component by using one Wiener process per chemical reaction. Subsequently, a number of authors have realised that the CLE (and SDEs in general) can be written in alternative, equivalent forms using arguments from stochastic analysis and multivariate statistics (Wilkinson, 2006; Allen *et al.*, 2008; Ullah and Wolkenhauer, 2009). These changes to the equation do not change the finite-dimensional distributions of its solution. As such, the alternative formulations do not arise from model reduction. These authors used this insight solely to reformulate the CLE with the minimum number of Wiener processes.

Our contribution in this chapter is a detailed exploration of this insight. We investigate the minimum number of Wiener processes in the CLE *in conjunction* with a state space reduction by removing linearly dependent variables through conservation laws exhibited by the chemical system. This will shed light on the structure of the CLE from a linear algebraic viewpoint. Secondly, we show how from the standard form CLE one Wiener process can be omitted for each pair of reversible reactions. If there are m_1 pairs of reversible reactions and m_2 irreversible reactions, then only $m_1 + m_2$ Wiener processes are needed, rather than $2m_1 + m_2$. We demonstrate that this simplification, in addition

to its appealing explanatory power, can significantly speed up numerical simulations. Adalsteinsson, McMillen, and Elston (2004) also use this formulation, but they did not point out that this formulation is different to the original one and did not compare the computational costs. Singer *et al.* (2009) derived this new form for a certain application from the Euler–Maruyama discretisation (see, for example, Kloeden and Platen, 1992) of the CLE, which is an approximative method and does not possess the same level of rigour as the one we follow.

We will initially intend to construct the CLE in such a way that its mean and variance match the mean and the variance of the CME at any time t . It will become clear that with our current knowledge this is an insurmountable task in the general case because the moment equations are not closed. Still, we will arrive at the standard, well-known form of the CLE and will use that as the best approximation available. All our alternative formulations will be equivalent to this standard form.

3.1 Gillespie’s derivation of the chemical Langevin equation

Gillespie (2000) set out to approximate the finite-dimensional distributions of the discrete Markov jump process of molecular counts by making two simplifying assumptions. His argument was the following.

Suppose that at time t the chemical system is in state $X(t) \in \mathbb{R}^n$. If the random variable $K_j(X, h)$ denotes the number of times reaction $\mathcal{R}_j$ occurs in a time interval of length h if the system is released from state X , then after time h has passed, the system

will be in state

$$X(t+h) = X(t) + \sum_{j=1}^m K_j(X(t), h) \nu_j. \quad (3.1.1)$$

Now assume that *Condition 1* holds:

1. h is small enough so that the change in the state during $[t, t+h]$ will be so small that none of the propensity functions a_j changes substantially,

$$a_j(X(s)) \approx a_j(X(t)), \quad \text{for all } s \in [t, t+h].$$

Recall that in the most typical cases $a_j(x)$ is a polynomial of degree not greater than 2.

In any reaction typically no molecular count changes by more than two. Hence this condition can be satisfied if the expected number of firings of a reaction channel is much smaller than the population size of the least populous species. This requirement can always be met if all molecular populations are sufficiently large. The assumption that the propensities remain approximately constant in the time interval implies that the random variables $K_1(X(t), h), \dots, K_m(X(t), h)$ are independent, and $K_j(X(t), h)$ is Poisson distributed with parameter $a_j(X(t))h$ for all j .

Then *Condition 2* is stipulated:

2. h is large enough so that the expected number of firings for each reaction channel $\mathcal{R}_j$, namely $E(K_j(X(t), h)) = a_j(X(t))h$, is much larger than 1.

This obviously runs counter to Condition 1. In cases where the two cannot be met simultaneously, Gillespie's approximation will fail. But large molecular populations help to satisfy this condition, just as with the previous one. In this case, the Poisson random variable $K_j(X(t), h)$ is well approximated by a normal random variable with matching mean and variance,

$$\mathcal{N}(a_j(X(t))h, a_j(X(t))h).$$

(Once again, we use the notation $\mathcal{N}(\mu, \sigma^2)$ for a normal distribution with mean μ and variance σ^2 .)

Thus the independent discrete Poisson random variables are replaced by the same number of independent but continuous normal random variables. It is well known that this distribution is a linear transformation of some other normal distribution:

$$\mathcal{N}(a_j(X(t))h, a_j(X(t))h) \sim a_j(X(t))h + \sqrt{a_j(X(t))} \mathcal{N}(0, h).$$

Substituting these approximations, Eq. (3.1.1) takes the form

$$X(t+h) = X(t) + \sum_{j=1}^m a_j(X(t))h \nu_{.j} + \sum_{j=1}^m \sqrt{a_j(X(t))} \nu_{.j} N_j \quad (3.1.2)$$

with independent normal random variables $N_1, \dots, N_m$ with mean zero and variance h . Here we keep t fixed and omit the dependence of N_j on t in our notation. Using the notion of *Wiener process*, clearly, when $h \rightarrow 0$, then Eq. (3.1.2) becomes nothing else but an n -variable Itô SDE

$$dX(t) = \sum_{j=1}^m a_j(X(t)) \nu_{.j} dt + \sum_{j=1}^m \sqrt{a_j(X(t))} \nu_{.j} dW_j(t), \quad (3.1.3)$$

which is called the *chemical Langevin equation*. This is how in Gillespie's derivation two approximative steps facilitated by two assumptions lead to an Itô SDE model.

It is worth noting that the approximation that followed from Condition 1 is the key idea of the τ -leap method (Gillespie, 2001), as we already indicated in Section 2.3.

3.2 Alternative motivation: comparison of moment equations for the discrete Markov process and an Itô diffusion

In the Langevin regime, one hopes to be able to construct an SDE such that the evolution equations for the state mean and the second moments are exactly Eqs. (2.2.3) and (2.2.4) (from Section 2.2.3) in order to get the same mean and covariance matrix as in the CME at any time t . Of course this is not equivalent to identical finite-dimensional distributions, but accurately describing the first two moments using an SDE can be useful for many practical applications.

The approach of matching the first two moments is implicit in Gillespie's derivation. In that case, as opposed to our approach, matching is only enforced locally, for each small time step. Global matching of the first two moments is more explicitly mentioned by, for instance, Wilkinson (2006), but it has not been explored in depth.

Formally, we are seeking $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $g : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times d}$ such that the solution to

$$dx(t) = f(x(t)) dt + g(x(t)) dW(t), \quad (3.2.1)$$

with a d -dimensional standard Wiener process W , has its first two moments given by Eqs. (2.2.3), (2.2.4).

We will now derive the ODEs that describe the evolution of the first two moments of Eq. (3.2.1). By taking the expectation on both sides of Eq. (3.2.1), it is easily seen that

$$\frac{d}{dt} \mathbb{E}(x(t)) = \mathbb{E}(f(x(t))). \quad (3.2.2)$$

To calculate the second moment for $x(t)$ it is enough to calculate $\mathbb{E}(x_i(t)x_k(t))$ for all

$i, k \in \{1, \dots, n\}$, $i \leq k$.

Proposition 3.2.1 For Eq. (3.2.1),

$$\begin{aligned} \frac{d}{dt} \mathbb{E}(x_i(t)x_k(t)) &= \mathbb{E}(f_i(x(t))x_k(t)) + \mathbb{E}(x_i(t)f_k(x(t))) \\ &+ \sum_{j=1}^d \mathbb{E}(g_{ij}(x(t))g_{kj}(x(t))). \end{aligned} \quad (3.2.3)$$

Proof The proof of this proposition is found in the appendix of Part I. ■

Comparing Eq. (2.2.3) with Eq. (3.2.2) suggests we should choose

$$f(x) = \nu a(x). \quad (3.2.4)$$

Once this is fixed, a comparison between Eqs. (2.2.4) and (3.2.3) suggests that

$$\sum_{j=1}^d g_{ij}(x)g_{kj}(x) = B_{ik}(x)$$

for all i and k is the most natural choice for g , that is,

$$g(x)g(x)^T = \nu \operatorname{diag}(a(x))\nu^T. \quad (3.2.5)$$

It is important to point out here that such a choice of f and g can only guarantee that the first two moments for the CME and the CLE are the same when all reactions are at most first order (under the law of mass action). In the case when there is at least one bimolecular reaction (when the moment equations are not closed), the mere formal matching of the two pairs of equations is not sufficient to match the first and second moments because Eqs. (2.2.4) and (3.2.3) are dependent on higher moments which we do not attempt to match and whose moment equations are in actual fact different.

It was shown in the previous section how Gillespie's argument gives the CLE (3.1.3), or this equivalent form:

$$dX(t) = \nu a(X(t)) dt + \sum_{j=1}^m \nu_{\cdot j} \sqrt{a_j(X(t))} dW_j(t). \quad (3.2.6)$$

Clearly, it satisfies both Eqs. (3.2.4) and (3.2.5) with $d = m$.

Corollary 3.2.2 *Under the law of mass action kinetics, if all reactions are at most first order, then the chemical Langevin equation (3.2.6) gives the same first and second moments (means and covariances) as the chemical master equation.*

This result has already appeared in Higham and Khanin (2008, Theorem 4.1), but we have used more succinct notation and therefore this derivation is more transparent. As we have just noted, in the general case this does not imply that the first two moments of the CLE are the same as those of the CME. Instead of further studying the relationship of the CME and the CLE, we accept Eq. (3.2.6) as the standard reference SDE model for chemical reaction systems.

3.3 The CLE as a parametric family of SDEs: alternative, weakly equivalent formulations

In the following, we derive alternative formulations of the CLE based on the following insight.

Proposition 3.3.1 *Different solutions g to the factorisation problem Eq. (3.2.5) all give chemical Langevin equations that have the same finite-dimensional distributions (in different terminology: which coincide in law).*

Beware that in its current formulation, this is not a rigorous statement. Both proofs we give here are incomplete and are intended to reveal the core idea and to cast light on why some form of this statement is true. Section 3.6 will study the difficulties associated with making this argument rigorous, and Proposition 3.6.8 of the concluding Section 3.6.4 will state and prove the coveted claim.

Proof We assume previous knowledge of a standard tool, the Kolmogorov forward equation, which is also known as the Fokker–Planck equation (Gardiner, 1985; Øksendal, 2007). For the solution g of Eq. (3.2.5), the probability density function $p_t(x_0, x)$ of a transition from x_0 to x in a time interval of length t evolves according to the partial differential equation

$$\begin{aligned} \frac{\partial p_t(x_0, x)}{\partial t} = & - \sum_{i=1}^n \frac{\partial \left(p_t(x_0, x) (\nu a(x))_i \right)}{\partial x_i} \\ & + \frac{1}{2} \sum_{i,k=1}^n \frac{\partial^2 \left(p_t(x_0, x) (g(x)g(x)^T)_{ik} \right)}{\partial x_i \partial x_k}. \end{aligned}$$

$p_t(x_0, x)$ evolves identically for all solutions g to (3.2.5), because the parameters in the Kolmogorov forward equation $\nu a(x)$ and $g(x)g(x)^T = \nu \operatorname{diag}(a(x))\nu^T$ are identical for any g , and $p_0(x_0, x) = \delta_{x_0}(x)$ (the Dirac delta function at x_0) does not depend on g . The reason why this proof is incomplete is that it requires that the transition measure of $X(t)$ has a density.

An alternative choice is to apply Theorem 8.4.3 of Øksendal (2007). The argument used there is based on the equivalence between solutions of an SDE (whose coefficients only depend on the state, not on time) and solutions of the *martingale problem*. The approach to study the weak existence and uniqueness of solutions to SDEs via the martingale problem was developed by Stroock and Varadhan (see e.g. Stroock and Varadhan, 1997), but it can be found in several other textbooks, such as Chung and Williams (1990); Karatzas and Shreve (1998). The martingale problem that corresponds to

$$dX(t) = \nu a(X(t)) dt + g(X(t)) dW(t) \tag{3.3.1}$$

is to find a measure P^{x_0} on the set of continuous functions equipped with the Borel σ -algebra $(C([0, \infty[, \mathbb{R}^n), \mathcal{B})$, such that for all twice continuously differentiable functions $h : \mathbb{R}^n \rightarrow \mathbb{R}$ with compact support, $h \in C_0^2(\mathbb{R}^n)$, $M(t) = h(\omega(t)) - \int_0^t Lh(\omega(s)) ds$ is

a martingale under P^{x_0} with respect to $\mathcal{B}_t$, the filtration generated by the continuous functions restricted to $[0, t]$, where ω is a general element of $C([0, \infty[, \mathbb{R}^n)$,

$$L = \sum_{i=1}^n (\nu a(x))_i \frac{\partial}{\partial x_i} + \frac{1}{2} \sum_{i,k=1}^n (g(x)g(x)^T)_{ik} \frac{\partial^2}{\partial x_i \partial x_k},$$

and $M(0) = h(x_0)$ P^{x_0} -almost surely. The correspondence between the SDE and the martingale problem (under certain regularity conditions) means the following. First, the probability measure on $(C([0, \infty[, \mathbb{R}^n), \mathcal{B})$ induced by the law of the Itô diffusion $(X(t))_{t \geq 0}$ solution of Eq. (3.3.1) solves the martingale problem. Vice versa, if there exists a solution to the martingale problem, then there exists a weak solution of the SDE (3.3.1). Uniqueness also carries over: there is a unique weak solution of the SDE (3.3.1) if and only if there is a unique solution of the associated martingale problem. The coefficients in the differential operator L of our martingale problem are benign, typically polynomials of degree not greater than two, so we can assume that the martingale problem is well posed, that is, there is a unique measure P^{x_0} that solves it. We will say more about this assumption in Section 3.6.

The rest of the argument is identical to that in our first sketch proof. The function g enters the martingale problem only via $g(x)g(x)^T = \nu \text{diag}(a(x))\nu^T$ in L , therefore different g all give the same martingale problem. This single martingale problem has a unique solution to which corresponds a unique weak solution of Eq. (3.3.1). ■

This proposition implies that the SDE (3.3.1) defined with different g , each of which satisfies Eq. (3.2.5), are equivalent in the sense that the distributions of their solutions are the same at any time t . It follows that *all* their moments will be identical at fixed times.

The main goal of Part I is to explore the different possibilities of how the SDE (3.3.1) can be parameterised with different g such that the multidimensional Itô diffusion pro-

cesses given by these different parameterisations all preserve the distribution of the standard CLE (3.2.6). Note that although the number of rows of g is fixed from the beginning to be the number of chemical species n , the number of columns, d , is not *a priori* specified. Our first line of enquiry will examine the minimum d for which the factorisation (3.2.5) is possible — this is also the minimum number of Wiener processes necessary to describe the Itô diffusion process. Then, using the mathematical framework we have developed, we will be able to construct and prove the validity of a reduced formulation of the original CLE where there is only one Wiener process associated with each pair of reversible reactions. This formulation may be considered as a more natural model of chemical reaction systems than the original model (3.2.6). We will also demonstrate that this reduced formulation can speed up numerical simulations considerably without compromising accuracy.

3.3.1 Gillespie’s original formulation

Construction 1 *Assuming that the number of firings of different reaction channels are independent in short time intervals and that the expected number of firings is large enough for each reaction channel, in his seminal paper Gillespie (2000) derived that g is of the form*

$$g(x) = \nu \operatorname{diag}(\sqrt{a_1(x)}, \dots, \sqrt{a_m(x)}). \quad \clubsuit$$

As we said, this will be our reference model and is also a special case of Eq. (3.2.5) with $d = m$. Here every independent Wiener process corresponds to one reaction channel. Hence the physical interpretation of this model is quite clear. Every variable is forced by as many Wiener processes as there are reaction channels which change its count.

Gillespie mentioned that this is not the only possible formulation, and other formulations with differing numbers of Wiener processes are possible. He referred to his former

work (Gillespie, 1996), where equations were laid down which if satisfied by both a g_1 and a g_2 then the two Langevin equations with either g_1 or g_2 would have increments with identical distributions. This is analogous to our Proposition 3.3.1.

3.3.2 The minimal formulation

In what follows we will often refer to $\dim(\text{Ker } \nu)$, the dimension of the (right) nullspace of ν , and to $\dim(\text{Ker } \nu^T)$, the dimension of the left nullspace of ν . Left nullvectors correspond to *conservation laws* in the reaction system, that is, to preserved linear combinations of different species' counts. Each right nullvector corresponds to a multiset of firings of reaction channels (a *multiset* is a set whose elements can be members of this set repeatedly, here we even allow negative multiplicities) such that if starting in state x all reactions occur the number of times given by the right nullvector, then the chemical system will eventually return to the original molecular counts x . However, note that reaction firings should be interpreted broadly, as some entries of these right nullvectors may be negative.

The focus of Part I is to explore different natural choices for the formulation of the CLE (3.3.1) and hence for the choice of d . A natural question to ask is what the minimum number of Wiener processes is in the CLE, or equivalently, what the minimum d is for which the factorisation of

$$B(x) = \nu \text{diag}(a(x)) \nu^T$$

in Eq. (3.2.5) is possible.

As $B(x)$ is a symmetric positive semi-definite square matrix for all x , it can be diagonalised by a change of basis with an orthonormal matrix $U(x)$ whose columns are

eigenvectors of $B(x)$:

$$B(x) = U(x)D(x)U(x)^T.$$

We partition the eigenvectors based on whether they belong to zero eigenvalue ($U_0(x)$) or some positive eigenvalue ($U_1(x)$) and arrange them such that $U(x) = [U_1(x) \ U_0(x)]$. Then there are $n - \dim(\text{Ker} B(x))$ nonzero eigenvalues, so $D(x)$ is of the form

$$D(x) = \begin{bmatrix} D_1(x) & 0 \\ 0 & 0 \end{bmatrix}$$

with a diagonal

$$D_1(x) \in \mathbb{R}^{(n - \dim(\text{Ker} B(x))) \times (n - \dim(\text{Ker} B(x)))}.$$

The construction for $g(x)$ is then

$$g(x) = U(x)D(x)^{1/2} = [U_1(x)D_1(x)^{1/2} \ 0],$$

or simply

$$g(x) = U_1(x)D_1(x)^{1/2} \in \mathbb{R}^{n \times (n - \dim(\text{Ker} B(x)))}.$$

Indeed, $g(x)g(x)^T = U(x)D(x)^{1/2}D(x)^{1/2}U(x)^T = B(x)$. This formulation shows that $d = n - \dim(\text{Ker} B(x))$ independent Wiener processes are enough to define Eq. (3.3.1).

This factorisation is minimal indeed, since the rank of $g(x)$ cannot be less than the rank of $B(x) = g(x)g(x)^T$, that is, $n - \dim(\text{Ker} B(x))$.

The next proposition shows that the number of columns of $g(x)$ is independent of the state x . In order to avoid digression, the proofs of the following two propositions are found in the appendix of Part I.

Proposition 3.3.2 *For every strictly positive x (it is enough that for all x and each reaction channel j , $a_j(x) > 0$ holds), $\dim(\text{Ker} B(x))$ is equal to the number of linearly independent conservation laws of the reaction network, $\dim(\text{Ker} \nu^T)$. In fact, a vector*

$y \in \mathbb{R}^n \setminus \{0\}$ is a (right) nullvector of $B(x)$ if and only if it is a left nullvector of the stoichiometric matrix ν .

The following proposition states that this construction reduces the number of Wiener processes compared to the m Wiener processes of Gillespie's construction by the dimension of the right nullspace of ν .

Proposition 3.3.3 $n - \dim(\text{Ker } \nu^T) = m - \dim(\text{Ker } \nu)$.

We summarise the results of this section.

Construction 2 *The previously described*

$$g(x) = U_1(x)D_1(x)^{1/2}$$

gives a chemical Langevin equation (3.3.1) with $n - \dim(\text{Ker } \nu^T) = m - \dim(\text{Ker } \nu)$ independent Wiener processes. Any CLE requires at least this many independent Wiener processes. ♣

Note that this result is an improvement over Gillespie (2000, Appendix B) and Wilkinson (2006, p 189) in that both texts claim that generally the number of Wiener processes d must be no less than n . We will return to the problem of state space reduction in Section 3.3.4, where we prove that there is an equivalent formulation of the CLE with $n - \dim(\text{Ker } \nu^T)$ states, and as we see here, $n - \dim(\text{Ker } \nu^T)$ Wiener processes.

The minimum number of Wiener processes needed is interesting for efficient numerical simulation (Kloeden and Platen, 1992). Notice that the solution in Construction 2 is not satisfactory since U_1 is dependent on x . Hence, in a numerical simulation scheme at each time step a new diagonalisation of $B(x)$ is required, which is computationally expensive.

As a first improvement, we propose another approach that results in a g of the same size, but potentially decreases the requirement for repeated computation at the cost of increased initial, one-off computation. A substantially different construction will be presented in Construction 4.

Let $W = [W_1 \ W_0] \in \mathbb{R}^{n \times n}$ be an orthogonal matrix such that the columns of $W_0 \in \mathbb{R}^{n \times \dim(\text{Ker } \nu^T)}$ form an orthonormal basis in the left nullspace of ν , $\text{Ker } \nu^T$, and the columns of $W_1 \in \mathbb{R}^{n \times (n - \dim(\text{Ker } \nu^T))}$ are an orthonormal basis in the orthogonal complement, the image space of ν , $\text{Im } \nu$. Let us define the square root $\bar{M} = \sqrt{M}$ of a square matrix $M \in \mathbb{R}^{k \times k}$ as any square matrix $\bar{M} \in \mathbb{R}^{k \times k}$ such that $\bar{M}\bar{M}^T = M$, if such an $\bar{M}$ exists.

Construction 3 *For notational brevity let $A(x) = \text{diag}(a(x))$. Then*

$$g(x) = W_1 \sqrt{W_1^T \nu A(x) \nu^T W_1}$$

gives a chemical Langevin equation (3.3.1) with $n - \dim(\text{Ker } \nu^T) = m - \dim(\text{Ker } \nu)$ independent Wiener processes. ♣

Proof We verify that $\hat{g}(x) = W \sqrt{W^T \nu A(x) \nu^T W}$ is an equally valid diffusion term (it satisfies Eq. (3.2.5)) and that the stated g is equivalent to $\hat{g}$. Note that $W^T \nu A(x) \nu^T W$ and $W_1^T \nu A(x) \nu^T W_1$ are symmetric positive semi-definite matrices, therefore their square root can be evaluated as for $B(x)$ earlier. Thus

$$\hat{g}(x) \hat{g}(x)^T = W W^T \nu A(x) \nu^T W W^T = \nu A(x) \nu^T$$

since W is orthogonal, so Eq. (3.2.5) is satisfied. Also,

$$\begin{aligned} W^T \nu A(x) \nu^T W &= \begin{bmatrix} W_1^T \nu \\ 0 \end{bmatrix} A(x) \begin{bmatrix} \nu^T W_1 & 0 \end{bmatrix} \\ &= \begin{bmatrix} W_1^T \nu A(x) \nu^T W_1 & 0 \\ 0 & 0 \end{bmatrix} \end{aligned}$$

shows that it is enough to use the top left block with $W_1^T \nu A(x) \nu^T W_1$. This is because when constructing $g(x)$ from this, on the left the columns of W_0 would be multiplied by zeros, and when constructing the CLE (3.3.1) on the right the last $\dim(\text{Ker } \nu^T)$ Wiener processes would be multiplied by zeros. Hence we can omit those. $\blacksquare$

This is an improvement over Construction 2 in that here the square root of a state-dependent $(n - \dim(\text{Ker } \nu^T)) \times (n - \dim(\text{Ker } \nu^T))$ matrix is used instead of an $n \times n$ matrix.

3.3.3 A general, state-independent reduction technique

In the previous section a practical constraint for numerical simulations was discussed. Constructions that require in each time step an eigendecomposition of a state-dependent matrix are computationally too costly. In the following, we develop a construction in which, to compute $g(x)$, only matrix products and taking the square root of a state-dependent diagonal matrix are required. This construction will give a CLE that generally may need more than $n - \dim(\text{Ker } \nu^T) = m - \dim(\text{Ker } \nu)$ independent Wiener processes, but certainly not more than m .

For a positive integer k , let I_k denote the $k \times k$ identity matrix. We say two nonzero vectors $y_1, y_2 \in \mathbb{R}^n \setminus \{0\}$ represent the same *direction*, if there is a $\lambda \in \mathbb{R} \setminus \{0\}$ such that $y_1 = \lambda y_2$.

