UPPSALA
UNIVERSITET

Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology 832

Multiscale Stochastic Simulation of Reaction-Transport Processes

Applications in Molecular Systems Biology

ANDREAS HELLANDER

ACTA
UNIVERSITATIS
UPSALIENSIS
UPPSALA
2011

ISSN 1651-6214
ISBN 978-91-554-8089-9
urn:nbn:se:uu:diva-152098

Dissertation presented at Uppsala University to be publicly examined in Room 2446, Polacksbacken, Lägerhyddsvägen 2D, Uppsala, Friday, June 10, 2011 at 10:15 for the degree of Doctor of Philosophy. The examination will be conducted in English.

Abstract
Hellander, A. 2011. Multiscale Stochastic Simulation of Reaction-Transport Processes. Applications in Molecular Systems Biology. Acta Universitatis Upsaliensis. *Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology* 832. 63 pp. Uppsala. ISBN 978-91-554-8089-9.

Quantitative descriptions of reaction kinetics formulated at the stochastic mesoscopic level are frequently used to study various aspects of regulation and control in models of cellular control systems. For this type of systems, numerical simulation offers a variety of challenges caused by the high dimensionality of the problem and the multiscale properties often displayed by the biochemical model.

In this thesis I have studied several aspects of stochastic simulation of both well-stirred and spatially heterogenous systems. In the well-stirred case, a hybrid method is proposed that reduces the dimension and stiffness of a model. We also demonstrate how both a high performance implementation and a variance reduction technique based on quasi-Monte Carlo can reduce the computational cost to estimate the probability density of the system.

In the spatially dependent case, the use of unstructured, tetrahedral meshes to sample realizations of the stochastic process is proposed. Using such meshes, we then extend the reaction-diffusion framework to incorporate active transport of cellular cargo in a seamless manner. Finally, two multilevel methods for spatial stochastic simulation are considered. One of them is a space-time adaptive method combining exact stochastic, approximate stochastic and macroscopic modeling levels to reduce the simualation cost. The other method blends together mesoscale and microscale simulation methods to locally increase modeling resolution.

Keywords: stochastic simulation, chemical master equation, reaction-diffusion master equation, unstructured mesh, active transport, hybrid methods, URDME

Andreas Hellander, Department of Information Technology, Division of Scientific Computing, Numerical Analysis, Box 337, Uppsala University, SE-75105 Uppsala, Sweden.

© Andreas Hellander 2011

ISSN 1651-6214
ISBN 978-91-554-8089-9
urn:nbn:se:uu:diva-152098 (http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-152098)

To my family

List of Papers

This thesis is based on the following papers, which are referred to in the text by their Roman numerals.

I A. Hellander and P. Lötstedt. Hybrid method for the chemical master equation, J. Comput. Phys., 227(1):100–122, 2007.

II E. Caulfield and A. Hellander. CellMC – a multiplatform model compiler for the Cell Broadband Engine and x86, Bioinformatics, 26(3):426–428, 2010.

III A. Hellander. Efficient computation of transient solutions of the chemical master equation based on uniformization and quasi–Monte Carlo, J. Chem. Phys., 128, 154110, 2008.

IV S. Engblom, L. Ferm, A. Hellander and P. Lötstedt. Simulation of stochastic reaction-diffusion processes on unstructured meshes, SIAM J. Sci. Comput., 31(3):1774–1797, 2009.

V A. Hellander and P. Lötstedt. Incorporating active transport of cellular cargo in stochastic mesoscopic models of living cells, Multiscale Model. Simul., 8(5):1691–1714, 2010.

VI L. Ferm, A. Hellander and P. Lötstedt. An adaptive algorithm for simulation of stochastic reaction-diffusion processes, J. Comput. Phys., 229(2):343–360, 2010.

VII A. Hellander, S. Hellander and P. Lötstedt. Coupling mesoscopic and microscopic simulation of stochastic reaction-diffusion processes in mixed dimensions, Technical report 2011-005, Department of Information Technology, Uppsala University, 2011. *(Submitted)*

Although not explicitly discussed in the comprehensive summary, the following papers are related to the contents of this thesis.

1. B. Drawert, S. Engblom and A. Hellander. URDME v1.1: User's manual, Technical report 2011-003, Department of Information Technology, Uppsala University, 2011.
2. L. Ferm, P. Lötstedt and A. Hellander. A hierarchy of approximations of the master equation scaled by a size parameter, J. Sci. Comp., 34(2):127–151, 2008.
3. M. Hegland, A. Hellander and P. Lötstedt. Sparse grids and hybrid methods for the chemical master equation, BIT Num. Math., 48(2):265–283, 2008

Reprints were made with permission from the publishers.

Contents

1 Introduction 9
2 Mesoscopic modeling and simulation of reaction-diffusion processes 13
2.1 Mesoscale description of well-stirred reaction networks 13
2.2 Numerical simulation methods for well-stirred systems 15
2.2.1 Direct solution metods 15
2.2.2 Stochastic simulation and Monte Carlo 16
2.2.3 Monte Carlo simulation of reaction networks using CellMC 18
2.2.4 Variance reduction based on quasi-Monte Carlo 19
2.3 Introducing spatial dependence 21
2.4 Simulating the RDME on unstructured meshes 23
2.5 Computational approaches to spatial stochastic simulation 27
2.6 Validity of the RDME 29
3 Beyond diffusive transport 31
4 Multilevel modeling and simulation of reaction-diffusion models . . 35
4.1 An hybrid method for well-stirred systems 35
4.2 A space-time adaptive method for spatial stochastic simulation 37
4.3 Coupling the microscopic and mesoscopic modeling levels . . . 41
5 Conclusions 45
6 Summary in Swedish 47
7 Author's contributions 51
8 Acknowledgment 53
Bibliography 55

1. Introduction

Quantitative modeling of biochemical reaction kinetics has proven itself useful in studying the behavior of diverse cellular phenomena, and has emerged as one of the most important tools in molecular systems biology. Based on prior knowledge or assumptions about its molecular interactions, the dynamic behavior of a biochemical network can be investigated *in silico*, providing a powerful complement to more traditional, experimental approaches.

A prominent theme in computational systems biology is the study of stochastic models of biochemical reaction networks. Many of the regulatory pathways in the living cell will, due to very low copy numbers of some of their constituents, need to operate reliably despite a significant level of intrinsic noise in the network. For example, the robustness of gene regulatory networks to noise has been a frequently studied topic and several mechanisms that have the potential of attenuating noise have been identified [107, 33, 106, 90].

When stochastic effects can be expected to play an important role for a system's dynamics, the classical, deterministic reaction-rate equations (RREs) are inadequate to describe the time evolution of the system. The RREs are formulated as a system of, generally non-linear, ordinary differential equations. They describe the time evolution of the concentration, or mean value, of a species and are as such posed on the macroscopic level. A better approach is to use a model on the so-called mesoscopic scale, where the discreteness of the species and the randomness of reactions are accounted for explicitly. The most popular mathematical framework that takes the effects of intrinsic noise into account is continuous time discrete space Markov processes. The Markov models are simple to formulate based on the description of a system in terms of elementary reactions, and several studies have highlighted the potential modeling benefit from using them over the deterministic models in a cell biology context [83, 6, 33, 109