Construction 4 Let s be the number of different directions given by the columns of ν . There exist matrices $J \in \mathbb{R}^{m \times s}$ and $V \in \mathbb{R}^{s \times m}$ such that $VA(x)V^T \in \mathbb{R}^{s \times s}$ is diagonal with only nonnegative entries and

$$g(x) = \nu J \sqrt{VA(x)V^T} \quad (3.3.2)$$

gives a chemical Langevin equation (3.3.1) with s independent Wiener processes, $m - \dim(\text{Ker } \nu) \leq s \leq m$. ♣

Proof Permute the columns of $\nu \in \mathbb{R}^{n \times m}$ such that $\nu = [\nu_1 \ \nu_2]$, where $\nu_1 \in \mathbb{R}^{n \times s}$ has one representative column vector for each direction given by the columns of ν . Then the columns that are left (ν_2) are each a constant multiple of one column in ν_1 . We permute the entries of $A(x)$ accordingly.

Let

$$\nu_2 = [\nu_1 v^{(1)} \ \dots \ \nu_1 v^{(m-s)}],$$

where for all i , $v^{(i)} \in \mathbb{R}^s$ has one nonzero entry.

Introducing $M = [v^{(1)} \ \dots \ v^{(m-s)}] \in \mathbb{R}^{s \times (m-s)}$, the definitions are

$$J = \begin{bmatrix} I_s \\ 0 \end{bmatrix} \in \mathbb{R}^{m \times s},$$

$$V = \begin{bmatrix} I_s & M \end{bmatrix} \in \mathbb{R}^{s \times m}.$$

First, partitioning $A(x)$ according to the sizes of blocks of V ,

$$\begin{aligned} VA(x)V^T &= \begin{bmatrix} I_s & M \end{bmatrix} \begin{bmatrix} A_1(x) & 0 \\ 0 & A_2(x) \end{bmatrix} \begin{bmatrix} I_s \\ M^T \end{bmatrix} \\ &= A_1(x) + MA_2(x)M^T \\ &= A_1(x) + \sum_{j=1}^{m-s} (A_2(x))_{jj} v^{(j)} v^{(j)T}, \end{aligned}$$

where the last step follows from

$$(MA_2(x)M^T)_{ik} = \sum_{j=1}^{m-s} (v^{(j)})_i (A_2(x))_{jj} (v^{(j)})_k^T.$$

Since $v^{(j)}$ has only one nonzero entry for all j , $\sum_{j=1}^{m-s} (A_2(x))_{jj} v^{(j)} v^{(j)T}$ is diagonal with only nonnegative entries, and consequently $VA(x)V^T$ is too. Hence $\sqrt{VA(x)V^T}$ exists trivially.

Second,

$$\nu JV = [\nu_1 \ \nu_2] \begin{bmatrix} I_s & M \\ 0 & 0 \end{bmatrix} = [\nu_1 \quad \nu_1 M] = [\nu_1 \ \nu_2] = \nu.$$

It follows

$$\begin{aligned} g(x)g(x)^T &= \nu J \sqrt{VA(x)V^T} \left(\nu J \sqrt{VA(x)V^T} \right)^T \\ &= \nu JV A(x) V^T J^T \nu^T = \nu A(x) \nu^T \end{aligned}$$


so Eq. (3.2.5) is satisfied. The actual form of g is

$$\begin{aligned} g(x) &= [\nu_1 \ \nu_2] \begin{bmatrix} \sqrt{A_1(x) + MA_2(x)M^T} \\ 0 \end{bmatrix} \\ &= \nu_1 \sqrt{A_1(x) + MA_2(x)M^T} \\ &= \nu_1 \sqrt{A_1(x) + \sum_{j=1}^{m-s} (A_2(x))_{jj} v^{(j)} v^{(j)T}}. \end{aligned} \quad \blacksquare$$

Corollary 3.3.4 *There is a formulation of the chemical Langevin equation (3.3.1) that is constructed from Gillespie's original CLE by omitting one independent Wiener process for each pair of reversible reactions and assigning to the retained Wiener process either respective stoichiometric vector multiplied by the square root of the sum of the two propensities. This is computationally inexpensive to numerically simulate. If m_1 is the number of pairs of reversible reactions, then in Gillespie's formulation there would be*

$2m_1$ Wiener processes for the reversible reactions, while in this formulation there would only be m_1 .

In fact, the result is slightly more general than this. Consider chemical systems with reactions



In both cases one independent Wiener process can be spared. Note that the reactions in these examples are at most bimolecular.

3.3.4 State space reduction

Another form of system-size reduction we have not discussed yet is the reduction of the number of variables. The conservation laws describe linear dependencies between the counts of molecular species. This can be used to express certain variables as functions of others. With $\dim(\text{Ker } \nu^T)$ linearly independent conservation laws it is possible to reduce the number of variables from n to $n - \dim(\text{Ker } \nu^T)$ without loss of accuracy.

To this end, we specify an invertible matrix $T \in \mathbb{R}^{n \times n}$ such that $T\nu$ will take over the role of ν . (For aesthetic reasons one may prefer $T \in \mathbb{Z}^{n \times n}$.) T is just a change of basis of the state space. To see this, multiply the CLE (3.3.1) with T from the left to get an equation in a new variable $z = Tx$:

$$d(Tx) = T\nu a(T^{-1}Tx) dt + Tg(T^{-1}Tx) dW(t),$$

or, by letting $\circ$ denote the composition of functions, and $\cdot$ multiplication (a special composition), we have

$$dz = (T \cdot \nu \cdot a \circ T^{-1})(z) dt + (T \cdot g \circ T^{-1})(z) dW(t).$$

We define T such that the last $\dim(\text{Ker } \nu^T)$ coordinates of the new state variable z are the conservation laws, which do not change at all.

We give T for Construction 1 first. Order the columns of $\nu \in \mathbb{R}^{n \times m}$ such that $\nu = [\nu_b \ \nu_c]$, where the columns of $\nu_b \in \mathbb{R}^{n \times (m - \dim(\text{Ker } \nu))}$ form a basis for $\text{Im } \nu$, and $\nu_c \in \mathbb{R}^{n \times \dim(\text{Ker } \nu)}$ is the collection of the rest of the column vectors. These are linearly dependent on columns of ν_b . Then, similarly to Construction 4, there are vectors

$$w^{(1)}, \dots, w^{(\dim(\text{Ker } \nu))} \in \mathbb{R}^{m - \dim(\text{Ker } \nu)},$$

and a matrix

$$R = [w^{(1)} \ \dots \ w^{(\dim(\text{Ker } \nu))}] \in \mathbb{R}^{(m - \dim(\text{Ker } \nu)) \times \dim(\text{Ker } \nu)}$$

such that $\nu_c = \nu_b R$. Define $\nu_b^\perp \in \mathbb{R}^{n \times \dim(\text{Ker } \nu^T)}$ such that its columns form a basis of the orthogonal complement space of $\text{Im } \nu$. The last preparatory step to the definition of T is to note the following fact, which is proved in the appendix of Part I.

Proposition 3.3.5 $\nu_b^T \nu_b$ is invertible.

Now let

$$T = \begin{bmatrix} (\nu_b^T \nu_b)^{-1} \nu_b^T \\ (\nu_b^\perp)^T \end{bmatrix}.$$

(To get an integer-valued T , we may put an appropriate diagonal matrix

$$D_0 \in \mathbb{Z}^{(n - \dim(\text{Ker } \nu^T)) \times (n - \dim(\text{Ker } \nu^T))}$$

in front of $(\nu_b^T \nu_b)^{-1} \nu_b^T$, and choose $\nu_b^\perp \in \mathbb{Z}^{n \times \dim(\text{Ker } \nu^T)}$.) T is invertible, as required: it is easy to verify that $T^{-1} = [\nu_b \ \nu_b^\perp ((\nu_b^\perp)^T \nu_b^\perp)^{-1}]$. Here $(\nu_b^\perp)^T \nu_b^\perp$ is invertible for the same reason as $\nu_b^T \nu_b$.

Hence

$$\begin{aligned}
T\nu &= \begin{bmatrix} (\nu_b^T \nu_b)^{-1} \nu_b^T \\ (\nu_b^\perp)^T \end{bmatrix} [\nu_b \ \nu_c] \\
&= \begin{bmatrix} I_{m-\dim(\text{Ker } \nu)} & (\nu_b^T \nu_b)^{-1} \nu_b^T \nu_c \\ 0 & 0 \end{bmatrix} \\
&= \begin{bmatrix} I_{m-\dim(\text{Ker } \nu)} & R \\ 0 & 0 \end{bmatrix}.
\end{aligned}$$

Therefore in no CLE formulation will the last $\dim(\text{Ker } \nu^T)$ variables be affected by the drift term $T\nu a(x)$. Since in Constructions 1 and 4 the first factor in $g(x)$ is ν , the last $\dim(\text{Ker } \nu^T)$ rows of the diffusion term $Tg(x)$ will vanish too. Consequently, the last $\dim(\text{Ker } \nu^T)$ variables of z are constant and can be omitted from a numerical simulation.

The same argument holds for Construction 3, using W_1 and W_0 instead of ν_b and $\nu_b^\perp$, respectively, in T . In the case of Construction 2, the state space reduction must precede the reduction of the number of Wiener processes. This method is very similar to Construction 3. For Construction 4, a finer partitioning of matrices ν, J, V is proposed. The detailed calculations are in the appendix of Part I. These considerations prove the following result.

Theorem 3.3.6 *For Constructions 1–4 a state space transformation is possible which reduces the number of variables from n to $n - \dim(\text{Ker } \nu^T) = m - \dim(\text{Ker } \nu)$ without changing the number of independent Wiener processes.*


3.4 Applications

We demonstrate the reduction of the number of independent Wiener processes in the CLE by three examples. The first example is a simple one merely to illustrate our

ideas, while the other two systems are models of biological systems of real interest: a Markov model for a human ether a-go-go related gene (HERG) encoded K^+ channel (Brennan *et al.*, 2009), and the Goldbeter–Koshland switch (Goldbeter and Koshland, 1981), which plays a vital role in many cellular pathways (Huang and Ferrell, 1996). In order to focus on the application of our main results, we will not carry out the fairly well-known state space reduction in any example.

3.4.1 A cyclical reaction system

Consider the following ring of $m = 3$ reactions with $n = 3$ species, $(A_1, A_2, A_3)^T$:



The indexing of reactions follows that of rate constants k_j . This specifies the order of columns in the stoichiometric matrix

$$\nu = \begin{bmatrix} -1 & 0 & 1 \\ 1 & -1 & 0 \\ 0 & 1 & -1 \end{bmatrix},$$

which has rank 2. The propensity vector function is just

$$a(x) = (k_1 x_1, k_2 x_2, k_3 x_3)^T.$$

Gillespie’s diffusion term (Construction 1) is

$$g^1(x) = \begin{bmatrix} -\sqrt{k_1 x_1} & 0 & \sqrt{k_3 x_3} \\ \sqrt{k_1 x_1} & -\sqrt{k_2 x_2} & 0 \\ 0 & \sqrt{k_2 x_2} & -\sqrt{k_3 x_3} \end{bmatrix}.$$

As there are no parallel stoichiometric vectors, Construction 4 cannot reduce the number of Wiener processes.

Constructions 2 and 3 can be computed analytically for such a small example. In Construction 2 finding the eigenvalues of the rank 2, 3×3 matrix requires the solution of a cubic equation (roots of the characteristic polynomial). But we know that one eigenvalue is zero and this reduces the problem to a quadratic. This gives $D(x)$. Finding the eigenvectors is done by solving a linear equation for each nonzero eigenvalue, and then the vectors need to be normalised to create $U_1(x)$.

The calculations giving Construction 3 can be coded in step-by-step instructions.

The orthogonal matrix W can be chosen as

$$W = [W_1 \ W_0] = \left[\begin{array}{cc|c} -1/\sqrt{2} & -1/\sqrt{6} & 1/\sqrt{3} \\ 1/\sqrt{2} & -1/\sqrt{6} & 1/\sqrt{3} \\ 0 & 2/\sqrt{6} & 1/\sqrt{3} \end{array} \right].$$

This is computed only once, therefore its computational cost is almost irrelevant. Then

$$W_1^T \nu A(x) \nu^T W_1 = \begin{bmatrix} 2a_1(x) + \frac{1}{2}a_2(x) + \frac{1}{2}a_3(x) & -\frac{\sqrt{3}}{2}a_2(x) + \frac{\sqrt{3}}{2}a_3(x) \\ -\frac{\sqrt{3}}{2}a_2(x) + \frac{\sqrt{3}}{2}a_3(x) & \frac{3}{2}a_2(x) + \frac{3}{2}a_3(x) \end{bmatrix}.$$

To take the square root of this or, in general, of a matrix

$$\begin{bmatrix} M_{11} & M_{12} \\ M_{12} & M_{22} \end{bmatrix},$$

we can compute the two eigenvalues as the roots of the quadratic characteristic polynomial. These are

$$\lambda_{1,2} = \frac{M_{11} + M_{22} \pm \sqrt{(M_{11} - M_{22})^2 + 4M_{12}^2}}{2}.$$

The corresponding normalised eigenvectors are

$$v_1 = \frac{1}{\sqrt{(\lambda_1 - M_{22})^2 M_{12}^{-2} + 1}} \begin{pmatrix} (\lambda_1 - M_{22})M_{12}^{-1} \\ 1 \end{pmatrix},$$

$$v_2 = \frac{1}{\sqrt{(\lambda_2 - M_{22})^2 M_{12}^{-2} + 1}} \begin{pmatrix} (\lambda_2 - M_{22})M_{12}^{-1} \\ 1 \end{pmatrix}.$$

Thus


$$g^3(x) = W_1 \begin{bmatrix} \sqrt{\lambda_1(x)}v_1(x) & \sqrt{\lambda_2(x)}v_2(x) \end{bmatrix}$$

is the product of a 3×2 and a 2×2 matrix, and the CLE requires 2 Wiener processes.

The construction that requires the least computation time hinges on how the cost of these computations compares to the cost of generating independent Wiener increments (that is, normal random variables).

3.4.2 A K^+ channel

We model the transformations of HERG encoded K^+ channels between three closed states (C_1, C_2, C_3), one open state (O) and one inactivation state (I) as $n = 5$ chemical species $(C_1, C_2, C_3, O, I)^T$ reacting through $m = 10$ reactions:



(For details, see Brennan *et al.* (2009) and references therein.) Thus the stoichiometric matrix is

$$\nu = \begin{bmatrix} -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & -1 & -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & -1 & 1 & 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & 0 & 1 & -1 & -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & -1 & -1 & 1 \end{bmatrix},$$

and the propensity vector function is

$$a(x) = (k_1x_1, k_2x_2, k_3x_2, k_4x_3, k_5x_3, k_6x_4, k_7x_4, k_8x_5, k_9x_5, k_{10}x_3)^T.$$

Gillespie's formulation (Construction 1) needs 10 Wiener processes with

$$g^1(x) = \nu \sqrt{\text{diag}(a(x))}.$$

The rank of the stoichiometric matrix ν is 4, which allows for a CLE specification with 4 Wiener processes. Thus the minimal solutions g^2 and g^3 from Constructions 2 and 3, respectively, are of the form

$$g^2(x) = U_1(x)D_1(x)^{1/2},$$

$$g^3(x) = W_1\sqrt{W_1^T\nu A(x)\nu^TW_1},$$

where $U_1(x), W_1$ are 5×4 , $D_1(x)$ and $\sqrt{W_1^T\nu A(x)\nu^TW_1}$ are 4×4 matrices, respectively. With the exception of W_1 , we could only compute either of these matrices analytically if we solved a quartic equation. To avoid this laborious task, one can use standard numerical computations that we do not present here.


On the other hand, Construction 4 gives a simple closed form diffusion term. Indeed, this is a straightforward example where the number of Wiener processes can be decreased by half, to 5, with

$$g^4(x) = \begin{bmatrix} -1 & 0 & 0 & 0 & 0 \\ 1 & -1 & 0 & 0 & 0 \\ 0 & 1 & -1 & 0 & 1 \\ 0 & 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 1 & -1 \end{bmatrix} \text{diag} \begin{pmatrix} \sqrt{a_1(x) + a_2(x)} \\ \sqrt{a_3(x) + a_4(x)} \\ \sqrt{a_5(x) + a_6(x)} \\ \sqrt{a_7(x) + a_8(x)} \\ \sqrt{a_9(x) + a_{10}(x)} \end{pmatrix}.$$

3.4.3 The Goldbeter–Koshland switch

This example studied by Goldbeter and Koshland (1981) is a system of covalent modifications facilitated by two converter enzymes, E_1 and E_2 . A typical example is a

phosphorylation–dephosphorylation system. It consists of the following $m = 6$ reactions:



with $n = 6$ chemical species, $(S, E_1, C_1, P, E_2, C_2)^T$. The corresponding stoichiometric matrix is

$$\nu = \begin{bmatrix} -1 & 1 & 0 & 0 & 0 & 1 \\ -1 & 1 & 1 & 0 & 0 & 0 \\ 1 & -1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 1 & 0 \\ 0 & 0 & 0 & -1 & 1 & 1 \\ 0 & 0 & 0 & 1 & -1 & -1 \end{bmatrix},$$

while the propensity vector function is given by

$$a(x) = (k_1 x_1 x_2, k_2 x_3, k_3 x_3, k_4 x_4 x_5, k_5 x_6, k_6 x_6)^T.$$

Gillespie’s formulation (Construction 1) with 6 Wiener processes is

$$g^1(x) = \nu \sqrt{\text{diag}(a(x))}.$$

However, the rank of the stoichiometric matrix ν is 3, which implies that only 3 Wiener processes are needed in the CLE. As with the K^+ channel, this can only be practically computed through numerical computation.

The closed form diffusion term from Construction 4 requires 4 Wiener processes.

Removing the stoichiometric vectors corresponding to reactions 2 and 5, we have

$$g^4(x) = \begin{bmatrix} -\sqrt{a_1(x) + a_2(x)} & 0 & 0 & \sqrt{a_6(x)} \\ -\sqrt{a_1(x) + a_2(x)} & \sqrt{a_3(x)} & 0 & 0 \\ \sqrt{a_1(x) + a_2(x)} & -\sqrt{a_3(x)} & 0 & 0 \\ 0 & \sqrt{a_3(x)} & -\sqrt{a_4(x) + a_5(x)} & 0 \\ 0 & 0 & -\sqrt{a_4(x) + a_5(x)} & \sqrt{a_6(x)} \\ 0 & 0 & \sqrt{a_4(x) + a_5(x)} & -\sqrt{a_6(x)} \end{bmatrix}.$$

These examples demonstrate cases in which the stoichiometric matrix is rank deficient and a reduction in the number of Wiener processes is possible. In Example 1 there were no parallel stoichiometric vectors, thus Construction 4 could not be deployed. In Examples 2 and 3 some Wiener processes could be spared for reversible reactions. These were also cases in which Constructions 2 and 3 could reduce the system size even further.

3.5 Simulations

In this section, we present computational benchmarking of numerical simulations of the examples described in Section 3.4. In addition to this, in order to demonstrate the theory we developed, we compare the numerically computed empirical means and variances from simulations that use different constructions for g in the CLE (3.3.1). As we have already stated in Proposition 3.3.1, all these different CLEs have the same finite-dimensional distributions and thus we expect all moments calculated with different g to agree (up to Monte Carlo sampling error).

3.5.1 A cyclical reaction system

For this example we chose rate constants to be $k_1 = k_2 = k_3 = 0.1$ and set the initial state to be $(100, 80, 100)^T$. Our numerical computations were carried out in MATLAB on

a desktop computer. We integrated using the Euler–Maruyama method (Kloeden and Platen, 1992) up to time 5 with a time step size 0.005 and generated 10^4 realisations for each simulated construction.

We simulated two different CLE formulations: the standard formulation (Construction 1) and Construction 3. The first needs 3 Wiener processes, while the latter only needs 2 Wiener processes. In our simulations of Construction 3, we used the explicit formula from Section 3.4.1 for the square root of the 2×2 matrix.

For this simple example, the running time required to generate the sample with g given by Construction 1 was 255 seconds, while with Construction 3 it was 256 seconds. This lack of computational improvement ought not to be surprising since the time saved by using one less Wiener process could be expected to be comparable to the time spent evaluating the complicated exact formula for the matrix square root.

	Construction 1	Construction 3
$E(X_1(5))$	98.41	98.52
$E(X_2(5))$	87.76	87.62
$E(X_3(5))$	93.83	93.86
$\text{Var}(X_1(5))$	50.17	51.19
$\text{Var}(X_2(5))$	47.37	47.49
$\text{Var}(X_3(5))$	47.80	48.46

Table 3.1: Comparison of the empirical means and variances in the cyclical reaction system at time 5 for CLE Constructions 1 and 3.

The results of the comparison of the simulated means and variances using the two different constructions are presented in Table 3.1. As one can see they agree very accurately, as expected.

3.5.2 A K^+ channel

We chose all rate constants to be $k_j = 0.1$ ($j \in \{1, \dots, 10\}$) and set the initial state to be $(100, 50, 100, 50, 100)^T$. We used the Euler–Maruyama integration scheme with a time step size 0.005 to generate 10^4 realisations up to time 5.

We simulated two different CLE constructions: the standard formulation (Construction 1) and Construction 4. The first needs 10 Wiener processes, while the second needs 5 Wiener processes.

For this example, the running time required to generate the sample with g given by Construction 1 was 455 seconds, while with Construction 4 it was 261 seconds. This is a saving of approximately 42%. This saving compares to a 50% decrease in the number of Wiener processes. Through the separate, preceding batch sampling of Wiener increments, it was established that the saving arose mainly from the decrease in the computational cost of matrix multiplications to compute the diffusion term. This observation accentuates the considerable benefit our reduction method can provide. As in Section 3.5.1 the means and the variances calculated using the two different constructions agreed (data not shown).

3.5.3 The Goldbeter–Koshland switch

We chose rate constants $k_1 = 0.05$, $k_2 = 0.1$, $k_3 = 0.1$, $k_4 = 0.01$, $k_5 = 0.1$, $k_6 = 0.1$ and set the initial state to be $(110, 100, 30, 30, 100, 30)^T$. We generated 10^4 realisations up to time 5 with the Euler–Maruyama method, with a time step size 0.005.

In our simulations, we compared the standard formulation (Construction 1) with the reduced one, Construction 4. Whereas the first requires 6 Wiener processes, the latter only needs 4.

The running time required to generate the sample for Construction 1 was 349 seconds, while for Construction 4 it was 254 seconds. Therefore the saving in computational time was approximately 27%, which is the result of a 33% reduction in the number of Wiener processes. Just as in the previous two cases, the means and the variances calculated using the two different constructions agreed (data not shown).

3.6 Criticism of the chemical Langevin equation

The forthcoming section explores issues with the CLE that have so far been swept under the carpet in this work. Specifically, we discuss the question of existence and uniqueness of solutions.

Let us make some simplifying but very classical assumptions. Suppose that all reactions are at most second order (at most two molecules interact in each) and that the kinetics follow the law of mass action. For reactions where two molecules of the same species $\mathcal{S}_i$ react, we stipulate that the propensity function is of the form $a_j(x) = k_r x_i^2$, instead of the combinatorially justified $k_r x_i(x_i - 1)$. This is more in line with the continuous nature of the CLE and causes considerable difference only at molecular numbers so low that are outside the normal application regime of the CLE. This assumption implies that $a_j(x) > 0$ even if $x_i < 2$, down to 0. From the modelling perspective $k_r x_i(x_i - 1)$ would be more accurate but it would still allow reactions to occur if $x_i \in]1, 2[$. Our choice will simplify the treatment of the behaviour at the boundary: we will not need to distinguish between molecular counts falling to 1 or to 0.

We will observe that the trajectories exit the nonnegative orthant with positive probability (Section 3.6.1). We will also find that in certain cases the state may explode in finite time, but identifying all such cases remains a challenging open problem (Sec-

tion 3.6.2). A sketch proof using a standard technique will be given as to how to define the trajectories of the stochastic process sequentially on the open positive orthant when the coefficients are only locally Lipschitz (Section 3.6.3), as they are in our case. We will revisit the fundamental result of this chapter, Proposition 3.3.1, to refine its proof in the light of these findings (Section 3.6.4). The analysis in this section, Section 3.6, is original and unpublished, except for Example 3.6.2.

For reference, we recapitulate that the original formulation of the CLE is the following:

$$dX(t) = \nu a(X(t)) dt + \sum_{j=1}^m \nu_{\cdot j} \sqrt{a_j(X(t))} dW_j(t). \quad (3.2.6)$$

This is the formulation we will use for the following arguments.

3.6.1 Species in the chemical Langevin equation can become negative

The starting observation of this line of argument is that the nonnegativity of the variables will be violated with positive probability at all time instances in the CLE model of even simple chemical systems. This has been noted in the literature only very recently (Wilkie and Wong, 2008; Szpruch and Higham, 2010), but we expose this phenomenon in greater generality.


The great concern is that we cannot have negative numbers under the square roots in (3.2.6), but we also expect our model to reflect the physical reality by assigning nonnegative values to molecular counts (or concentrations).

In practice (that is, in computer simulations) this problem does not necessarily appear. This is because the probability of negativity is small so it may not be encountered in the finite sample that the computer generates. But in theory — particularly when the

weak solution is considered instead of the strong solution, that is, the trajectories that the simulations generate — this means that the CLE is seriously flawed and it ought to be replaced by an improved SDE model.

We give three increasingly general cases of how (and why) trajectories will leave the nonnegative orthant. The first has been in the folklore for some time, the second was discussed by Wilkie and Wong (2008) and Szpruch and Higham (2010), and the third is a new and unpublished finding.

Example 3.6.1 The case in which the problem can be seen most easily is the *constitutive production* of a species, say $\mathcal{S}_i$:



Let us assume that this is reaction $\mathcal{R}_1$ among m reactions. Its propensity is $a_1(x) = k_1$.

Then, with ν_i denoting the i th row of ν , the row corresponding to $\mathcal{S}_i$ in the CLE is

$$dX_i(t) = \nu_i \cdot a(X(t)) dt + \sqrt{k_1} dW_1(t) + \sum_{j=2}^m \nu_{ij} \sqrt{a_j(X(t))} dW_j(t). \quad (3.6.1)$$


Here $\sqrt{k_1} dW_1(t)$ is an additive noise term. Hence, for any $t > 0$ and $M > 0$,

$$\mathbb{P} \left(\int_0^t \sqrt{k_1} dW_1(s) < -M \right) > 0.$$

Therefore no matter what the other terms are on the right-hand side of Eq. (3.6.1), this reaction in itself can drive X_i negative in any short time interval with positive probability. For instance, it makes no difference what the drift term is; it can only shift the distribution of X_i upwards but it will not prevent it from having an infinitely long tail towards negative values. ■

Example 3.6.2 Wilkie and Wong (2008) and Szpruch and Higham (2010) gave *reversible isomerisation* as an example. This example has two molecular species and two

reactions:



In this case the total number of molecules is conserved: $X_T = X_1(t) + X_2(t)$. The row that corresponds to $\mathcal{S}_1$ in the CLE is

$$\begin{aligned} dX_1(t) = & -k_1 X_1(t) dt + k_2 X_2(t) dt \\ & - \sqrt{k_1 X_1(t)} dW_1(t) + \sqrt{k_2 X_2(t)} dW_2(t), \end{aligned}$$

or after substituting the above conservation law,



$$\begin{aligned} dX_1(t) = & -k_1 X_1(t) dt + k_2 (X_T - X_1(t)) dt \\ & - \sqrt{k_1 X_1(t)} dW_1(t) + \sqrt{k_2 (X_T - X_1(t))} dW_2(t). \end{aligned}$$

When $X_1(t)$ is close to zero, then $X_2(t) = X_T - X_1(t)$ will be close to X_T . If $X_1(t)$ is below some fixed value, then $X_2(t)$ is bounded from below. Similarly to Example 3.6.1, it is enough to consider the diffusion terms for our argument. Under the assumption that $k_2(X_T - X_1(t))$ is bounded from below, $\sqrt{k_2(X_T - X_1(t))} dW_2(t)$ will be similar to the additive noise term of Example 3.6.1. The same argument applied here proves that this noise term on its own can drive $X_1(t)$ negative.

When $X_1(t)$ is close to the other end of the $[0, X_T]$ interval, close to X_T , a similar argument with the roles of X_1 and X_2 reversed proves that X_2 can become negative with positive probability.

This example is mathematically equivalent to the ion channel model that Dangerfield, Kay, and Burrage (2010) studied. The authors proposed to replace the CLE model for this system with the Wright–Fisher model from population genetics because that preserves the interval constraint on the state. ■

Example 3.6.3 The same phenomenon is present in the most typical reactions, like



For the second reaction (the argument holds for the first too), the row of the CLE for $\mathcal{S}_i$ is

$$\begin{aligned} dX_i(t) = & \nu_i a(X(t)) dt + \sqrt{k_1 X_{h_1}(t) X_{h_2}(t)} dW_1(t) \\ & + \sum_{j=2}^m \nu_{ij} \sqrt{a_j(X(t))} dW_j(t). \end{aligned} \quad (3.6.2)$$

For this argument, we suppose that $\mathcal{S}_{h_1}$ and $\mathcal{S}_{h_2}$ do not appear on the reactant side of any other reaction. Again, for any $t > 0$ and $M > 0$,

$$\mathbb{P} \left(\int_0^t 1 dW_1(s) < -M \right) > 0.$$

As there is an independent Wiener process W_ℓ associated with each reaction channel $\mathcal{R}_\ell$, the event

$$\left\{ \begin{array}{l} \text{for all } \ell \in \{2, 3, \dots, m\} \text{ and } u \in [0, t], \\ \int_0^u \nu_{i\ell} a_\ell(X(s)) ds + \int_0^u \nu_{i\ell} \sqrt{a_\ell(X(s))} dW_\ell(s) \text{ is bounded} \end{array} \right\} \quad (3.6.3)$$

has positive probability (we have just stipulated that none of the $a_\ell(x)$ is dependent on x_{h_1} or x_{h_2}). Therefore, if M is large enough, then on the event $\{\int_0^t 1 dW_1(s) < -M\}$ the first noise term $\sqrt{k_1 X_{h_1}(t) X_{h_2}(t)} dW_1(t)$ can dominate the right-hand side of Eq. (3.6.2). Thus with positive probability X_i will *decrease*. As it approaches zero, as neither $\mathcal{S}_{h_1}$ nor $\mathcal{S}_{h_2}$ is used on the reactant side in any other reaction, the propensity $a_1(X(u))$ will remain bounded from below:

$$a_1(X(u)) = k_1 X_{h_1}(u) X_{h_2}(u) \geq k_1 X_{h_1}(0) X_{h_2}(0) > 0. \quad (3.6.4)$$

So the term that is forcing X_i towards zero will not vanish and there will be nothing to prevent X_i becoming negative. By requiring low enough bounds in Eq. (3.6.3) for each reaction channel it can also be guaranteed that with positive probability neither $X_{h_1}(u)$ nor $X_{h_2}(u)$ decreases by so much when they are on the product side of reactions in $\{\mathcal{R}_2, \mathcal{R}_3, \dots, \mathcal{R}_m\}$ that Eq. (3.6.4) is violated. This completes the argument. ■

Why certain variables become negative with a positive probability can be understood from Section 3.1, the derivation of the CLE by Gillespie (2000). That argument replaced nonnegative Poisson random variables with normal random variables, for each reaction channel separately. However, the left tail of the normal distribution is unbounded. In Example 3.6.3 the drift pushes X_i towards ever higher positive values but the Wiener noise effectively makes the reaction channel reversible (although the reverse direction will have a small probability). As the reaction propensity is only dependent on the left-hand side of the reaction, the right-hand side can become negative due to this backwards-acting noise by creating $\mathcal{S}_{h_1}$ and $\mathcal{S}_{h_2}$ from $\mathcal{S}_i$.

Note that this phenomenon is inherently multidimensional. In the example X_i becomes negative because its noise is dictated by the other molecular counts (X_h , or X_{h_1} and X_{h_2}).

Therefore it will not be surprising what Wilkie and Wong (2008) proposed in order to preserve nonnegativity in the CLE. They recommend that for each reaction, the corresponding noise term should be kept only in the equations of species that are reactants of that reaction. The noise term should be omitted from the equations of species that appear only on the product side of the reaction. However, this leads to kinetics where mass is not conserved in individual reactions. We think this is too high a price to pay. The authors are forced to reinterpret the meaning of this noise which only affects the reactant but not the product side; they stipulate that it is due to the temporary

unavailability or the release of reactant molecules by clustering, caging, adsorption or similar processes. This modified CLE, even if nonnegativity preserving, is no longer a quantitative model of intrinsic noise.

Numerical analysts are in a disadvantaged situation to recognise this negativity issue with the CLE. If a simulation runs into negative territory, it is typically attributed to the time step being too long. The observation underlying this is that the reaction intensities converge to zero with the decreasing number of reactants ($\mathcal{S}_h$, $\mathcal{S}_{h_1}$ or $\mathcal{S}_{h_2}$). What we have just shown is that it is not one of these molecules, but $\mathcal{S}_i$ whose count may become negative. The logically incorrect step, as we can see it now, is to prematurely deduce that if one of the species' counts becomes negative, it must have been caused by discretisation error. The following quotation from Sotiropoulos *et al.* (2009, p 472) is a typical example which suggests that many users of the CLE have an incomplete view of the negativity issue:

‘Every time the integration failed, we decreased the time step [...]; 5×10^{-5} s was the first time step for which the integration did not fail, i.e., species populations did not attain negative values. [...] In the remainder of the present work, whenever we state that the integration fails for larger time steps, it is assumed that we followed a similar approach to arrive at the selected time step.’

We verified this claim and found that in each of their three examples, with step sizes that they had found too long, negativity arose immediately in the first time step. In each case it was caused by the drift and not the diffusion term. In this case the authors were right, and the discretisation error was responsible for negativity. However, we suspect that in the CLE modelling framework reversible reactions, like Example 3.6.2 or the complex formation in the simple enzymatic reaction associated with Michaelis and Menten, often

drive scarce molecular species negative due to the reason we discussed. It is important that users of the CLE are aware of this possibility.

3.6.2 Problems to be addressed in any proof of existence and uniqueness

As we have just seen, it is a fundamental property of the CLE that its trajectories can exit the nonnegative orthant. In this section we conduct a survey into what challenges there are which one needs to overcome in order to guarantee a unique solution for the CLE on an infinite time horizon.

Textbooks typically give the following set of sufficient conditions for the existence and uniqueness of strong solutions for SDEs. (For our purposes the possibility of time-dependence of the coefficients will be ignored.)

Theorem 3.6.4 *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $g : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times d}$ be measurable functions satisfying*

$$|f(x)| + |g(x)| \leq C(1 + |x|) \quad (3.6.5)$$

for all $x \in \mathbb{R}^n$ and some constant C (where $|g(x)|^2 = \sum_{i,j} |g_{ij}(x)|^2$) and such that

$$|f(x) - f(y)| + |g(x) - g(y)| \leq D|x - y| \quad (3.6.6)$$

for all $x, y \in \mathbb{R}^n$ and some constant D . Let Z be a random variable taking values in $\mathbb{R}^n$ which is independent of the σ -algebra $\mathcal{F}_\infty$ generated by $(W(t))_{t \geq 0}$ and such that

$$\mathbb{E}(|Z|^2) < \infty.$$

Then the SDE

$$dX(t) = f(X(t)) dt + g(X(t)) dW(t), \quad X(0) = Z$$

has a unique t -continuous solution $X(t, \omega)$ with the property that for all $t \geq 0$, $X(t, \omega)$ is adapted to the filtration $(\mathcal{F}_t^Z)$ generated by Z and $(W(s))_{0 \leq s \leq t}$ and

$$\mathbb{E} \left(\int_0^t |X(s)|^2 ds \right) < \infty.$$


(This is Theorem 5.2.1 in Øksendal (2007), but it can also be found as Theorem 1.1 in Chapter 5 of Friedman (1975); Theorem 2.9 in Chapter 5 of Karatzas and Shreve (1998). See also Theorem 10.6 in Chung and Williams (1990).)

In fact, there is redundancy in these conditions: the global Lipschitz property (3.6.6) implies at most linear growth (3.6.5). Instead, it is enough to require local Lipschitzness and linear growth.

The CLE does not satisfy the Lipschitz condition (3.6.6): its diffusion term $g(x) = \nu \text{diag}(\sqrt{a(x)})$ is not Lipschitz, only locally Lipschitz on the open, strictly positive orthant $]0, \infty[^n$ because of the square roots. If there are reactions involving two molecules of the same species (second-order or higher), particularly a dimerisation with propensity $a_j(x) = k_j x_i^2$, then the Lipschitz condition will also be violated by $f(x) = \nu a(x)$ as $x_i \rightarrow \infty$.

The violation of Condition (3.6.5) can cause the state to explode in finite time. We can even give a simple chemical example.

Example 3.6.5 Consider the following chemical reaction system:



The corresponding CLE is

$$dX_1(t) = k_1 X_1(t)^2 dt + \sqrt{k_1} |X_1(t)| dW_1(t), \quad X_1(0) = a > 0.$$

If we take the expectation of both sides and look at the remaining ODE (the reaction rate equation), then by Jensen's inequality

$$d\mathbb{E}(X_1(t)) = k_1 \mathbb{E}(X_1(t)^2) dt \geq k_1 \mathbb{E}(X_1(t))^2 dt.$$

The ODE


$$\frac{d}{dt}x(t) = k_1x(t)^2, \quad x(0) = a > 0$$

can be solved explicitly:

$$x(t) = \frac{a}{1 - ak_1t},$$

and this blows up by $t = \frac{1}{ak_1}$. Since $E(X_1(t))$ increases at least as rapidly as $x(t)$, and their initial conditions are identical, so it too must blow up by that time. ■

Obviously, such a reaction is unlikely to form part of a reaction system. However, in general it cannot be expected that some conservation law (e.g. the conservation of mass or the conservation of the number of atoms of certain elements) holds for each reaction channel, certainly not in practical biochemical applications. The reason for this is that very often only the counts and the reactions of macromolecules (such as nucleic acids, proteins) are modelled, simple molecules and building blocks of complex molecules (e.g. ATP, nucleotides, amino acids) are assumed to be abundant enough not to be a limiting factor. Thus in applications one often writes reactions like



In complex biochemical networks with such non-conserving reactions it is unclear whether there is a possibility for runaway production of a molecular species as in Example 3.6.5. One would like to give as general conditions as possible under which the existence of a solution of the CLE that does not explode in finite time can be guaranteed.

At the most conservative end of the spectrum, one can impose the conservation of mass for each reaction channel. In this case the state space is the compact set given by the intersection of the stoichiometric subspace with $[0, \infty]^n$, and state explosion is not possible. One is interested in how much this condition can be relaxed such that

state explosion still cannot occur in finite time. In many models we do expect that variables will ultimately tend to infinity. In a cell cycle model DNA (and essentially all constituents of a cell) will need to be doubled by the end of the cell cycle. If cell division is not incorporated into the model, then the biochemical processes are expected to generate exponentially increasing masses of each molecular species.

In the deterministic reaction rate equation setting, with law of mass action kinetics, trajectories starting from a nonnegative initial state remain nonnegative. If the chemical reaction system is *weakly reversible* (the definition of this classical notion would require some prerequisites, therefore we choose to direct the reader to Feinberg, 1987), then there is a bounded subset of the positive orthant $S \subset]0, \infty[^n$ such that for every strictly positive initial state the trajectory will enter S and will remain inside S for all time after a certain threshold that is dependent on the initial state (August and Barahona, 2010). This means that weak reversibility rules out finite time explosion in the deterministic model under the law of mass action.

The development of comparably general results for the stochastic case is left for future research (Chapter 10).

3.6.3 A strong existence and uniqueness proof for the chemical Langevin equation on the strictly positive orthant

We now consider how to define strong solutions of the CLE taking into account that close to the boundary the diffusion term satisfies Ineq. (3.6.6) only locally. We have already indicated that for large values of the variables, the drift term too may only satisfy a local Lipschitz condition. For this argument we require that the initial state is

a strictly positive deterministic vector, $x_0 \in]0, \infty[^n$.

One may take positive sequences $(\varepsilon_k) \searrow 0$ and $(K_k) \nearrow \infty$ (that is, the sequence (ε_k) is monotonically decreasing and converges to 0, and (K_k) monotonically increases to infinity) such that $\varepsilon_1 < (x_0)_i < K_1$ for all $i \in \{1, \dots, n\}$. For each element ω of the probability space, we define the (continuous) path of the stochastic process up to stopping times

$$\tau_k(\omega) := \inf\{t \geq 0 \mid X(t, \omega) \notin]\varepsilon_k, K_k[^n \text{ given } X(0, \omega) = x_0\}$$

sequentially. This can be achieved by replacing $g(x) = \nu \operatorname{diag}(\sqrt{a(x)})$ with some g^k such that for each $k \in \mathbb{N}$ their restrictions to $] \varepsilon_k, K_k[^n$ coincide:

$$g \restriction] \varepsilon_k, K_k[^n = g^k \restriction] \varepsilon_k, K_k[^n,$$

$g^k(x)$ is zero whenever at least one coordinate of x is nonpositive and g^k is globally Lipschitz. This implies that g^k satisfies the linear growth condition (3.6.5). Similarly, f can be modified outside $] \varepsilon_k, K_k[^n$ (more specifically, for all $i \in \{1, \dots, n\}$ for $x_i > K_k$) such that f^k meets conditions (3.6.5) and (3.6.6). Then Theorem 3.6.4 can be applied to the SDEs

$$dX^k(t) = f^k(X^k(t)) dt + g^k(X^k(t)) dW(t).$$

For each ω , the strong solutions $X^k(t)$ of these SDEs have the property that for $k_1 < k_2$,

$$(X^{k_1}(t, \omega))_{0 \leq t \leq \tau_{k_1}(\omega)} = (X^{k_2}(t, \omega))_{0 \leq t \leq \tau_{k_1}(\omega)},$$

that is, $X^{k_2}(t)$ extends $X^{k_1}(t)$, but it coincides with the trajectory we wish to define as the solution of the CLE up to τ_{k_2} , which is greater than τ_{k_1} . This is a standard technique, e.g. compare with Section 5.3 in Durrett (1996).

The sequence of stopping times $\tau_k(\omega)$ is strictly increasing. Therefore it has some limit $\tau(\omega)$, either finite or infinite. We define the trajectory of the CLE, $X(t)$, up to

but excluding $\tau(\omega)$. If $\tau(\omega) = \infty$, then the trajectory thus constructed is defined for all $t \geq 0$. If $\tau(\omega)$ is finite, then the trajectory may have a limit as $t \rightarrow \tau(\omega)$, which limit may have a coordinate that is zero (the trajectory reached the boundary), or that has diverged to infinity (finite time explosion).

Let us summarise what we found in this section. The CLE, the standard SDE model for chemical reaction kinetics, poses several challenges to a rigorous analysis of the existence and uniqueness of its solutions. Its diffusion coefficient satisfies the Lipschitz condition on the positive orthant $]0, \infty[^n$ only locally. In certain chemical systems the state may explode in finite time but in this it does not differ from the deterministic reaction rate equation model. Neither the drift nor the diffusion coefficient vanishes on the boundary. Therefore extending them to be zero on the entire $\mathbb{R}^n$ outside the nonnegative orthant is not a continuous extension. This is of importance for a proof of Proposition 3.3.1: either we aim to prove the existence and uniqueness of a solution to the Kolmogorov forward equation or to the martingale problem, most references give uniqueness results which require that the coefficients are continuous on the entire $\mathbb{R}^n$.

3.6.4 Revisiting Proposition 3.3.1: the proof of the weak equivalence of our alternative formulations

Lastly, we will reconsider the statement which was at the heart of the development of different formulations of the CLE. This is motivated by the more accurate understanding that has been developed over the past few pages of the subtle issues with defining the solution of the CLE. The second proof of Proposition 3.3.1 used the martingale problem formulation, and implicitly assumed that the process is defined on the entire $\mathbb{R}^n$. Now we can make a rigorous argument that takes into account that the coefficients of the CLE are only defined on the nonnegative orthant. The coefficients are only locally Lipschitz

on the strictly positive orthant; for example, on sets of the form $] \varepsilon, K[^n$ ($0 < \varepsilon < K$) they are bounded and Lipschitz.

For this argument we will rely on Chapter 1 of Pinsky (1995). The main result that we will need to invoke is in its Section 1.13, *The generalized martingale problem on* $D \subseteq \mathbb{R}^n$. Notice, first, that here the martingale problem is defined on a subset of $\mathbb{R}^n$. Second, the word *generalised* refers to the fact that this setup will allow the explosion of the corresponding stochastic process. We need to introduce some new definitions now.

Let $D \subseteq \mathbb{R}^n$ be a *domain* (an open, connected set), and let $(D_k)_{k \in \mathbb{N} \setminus \{0\}}$ be an increasing sequence of bounded domains in $\mathbb{R}^n$ such that for all k , for the closure of D_k , $\bar{D}_k \subseteq D_{k+1}$ holds, and $D = \bigcup_{k=1}^{\infty} D_k$. Let $\hat{D} := D \cup \{\Delta\}$ denote the *one-point compactification* of D : if D is bounded, then Δ is identified with ∂D , the boundary of D ; if D is unbounded, then Δ is identified with ∂D extended with the ‘point at infinity’. With the introduction of an appropriate metric on $\hat{D}$ (Pinsky, 1995, Section 1.12), let $C([0, \infty[, \hat{D})$ denote the space of continuous functions ω from $[0, \infty[$ to $\hat{D}$ with the topology of uniform convergence (with respect to the metric on $\hat{D}$) on bounded intervals. The choice of the trajectories themselves as elementary events (the so-called *canonical* construction) is already manifest. Similarly to the previous section, the *first exit time* from D_k will be required:

$$\tau_k(\omega) := \inf\{t \geq 0 \mid \omega(t) \notin D_k\}.$$

Define $\tau_D := \lim_{k \rightarrow \infty} \tau_k$. Now

$$\hat{\Omega}_D := \{\omega \in C([0, \infty[, \hat{D}) \mid \text{either } \tau_D(\omega) = \infty$$

$$\text{or } \tau_D(\omega) < \infty \text{ and for all } t > \tau_D(\omega), \omega(t) = \Delta\}.$$

Let $\hat{\mathcal{B}}^D$ denote the Borel σ -algebra on $\hat{\Omega}_D$ and define the filtration $\hat{\mathcal{B}}_t^D := \sigma((\omega(s))_{0 \leq s \leq t})$.

Then we have $\hat{\mathcal{B}}^D = \sigma((\hat{\mathcal{B}}_t^D)_{0 \leq t < \infty}) = \sigma((\hat{\mathcal{B}}_{\tau_k}^D)_{k \in \mathbb{N} \setminus \{0\}})$.

Let S_n denote the set of real symmetric positive semidefinite $n \times n$ matrices. With $f : D \rightarrow \mathbb{R}^n$ and $B : D \rightarrow S_n$, we define the following differential operator on D :

$$L = \sum_{i=1}^n f_i(x) \frac{\partial}{\partial x_i} + \frac{1}{2} \sum_{i,k=1}^n B_{ik}(x) \frac{\partial^2}{\partial x_i \partial x_k}.$$

For each $k \in \mathbb{N} \setminus \{0\}$, let ψ_k be a smooth function: $\psi_k \in C^\infty(\mathbb{R}^n, \mathbb{R})$, satisfying $\psi_k \upharpoonright D_k = 1$, $\psi_k \upharpoonright \mathbb{R}^n \setminus D_{k+1} = 0$ and $0 \leq \psi_k \leq 1$. Define $f^k : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $B^k : \mathbb{R}^n \rightarrow S_n$ by

$$f^k = \psi_k f,$$

$$B^k = \psi_k B + (1 - \psi_k) I_n.$$

We define L_k on $\mathbb{R}^n$ similarly to L , with f^k and B^k in the place of f and B . For $s, t \in \mathbb{R}$, let $s \wedge t := \min\{s, t\}$. Following these preparatory steps, we are ready to define the generalised martingale problem on D and state important results about it.

The generalised martingale problem for L on D For each $x_0 \in \hat{D}$, find a probability measure P^{x_0} on $(\hat{\Omega}_D, \hat{\mathcal{B}}^D)$ such that

1. $P^{x_0}(\omega(0) = x_0) = 1$;
2. $h(\omega(t \wedge \tau_k)) - \int_0^{t \wedge \tau_k} Lh(\omega(s)) ds$ is a martingale with respect to $(\hat{\Omega}_D, \hat{\mathcal{B}}_t^D, P^{x_0})$ for all $k \in \mathbb{N} \setminus \{0\}$ and all $h \in C^2(D)$.

Theorem 3.6.6 (Pinsky, 1995, Theorem 1.13.1) *Let $f : D \rightarrow \mathbb{R}^n$ and $B : D \rightarrow S_n$ be locally bounded and measurable on D and assume that B is continuous on D and that*

$$\sum_{i,k=1}^n B_{ik}(x) y_i y_k > 0, \text{ for all } x \in D \text{ and } y \in \mathbb{R}^n \setminus \{0\}. \quad (3.6.7)$$

Then there exists a unique solution $\{P^{x_0} \mid x_0 \in \hat{D}\}$ to the generalised martingale problem for L on D . Letting $\{P_{(k)}^{x_0} \mid x_0 \in \mathbb{R}^n\}$ denote the unique solution to the martingale problem

on $\mathbb{R}^n$ for L_k with coefficients f^k and B^k , then $P^{x_0} \upharpoonright \hat{\mathcal{B}}_{\tau_k}^D = P_{(k)}^{x_0} \upharpoonright \mathcal{B}_{\tau_k}$ for all $k \in \mathbb{N} \setminus \{0\}$ and all $x_0 \in D$ (where $\mathcal{B}_{\tau_k}$ has been identified with $\hat{\mathcal{B}}_{\tau_k}^D$ in the obvious way). The family $\{P^{x_0} \mid x_0 \in D\}$ possesses the Feller property and the family $\{P^{x_0} \mid x_0 \in \hat{D}\}$ possesses the strong Markov property. Furthermore, the sub-probability measure $P^{x_0}(\omega(t) \in \cdot)$ on D possesses a density $p_t(x_0, x)$ and for any finite stopping time σ ,

$$P^{x_0}(\omega(\sigma + t) \in B \mid \hat{\mathcal{B}}_{\sigma}^D) = P^{\omega(\sigma)}(\omega(t) \in B)$$

for all $t \geq 0$ and $B \subseteq \mathbb{R}^n$ Borel. If

$$\lim_{k \rightarrow \infty} P_{(k)}^{x_0}(\tau_k \leq t) = 0 \text{ for all } t > 0 \text{ and } x_0 \in D,$$

then $\{P^{x_0} \mid x_0 \in D\}$ in fact solves the martingale problem for L on D .

The following is an immediate consequence.

Theorem 3.6.7 (Pinsky, 1995, Theorem 1.13.2) *Let $D_P, D_Q \subseteq \mathbb{R}^n$ be domains. Let $f^P : D_P \rightarrow \mathbb{R}^n$ and $B^P : D_P \rightarrow S_n$ be locally bounded and measurable. With these coefficients, define L_P in the usual way. Let $f^Q : D_Q \rightarrow \mathbb{R}^n$, $B^Q : D_Q \rightarrow S_n$ and L_Q satisfy the respective conditions. Let $\{P^{x_0} \mid x_0 \in \hat{D}_P\}$ solve the generalised martingale problem for L_P on D_P and let $\{Q^{x_0} \mid x_0 \in \hat{D}_Q\}$ solve the generalised martingale problem for L_Q on D_Q . Let $U \subseteq D_P \cap D_Q$ be a domain on which the coefficients of L_P and L_Q coincide. Assume that these coefficients, when restricted to U , satisfy the conditions of Theorem 3.6.6. Then for all $x_0 \in U$, $P^{x_0} \upharpoonright \mathcal{B}_{\tau_U} = Q^{x_0} \upharpoonright \mathcal{B}_{\tau_U}$. ($\mathcal{B}_{\tau_U}$ has been identified with $\hat{\mathcal{B}}_{\tau_U}^{D_P}$ and $\hat{\mathcal{B}}_{\tau_U}^{D_Q}$ in the obvious way.)*

These two theorems will be directly applied to prove our equivalence statement. An issue that needs to be dealt with is that the diffusion matrix of the original CLE (3.2.6), $B(x) = \nu \text{diag}(a(x))\nu^T$, does not generally satisfy Condition (3.6.7) due to the existence of conservation laws (compare with Proposition 3.3.2). This failure to meet the positive

definiteness criterion is called the *degenerate* case in the theory of martingale problems (Stroock and Varadhan, 1997). In order to rectify this issue, we will carry out a state space reduction on the CLE (Section 3.3.4), and apply Theorem 3.6.6 to a reduced, $(n - \dim(\text{Ker } \nu^T))$ -dimensional formulation.

It is clear that if a basis transformation is applied to this reduced CLE formulation whose corresponding generalised martingale problem has a unique solution, or the state space of this reduced CLE is embedded into a higher dimensional space, the solution will still be essentially the same. Therefore we may state our equivalence result allowing for such contingencies. In fact, it is vital to do so; otherwise even the equivalence of the original CLE (3.2.6) and its $(n - \dim(\text{Ker } \nu^T))$ -dimensional reduced form would not be covered.

Proposition 3.6.8 *Let $x_0 \in]0, \infty[^n$. Let $N \in \mathbb{N}$, $N \geq n - \dim(\text{Ker } \nu^T)$. Assume that there exist invertible matrices*

$$T = \begin{bmatrix} T_1 \\ T_0 \end{bmatrix} \in \mathbb{R}^{n \times n}, \quad S = \begin{bmatrix} S_1 \\ S_0 \end{bmatrix} \in \mathbb{R}^{N \times N}$$

such that $T_1 \in \mathbb{R}^{(n - \dim(\text{Ker } \nu^T)) \times n}$, $T_0 \in \mathbb{R}^{\dim(\text{Ker } \nu^T) \times n}$, $S_1 \in \mathbb{R}^{(n - \dim(\text{Ker } \nu^T)) \times N}$, $S_0 \in \mathbb{R}^{(N - n + \dim(\text{Ker } \nu^T)) \times N}$ and $T_0 \nu = 0$. Consider functions

$$b : S_1^{\leftarrow}(T_1(]0, \infty[^n)) \rightarrow \mathbb{R}^N, \text{ and}$$

$$\sigma : S_1^{\leftarrow}(T_1(]0, \infty[^n)) \rightarrow \mathbb{R}^{N \times d}$$

(here $S_1^{\leftarrow}(T_1(]0, \infty[^n)) \subseteq \mathbb{R}^N$ is the preimage of $T_1(]0, \infty[^n)$ under S_1) which are locally Lipschitz, satisfy the linear growth conditions (3.6.5) and for which $S_0 b = 0$ and $S_0 \sigma = 0$ hold. Let $(E_k)_{k \in \mathbb{N} \setminus \{0\}}$ be an increasing sequence of bounded domains in $]0, \infty[^n$ such that $x_0 \in E_1$, for all k , $\bar{E}_k \subseteq E_{k+1}$ holds, and $\bigcup_{k=1}^{\infty} E_k =]0, \infty[^n$. With these and a d -

dimensional standard Wiener process $\widetilde{W}$, consider the SDE

$$d\xi(t \wedge \tau_\xi) = b(\xi(t \wedge \tau_\xi)) dt + \sigma(\xi(t \wedge \tau_\xi)) d\widetilde{W}(t), \quad (3.6.8)$$

$$\xi(0) = \xi_0 \in S_1^{\leftarrow}(T_1([0, \infty[^n]),$$

which is stopped at

$$\tau_\xi := \lim_{k \rightarrow \infty} \inf \left\{ t \geq 0 \mid \xi(t) \notin S_1^{\leftarrow}(T_1(E_k)) \right\}.$$

If

$$S_1 \xi_0 = T_1 x_0, \quad (3.6.9)$$

$$S_1 b \left(S^{-1} \begin{pmatrix} \cdot \\ S_0 \xi_0 \end{pmatrix} \right) = T_1 \nu a \left(T^{-1} \begin{pmatrix} \cdot \\ T_0 x_0 \end{pmatrix} \right), \quad (3.6.10)$$

$$S_1 \sigma \left(S^{-1} \begin{pmatrix} \cdot \\ S_0 \xi_0 \end{pmatrix} \right) \sigma \left(S^{-1} \begin{pmatrix} \cdot \\ S_0 \xi_0 \end{pmatrix} \right)^T S_1^T = T_1 \nu \operatorname{diag} \left(a \left(T^{-1} \begin{pmatrix} \cdot \\ T_0 x_0 \end{pmatrix} \right) \right) \nu^T T_1^T, \quad (3.6.11)$$

(Eqs. (3.6.10) and (3.6.11) are required to hold on $T_1([0, \infty[^n)$), then the stopped SDE (3.6.8) has essentially the same finite-dimensional distributions as the CLE (3.2.6) started from x_0 and stopped at

$$\tau_{[0, \infty[^n} = \lim_{k \rightarrow \infty} \inf \left\{ t \geq 0 \mid X(t) \notin E_k \right\}.$$

(In different terminology: the two have essentially the same law.) Here essentially refers to that the finite-dimensional distributions of Eqs. (3.2.6) and (3.6.8) are bijectively connected by linear transformations: for any $\ell \in \mathbb{N}$, $t_1, \dots, t_\ell \geq 0$ and Borel sets $B_1, \dots, B_\ell \subseteq \mathbb{R}^{n - \dim(\operatorname{Ker} \nu^T)}$,

$$\begin{aligned} \mathbb{P}(S_1 \xi(t_1 \wedge \tau_\xi) \in B_1, \dots, S_1 \xi(t_\ell \wedge \tau_\xi) \in B_\ell) &= \\ &= \mathbb{P}(T_1 X(t_1 \wedge \tau_{[0, \infty[^n}) \in B_1, \dots, T_1 X(t_\ell \wedge \tau_{[0, \infty[^n}) \in B_\ell). \end{aligned}$$

Proof Theorem 3.6.4 and considerations in Section 3.6.3 give that both the stopped SDE (3.6.8) and the stopped CLE (3.2.6) have unique strong solutions. The state space reduction discussed in Section 3.3.4 can be applied to Eq. (3.6.8) with matrix S and to the CLE (3.2.6) with matrix T . The conditions $S_0 b = 0$, $S_0 \sigma = 0$ and $T_0 \nu = 0$ imply that certain entries of the new variables, namely $S_0 \xi_0$ and $T_0 x_0$, will be constant. They still influence the propensity functions, but they can be treated as fixed parameters. Thereby both Eqs. (3.6.8) and (3.2.6) have been reduced to $(n - \dim(\text{Ker } \nu^T))$ -dimensional SDEs. To each corresponds a generalised martingale problem on the domain $D = T_1([0, \infty[^n) \subseteq \mathbb{R}^{n - \dim(\text{Ker } \nu^T)}$, and each has a solution provided by the corresponding strong solution. By Eqs. (3.6.10) and (3.6.11), the two generalised martingale problems are the same; its coefficient f equals Eq. (3.6.10), and B equals Eq. (3.6.11).

Now we check that the conditions of Theorems 3.6.6 and 3.6.7 are satisfied. The right-hand sides of Eqs. (3.6.10) and (3.6.11) show that the local boundedness, measurability and continuity conditions are met.

To check Condition (3.6.7), pick any $y \in \mathbb{R}^{n - \dim(\text{Ker } \nu^T)} \setminus \{0\}$ and consider the right-hand side of Eq. (3.6.11). The propensity a is evaluated in $T^{-1}(T_1([0, \infty[^n))_{T_0 x_0})$, that is, in a general element of $]0, \infty[^n$, so all its coordinates are positive. For the sake of simplicity, let them be denoted by $\tilde{a}_1, \dots, \tilde{a}_m$. Since T is invertible, the rows of T_1 are linearly independent. Therefore $y^T T_1 \neq 0$. Also, $\nu^T T_1^T y \neq 0$, because if $T_1^T y \in \text{Ker } \nu^T$ was the case, then $y^T T_1$ would be generated by the rows of T_0 (they form a basis in $\text{Ker } \nu^T$), which contradicts that the rows of T_1 and T_0 combined are linearly independent. From $\nu^T T_1^T y \neq 0$ it follows that $\text{diag}(\sqrt{\tilde{a}_1}, \dots, \sqrt{\tilde{a}_m}) \nu^T T_1^T y \neq 0$, because the $\tilde{a}_1, \dots, \tilde{a}_m$ are all positive. Thus

$$0 < \left| \text{diag}(\sqrt{\tilde{a}_1}, \dots, \sqrt{\tilde{a}_m}) \nu^T T_1^T y \right|^2 = y^T T_1 \nu \text{diag}(\tilde{a}_1, \dots, \tilde{a}_m) \nu^T T_1^T y,$$

so B is positive definite everywhere on $T_1([0, \infty[^n)$. Consequently, the two theorems

apply. In Theorem 3.6.7, we can choose $U = D_P = D_Q = T_1([0, \infty[^n)$. This gives that the reduced, stopped SDE (3.6.8) has the same finite-dimensional distributions as the reduced, stopped CLE (3.2.6).

The original stopped SDEs can be retrieved by appending $S_1\xi_0$ (or T_1x_0) to the reduced state vector of Eq. (3.6.8) (of Eq. (3.2.6), respectively) and applying S^{-1} (T^{-1} , respectively). ■

We required local Lipschitz and linear growth conditions for b and σ in order to ensure the existence of strong solutions. This is because the application we have in mind is the numerical simulation of trajectories. Without the linear growth condition, finite time explosion becomes possible. We could have substituted these conditions with the weaker local boundedness and continuity conditions from Theorem 3.6.6. In that case, the existence of a solution would have followed from Theorem 3.6.6 and the inverse linear transformation under S and T of the weak solution. To identify the solution on $D = T_1([0, \infty[^n)$, one can choose $D_k = T_1([\varepsilon_k, K_k[^n)$ for all $k \in \mathbb{N} \setminus \{0\}$ in Theorem 3.6.6, similarly to Section 3.6.3.

We proved our equivalence result for the CLE with mass action kinetics, but generalisations to wider classes of propensity functions a are possible.

We quickly demonstrate the use of Proposition 3.6.8 in practice by showing the equivalence of the four constructions developed earlier in this chapter. Recall that the variability in formulations hinged on how $g : [0, \infty[^n \rightarrow \mathbb{R}^{n \times d}$ in the SDE

$$dX(t) = \nu a(X(t)) dt + g(X(t)) dW(t) \quad (3.3.1)$$

was chosen under the condition $g(x)g(x)^T = \nu \text{diag}(a(x))\nu^T$. We sometimes used the

shorthand $A(x) = \text{diag}(a(x))$. Consider Constructions 1, 3 and 4:

$$g^1(x) = \nu \sqrt{\text{diag}(a(x))},$$

$$g^3(x) = W_1 \sqrt{W_1^T \nu A(x) \nu^T W_1},$$

$$g^4(x) = \nu J \sqrt{V A(x) V^T}.$$

Construction 1 is the original CLE (3.2.6), the reference to which Constructions 3 and 4 will be compared; the latter ones will be used as Eq. (3.6.8). Here $W = [W_1 \ W_0] \in \mathbb{R}^{n \times n}$ was an orthogonal matrix such that the columns of $W_0 \in \mathbb{R}^{n \times \dim(\text{Ker } \nu^T)}$ form an orthonormal basis in $\text{Ker } \nu^T$, and the columns of $W_1 \in \mathbb{R}^{n \times (n - \dim(\text{Ker } \nu^T))}$ are an orthonormal basis in $\text{Im } \nu$. We can pick these as the basis transformation matrices in all three cases. Note how simple this example is: all three processes live in the same n -dimensional space, the basis transformation is only required to meet the positive definiteness condition (3.6.7) in a lower dimensional space. With

$$T_1 = S_1 = W_1^T \text{ and } T_0 = S_0 = W_0^T,$$

$T_0 \nu = W_0^T \nu = 0$, $S_0 \nu a(x) = W_0^T \nu a(x) = 0$, and both $S_0 g^3(x) = W_0^T W_1 \sqrt{W_1^T \nu A(x) \nu^T W_1} = 0$ and $S_0 g^4(x) = 0$ are satisfied. Obviously, we choose $x_0 = \xi_0$, so Eq. (3.6.9) holds. Eq. (3.6.10) is satisfied because all our constructions use the same drift term. Eq. (3.6.11) follows from that g is always chosen to satisfy $g(x)g(x)^T = \nu \text{diag}(a(x))\nu^T$.

Therefore the conditions of Proposition 3.6.8 are all satisfied. The solution to the generalised martingale problem can be lifted up into the original space by multiplication with $T^{-1} = S^{-1} = W$ from the left. So the three SDE constructions, each stopped when it leaves the positive orthant, have the same finite-dimensional distributions.

The case of comparing Construction 2,

$$g^2(x) = U_1(x) D_1(x)^{1/2},$$

to Construction 1 calls for the application of

$$T_1 = S_1 = U_1(x)^T \text{ and } T_0 = S_0 = U_0(x)^T$$

as basis transformation matrices. However, these are now state dependent. $T_0\nu = U_0(x)^T\nu = 0$ due to Proposition 3.3.2, $S_0\nu a(x) = U_0(x)^T\nu a(x) = 0$, and $S_0g^2(x) = U_0(x)^TU_1(x)D_1(x)^{1/2} = 0$ all hold for any $x \in]0, \infty[^n$. Eqs. (3.6.9), (3.6.10) and (3.6.11) hold as in the case of Constructions 3 and 4. It is also important to verify that the latter two equations provide valid coefficients for the generalised martingale problem. Local boundedness, measurability and continuity (in Theorem 3.6.6) follow from the well-known continuity of the eigendecomposition (that gives $U(x)$ and $D(x)$) as a function of the matrix. Hence Proposition 3.6.8 applies.

It is worth mentioning here that the argument of Section 3.6.3 about pathwise existence is directly applicable to Construction 4 as well. For Constructions 2 and 3, one also needs that taking the square root via the eigendecomposition of the diffusion matrix $\nu \text{diag}(a(x))\nu^T$ (or of $W_1^T\nu A(x)\nu^TW_1$, respectively) is not only a continuous but a locally Lipschitz function of the state. Since the propensity functions are polynomials with degree not greater than 2, the conditions of Theorem 6.1.2 of Friedman (1975) or Theorem 5.2.3 of Stroock and Varadhan (1997) are satisfied in $]0, \infty[^n$ and these give the locally Lipschitz property as required.

We conclude by reiterating that the greatest deficiency of the CLE stems from the fact that its diffusion coefficient does not vanish on the boundary — trajectories may leave the nonnegative orthant with positive probability. There may exist a natural, consistent way to define a boundary behaviour that prevents the state from leaving the nonnegative orthant (e.g. reflection in a prespecified direction). Still, this fundamental problem may suggest that the entire CLE model needs to be reconsidered and a new SDE approximation to the discrete Markov jump process is needed for a satisfactory

continuous stochastic reaction kinetics model.

Chapter 4

Summary of the results of Part I

In Part I, we have given a thorough analysis of the form of the chemical Langevin equation (CLE). We proved that the finite-dimensional distribution of the Itô diffusion process described by the CLE does not change if the diffusion term $g(x) = \nu \sqrt{\text{diag}(a(x))}$ of the standard form is replaced by another one as long as it satisfies (3.2.5), $g(x)g(x)^T = \nu \text{diag}(a(x))\nu^T$. We explored different ways how the factorisation of the right-hand side can be carried out.

Via the combination of the factorisation with the minimum number of columns in g (Section 3.3.2) with a state space reduction by the removal of dependent variables (Section 3.3.4), we showed that the CLE can be given in a form where there are as many independent Wiener processes as there are linearly independent variables. This is also the number of independent stoichiometric vectors, that is, the dimension of the linear space spanned by the update vectors, rank ν . Neither the number of variables nor the number of Wiener processes can be further reduced without loss of accuracy of the system description. The number of Wiener processes can be reduced independently from the state space reduction: one can carry out just one or the other or both. Indeed, the

state space reduction was achieved by multiplying the state x and the functions f and g by a matrix T from the left and by compensating for this by ‘fitting’ both f and g with an ‘input converter’ T^{-1} , which are ‘external’ changes. The reduction in the number of Wiener processes relies on what happens with the entirety of $g(x)g(x)^T$, that is, on the internal structure of $g(x)$.

Due to the stoichiometric constraints, after its release from the initial state, a chemical reaction model can only move within a $(\text{rank } \nu)$ -dimensional affine subspace of the full n -dimensional state space. The state space reduction reflects this geometric constraint. One of our contributions is that we found that this many, $\text{rank } \nu$, independent Wiener processes are sufficient (and also necessary) to describe the distribution of the process given by the CLE at any time instance t . The eigendecomposition used for Construction 2 is apparently just an algorithm to disentangle the directions (locally) in which the m Wiener processes of the standard CLE fluctuate to $\text{rank } \nu$ orthogonal directions.

Another contribution of this work is that we showed using the same mathematical framework that in the case where there are m_1 pairs of reversible chemical reactions and m_2 irreversible reactions there is another, transparent formulation of the CLE with only $m_1 + m_2$ Wiener processes, whereas the standard form uses $2m_1 + m_2$. This new form can be seen to be more intuitive than the standard approach. At the heart of this construction is a transformation in which the two Gaussian noise processes that correspond to the two directions of a reversible reaction are replaced by a single one with variance equal to the sum of the two variances.

It is important to distinguish our system-size reduction methods from model reduction techniques, such as time scale separation applied to multiscale systems (Sotiropoulos *et al.*, 2009; Kang *et al.*, 2010). Ours are not approximations but transformations of the

CLE which avoid information loss about the statistical properties of the CLE model of the kinetics of the chemical system. In computer jargon, ours are ‘*lossless compression*’ methods (of the information content of the diffusion matrix $B(x)$) as opposed to model reduction approaches that are ‘*lossy*’.

We illustrated these ideas by considering alternative forms of the CLE for a HERG ion channel model and the Goldbeter–Koshland switch. We showed that considerable savings in running time can be achieved when using the reduced form of the CLE for numerical simulation. We believe that all software implementations of the CLE should include this reduced form. This would only require a small change in code and would accelerate simulation without changing the statistical properties of the generated stochastic process. The CLE is an important tool for the analysis and simulation of multiscale chemical reaction systems and it is vital to choose its most appropriate or most efficient formulation according to the requirements of the application.

The problem with the boundary behaviour of the CLE underscores the importance of using the CLE only in the appropriate regime, where all species are sufficiently abundant, and the state is far from the boundary. The genuine solution would be the development of a nonnegativity preserving SDE model for chemical reaction kinetics.

Chapter A

Appendix. Proofs deferred from Chapter 3

A.1 Proof of Proposition 3.2.1

Proof For ease of notation we will drop the time variable t from $x(t)$. We apply the multidimensional Itô's formula. This states that when substituting time t and a diffusion process $x(t)$ into a function $u(t, y) : \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R}$,

$$du(t, x) = \frac{\partial u(t, x)}{\partial t} dt + \sum_{i=1}^n \frac{\partial u(t, x)}{\partial y_i} dx_i + \frac{1}{2} \sum_{i,k=1}^n \frac{\partial^2 u(t, x)}{\partial y_i \partial y_k} dx_i dx_k$$

holds, where the rules for computing $dx_i dx_k$ are $dt dt = dt dW_{j,t} = dW_{j,t} dt = 0$, $dW_{j,t} dW_{j',t} = \delta_{jj'} dt$ (for all $j, j' \in \{1, \dots, d\}$, $\delta_{jj'}$ is the Kronecker delta) (Øksendal, 2007), where these are approximations that hold in an $o(dt)$ sense. Applying the formula

with $u(t, y) = y_i y_k$ gives

$$\begin{aligned}
d(x_i x_k) &= 0 + (x_k dx_i + x_i dx_k) + \frac{1}{2} (dx_i dx_k + dx_k dx_i) \\
&= (x_k dx_i + x_i dx_k) \\
&\quad + \frac{1}{2} 2 \left(f_i(x) dt + \sum_{j=1}^d g_{ij}(x) dW_j(t) \right) \\
&\quad \times \left(f_k(x) dt + \sum_{j'=1}^d g_{kj'}(x) dW_{j'}(t) \right) \\
&= (x_k dx_i + x_i dx_k) \\
&\quad + \sum_{j=1}^d \sum_{j'=1}^d g_{ij}(x) g_{kj'}(x) dW_j(t) dW_{j'}(t) \\
&= \left(x_k f_i(x) dt + x_k \sum_{j=1}^d g_{ij}(x) dW_j(t) \right. \\
&\quad \left. + x_i f_k(x) dt + x_i \sum_{j=1}^d g_{kj}(x) dW_j(t) \right) \\
&\quad + \sum_{j=1}^d g_{ij}(x) g_{kj}(x) dt.
\end{aligned}$$

Taking the expectation on both sides yields

$$dE(x_i x_k) = E(x_k f_i(x)) dt + E(x_i f_k(x)) dt + \sum_{j=1}^d E(g_{ij}(x) g_{kj}(x)) dt,$$

which is just another form of Eq. (3.2.3). ■

A.2 Proof of Proposition 3.3.2

Proof If $y \in \mathbb{R}^n \setminus \{0\}$ is a left nullvector of ν , $y^T \nu = 0$, then it is trivially a nullvector of $B(x)$:

$$B(x)y = \nu \operatorname{diag}(a(x)) \nu^T y = 0.$$

In the other direction, if $B(x)y = 0$, we first prove that y is a left nullvector of the factor in Gillespie's factorisation,

$$\sigma(x) = \nu \operatorname{diag}(\sqrt{a_1(x)}, \dots, \sqrt{a_m(x)}).$$

Indeed, $0 = B(x)y = \sigma(x)\sigma(x)^T y$, hence $0 = y^T \sigma(x)\sigma(x)^T y = |\sigma(x)^T y|^2 = |y^T \sigma(x)|^2$, so $y^T \sigma(x) = 0$ and y is a left nullvector to $\sigma(x)$.

The left nullvectors of $\sigma(x)$ and ν are the same, since $\sqrt{a_1(x)}, \dots, \sqrt{a_m(x)}$ are all positive by assumption. Therefore y is a left nullvector of ν , as claimed. ■

A.3 Proof of Proposition 3.3.3

Proof The column rank of $\nu \in \mathbb{R}^{n \times m}$ is just $\dim(\text{Im } \nu)$. It is well known that

$$\dim(\text{Im } \nu) + \dim(\text{Ker } \nu) = m.$$

The row rank of ν is the column rank of ν^T , or $\dim(\text{Im } \nu^T)$. Similarly,

$$\dim(\text{Im } \nu^T) + \dim(\text{Ker } \nu^T) = n.$$

It is also well known that the column and row ranks are always equal. Therefore

$$m - \dim(\text{Ker } \nu) = \text{rank } \nu = n - \dim(\text{Ker } \nu^T). \quad \blacksquare$$

A.4 Proof of Proposition 3.3.5

Proof By its definition, the columns of

$$\nu_b \in \mathbb{R}^{n \times (m - \dim(\text{Ker } \nu))} = \mathbb{R}^{n \times (n - \dim(\text{Ker } \nu^T))}$$

are linearly independent. Using this, we will prove that the kernel of

$$\nu_b^T \nu_b \in \mathbb{R}^{(n - \dim(\text{Ker } \nu^T)) \times (n - \dim(\text{Ker } \nu^T))}$$

is trivial, hence $\nu_b^T \nu_b$ is full rank.

Let $y \in \mathbb{R}^{n - \dim(\text{Ker } \nu^T)} \setminus \{0\}$ be a nullvector of $\nu_b^T \nu_b$, $\nu_b^T \nu_b y = 0$. Then $0 = y^T \nu_b^T \nu_b y = |\nu_b y|^2$, so $\nu_b y = 0$. By the linear independence of the columns of ν_b , y is zero. ■

A.5 State space reduction for Construction 4

For Construction 4, a finer partitioning of matrices ν, J, V is proposed. Let us order the columns of $\nu \in \mathbb{R}^{n \times m}$ such that $\nu = [\nu_1 \ \nu_2 \ \nu_3 \ \nu_4]$, where the columns of $\nu_1 \in \mathbb{R}^{n \times (m - \dim(\text{Ker } \nu))}$ form a basis for $\text{Im } \nu$; ν_3 is the collection of the column vectors that are constant multiples of any single column of ν_1 ; the columns of ν_2 represent all the directions specified by columns of ν that are distinct to directions of the columns of ν_1 (columns of ν_2 are linearly dependent on columns of ν_1 , they are a linear combination of more than one); and finally ν_4 is the collection of the column vectors that are constant multiples of any single column of ν_2 . Let the sizes of these matrices define r_2, r_3 and r_4 such that $\nu_2 \in \mathbb{R}^{n \times r_2}, \nu_3 \in \mathbb{R}^{n \times r_3}, \nu_4 \in \mathbb{R}^{n \times r_4}$. Obviously, $m - \dim(\text{Ker } \nu) + r_2 = s$, and $r_2 + r_3 + r_4 = \dim(\text{Ker } \nu)$. The entries of $A(x)$ are permuted accordingly, and then $A(x)$ is partitioned into blocks.

Such a choice specifies the matrices

$$\begin{aligned} R &\in \mathbb{R}^{(m - \dim(\text{Ker } \nu)) \times r_2}, \\ M_3 &= [v^{(1)} \ \dots \ v^{(r_3)}] \in \mathbb{R}^{(m - \dim(\text{Ker } \nu)) \times r_3}, \\ M_4 &= [w^{(1)} \ \dots \ w^{(r_4)}] \in \mathbb{R}^{r_2 \times r_4}, \end{aligned}$$

such that $\nu_2 = \nu_1 R, \nu_3 = \nu_1 M_3, \nu_4 = \nu_2 M_4$, and all $v^{(i)}$ and $w^{(k)}$ have only one nonzero

entry each. Then let

$$J = \begin{bmatrix} I_{m-\dim(\text{Ker } \nu)} & 0 \\ 0 & I_{r_2} \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \in \mathbb{R}^{m \times s},$$

$$V = \begin{bmatrix} I_{m-\dim(\text{Ker } \nu)} & 0 & M_3 & 0 \\ 0 & I_{r_2} & 0 & M_4 \end{bmatrix} \in \mathbb{R}^{s \times m},$$

J having first r_3 then r_4 rows of zeros.

The construction is again as in Eq. (3.3.2). For the sake of notational clarity, let

$$C_1(x) = A_1(x) + \sum_{j=1}^{r_3} (A_3(x))_{jj} v^{(j)} v^{(j)T} \in \mathbb{R}^{(m-\dim(\text{Ker } \nu)) \times (m-\dim(\text{Ker } \nu))},$$

$$C_2(x) = A_2(x) + \sum_{j=1}^{r_4} (A_4(x))_{jj} w^{(j)} w^{(j)T} \in \mathbb{R}^{r_2 \times r_2}.$$

Then

$$VA(x)V^T = \begin{bmatrix} C_1(x) & 0 \\ 0 & C_2(x) \end{bmatrix}$$

is diagonal, $\nu JV = \nu$ and Eq. (3.2.5) hold. Defining T with ν_1 in the role of ν_b ,

$$g(x) = \nu_1 \begin{bmatrix} \sqrt{C_1(x)} & R\sqrt{C_2(x)} \end{bmatrix},$$

$$Tg(x) = \begin{bmatrix} \sqrt{C_1(x)} & R\sqrt{C_2(x)} \\ 0 & 0 \end{bmatrix},$$

whose nonzero blocks together are in $\mathbb{R}^{(m-\dim(\text{Ker } \nu)) \times s}$, as required.

Part II

Optimal experiment design for discrimination between rival biochemical models

Chapter 5

Motivation for the mathematical design of experiments

In Part I our focus was on the development, justification and understanding of a certain mathematical modelling framework for biochemical reaction kinetics. In Part II our attention will turn to choosing the best model for a particular biological system within a simpler, fixed modelling framework, the deterministic reaction rate equation. We note here that the methods we will develop apply to other nonlinear ODE models too.

A successful modelling effort is necessarily an iteration between changing the model and testing the new one by experimentation. Testing the accuracy and fitness for purpose of a mathematical description of any physical process should be done by confronting it with experimental data. At the same time, models should inform the design of new experiments.

Traditionally, experiments have been designed using heuristic approaches: experience, intuition, or simple causal analyses. Evidently, such heuristically designed experiments are not always maximally informative, a great impediment given the labour, mon-

etary cost and effort expended on the implementation of experiments or the development of new measurement techniques. As a result, it is becoming increasingly necessary to systematically design more informative experiments, in order for the iterative modelling process to result in reliable models at the lowest possible cost.

Our view on the computational modelling of biological networks is that the iterative modelling procedure should have three stages: model identification (model fitting), model discrimination (in which a new experiment is designed) and model invalidation (using the new experimental data). Note that we talk about invalidation and not about validation; we will explain this shortly. The new experiment may give new insights that can lead to the modification of the best model, thus closing the flow-chart of this iterative process into a cycle. All three tasks present serious challenges, and remain important areas of research.

5.1 Previous work on optimal experiment design for biochemical reaction systems

5.1.1 Experiment design for system identification

To date, the majority of studies addressing experiment design for biological reaction systems have considered system identification. In this area, experiments are designed such that the resulting data are most informative about model structure or parameter values — see, for example, Barrett and Palsson (2006) and Feng and Rabitz (2004); Gunawan, Gadkar, and Doyle III (2006); Bandara *et al.* (2009), respectively. Several groups studied statistically orientated frameworks for optimal structure identification (Casey *et al.*, 2007) or for parameter identification (Casey *et al.*, 2007; Balsa-Canto

et al., 2008). These approaches aim to find the weighted least squares of differences between data and model prediction and make use of the *Fisher information matrix* and the associated notions of *A*-, *D*- and *E-optimality*. In this framework, Yue *et al.* (2008) examined optimally designed parameter estimation methods that are robust to model uncertainties (*robust experiment design*).

5.1.2 Experiment design for model discrimination

In numerous practical situations, accumulated biological knowledge about the system of interest can constrain the set of plausible model structures. In such a case, several different model structures may be proposed based on corresponding biological hypotheses. Experiment design in this context aims for the effective discrimination between these well-defined alternative models. In particular, several mathematical models of similar complexity, each corresponding to a different network topology, can fit the available set of experimental data and describe the behaviour of this system within error bounds reflecting uncertainty in experimental environment and inaccuracies of measurements. For instance, network determination algorithms (such as the ones developed by August and Papachristodoulou, 2009a) yield such alternative models. Discriminatory experiment design can then be used to differentiate between the models.

Various aspects of model discriminatory experiment design have been addressed in the literature. Bardsley, Wood, and Melikhova (1996) investigated the problem of how measurements should be spaced in time to perform an optimally discriminating experiment between two models, and how many of them are required. More specifically, they compared different patterns of measurement spacings (geometric versus uniform spacing of points in time, or uniform spacing on the axis that corresponds to the state). Chen and Asprey (2003) developed statistical approaches to parameter estimation, the assess-

ment of model fit and model discrimination, assuming that the response variables are uncertain. In this framework, model discrimination is based on a Bayesian approach, which assigns prior probabilities to each model, updates these after each experiment and chooses the model with the likelihood that has become sufficiently large compared to others. An alternative frequentist method uses repeated hypothesis tests to reject models one by one. Donckels *et al.* (2009) separated the uncertainty of the model predictions and the uncertainty of the measurements and used these to design the next experiment such that it is most informative. As opposed to the traditional approach, here the expected information content of the newly designed experiment is also taken into account (*anticipatory design*) in order to assess the uncertainties more accurately. Kreutz and Timmer (2009) gave a review of approaches to parameter estimation and model discrimination (discussing the Akaike Information Criterion, the likelihood ratio test and alternative forms of the sum of squared differences between two models' outputs). They also discussed relevant classical statistical aspects of experiment design, such as randomisation, replication and confounding.

Tidor and co-workers (Apgar *et al.*, 2008) developed dynamic model-based controllers that drive the output along a prescribed target trajectory (usually a constant output). If such a control input signal achieves the required output trajectory in an experiment, then the model is more accurate than another model which gives a different output trajectory for this particular input. Kremling *et al.* (2004) presented three methods for optimal pairwise discriminating experiment design, and compared them on a test example. Their first method compares combinations of certain initial input levels and subsequent changes in input in order to determine which combination will lead to the largest difference in the outputs. Their second method replaces models with their linearised counterparts in order to find a sinusoidal input with a frequency that maximises

the difference between phase shifts of the two models. Their third method follows the work by Chen and Asprey (2003), and aims to find an input profile that brings the output responses of the two models as far apart as possible. The distance is measured by a weighted objective function. The weighting is set up such that if the measurement error of an output variable is large, then the difference of these outputs contributes less to the weighted objective function. The authors concluded that the most appropriate method strongly depends on the possible ways to stimulate the system and the quality of the measurements.

5.1.3 Closing the cycle: model invalidation

The findings of model discriminatory experiments feed into model invalidation procedures (Anderson and Papachristodoulou, 2009). The experiments need to be carefully designed and implemented to produce new data that show system behaviour that cannot be represented by a particular model. This can be used to invalidate a seemingly good but incorrect model.

Note that, logically, one can never validate a model. At best, a model will be capable of explaining all available data and can be tested against some of its predictions. Therefore, narrowing down on the correct model can only be done from the other direction through invalidation, by systematically ruling out incorrect models.

5.2 Problem formulation

5.2.1 Overview

Our work addresses the problem of model discrimination. We present a framework for defining and designing optimally discriminating experiments, that is, experiments that

are the best (in some mathematically defined but practically meaningful way) at discriminating between rival models.

There are cases when it is difficult or even impossible to distinguish between rival models due to the incomplete observability of their internal states. Tests exist to identify such cases (August and Papachristodoulou, 2009b). Even when model discrimination is possible, it can be expected to be difficult as the starting assumption is that the rival models both describe all available data well.

Our key principle is to maximise the difference between the outputs of two different models, in particular, the L_2 -norm of the output difference. Although similar in principle, our investigation follows a direction distinctly different from and more practical than the work by Chen and Asprey (2003): we use deterministic models that do not take account of measurement noise directly. Instead, we try to make the outputs of the two models as distant as possible to ensure that even a noisy measurement has a good chance of discriminating between them.

We discuss three approaches to achieve this goal. In the first approach, the *Initial condition design for model discrimination*, we find the initial state of the system which results in the most discriminating output between the two examined models. This is a result by Papachristodoulou and El-Samad (2007).

The second method, *Design of structural changes for model discrimination*, combines optimal initial condition choice with optimal systemic modifications. The latter reflects the assumption that in the experiment it is possible, for example, to up- or downregulate the expression of certain genes, either through genetic manipulations or other techniques, such as RNA interference (RNAi) technology. The gene product may be an enzyme whose concentration is not explicitly modelled but is reflected in a chemical rate constant, or some protein which exists in (possibly various) phosphorylated and dephosphorylated

forms such that the sum of their concentrations is constant. In our mathematical model this means a free choice in some parameter values within given intervals.

The third method, *Input design for model discrimination*, is a novel one. It assumes the possibility for external stimulation during the experiment and searches for the best such stimulus from a set of allowable stimuli. This approach is reminiscent of, but different from, the second method in Kremling *et al.* (2004) — there, the difference between phase shifts is maximised, whereas in our method the difference between amplitudes is maximised.

In all three approaches, the problems are cast in an optimisation framework and the *sum of squares* (SOS) technique (Parrilo, 2000) is used for the experiment design, allowing us to treat the nonlinear system descriptions directly. The theoretical results are demonstrated by the application of each method to a discrimination problem for two models of signal processing for chemotaxis in *Dictyostelium discoideum* amoebæ.

5.2.2 Mathematical formalisation of the problem

In the following we consider different models describing the same biological system by a set of autonomous ODEs. This can be the reaction rate equation using mass action kinetics, but other formalisms are possible, like generalised mass action or Michaelis–Menten kinetics. In general, the i th model takes the form

$$\begin{aligned}\frac{dx_i(t)}{dt} &= f_i(x_i(t)) + g_i(x_i(t))u(t), \\ y_i(t) &= h_i(x_i(t)),\end{aligned}\tag{5.2.1}$$

where

$u: [0, \infty[\rightarrow \mathbb{R}^q$ denotes the input function,

$x_i: [0, \infty[\rightarrow \mathbb{R}^{n_i}$ the state, and

$y_i: [0, \infty[\rightarrow \mathbb{R}^\ell$ the output function.

The input functions (and the output functions, respectively) are of the same dimension for each model. Function g_i is matrix valued, with size $n_i \times q$. The structure of the functions f_i , g_i and h_i will depend on the modelling framework in use to describe the biological system, but we assume that all of them are smooth. We will use the shorthand $\dot{x}_i(t) := \frac{dx_i(t)}{dt}$ for the differential of a function with respect to time.

The output function represents measurements an experimenter obtains from the system, and the input function represents the stimuli or perturbations the experimenter introduces to the system during the experiment. For mathematical simplicity we assume that the input does not affect the output directly.

In this work our aim is to discriminate between two models of the form Eq. (5.2.1) ($i \in \{1, 2\}$), which have n_1 and n_2 state variables, respectively. As these two models represent the same underlying biological system, we require that they both admit the same steady states and fit already available experimental data. Our aim is to design the next experiment that will facilitate their discrimination. A natural way to formulate the discrimination problem is to concatenate the two models and generate the difference between their outputs: with

$$x(t) = \begin{pmatrix} x_1(t) \\ x_2(t) \end{pmatrix}, \quad f(x) = \begin{pmatrix} f_1(x_1) \\ f_2(x_2) \end{pmatrix}, \quad g(x) = \begin{bmatrix} g_1(x_1) \\ g_2(x_2) \end{bmatrix},$$

$$h(x) = h_1(x_1) - h_2(x_2), \quad y(t) = y_1(t) - y_2(t),$$

the concatenated system is

$$\begin{aligned}\dot{x}(t) &= f(x(t)) + g(x(t))u(t), \\ y(t) &= h(x(t)).\end{aligned}\tag{5.2.2}$$

We call an experiment *optimal*, if the difference between the outputs of the two models ($y_1 - y_2$) is maximal over a set of experimental perturbations of bounded ‘size’. In technical terms, we aim to pick the best point in a set of allowable perturbations of the initial state conditions (*Initial condition design for model discrimination*), the set of some admissible parameter changes and the set of common initial conditions (*Design of structural changes for model discrimination*), or the set of allowable inputs u (*Input design for model discrimination*) in order to maximise the L_2 -distance between the outputs of the two rival models:

$$\|y\|_2 = \|y_1 - y_2\|_2 = \left(\int_0^\infty \sum_{k=1}^\ell (y_{1k}(t) - y_{2k}(t))^2 dt \right)^{\frac{1}{2}}.$$

To help interpretation, we implement a change of coordinates that places the investigated steady state to zero in both models. We assume that the outputs are identical in this common steady state, now the origin: $h_1(0) = h_2(0)$. Throughout this work it is also assumed that the examined steady state is asymptotically stable in both models in Eq. (5.2.2).

Since experiments must be implemented in finite time, we require that the designed input u has compact support. For convenience, we sometimes relax this requirement and only assume that u is ‘very small’ after a certain time. Clearly, since there is only one experimental setup in reality, the input u must be identical for the two models.

In the case of linear systems, the description of the concatenated system (5.2.2)

becomes

$$\begin{aligned}\dot{x}(t) &= Ax(t) + Bu(t), \\ y(t) &= Cx(t),\end{aligned}\tag{5.2.3}$$

with

$$\begin{aligned}x(t) &= \begin{pmatrix} x_1(t) \\ x_2(t) \end{pmatrix} \in \mathbb{R}^{n_1+n_2}, & A &= \begin{bmatrix} A_1 & 0 \\ 0 & A_2 \end{bmatrix} \in \mathbb{R}^{(n_1+n_2) \times (n_1+n_2)}, \\ B &= \begin{bmatrix} B_1 \\ B_2 \end{bmatrix} \in \mathbb{R}^{(n_1+n_2) \times q}, & C &= [C_1 \quad -C_2] \in \mathbb{R}^{\ell \times (n_1+n_2)}.\end{aligned}\tag{5.2.4}$$

We assume that all eigenvalues of both A_1 and A_2 have negative real parts (we call these matrices *Hurwitz*), thus they define asymptotically stable systems. This makes A Hurwitz too.

Chapter 6

Initial condition design and its extension to the design of optimal structural changes for model discrimination

6.1 Initial condition design for model discrimination

Many biological experiments drive a cellular system into an informative out-of-equilibrium state (e.g. heat shock, osmotic shock, chemical stimulus) and then glean information from the patterns of return to equilibrium in the absence of an input. In an optimisation formulation this amounts to searching for normalised initial condition $x_1(0) = x_2(0)$ for the two models between which one wishes to discriminate that maximises the output

difference $\|y\|_2$ — where y is defined in Eq. (5.2.2) — for the unforced system ($u = 0$). Here we assume that the two alternative model representations of the system are written in terms of the same chemical species, thus $n_1 = n_2 =: n$.

6.1.1 Linear case

If $x_1(0)$ is not required to be equal to $x_2(0)$, then the solution can be borrowed from standard results in linear systems theory. In particular, the optimal direction for the initial value of Eq. (5.2.3) can be found by the following procedure.

1. Find a symmetric, positive semidefinite matrix $P \in \mathbb{R}^{2n \times 2n}$ that solves the so-called *observability Lyapunov equation*

$$A^T P + P A + C^T C = 0.$$

The solution P is called the *observability gramian* (Dullerud and Paganini, 2000).

It is known to satisfy

$$P = \int_0^\infty e^{A^T t} C^T C e^{A t} dt.$$

2. Find the normalised eigenvector $\bar{x} \in \mathbb{R}^{2n}$ ($|\bar{x}| = 1$) corresponding to the largest eigenvalue $\bar{\lambda}$ of P , that is, for $\bar{\lambda}$ find $\bar{x}$ such that

$$P \bar{x} = \bar{\lambda} \bar{x}.$$

Indeed, the direction $x(0) = \bar{x}$ gives the maximum output L_2 -norm (output *energy*), since the output is $y(t) = C e^{A t} x(0)$ and the output energy is given by

$$\|y\|_2^2 = x(0)^T P x(0). \tag{6.1.1}$$

However, this computation is not satisfactory since an experimentally meaningful initial condition should satisfy $x_1(0) = x_2(0) = \hat{x} \in \mathbb{R}^n$. To enforce this condition, we can partition P into blocks of size $n \times n$,

$$P = \begin{bmatrix} P_{11} & P_{12} \\ P_{12}^T & P_{22} \end{bmatrix}.$$

With this decomposition, the optimal initial state is the unit norm eigenvector $\hat{x}$ corresponding to the largest eigenvalue of the matrix

$$R = P_{11} + P_{22} + P_{12} + P_{12}^T.$$

To see this, substitute $x(0) = \begin{pmatrix} \hat{x} \\ \hat{x} \end{pmatrix} \in \mathbb{R}^{2n}$ in Eq. (6.1.1) to get:

$$\|y\|_2^2 = \hat{x}^T R \hat{x}.$$

Hence $\|y\|_2^2$ is maximised exactly when $\hat{x}$ is the eigenvector corresponding to the largest eigenvalue of R .

6.1.2 Nonlinear case

The ideas behind model discrimination in the linear case can be generalised for application to nonlinear systems. However, we cannot explicitly compute the exact difference in the outputs of the two rival models $\|y\|_2$. Our approach avoids simulations and concentrates on finding an upper bound on $\|y\|_2$ using *storage functions* (Willems, 1972a,b) and *sum of squares* algorithmic relaxations of the resulting optimisation problem.

To determine an upper bound on $\|y\|_2$ for the nonlinear system (5.2.2), suppose there

exists a continuously differentiable function $S : \mathbb{R}^{2n} \rightarrow \mathbb{R}$ satisfying

$$S \geq 0, \quad S(0) = 0, \quad \text{and} \quad (6.1.2)$$

$$\text{for all } x \in D, \quad -\frac{\partial S(x)}{\partial x} f(x) - h(x)^T h(x) \geq 0, \quad (6.1.3)$$

where D is a neighbourhood of the steady state defined by

$$D = \left\{ \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \in \mathbb{R}^{2n} \mid |x_1| \leq \alpha, |x_2| \leq \alpha \right\} \quad (6.1.4)$$

for some $\alpha > 0$. Here we assume that D does not include states which are not physically meaningful, and the value of α will ensure this. This implies that the system is *dissipative* with supply rate $-h(x)^T h(x)$. Suppose that the system is released from an initial state $x(0)$ inside the largest level set of S that fits into D , so that $|x_1(0)| = |x_2(0)| = \beta \leq \alpha$. (Note that here in the nonlinear case assuming that the initial state is normalised to unit length would restrict generality.) In this case, integrating condition (6.1.3) and using $\frac{\partial S}{\partial x} f(x) = \frac{dS}{dt}$, we get

$$\int_0^T h(x(t))^T h(x(t)) dt \leq S(x(0)) - S(x(T)).$$

If we let $T \rightarrow \infty$, then

$$\begin{aligned} \|y\|_2^2 &= \int_0^\infty h(x(t))^T h(x(t)) dt \\ &\leq S(x(0)) - \lim_{T \rightarrow \infty} S(x(T)) \leq S(x(0)), \end{aligned} \quad (6.1.5)$$

by the nonnegativity of S . This implies that $\|y\|_2^2 \leq S(x(0))$, since condition (6.1.3) is valid within the whole region D and level sets of S are invariant. Hence we have found a way to bound $\|y\|_2^2$, which requires constructing the function S .

It is worth noting that the result from the linear and nonlinear cases have a similar purpose. Whereas in the linear case the result is rooted in a Lyapunov equality and

provides optimal solution, in the general nonlinear case one has to be content with an estimate given by inequality (6.1.5).

A condition missing from the above construction is that the two system models must be released from the same initial state. Hence, following from our discussion in the linear case, one has to construct an appropriately modified S , $\hat{S}(\hat{x}) := S(\begin{pmatrix} \hat{x} \\ \hat{x} \end{pmatrix})$ ($\hat{x} \in \mathbb{R}^n$). In the linear case, the desired initial conditions correspond to those that maximise the quadratic form $\hat{x}^T R \hat{x}$. That is, the optimal direction was that of the eigenvector corresponding to the largest eigenvalue of matrix R . This is also exactly the direction corresponding to the smallest semi-axis of the ellipsoid $\hat{x}^T R \hat{x} = r$ for some $r > 0$. In other words, we were looking for

$$\inf \left\{ \gamma > 0 \mid \{ \hat{x} \in \mathbb{R}^n \mid |\hat{x}| = \beta, \hat{x}^T R \hat{x} \geq \gamma \} = \emptyset \right\}.$$

Similarly to the linear case, in the nonlinear case we will also use a geometric argument to achieve initial condition design. Here, $\hat{S}(\hat{x})$ plays the role of the quadratic form $\hat{x}^T R \hat{x}$, and one can now decrease $\gamma > 0$ from infinity until the shrinking level set $\{ \begin{pmatrix} \hat{x} \\ \hat{x} \end{pmatrix} \in \mathbb{R}^{2n} \mid \hat{S}(\hat{x}) = \gamma \}$ touches $D' = \{ \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \in \mathbb{R}^{2n} \mid |x_1| \leq \beta, |x_2| \leq \beta \}$. Therefore, we need to solve the following optimisation problem:

$$\begin{aligned} &\text{minimise} && \gamma, \\ &\text{such that} && \text{for all } \hat{x} \in \mathbb{R}^n, |\hat{x}| = \beta : \\ &&& \hat{S}(\hat{x}) - \gamma \leq 0. \end{aligned} \tag{6.1.6}$$

The sequence of results presented so far asserts that the presence of a function S with the properties delineated above provides an upper bound on the difference in output energies of two rival models.

This information can be exploited to generate experimental initial conditions that drive the system towards this bound. These methods, however, do not prescribe how

one should go about finding such a function. Constructing a nonnegative function is in general a difficult problem. However, advances in the theory of sum of squares provide a computationally tractable way to relax this problem (Parrilo, 2000). In a nutshell, instead of searching for a general nonnegative function, we can constrain our search to functions that can be parameterised as sums of squares of polynomials. Within this class, the problem can be solved through semidefinite programming, with a polynomial time algorithm (see Appendix B, Sections B.1 and B.2).

Therefore, our strategy to find a near optimal initial state for the nonlinear model discrimination is a two-step process. First, we construct an SOS function S that satisfies conditions (6.1.2) and (6.1.3). In the second step, we search for the direction in which $\hat{S}$ is maximal, that is, we solve the optimisation problem (6.1.6). The details are given in Appendix B, in Section B.2.

6.2 Design of structural changes for model discrimination

A class of experiments is based on the introduction of internal changes (genetic and biochemical manipulations) to the system. To mirror such experiments, we develop a methodology to choose numerical changes of parameters in the system that maximise the difference between the outputs of two rival models of its internal structure. Since two such models are different, they do not necessarily have the same number of parameters. Therefore the design concentrates on the parameters that the models have in common, which we denote by p_j ($j \in \{1, \dots, m\}$). We assume that their values can be chosen within closed intervals $[a_j, b_j]$ (where $b_j \geq a_j \geq 0$ for all j), that is, $p \in \Pi$, where $\Pi := \prod_{j=1}^m [a_j, b_j] \subset \mathbb{R}^m$. We rewrite Eq. (5.2.2) to underline the dependence of the

model on those parameters as

$$\dot{x}(t) = f(x(t), p) + g(x(t), p)u(t),$$

$$y(t) = h(x(t)),$$

with

$$x(t) = \begin{pmatrix} x_1(t) \\ x_2(t) \end{pmatrix}, \quad f(x, p) = \begin{pmatrix} f_1(x_1, p) \\ f_2(x_2, p) \end{pmatrix}, \quad g(x, p) = \begin{bmatrix} g_1(x_1, p) \\ g_2(x_2, p) \end{bmatrix},$$

$$h(x) = h_1(x_1) - h_2(x_2), \quad y(t) = y_1(t) - y_2(t),$$

and again assume that $n_1 = n_2 = n$, $u = 0$, $f_1(0, p) = f_2(0, p) = 0$ (for every $p \in \Pi$),

$h_1(0) = h_2(0)$ and let

$$D = \left\{ \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \in \mathbb{R}^{2n} \mid |x_1| \leq \alpha, |x_2| \leq \alpha \right\}.$$

The steady state of either model may change with changing parameter p . Therefore the assumption $f_1(0, p) = f_2(0, p) = 0$ should be interpreted as a change of coordinates that shifts the steady state of each model to the origin individually for each p . We are not interested in how far the two equilibria shift *per se*, which is an algebraic problem, instead we are interested in the difference in their dynamic responses. This would reflect a situation in which a change in parameters would not be reflected in a significant change in the steady state but which could result in a substantial difference in the dynamics of the system.

As with the previous method, our methodology will rely on the construction of an appropriate function S that sets an upper bound on the difference between the outputs of the two models, followed by a computationally efficient formulation for the construction of this function using SOS.

For the above system, suppose that there exists a function $S : \mathbb{R}^{2n} \times \Pi \rightarrow \mathbb{R}$ which is sufficiently smooth and satisfies

$$\begin{aligned} S &\geq 0, \quad \text{for all } p \in \Pi, \quad S(0, p) = 0, \quad \text{and} \\ \text{for } (x, p) &\in D \times \Pi, \quad -\frac{\partial S(x, p)}{\partial x} f(x, p) - h(x)^T h(x) \geq 0. \end{aligned} \quad (6.2.1)$$

Then

$$\begin{aligned} \|y\|_2^2 &= \int_0^\infty h(x(t))^T h(x(t)) \, dt \\ &\leq S(x(0), p) - \lim_{T \rightarrow \infty} S(x(T), p) \leq S(x(0), p) \end{aligned} \quad (6.2.2)$$

if the system is released from an initial state $(x(0), p) \in D \times \Pi$ where $x(0)$ is in a level set of S entirely contained in D , $|x_1(0)| = |x_2(0)| = \beta \leq \alpha$. The last inequality in (6.2.2) holds since $S(x, p) \geq 0$.

The computational relaxation and implementation of the search for the function S is presented in Appendix B, Section B.3. Once this function has been constructed, one can extract the optimal point $\hat{x}$ and parameter point $\hat{p}$ that maximises the difference between the measured outputs of the two models.

Chapter 7

Input design for model discrimination

A powerful approach to discriminate between two plausible models of a molecular biological system is to design an experimental input that maximally differentiates between the dynamical behaviours of the outputs of the two models. If this input generates qualitatively different patterns in the model outputs, then one can subject the actual physical system to this designed input and then eliminate the model which differs from this pattern. We study this general form of the optimally discriminating input problem:

$$\text{Given} \quad \begin{pmatrix} \dot{x}_1(t) \\ \dot{x}_2(t) \end{pmatrix} = \begin{pmatrix} f_1(x_1(t)) \\ f_2(x_2(t)) \end{pmatrix} + \begin{bmatrix} g_1(x_1(t)) \\ g_2(x_2(t)) \end{bmatrix} u(t),$$

$$y = h_1(x_1(t)) - h_2(x_2(t)),$$

$$x_1(0) = x_2(0) = 0,$$

$$\text{find} \quad \text{the input } u, \quad \|u\|_2 \leq 1,$$

$$\text{such that} \quad \|y\|_2 \text{ is maximal.}$$

The *system gain* (or *L_2 -gain*) is the quantity given by

$$\sup\{\|y\|_2 \mid \|u\|_2 \leq 1\} = \sup\{\|y\|_2 \mid \|u\|_2 = 1\}.$$

We assume that the input (or in our numerical example, the perturbation of the input from a basal value) is of unit L_2 -norm (in other words, is of unit energy) at most. This assumption can be made without loss of generality as one can scale the equations accordingly, depending on the amount of input (ligand) available and the constraints of the system under study. Our goal is to maximise the difference between the two model outputs over a transient period after application of the new input. Recall that the two models describe currently available data equally well. Hence it can be assumed that for the same basal input they have the same pre-stimulus steady states and the same outputs.

Solving this optimisation problem in order to generate the maximally informative input is computationally challenging. In fact, even the first order condition of optimality is a $2(n_1 + n_2)$ -variable ODE with boundary conditions at both ends of the time interval (for $n_1 + n_2$ variables the boundary conditions are at the beginning of the interval, for the remaining $n_1 + n_2$ variables they are at the end; Bryson and Ho, 1975).

For that reason, our strategy is an approximative one. We approximate the maximally discriminating input with the maximally discriminating input of the linearisation of system (5.2.2) in Section 7.1. Then in Section 7.2 we assess the suitability of this (possibly suboptimal) input for the nonlinear system by comparing the $\|y\|_2$ value it achieves when fed into the nonlinear system to the greatest $\|y\|_2$ value over the set of all possible inputs (the latter is just the L_2 -gain of system (5.2.2)). This gain will again be computed approximately using an SOS decomposition approach.

The benefits of this strategy reside in the fact that we can use established, simple methods to find an input that gives the maximal L_2 -norm output for the linearised

system. Although this input may well be suboptimal, its effectiveness can be assessed. As we have just said, this input can be applied to the nonlinear system, and a comparison made how the realised output L_2 -norm compares to the optimal, maximally discriminating L_2 -gain.

7.1 Designing an input profile using linearisation

Designing an input profile for optimal discrimination using linearisation is more easily addressed in the frequency domain than in the time domain. The standard theory, which we will summarise here, treats single-input single-output systems ($q = \ell = 1$), but a straightforward generalisation to multiple-input multiple-output systems exists (Zhou and Doyle, 1997), which we will also use in our case study (Section 8.3). From a state space description, the transformation to the frequency domain is done by taking the Laplace transform of the state, input and output functions: for a function f in the time domain its *Laplace transform* is

$$\hat{f}(s) = \int_0^\infty f(t)e^{-st} dt.$$

For the linear system (5.2.3), if $x(0) = 0$, for the transformed functions the input–output relationship becomes

$$\hat{y}(s) = \hat{G}(s)\hat{u}(s),$$

where

$$\hat{G}(s) = C(sI - A)^{-1}B$$

is called the *transfer function* of the system (Doyle, Francis, and Tannenbaum, 1990; Dullerud and Paganini, 2000). Here the matrices A , B and C are defined in Eqs. (5.2.4),

and I is the identity matrix. The ∞ -norm of the transfer function $\hat{G}$ is defined as

$$\|\hat{G}\|_\infty := \sup_{\omega \in \mathbb{R}} |\hat{G}(i\omega)|,$$

where $i \in \mathbb{C}$ is the imaginary unit: $i^2 = -1$. The importance of this notion is that the system gain is exactly the same as the ∞ -norm of the transfer function:

$$\sup\{\|y\|_2 \mid \|u\|_2 = 1\} = \|\hat{G}\|_\infty.$$

This means that in order to find the input that corresponds to the induced L_2 -norm gain, one needs to find where the *Bode magnitude plot* peaks, the graph of the function $\omega \mapsto |\hat{G}(i\omega)|$ ($\omega \in \mathbb{R}$). In the most commonly used form the Bode magnitude plot uses decibel units on the vertical axis, hence it is defined by

$$\omega \mapsto 20 \log_{10} |\hat{G}(i\omega)| = 20 \log_{10} |C(i\omega I - A)^{-1}B|$$

for each frequency $\omega \in \mathbb{R}$. The scale on the horizontal axis is also logarithmic usually.

We will present such a plot for our case study in Section 8.3 (Figure 8.3).

When the frequency at which $|G(i\omega)|$ peaks is ω_0 (that is, when $\|\hat{G}\|_\infty = |\hat{G}(i\omega_0)|$), then the optimal input in the frequency domain is

$$\hat{u}(i\omega) = \frac{1}{2}(\delta_{\omega_0}(\omega) + \delta_{-\omega_0}(\omega))$$

for $\omega \in \mathbb{R}$ (here δ stand for Dirac delta functions). This optimal input in the time domain is $u(t) = \cos(\omega_0 t)$, but beware that it is not a permissible input: $\|\cos(\omega_0 t)\|_2 = \infty$.

Instead, for any appropriately small $\varepsilon \in]0, \omega_0[$, we can approximate the optimal input in the frequency domain with

$$\hat{u}_\varepsilon(i\omega) = \begin{cases} \sqrt{\pi/2\varepsilon}, & \text{if } |\omega - \omega_0| < \varepsilon \text{ or } |\omega - (-\omega_0)| < \varepsilon, \\ 0, & \text{otherwise.} \end{cases}$$

(See Doyle *et al.*, 1990) In the time domain this input function is

$$u_\varepsilon(t) = Z \cos(\omega_0 t) \frac{\sin(\varepsilon t)}{t} = Z\varepsilon \cos(\omega_0 t) \text{sinc}(\varepsilon t) \quad (7.1.1)$$

with Z normalising constant which ensures that the L_2 -norm of $u_\varepsilon(t)$ is now 1.

7.2 Obtaining an upper bound on the L_2 -gain of the system to assess input performance

To assess the near-optimal input designed using linearisation, we can compare its performance in driving the difference of the outputs of the two rival nonlinear models to the largest achievable difference, the L_2 -gain. We can obtain an upper bound for the L_2 -gain by constructing an appropriate storage function S . To do so, given the appropriately normalised system (5.2.2) and $\varepsilon > 0$, we assume that the trajectories with input u with $\|u\|_2 = \varepsilon$ and initial condition 0 (the steady state) remain in a region D around this initial state for all time. If there exists a $\gamma > 0$ and a continuously differentiable function $S : \mathbb{R}^{n_1+n_2} \rightarrow \mathbb{R}$ satisfying

$$S \geq 0, \quad S(0) = 0, \quad \text{and} \quad (7.2.1)$$

$$\text{for all } x \in D, \quad -\frac{\partial S}{\partial x} f(x) - y^T y + \gamma u^T u \geq 0, \quad (7.2.2)$$

then

$$\frac{\|y\|_2^2}{\|u\|_2^2} \leq \gamma.$$

In other terms, γ is the desired upper bound on the maximum difference in the output of the rival nonlinear models. To see this, integrating condition (7.2.2) from 0 to T leads to

$$\begin{aligned} \int_0^T (-y(t)^T y(t) + \gamma u(t)^T u(t)) \, dt &\geq S(x(T)) - S(x(0)) \\ &= S(x(T)) \geq 0. \end{aligned}$$

Therefore, for $T \rightarrow \infty$, if $x \in D$ for the whole time, we obtain

$$\gamma \geq \frac{\int_0^\infty y(t)^T y(t) dt}{\int_0^\infty u(t)^T u(t) dt}.$$

Here again, obtaining such a function S that provides the upper bound is difficult.

The task of finding this bound can be relaxed to solving an SOS programme and its subsequent solution using semidefinite programming (see Appendix B, Section B.4).

Chapter 8

An application: signal sensing in *Dictyostelium discoideum*

Perfect adaptation is a critical feature of many cellular signalling networks — it allows a cell to respond to a stimulus, but to resensitise itself so that further increases in stimulus can be detected. Adaptation is commonly used in sensory and other signalling networks to expand the input range that a circuit is able to sense, to more accurately detect changes in the input and to maintain homeostasis in the presence of perturbations. One of the earliest examples of cellular networks exhibiting perfect adaptation is chemotaxis, which we use as a test case to illustrate our algorithms. Specifically, we use the chemotactic response in the social amoeba *Dictyostelium discoideum*. Under starvation, *Dictyostelium* secretes cyclic adenosine monophosphate (cAMP) thus attracting other *Dictyostelium* amoebæ to aggregate and form a multicellular slug and then a fruiting body, which produces spores. Spores are inactive cells that are capable of becoming active when food is abundant. Experiments indicate that a step input of chemoattractant triggers a transient response, after which the chemosensory mechanism returns to its

pre-stimulus values (to its steady state), indicating perfect adaptation (Van Haastert and Van der Heijden, 1983).




Figure 8.1: Two models of the signal sensing system of the *Dictyostelium* amoeba.

At least two different simple models can describe the adaptation mechanism observed when an amoeba encounters the chemoattractant cAMP (Figure 8.1; Levchenko and Iglesias, 2002). In both models, a chemotaxis response regulator R becomes active (R^*) through the action of an activator enzyme A when a cAMP ligand S appears. However, the deactivation mechanism determined through the interaction of an inhibiting molecule I in the two models is different.

Since the sum of the concentrations of the active and inactive response regulators in the two models is constant, we can write $R_T = R^*(t) + R(t)$. Consequently, $\dot{R}^*$ can be derived under mass action kinetics as

$$\begin{aligned}\dot{R}^* &= -k_{-r} I R^* + k_r A R \\ &= -(k_{-r} I + k_r A) R^* + k_r A R_T,\end{aligned}$$

with activation and deactivation rate constants k_r and k_{-r} .

In Model 1 both molecules, A and I , are regulated by the external signal, which is proportional to cAMP concentration S . With rate constants k_a , k_{-a} , k_{-i} and k_{i_1} , the dynamics of A and I are given by

$$\begin{aligned}\dot{A} &= -k_{-a}A + k_aS, \\ \dot{I} &= -k_{-i}I + k_{i_1}S.\end{aligned}$$

In Model 2 the inhibitory molecule I is activated through the indirect action of activator A instead of direct activation by sensing of ligand binding, giving

$$\begin{aligned}\dot{A} &= -k_{-a}A + k_aS, \\ \dot{I} &= -k_{-i}I + k_{i_2}A,\end{aligned}$$

where k_{i_2} is a rate constant. The equations for A are obviously identical in both models. The parameter values used are given in Table 8.1. Simple manipulations show that the steady state value for R^* in Model 1 is given by

$$\bar{R}^* = \frac{k_r \frac{k_a}{k_{-a}} R_T}{k_r \frac{k_a}{k_{-a}} + k_{-r} \frac{k_{i_1}}{k_{-i}}},$$

while the steady state value for R^* in Model 2 is given by

$$\bar{R}^* = \frac{k_r R_T}{k_r + k_{-r} \frac{k_{i_2}}{k_{-i}}}.$$

Both are independent of the stimulus, explaining perfect adaptation. The two models share the same unique steady state if $k_{i_2} = k_{i_1} k_{-a} / k_a$, a condition we impose.

8.1 Initial condition design for model discrimination

For the initial condition discriminating design, we set the input to a basal level of $S = S_0$, and assume that all three concentrations, A , I and R^* , can be measured, so the output

Parameter	k_r	k_{-r}	k_a	k_{-a}	k_{i_1}	k_{i_2}	k_{-i}	S_0	R_T
Value	1	1	3	2	1	2/3	0.1	0.2	23/30

Table 8.1: Parameter values of the two models in the *Initial condition design for model discrimination* and *Input design for model discrimination* cases.

is (A, I, R^*) . The most discriminating initial state (A, I, R^*) can be found based on the linearisation of the system around its steady state using the main linear case result. The common unit length initial state which provides the direction of the perturbation from equilibrium to maximise $\|y_1 - y_2\|_2$ is then given by $x_1(0) = x_2(0) = (1, 0, 0)$, where x_i ($i \in \{1, 2\}$) are the state vectors of Model 1 and 2, respectively.

By applying the analogous results from the nonlinear case, the unit norm direction that maximises the above function is also found to be $x_1(0) = x_2(0) = (1, 0, 0)$, which illustrates that at least in this example, linearisation can be capable of providing the correct information at a lower computational cost. Figure 8.2 compares the evolution of the states in Models 1 and 2, (a) from a common arbitrary initial state, and (b) from the common initial state generated by the nonlinear method by taking the differences between the states in the two rival models.




Figure 8.2: Difference between state variables of the two rival models (states of Model 1 minus states of Model 2) in the *Initial condition design for model discrimination* and *Design of structural changes for model discrimination* cases.

Simulation results for the difference between Models 1 and 2 when started (a) from an arbitrarily perturbed initial condition $(0.5774, 0.5774, 0.5774)$, (b) from the ‘best’ unit-norm perturbation of the initial condition $(1, 0, 0)$, and (c) with the optimal parameter changes from the corresponding optimally perturbed initial state $(1, 0, 0)$. The corresponding $\|y_1 - y_2\|_2$ values are 0.420, 0.729 and 0.747, respectively.

8.2 Design of structural changes for model discrimination

In the case of the design of the most informative structural and parameter perturbations, the equilibrium is dependent on the particular choice of parameters. Therefore, we illustrate in detail how to change coordinates in order to translate the equilibrium to zero. Let the state variables for Model 1 be $x_{11} = A$, $x_{12} = I$ and $x_{13} = R^*$ and those for Model 2 be $x_{21} = A$, $x_{22} = I$ and $x_{23} = R^*$. Assume that in both models the parameters that can be modified before the experiment are R_T , the total chemotaxis response regulator concentration, and $p = k_r$, the response regulator activation rate constant. Parameter values and the intervals of values that can be achieved are given in Table 8.2. The dynamics of the two models are given by

$$\begin{aligned}\dot{x}_{11} &= -k_{-a}x_{11} + k_a S_0, \\ \dot{x}_{12} &= -k_{-i}x_{12} + k_{i_1}S_0, \\ \dot{x}_{13} &= -(px_{11} + k_{-r}x_{12})x_{13} + pR_Tx_{11}, \\ \dot{x}_{21} &= -k_{-a}x_{21} + k_a S_0, \\ \dot{x}_{22} &= -k_{-i}x_{22} + k_{i_2}x_{21}, \\ \dot{x}_{23} &= -(px_{21} + k_{-r}x_{22})x_{23} + pR_Tx_{21}.\end{aligned}$$

Parameter	k_r	k_{-r}	k_a	k_{-a}	k_{i_1}	k_{i_2}	k_{-i}	S_0	R_T
Range or value	[0.5, 1.5]	1	3	2	1	2/3	0.1	0.2	[0.5, 3.0]

Table 8.2: Parameter ranges of the two models in the *Design of structural changes for model discrimination* case.

The steady states for these equations are given by

$$\begin{aligned} x_{11}^* &= x_{21}^* = \frac{k_a S_0}{k_{-a}}, & x_{12}^* &= \frac{k_{i_1} S_0}{k_{-i}}, & x_{22}^* &= \frac{k_{i_2}}{k_{-i}} \frac{k_a S_0}{k_{-a}}, \\ x_{13}^* &= \frac{p R_T k_a k_{-i}}{k_{-r} k_{i_1} k_{-a} + p k_a k_{-i}}, & x_{23}^* &= \frac{p R_T k_{-i}}{k_{-r} k_{i_2} + p k_{-i}}. \end{aligned}$$

Let $\tilde{x}_{ik} = x_{ik} - x_{ik}^*$, $i \in \{1, 2\}$, $k \in \{1, 2, 3\}$. We can now perform a change of coordinates to make the origin the equilibrium point of:

$$\begin{aligned} \dot{\tilde{x}}_{11} &= -k_{-a} \tilde{x}_{11}, \\ \dot{\tilde{x}}_{12} &= -k_{-i} \tilde{x}_{12}, \\ \dot{\tilde{x}}_{13} &= - \left(p \tilde{x}_{11} + \frac{p k_a S_0}{k_{-a}} + k_{-r} \tilde{x}_{12} + \frac{k_{-r} k_{i_1} S_0}{k_{-i}} \right) \tilde{x}_{13} \\ &\quad + \frac{p R_T k_{-r} k_{i_1} k_{-a}}{k_{-r} k_{i_1} k_{-a} + p k_a k_{-i}} \tilde{x}_{11} - \frac{k_{-r} p R_T k_a k_{-i}}{k_{-r} k_{i_1} k_{-a} + p k_a k_{-i}} \tilde{x}_{12}, \\ \dot{\tilde{x}}_{21} &= -k_{-a} \tilde{x}_{21}, \\ \dot{\tilde{x}}_{22} &= -k_{-i} \tilde{x}_{22} + k_{i_2} \tilde{x}_{21}, \\ \dot{\tilde{x}}_{23} &= - \left(p \tilde{x}_{21} + \frac{p k_a S_0}{k_{-a}} + k_{-r} \tilde{x}_{22} + \frac{k_{-r} k_{i_2} k_a S_0}{k_{-a} k_{-i}} \right) \tilde{x}_{23} \\ &\quad + \frac{p R_T k_{-r} k_{i_2}}{k_{-r} k_{i_2} + p k_{-i}} \tilde{x}_{21} - \frac{k_{-r} p R_T k_{-i}}{k_{-r} k_{i_2} + p k_{-i}} \tilde{x}_{22}. \end{aligned} \tag{8.2.1}$$

In order to discriminate between Models 1 and 2, we solve the optimisation programmes given in Section B.3 of Appendix B, as explained in Section 6.2. We allow parameters R_T and k_r to vary. We obtain that $\tilde{x}_1(0) = \tilde{x}_2(0) = (1, 0, 0)$ are the initial conditions, and $R_T = 3$ and $k_r = p = 1.5$ are the values of the parameters that have maximal discriminating power between the two models. The increased difference between states can be seen in Figure 8.2c. This result means that one needs to overexpress the total number of chemotaxis response regulators and increase their rate of activation in order to see a large difference between the two models.

8.3 Input design for model discrimination

To discriminate between the two models based on an optimally chosen input profile, we first obtained an upper bound on the L_2 -gain of the difference system from input S to output (A, I, R^*) using the algorithm in Section B.4 of Appendix B. This bound was about 0.477, a value also close to that determined through the linearisation of the system (0.4766, see the Bode plot shown in Figure 8.3).




Figure 8.3: Bode magnitude plots of the linearised difference system of the two models with output (A, I, R^*) , I and R^* .

The difference between the R^* values of the two models is marginal compared to the I values. The L_2 -gains for outputs (A, I, R^*) , I and R^* are 0.4766, 0.4762 and 0.02038, respectively. At this resolution one cannot see a difference between the cases when the output is the full state (A, I, R^*) or I only. The shape of the Bode magnitude plot for single outputs I and R^* is similar, with only slightly different critical frequencies (0.4472 and 0.3853, respectively). The critical frequency corresponding to completely observed state is 0.4470.

In order to evaluate the performance of our algorithm, we simulated the original,

nonlinear system subjected to different inputs between 0 and 60 time units including a constant (step function), sine, cosine, sine with an exponentially increasing then decreasing multiplier, sinc, the function given by Eq. (7.1.1), or a square wave function (see Figure 8.4). As described earlier, for periodic inputs, the period of the input was determined by finding the frequency corresponding to the maximum amplification in the Bode magnitude plot of the linearised system (Figure 8.3).

Table 8.3 summarises the results for different input perturbation functions, input perturbation energies ($\|u\|_2$) and output variables using simulations of the nonlinear system. Since a cosine input and the input function given by Eq. (7.1.1) ($u_\varepsilon(t) = Z \cos(\omega_0 t) \frac{\sin(\varepsilon t)}{t}$) gave indistinguishable results (because themselves are indistinguishable), their rows were merged into one. The top and bottom parts of the table show values for different input perturbation energies, the various columns for different output variables. The applied frequency, where relevant, was always the critical frequency corresponding to the particular output function. Figure 8.5 is a graphical representation of the same data.

Figure 8.6 compares the differences in state variables between the two models for three typical input perturbations. Interestingly, our results indicate that for all inputs used, discrimination between the two models should be accomplished by measuring output I rather than output R^* (also seen in Table 8.3). In the first plot, the basal input S_0 is perturbed by a step function between 0 and 60 time units (Figure 8.4a). In the second plot, the system is injected with a sine function (Figure 8.4b). In the third plot, the input is a square wave function (Figure 8.4c), a caricature of the sine function with preserved period that can be realised in practice more easily. The sine input yields a visibly larger difference than the step function between the R^* values of the two models. The square wave function produces a similarly good result. One has to note that one measurement may not be enough for the discrimination, but a series of measurements

$\ u\ _2 = 1$	(A, I, R^*)	I	R^*
Sine	0.472	0.472	0.0195
Sine w. exp. mult.	0.475	0.474	0.0209
Cosine or (7.1.1)	0.473	0.472	0.0200
Square wave	0.451	0.450	0.0182
Sinc	0.412	0.412	0.0153
Step	0.198	0.197	0.0070
$\ u\ _2 = 0.01$	(A, I, R^*)	I	R^*
Sine	0.476	0.481	0.0203
Sine w. exp. mult.	0.467	0.467	0.0195
Cosine or (7.1.1)	0.457	0.456	0.0191
Square wave	0.441	0.441	0.0184
Sinc	0.396	0.396	0.0173
Step	0.198	0.198	0.0085

Table 8.3: Achievable output differences for different input profiles.

Numerical estimates of $\max_u \|y_1 - y_2\|_2 / \|u\|_2$ with different inputs for the *Dictyostelium* models by simulating the nonlinear system. Here maximisation is over different frequencies for inputs where frequency makes sense. Output is either all states or I or R^* .

may be needed.

Perhaps the most notable outcome of the input design is that sinusoidal input perturbations generate the best L_2 -gains and are therefore superior to a step function for discriminating between rival chemotaxis models. Square wave stimulation is achievable in the reality of a laboratory. This is important since step inputs are usually used in experiments, often at the exclusion of other input signals. Our studies demonstrate how more dynamic inputs, in this case an oscillating input (on a finite time interval), might

be necessary to delineate subtle features of underlying network topologies.

The optimisation problems in all three cases were solved on a desktop computer. The most challenging was the first SOS programme for the choice of optimal initial state and parameter values with eight variables (three state variables for each model and two common parameters). Numerical methods will need to be improved in order to deal with SOS programmes resulting from the analysis of more complex systems biological models.




Figure 8.4: Three distinctive input profiles $S = S_0 + u$.

Basal input S_0 perturbed until time 60 by (a) a step function, (b) a sine function, or (c) a square wave function. $\|u\|_2 = 1$ in all cases. In (b) and (c) the frequency is 0.3853, the critical frequency corresponding to the single output R^* .




Figure 8.5: Bar chart of data in Table 8.3: achievable output differences for different input profiles.




Figure 8.6: Difference between state variables of the two rival models (states of Model 1 minus states of Model 2) in the *Input design for model discrimination* case.

Simulation results for Models 1 and 2 with (a) a step, (b) sine, or (c) square wave perturbation of the basal input S_0 until time 60 (the inputs in Figure 8.4). The output variable (on which the choice of the optimal frequency depends) is R^* and $\|u\|_2 = 1$. The corresponding $\|y_1 - y_2\|_2$ values are 0.0070, 0.0195 and 0.0182, respectively. Note that the input signal $S_0 + u$ is not included in the figure.

Chapter 9

Summary of the results of Part II

In Part II we developed methods for designing experiments to effectively discriminate between different models of a biological system. These methods are tailored to generate maximally informative data that can be used to invalidate models of gene regulatory pathways by ruling out certain connectivities in their underlying biochemical reaction networks. We approached the problem in a unified framework, developing methodologies for initial condition design (see also Papachristodoulou and El-Samad, 2007), for parameter modifications and for the design of dynamic stimulus profiles. These types of manipulations cover a large spectrum of what is experimentally feasible and this is what has motivated our formulation of the problem and the approach to its investigation.

If the field of systems biology is to accelerate the pace of biological discovery, rigorous mathematical methods should be developed to link computational models of biological networks to experimental data in tight rounds of analysis and synthesis. Any informative model should be analysed in light of existing data, but it should also be able to synthesise new experiments that further delineate the features of the underlying biological system. Despite many notable examples demonstrating the success of this iterative procedure,

progress has been slow due to the *ad hoc* nature of its implementation: the iterations between the development of models and the production of data is still mostly guided by the intuition of the modellers, and no rigorous algorithms exist to render this process more systematic and less biased. We believe that the work presented here constitutes an important step in this direction. By design, our formulation of the problem is of sufficient generality to accommodate many experiment design procedures, and is cast in a natural optimisation framework. Acknowledging that optimality of experiment designs must always be balanced with biological and other practical constraints, our formalism allows for the incorporation of limitations and constraints as dictated by the specific biological context.

We illustrated the applicability of our algorithms using two possible and widely accepted simplified models of the adaptation mechanism in *Dictyostelium discoideum* chemotaxis. Evidently, these models do not capture the full complexity of the biological circuit responsible for chemotactic behaviour. The models, however, illustrate the core circuit topologies that are sufficient to implement perfect adaptation in the system. These optimal experiment design methods were applied in practice to invalidate models of the chemotaxis pathway in *Rhodobacter sphaeroides* (Roberts *et al.*, 2009). There, the combination of a square wave profile stimulation and protein overexpression was necessary in the most challenging model discrimination problem. This demonstrates the practical demand for sophisticated experiment design techniques.

The recipe for model discrimination that we propose involves collecting mostly time series data. Every new time point at which measurements are made increases the labour and monetary cost of experiments, thus one must carefully balance the number of time points collected against the cost, and consider where along a time series to place observations. Our methods naturally offer insight into this question by predicting the time

series that the experiment will likely measure. Furthermore, if the optimal experiment is such that a differentiating dynamical phenotype only emerges several hours after a perturbation, our methods can be easily modified to balance optimality with practically measurable dynamics.

Finally, many commonly used perturbations (genetic or environmental) lead to either extreme stress responses that put a cell in a modified physiological state or kills it, or lead to quiescent states that do not have much measurable information about the underlying regulatory network. Experiments that generate less catastrophic failures of cellular networks under study, while being maximally informative, hold great promise for the study of biological networks. However, finding these perturbations is a nontrivial task. Model-based design of experiments will undeniably be instrumental for that, ultimately leading to many important biological discoveries.

Chapter B

Appendix. Methods deferred from Chapters 6 and 7

B.1 Sum of squares decompositions

Here we present the *sum of squares* (SOS) formalism which is used to relax and solve the optimisation problems posed by the various approaches for model discrimination considered in Part II.

Denote by $\mathbb{R}[y]$ the ring of polynomials in variables $y = (y_1, \dots, y_n)$ with real coefficients. $p(y) \in \mathbb{R}[y]$ is called *nonnegative* if and only if $p(y) \geq 0$ for all y . It is a *sum of squares* if there exist other polynomials $p_i(y) \in \mathbb{R}[y]$ ($i \in \{1, \dots, M\}$) such that $p(y) = \sum_{i=1}^M p_i(y)^2$.

Obviously, such a polynomial is nonnegative, but the converse is not always true (Parrilo, 2000). In fact, testing if $p(y) \geq 0$ is NP-hard (Murty and Kabadi, 1987), but testing if $p(y)$ is a sum of squares is equivalent to a *semidefinite programme* (SDP) (Parrilo, 2000), a convex optimisation problem for which there are algorithms that can

solve it with polynomial time complexity.

The MATLAB toolbox SOSTOOLS (Prajna, Papachristodoulou, and Parrilo, 2002) can be used to formulate this SDP which can be solved using SDP solvers such as SeDuMi (Sturm, 1999) or SDPT3 (Toh, Tütüncü, and Todd, 1999).

B.2 SOS programmes for initial condition design

The strategy outlined in the section on *Initial condition design for model discrimination* relies on the construction of a function S satisfying the nonnegativity conditions given by (6.1.2) and (6.1.3). As it was pointed out in Section B.1, constructing a nonnegative S is difficult. We therefore relax nonnegativity to the existence of an SOS decomposition and solve the problem through semidefinite programming. An SOS programme that can be used to generate S is

$$\begin{array}{ll}
\text{given} & f_1, f_2, \text{ and the set description } D, \\
\text{find} & S(x), \sigma_1(x_1), \sigma_2(x_2) \text{ all SOS,} \\
\text{s.t.} & S(0) = 0, \\
& -\frac{\partial S}{\partial x} f(x) - h^T h + \sigma_1(x_1)(|x_1|^2 - \alpha^2) \\
& + \sigma_2(x_2)(|x_2|^2 - \alpha^2) \text{ is SOS.}
\end{array}$$

The last constraint ensures that $-\frac{\partial S}{\partial x} f(x) - h^T h \geq 0$ when $x \in D$, since the multipliers $\sigma_1(x_1)$ and $\sigma_2(x_2)$ are SOS.

The solution S is not unique, but a heuristic to find the ‘best’ S is to optimise over the decision variables in the SOS description for S (by minimising the trace of the Jacobian of S at the origin), so that the resulting S has sub-level sets that have maximal area.

In the second step, we solve the SOS relaxation of the optimisation problem (6.1.6) to get the initial state $x_1(0) = x_2(0) = \hat{x}$:

$$\begin{array}{ll}
\text{given} & \hat{S}, \\
\text{minimise} & \gamma, \text{ and} \\
\text{find} & p(\hat{x}) \text{ polynomial,} \\
\text{s.t.} & -\hat{S}(\hat{x}) + \gamma + p(\hat{x})(|\hat{x}|^2 - \beta^2) \text{ is SOS.}
\end{array}$$

The point $\hat{x}$ can be obtained from the dual solution of this semidefinite programme, using SOSTOOLS.

B.3 SOS programmes for optimal structural design

The search for a function $S(x, p) \geq 0$ such that conditions (6.2.1) hold in $D \times \Pi$ can be formulated as

$$\begin{array}{ll}
\text{given} & f_1, f_2, \ b_i \geq a_i \geq 0 \ (i \in \{1, \dots, m\}), \\
\text{find} & S(x, p), \sigma_1(x, p), \sigma_2(x, p), v_1(x, p), \dots, v_m(x, p) \text{ all SOS,} \\
\text{s.t.} & S(0, p) = 0, \text{ and} \\
& -\frac{\partial S(x, p)}{\partial x} f(x, p) - h^T(x)h(x) \\
& \quad + \sigma_1(x, p)(|x_1|^2 - \alpha^2) \\
& \quad + \sigma_2(x, p)(|x_2|^2 - \alpha^2) \\
& \quad + \sum_{i=1}^m v_i(x, p)(p_i - a_i)(p_i - b_i) \text{ is SOS.}
\end{array}$$

As in the initial state design case, we would like to maximise the difference given by y . Although we cannot achieve this goal directly, we can maximise the approximation of $\|y\|_2^2$ given by $S(x(0), p)$ (see (6.2.2)), where we require that $x_1(0) = x_2(0) = \hat{x}$

and $|\hat{x}| = \beta$. We introduce the modified S , $\hat{S}(\hat{x}, p) := S(\begin{pmatrix} \hat{x} \\ \hat{x} \end{pmatrix}, p)$ ($\hat{x} \in \mathbb{R}^n$). Similarly, $\hat{f}(\hat{x}, p) := f(\begin{pmatrix} \hat{x} \\ \hat{x} \end{pmatrix}, p)$. The problem is then the following.

$$\begin{array}{ll}
\text{Given} & \hat{S}, \\
\\
\text{minimise} & \gamma, \text{ and} \\
\\
\text{find} & r(\hat{x}, p) \text{ polynomial, and } w_1(\hat{x}, p), \dots, w_m(\hat{x}, p) \text{ all SOS,} \\
\\
\text{s.t.} & -\hat{S}(\hat{x}, p) + \gamma + r(\hat{x})(|\hat{x}|^2 - \beta^2) \\
& + \sum_{i=1}^m w_i(\hat{x}, p)(p_i - a_i)(p_i - b_i) \text{ is SOS.}
\end{array}$$

Exactly as in the initial state design case, the point $\hat{x}$ can be obtained from the dual solution, using SOSTOOLS.

B.4 SOS programmes for optimal input design

In the *Input design for model discrimination*, we should first note that it may occasionally be the case that the set of inputs considered will lead to a system trajectory outside the region where S is constructed. This case can be ruled out by solving a related reachability problem (Papachristodoulou, 2005). Here, we assume that the containment of the trajectory in D has been ensured, and describe how to obtain an estimate of the L_2 -gain of the system.

To construct a function S which satisfies the conditions shown in Section 7.2, we use the SOS framework as follows. Condition (7.2.2) can be satisfied by searching for SOS multipliers $\sigma_1(x)$ and $\sigma_2(x)$ such that

$$\begin{aligned}
& -\frac{\partial S}{\partial x} f(x) - y^T y + \gamma u^T u + \sigma_1(x)(|x_1|^2 - \alpha^2) \\
& + \sigma_2(x)(|x_2|^2 - \alpha^2) \text{ is SOS,}
\end{aligned}$$

where $\alpha > 0$ is used to define the region D by

$$D = \left\{ \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \in \mathbb{R}^{n_1+n_2} \mid |x_1| \leq \alpha, |x_2| \leq \alpha \right\}.$$

This condition guarantees that $-\frac{\partial S}{\partial x}f(x) - y^Ty + \gamma u^Tu \geq 0$ for $x \in D$. The rest of the conditions can also be easily enforced in an SOS programming framework.

Consequently, the overall SOS programme for constructing S takes the form:

given	f_1, f_2, α , and the set description D ,
minimise	γ , and
find	$S(x), \sigma_1(x), \sigma_2(x)$ all SOS,
s.t.	$S(0) = 0$, and
	$-\frac{\partial S}{\partial x}f(x) - y^Ty + \gamma u^Tu + \sigma_1(x)(x_1 ^2 - \alpha^2)$
	$+ \sigma_2(x)(x_2 ^2 - \alpha^2)$ is SOS.

Chapter 10

Outlook and conclusion

10.1 Future directions in the modelling of chemical reaction kinetics with continuous stochastic models

Biochemical processes in individual cells are known to be multiscale: they occur on widely varying time scales, and the sizes of the molecular species' populations range over several orders of magnitude. One of the great challenges facing mathematical modelling is to reflect this multiscale property. The triad of the discrete Markov jump process, the CLE and the reaction rate equation provide accurate and economical descriptions for three complementary regimes of molecular population sizes. Each has its unique niche and each is essential in its own right for a comprehensive picture of biochemical system dynamics. Part I concerned the middle of these three regimes.

The most imminent task relating to the CLE was discussed in detail in Section 3.6: its trajectories can leave the nonnegative orthant with positive probability and those

trajectories cannot be continued such that they still satisfy the CLE. The discovery of this defect creates an urgent need for a new, nonnegative SDE model for reaction kinetics to replace the classical CLE.

We can envisage two routes to this end: either by modifying the boundary behaviour of the stochastic process to enforce nonnegativity whenever a variable is about to become negative, or by revisiting Gillespie’s argument that gave the CLE (Gillespie, 2000). It might be the case that there is no better diffusion approximation to our discrete Markov jump process than the CLE. All approximations have shortcomings and problematic boundary behaviour might be unavoidable, especially since a diffusion approximation can be expected to break down at the boundary. These questions are currently an active area of research.

Another open problem concerning the existence of solutions of the CLE was presented at the end of Section 3.6.2. Namely, providing sufficient conditions to avoid finite time explosion in the CLE that are more general than the compactness of the state space. This is an interesting direction for possible future research, especially given a nontrivial result for the deterministic analogue (August and Barahona, 2010).

The existence of multiple scales within the same chemical system calls for models that are themselves multiscale, and which treat faster and slower reactions, more and less abundant species differently. Arguably, it is rarely the case that the assumptions in Gillespie’s derivation of the CLE hold (Section 3.1), that is, that all discrete variables can be simultaneously and uniformly approximated with continuous ones. In fact, there are *hybrid* simulation strategies, in which the slowly changing variables follow the discrete stochastic process, and the fast variables obey, for instance, either the reaction rate equation or the τ -leap method or the CLE. Wilkinson (2006, Section 8.4) gives an overview of such algorithms and references to relevant primary sources.

The model reduction and analysis of multiscale biochemical systems is currently a very active area of research. It is heavily influenced by the chemical engineering and control engineering literature; the first methods were developed for deterministic ODE models, often for applications in the petrochemical industry. Subsequently, such techniques found extensive use in the mathematical biology community. Most recently, and we will point out just some of the promising directions, stochastic analogues of methods for deterministic systems are being developed. The partial equilibrium assumption or quasi-steady state approximation is such an example, which has already been applied to the discrete Markov jump process (Haseltine and Rawlings, 2002; Rao and Arkin, 2003; Goutsias, 2005; Chevalier and El-Samad, 2009) and directly to the CLE (Lan *et al.*, 2008; Sotiropoulos *et al.*, 2009). Its application to speed up the SSA gave the *slow-scale SSA* (Cao, Gillespie, and Petzold, 2005) and the *nested SSA* (E, Liu, and Vanden-Eijnden, 2005). The differences between the two approaches were discussed by both sets of authors (Gillespie, Petzold, and Cao, 2007; E, Liu, and Vanden-Eijnden, 2007).

A more sophisticated, although laborious technique to reduce multiscale, discrete space Markov jump process models of chemical reaction systems was developed by Ball *et al.* (2006). Their starting point, the discrete Markov process, was written with a stochastic equations formalism with independent Poisson processes for each reaction channel. Scaling constants were introduced for each molecular species, each reaction channel, and time. With the fine control of all these scaling parameters, the authors could approximate the different variables and reaction channels with diffusion approximations (SDEs) or even continuous deterministic processes (integral equations), depending on the inherent scaling properties of the system. However, they could not provide rules for the appropriate choice of scaling constants, which is a great hindrance to the application

of this model reduction technique. This shortcoming is addressed in Kang and Kurtz (2010).

As we have just pointed it out, in the derivation of the CLE, passing to a continuous limit in each variable uniformly is a questionable approach. Kang, Kurtz, and Popovic (2010) are studying diffusion approximations for multiscale chemical reaction models to develop a better understanding of this regime.

After a thorough analysis of multiscale deterministic chemical reaction models (Lee and Othmer, 2010a), Lee and Othmer (2010b) are working on a similar treatment of stochastic models, independently of the group around Kurtz.

10.2 A different optimal experiment design problem for model discrimination based on the Hankel operator

In Chapter 7 of Part II, in our discussion of optimal experimental stimulation design, our approach was to observe and stimulate the biological system from time zero, applying input and measuring output simultaneously. We briefly describe a different concept, in which input is applied only up to a certain time, say zero, which is also the time when the output measurement starts.

In the state space description, one can use the input–output operator $G: L_2(\cdot) - \infty, \infty) \rightarrow L_2(\cdot) - \infty, \infty)$, for which $y = Gu$. This operator is the time domain equivalent of the transfer function $\hat{G}$ we used for the input design.

The *Hankel operator* $\Gamma_G: L_2(\cdot) - \infty, 0] \rightarrow L_2([0, \infty)$ of G is defined by

$$\Gamma_G = P_+ G|_{L_2(\cdot) - \infty, 0]} ,$$

where $P_+ : L_2(]-\infty, \infty[) \rightarrow L_2([0, \infty[)$ is the projection by truncation. So the Hankel operator connects input

$$u :]-\infty, 0] \rightarrow \mathbb{R}^q$$

with output

$$y : [0, \infty[\rightarrow \mathbb{R}^\ell.$$

Equivalently, $\Gamma_G = \Psi_o \Psi_c$, where

- Ψ_c is the *controllability operator*, which maps u to $x(0)$,
- Ψ_o is the *observability operator*, which maps $x(0)$ to y with no input after time 0.

Indeed, as the support of input u is before 0, it only affects future output y through the state at time 0, $x(0)$. The introductory results of this theory are about the norm of the Hankel operator for linear systems (Dullerud and Paganini, 2000), which can potentially serve as a starting point for interesting investigations of optimal experiment design in this novel setting.

10.3 Conclusion

One of the characteristic aspects of contemporary biology is a strive for a systemic understanding of the spatio-temporal dynamics of intracellular networks of molecular reactions. In this effort, complementing experimentation, the contribution of mathematical modelling is crucial; in essence, having an accurate mathematical model that is amenable to analysis or to numerical simulation is what is considered as the understanding of a reaction network.

This thesis contributes to two stages of the mathematical modelling process for biochemical networks: to using the right modelling framework and to finding the right

model within a specific framework for a biological system. Part I considered what the best way is of modelling biochemical reaction kinetics with continuous stochastic models. We found that the moment equations give new justification for the standard SDE model, the CLE. We explored different formulations of this equation and thus we gained a better understanding of the geometry of the state space and discovered a way to accelerate the numerical simulation of the equation. We also gave a detailed analysis of the Achilles heel of the CLE, the negativity issue. We hope that armed with these observations, the development of a novel, nonnegativity preserving SDE model will become possible.

Part II addressed a more practical task: designing experiments that differentiate between similarly accurate models of a reaction network in order to rule out incorrect ones. This will improve the accuracy of our knowledge of the underlying biology and of the biochemical processes involved. We found that in the case of externally excitable systems, the optimal stimulation for model discrimination, perhaps surprisingly, is not a single impulse (step function input profile) but periodically repeated stimulation (sinusoidal or square wave function input). This result has already been applied in practice with resounding success.

We believe that both lines of our inquiries have important practical implications. Further work is required to disseminate these results to potential beneficiaries. With these considerations we bring our thesis to an end.

Acknowledgements

I would like to thank my two supervisors, Prof. Alison Etheridge and Dr. Antonis Papatristodoulou for their guidance during my doctoral studies. I am most grateful to Antonis for the long discussions we had, for the encouragement, and for his various advice on career planning. I thank Alison for stopping me — she told me I had enough results for a doctorate. I received much feedback on my work from both my supervisors. Both Alison and Antonis provided enormous support by sending me to conferences and introducing me to some of the leading figures of my research area. Dr. Martin Kolb deserves a special mention for recommending Pinsky's book.

I wish to acknowledge financial support from the Engineering and Physical Sciences Research Council through the Life Sciences Interface Doctoral Training Centre. The stipend gave me priceless intellectual freedom — I wonder if I will ever have this again. I hope I was able to give something back not only through my work but also through my non-academic activities.

I thank my first academic home at the University of Oxford, the Life Sciences Interface Doctoral Training Centre, for the teaching and the good company.

I thank the Department of Statistics and especially the great community in the Peter Medawar Building. Being here gave me a sense of belonging and it was always a pleasure to hang out with them in the coffee breaks. I am grateful to the IT support staff of the Department: their methodical work meant that I was never distracted by having to do

computer maintenance.

My secondary allegiance to the Control Group of the Department of Engineering Science brought me a broader perspective and a few good friends.

I greatly enjoyed being a member of Keble College. I made many friends there and had very good let's-solve-the-world's-problems-type conversations. Meals in the magnificent dining hall will always remain a defining experience of my time at Oxford.

Finally, I thank the University of Oxford for all the fascinating people I met here: the very people who make up this institution and the invited speakers who bring the world to Oxford with them.

List of abbreviations

ATP adenosine triphosphate

cAMP cyclic adenosine monophosphate

CLE chemical Langevin equation

CME chemical master equation

DNA deoxyribonucleic acid

Eq. equation

FSP finite state projection

gfp green fluorescent protein

HERG human ether a-go-go related gene

IPTG isopropyl β -D-thiogalactopyranoside

mRNA messenger ribonucleic acid

ODE ordinary differential equation

RNAi ribonucleic acid interference

SDE stochastic differential equation

SDP semidefinite programme

SOS sum of squares

SSA stochastic simulation algorithm

s.t. subject to

Bibliography

- D. Adalsteinsson, D. McMillen, and T. Elston. Biochemical Network Stochastic Simulator (BioNetS): software for stochastic modeling of biochemical networks. *BMC Bioinformatics*, 5 (2004) 24.
- B. Alberts, A. Johnson, J. Lewis, M. Raff, K. Roberts, and P. Walter. Molecular biology of the cell. Garland Science, New York, Abingdon, Fifth edition, 2008.
- E. J. Allen, L. J. S. Allen, A. Arciniega, and P. E. Greenwood. Construction of equivalent stochastic differential equation models. *Stochastic Analysis and Applications*, 26 (2008) 274–297.
- D. F. Anderson. A modified next reaction method for simulating chemical systems with time dependent propensities and delays. *The Journal of Chemical Physics*, 127 (2007) 214107.
- J. Anderson and A. Papachristodoulou. On validation and invalidation of biological models. *BMC Bioinformatics*, 10 (2009) 132.
- J. F. Apgar, J. E. Toettcher, D. Endy, F. M. White, and B. Tidor. Stimulus design for model selection and validation in cell signaling. *PLoS Computational Biology*, 4 (2008) e30.
- A. Arkin, J. Ross, and H. H. McAdams. Stochastic kinetic analysis of developmental

- pathway bifurcation in phage λ -infected *Escherichia coli* cells. *Genetics*, 149 (1998) 1633–1648.
- A. Auger, P. Chatelain, and P. Koumoutsakos. *R*-leaping: accelerating the stochastic simulation algorithm by reaction leaps. *The Journal of Chemical Physics*, 125 (2006) 084103.
- E. August and M. Barahona. Solutions of weakly reversible chemical reaction networks are bounded and persistent. In Proceedings of the 11th IFAC Symposium on Computer Applications in Biotechnology, 2010.
- E. August and A. Papachristodoulou. Efficient, sparse biological network determination. *BMC Systems Biology*, 3 (2009a) 25.
- E. August and A. Papachristodoulou. A new computational tool for establishing model parameter identifiability. *Journal of Computational Biology*, 16 (2009b) 875–885.
- K. Ball, T. G. Kurtz, L. Popovic, and G. Rempala. Asymptotic analysis of multiscale approximations to reaction networks. *The Annals of Applied Probability*, 16 (2006) 1925–1961.
- E. Balsa-Canto, A. A. Alonso, and J. R. Banga. Computational procedures for optimal experimental design in biological systems. *IET Systems Biology*, 2 (2008) 163–172.
- S. Bandara, J. P. Schlöder, R. Eils, H. G. Bock, and T. Meyer. Optimal experimental design for parameter estimation of a cell signaling model. *PLoS Computational Biology*, 5 (2009) e1000558.
- W. G. Bardsley, R. M. W. Wood, and E. M. Melikhova. Optimal design: a computer program to study the best possible spacing of design points. *Computers & Chemistry*, 20 (1996) 145–157.

- C. L. Barrett and B. O. Palsson. Iterative reconstruction of transcriptional regulatory networks: an algorithmic approach. *PLoS Computational Biology*, 2 (2006) 429–438.
- J. M. Bower and H. Bolouri (editors). Computational modeling of genetic and biochemical networks. MIT Press, Cambridge, 2004.
- T. Brennan, M. Fink, and B. Rodriguez. Multiscale modelling of drug-induced effects on cardiac electrophysiological activity. *European Journal of Pharmaceutical Sciences*, 36 (2009) 62–77.
- M. Brenner and J.-i. Tomizawa. Quantitation of ColE1-encoded replication elements. *Proceedings of the National Academy of Sciences USA*, 88 (1991) 405–409.
- A. E. Bryson and Y.-C. Ho. Applied optimal control: optimization, estimation and control. Taylor & Francis, London, 1975.
- K. Burrage, J. Hancock, A. Leier, and D. V. Nicolau Jr. Modelling and simulation techniques for membrane biology. *Briefings in Bioinformatics*, 8 (2007) 234–244.
- Y. Cao, D. T. Gillespie, and L. R. Petzold. The slow-scale stochastic simulation algorithm. *The Journal of Chemical Physics*, 122 (2005) 014116.
- F. P. Casey, D. Baird, Q. Feng, R. N. Gutenkunst, J. J. Waterfall, C. R. Myers, K. S. Brown, R. A. Cerione, and J. P. Sethna. Optimal experimental design in an epidermal growth factor receptor signalling and down-regulation model. *IET Systems Biology*, 1 (2007) 190–202.
- A. Chatterjee, D. G. Vlachos, and M. A. Katsoulakis. Binomial distribution based τ -leap accelerated stochastic simulation. *The Journal of Chemical Physics*, 122 (2005) 024112.

- B. H. Chen and S. P. Asprey. On the design of optimally informative dynamic experiments for model discrimination in multiresponse nonlinear situations. *Industrial & Engineering Chemistry Research*, 42 (2003) 1379–1390.
- M. W. Chevalier and H. El-Samad. A rigorous framework for multiscale simulation of stochastic cellular networks. *The Journal of Chemical Physics*, 131 (2009) 054102.
- K. L. Chung and R. J. Williams. Introduction to stochastic integration. Birkhäuser, Boston, second edition, 1990.
- A. Cornish-Bowden. Fundamentals of enzyme kinetics. Portland Press, London, 2004.
- C. E. Dangerfield, D. Kay, and K. Burrage. Stochastic models and simulation of ion channel dynamics. *Procedia Computer Science*, 1 (2010) 1581–1590.
- S. Di Talia, J. M. Skotheim, J. M. Bean, E. D. Siggia, and F. R. Cross. The effects of molecular noise and size control on variability in the budding yeast cell cycle. *Nature*, 448 (2007) 947–951.
- B. M. Donckels, D. J. De Pauw, B. De Baets, J. Maertens, and P. A. Vanrolleghem. An anticipatory approach to optimal experimental design for model discrimination. *Chemometrics and Intelligent Laboratory Systems*, 95 (2009) 53–63.
- J. Doyle, B. Francis, and A. Tannenbaum. Feedback control theory, chapter 2. Macmillan Publishing Co., New York, 1990. 11–26.
- B. Drawert, M. J. Lawson, L. Petzold, and M. Khammash. The diffusive finite state projection algorithm for efficient simulation of the stochastic reaction-diffusion master equation. *The Journal of Chemical Physics*, 132 (2010) 074101.

- G. E. Dullerud and F. G. Paganini. A Course in Robust Control Theory. A convex approach. Springer, New York, 2000.
- R. Durrett. Stochastic calculus. A practical introduction. CRC Press, Boca Raton, second edition, 1996.
- W. E. D. Liu, and E. Vanden-Eijnden. Nested stochastic simulation algorithm for chemical kinetic systems with disparate rates. *The Journal of Chemical Physics*, 123 (2005) 194107.
- W. E. D. Liu, and E. Vanden-Eijnden. Response to “Comment on ‘Nested stochastic simulation algorithm for chemical kinetic systems with disparate rates’ [J. Chem. Phys. 123, 194107 (2005)]”. *The Journal of Chemical Physics*, 126 (2007) 137102.
- M. B. Elowitz and S. Leibler. A synthetic oscillatory network of transcriptional regulators. *Nature*, 403 (2000) 335–338.
- M. B. Elowitz, A. J. Levine, E. D. Siggia, and P. S. Swain. Stochastic gene expression in a single cell. *Science*, 297 (2002) 1183–1186.
- R. Erban and S. J. Chapman. Stochastic modelling of reaction–diffusion processes: algorithms for bimolecular reactions. *Physical Biology*, 6 (2009) 046001.
- D. Fange, O. G. Berg, P. Sjöberg, and J. Elf. Stochastic reaction-diffusion kinetics in the microscopic limit. *Proceedings of the National Academy of Sciences USA*, 107 (2010) 19820–19825.
- M. Feinberg. Chemical reaction network structure and the stability of complex isothermal reactors—I. The Deficiency Zero and Deficiency One Theorems. *Chemical Engineering Science*, 42 (1987) 2229–2268.

- D. Fell. Understanding the control of metabolism. Portland Press, London, 1997.
- W. Feller. An introduction to probability theory and its applications, volume I. John Wiley & Sons, New York, Chichester, second edition, 1957.
- X.-j. Feng and H. Rabitz. Optimal identification of biochemical reaction networks. *Biophysical Journal*, 86 (2004) 1270–1281.
- A. Friedman. Stochastic differential equations and applications, volume I. Academic Press, New York, 1975.
- C. W. Gardiner. Handbook of stochastic methods for physics, chemistry and the natural sciences. Springer Series in Synergetics. Springer–Verlag, Berlin, Heidelberg, second edition, 1985.
- M. A. Gibson and J. Bruck. Efficient exact stochastic simulation of chemical systems with many species and many channels. *Journal of Physical Chemistry A*, 104 (2000) 1876–1889.
- D. T. Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics*, 22 (1976) 403–434.
- D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry*, 81 (1977) 2340–2361.
- D. T. Gillespie. A rigorous derivation of the chemical master equation. *Physica A*, 188 (1992) 404–425.
- D. T. Gillespie. The multivariate Langevin and Fokker-Planck equations. *The American Journal of Physics*, 64 (1996) 1246–1257.

- D. T. Gillespie. The chemical Langevin equation. *The Journal of Chemical Physics*, 113 (2000) 297–306.
- D. T. Gillespie. Approximate accelerated stochastic simulation of chemically reacting systems. *The Journal of Chemical Physics*, 115 (2001) 1716–1733.
- D. T. Gillespie, L. R. Petzold, and Y. Cao. Comment on “Nested stochastic simulation algorithm for chemical kinetic systems with disparate rates” [J. Chem. Phys. 123, 194107 (2005)]. *The Journal of Chemical Physics*, 126 (2007) 137101.
- A. Goldbeter and D. E. Koshland. An amplified sensitivity arising from covalent modification in biological systems. *Proceedings of the National Academy of Sciences USA*, 78 (1981) 6840–6844.
- C. A. Gómez-Uribe and G. C. Verghese. Mass fluctuation kinetics: Capturing stochastic effects in systems of chemical reactions through coupled mean-variance computations. *The Journal of Chemical Physics*, 126 (2007) 024109.
- J. Goutsias. Quasiequilibrium approximation of fast reaction kinetics in stochastic biochemical systems. *The Journal of Chemical Physics*, 122 (2005) 184102.
- R. Gunawan, K. G. Gadkar, and F. J. Doyle III. Methods to identify cellular architecture and dynamics from experimental data. In Z. Szallasi, J. Stelling, and V. Periwal (editors), *System modeling in cellular biology: from concepts to nuts and bolts*. MIT Press, Cambridge, 2006. 221–242.
- E. L. Haseltine and J. B. Rawlings. Approximate simulation of coupled fast and slow reactions for stochastic chemical kinetics. *The Journal of Chemical Physics*, 117 (2002) 6959–6969.

- J. Hattne, D. Fange, and J. Elf. Stochastic reaction-diffusion simulation with MesoRD. *Bioinformatics*, 21 (2005) 2923–2924.
- R. Heinrich and S. Schuster. The regulation of cellular systems. Chapman & Hall, New York, 1996.
- D. J. Higham and R. Khanin. Chemical Master versus Chemical Langevin for first-order reaction networks. *The Open Applied Mathematics Journal*, 2 (2008) 59–79.
- C.-Y. F. Huang and J. E. Ferrell. Ultrasensitivity in the mitogen-activated protein kinase cascade. *Proceedings of the National Academy of Sciences USA*, 93 (1996) 10078–10083.
- H.-W. Kang and T. G. Kurtz. Separation of time-scales and model reduction for stochastic reaction networks, 2010. Manuscript in preparation.
- H.-W. Kang, T. G. Kurtz, and L. Popovic. Diffusion approximations for multiscale chemical reaction models, 2010. Manuscript in preparation.
- I. Karatzas and S. E. Shreve. Brownian motion and stochastic calculus. Springer, New York, second edition, 1998.
- S. Karlin. A first course in stochastic processes. Academic Press, New York, 1966.
- P. E. Kloeden and E. Platen. Numerical solution of stochastic differential equations. Springer, Berlin, 1992.
- A. Kremling, S. Fischer, K. Gadkar, F. J. Doyle, T. Sauter, E. Bullinger, F. Allgöwer, and E. D. Gilles. A benchmark for methods in reverse engineering and model discrimination: problem formulation and solutions. *Genome Research*, 14 (2004) 1773–1785.

- C. Kreutz and J. Timmer. Systems biology: experimental design. *FEBS Journal*, 276 (2009) 923–942.
- T. G. Kurtz. Strong approximation theorems for density dependent Markov chains. *Stochastic Processes and their Applications*, 6 (1978) 223–240.
- Y. Lan, T. C. Elston, and G. A. Papoian. Elimination of fast variables in chemical Langevin equations. *The Journal of Chemical Physics*, 129 (2008) 214115.
- C. H. Lee and H. G. Othmer. A multi-time-scale analysis of chemical reaction networks: I. Deterministic systems. *Journal of Mathematical Biology*, 60 (2010a) 387–450.
- C. H. Lee and H. G. Othmer. A multi-time-scale analysis of chemical reaction networks: II. Stochastic systems, 2010b. Manuscript in preparation.
- A. Levchenko and P. A. Iglesias. Models of eukaryotic gradient sensing: application to chemotaxis of amoebae and neutrophils. *Biophysical Journal*, 82 (2002) 50–63.
- D. A. McQuarrie. Stochastic approach to chemical kinetics. *Journal of Applied Probability*, 4 (1967) 413–478.
- B. Mélykúti, E. August, A. Papachristodoulou, and H. El-Samad. Discriminating between rival biochemical network models: three approaches to optimal experiment design. *BMC Systems Biology*, 4 (2010) 38.
- B. Mélykúti, K. Burrage, and K. C. Zygalakis. Fast stochastic simulation of biochemical reaction systems by alternative formulations of the chemical Langevin equation. *The Journal of Chemical Physics*, 132 (2010) 164109.
- B. Munsky and M. Khammash. The finite state projection algorithm for the solution of the chemical master equation. *The Journal of Chemical Physics*, 124 (2006) 044104.

- J. D. Murray. Mathematical biology, volume 1–2. Springer–Verlag, New York, 2003.
- K. G. Murty and S. N. Kabadi. Some NP-complete problems in quadratic and nonlinear programming. *Mathematical Programming*, 39 (1987) 117–129.
- B. Øksendal. Stochastic differential equations. An introduction with applications. Springer, New York, sixth edition, 2007.
- A. Papachristodoulou. Scalable analysis of nonlinear systems using convex optimization. Ph.D. thesis, Control and Dynamical Systems, California Institute of Technology, Pasadena, CA, USA, 2005.
- A. Papachristodoulou and H. El-Samad. Algorithms for discriminating between biochemical reaction network models: towards systematic experimental design. In Proceedings of the 2007 American Control Conference. 2714–2719.
- P. A. Parrilo. Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization. Ph.D. thesis, California Institute of Technology, Pasadena, CA, USA, 2000.
- J. Paulsson and M. Ehrenberg. Noise in a minimal regulatory network: plasmid copy number control. *Quarterly Reviews of Biophysics*, 34 (2001) 1–59.
- R. G. Pinsky. Positive harmonic functions and diffusion. Cambridge University Press, Cambridge, 1995.
- S. Prajna, A. Papachristodoulou, and P. A. Parrilo. SOSTOOLS – Sum of Squares Optimization Toolbox, User’s Guide, 2002. Software available at <http://www.eng.ox.ac.uk/control/sostools>.

- C. V. Rao and A. P. Arkin. Stochastic chemical kinetics and the quasi-steady-state assumption: Application to the Gillespie algorithm. *The Journal of Chemical Physics*, 118 (2003) 4999–5010.
- M. A. J. Roberts, E. August, A. Hamadeh, P. K. Maini, P. E. McSharry, J. P. Armitage, and A. Papachristodoulou. A model invalidation-based approach for elucidating biological signalling pathways, applied to the chemotaxis pathway in *R. sphaeroides*. *BMC Systems Biology*, 3 (2009) 105.
- A. Singer, R. Erban, I. G. Kevrekidis, and R. R. Coifman. Detecting intrinsic slow variables in stochastic dynamical systems by anisotropic diffusion maps. *Proceedings of the National Academy of Sciences USA*, 106 (2009) 16090–16095.
- A. Singh and J. P. Hespanha. Lognormal moment closures for biochemical reactions. In Proceedings of the 45th IEEE Conference on Decision and Control 2006. 2063–2068.
- V. Sotiropoulos, M.-N. Contou-Carrere, P. Daoutidis, and Y. N. Kaznessis. Model reduction of multiscale chemical Langevin equations: a numerical case study. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 6 (2009) 470–482.
- F. St-Pierre and D. Endy. Determination of cell fate selection during phage lambda infection. *Proceedings of the National Academy of Sciences USA*, 105 (2008) 20705–20710.
- D. W. Stroock and S. R. S. Varadhan. Multidimensional diffusion processes. Springer, New York, 1997.
- J. F. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optimization Methods and Software*, 11–12 (1999) 625–653. Software available at <http://sedumi.ie.lehigh.edu>.

- Z. Szallasi, J. Stelling, and V. Periwal (editors). System modeling in cellular biology: from concepts to nuts and bolts. MIT Press, Cambridge, 2006.
- L. Szpruch and D. J. Higham. Comparing hitting time behavior of Markov jump processes and their diffusion approximations. *Multiscale Modeling and Simulation*, 8 (2010) 605–621.
- K. Takahashi, S. Tănase-Nicola, and P. R. ten Wolde. Spatio-temporal correlations can drastically change the response of a MAPK pathway. *Proceedings of the National Academy of Sciences USA*, 107 (2010) 2473–2478.
- T. Tian and K. Burrage. Binomial leap methods for simulating stochastic chemical kinetics. *The Journal of Chemical Physics*, 121 (2004) 10356–10364.
- K.-C. Toh, R. H. Tütüncü, and M. J. Todd. SDPT3 - a MATLAB software package for semidefinite-quadratic-linear programming, 1999. Software available at <http://www.math.nus.edu.sg/~mattohk/sdpt3.html>.
- R. Tomioka, H. Kimura, T. J. Kobayashi, and K. Aihara. Multivariate analysis of noise in genetic regulatory networks. *Journal of Theoretical Biology*, 229 (2004) 501–521.
- M. Ullah and O. Wolkenhauer. Investigating the two-moment characterisation of sub-cellular biochemical networks. *Journal of Theoretical Biology*, 260 (2009) 340–352.
- P. J. M. Van Haastert and P. R. Van der Heijden. Excitation, adaptation, and deadaptation of the cAMP-mediated cGMP response in *Dictyostelium discoideum*. *The Journal of Cell Biology*, 96 (1983) 347–353.
- N. G. van Kampen. Stochastic processes in physics and chemistry. North-Holland Personal Library. Elsevier, Amsterdam, third edition, 2007.

- J. S. van Zon and P. R. ten Wolde. Green's-function reaction dynamics: A particle-based approach for simulating biochemical networks in time and space. *The Journal of Chemical Physics*, 123 (2005) 234910.
- J. Wilkie and Y. M. Wong. Positivity preserving chemical Langevin equations. *Chemical Physics*, 353 (2008) 132–138.
- D. J. Wilkinson. Stochastic modelling for systems biology. Mathematical and Computational Biology Series. CRC Press, Boca Raton, 2006.
- D. J. Wilkinson. Stochastic modelling for quantitative description of heterogeneous biological systems. *Nature Reviews Genetics*, 10 (2009) 122–133.
- J. C. Willems. Dissipative dynamical systems, part I: general theory. *Archive for Rational Mechanics and Analysis*, 45 (1972a) 321–351.
- J. C. Willems. Dissipative dynamical systems, part II: linear systems with quadratic supply rates. *Archive for Rational Mechanics and Analysis*, 45 (1972b) 352–393.
- O. Wolkenhauer, M. Ullah, W. Kolch, and K.-H. Cho. Modeling and simulation of intracellular dynamics: choosing an appropriate framework. *IEEE Transactions on NanoBioscience*, 3 (2004) 200–207.
- H. Yue, M. Brown, F. He, J. Jia, and D. B. Kell. Sensitivity analysis and robust experimental design of a signal transduction pathway system. *International Journal of Chemical Kinetics*, 40 (2008) 730–741.
- K. Zhou and J. C. Doyle. Essentials of robust control. Prentice Hall, New York, 1997.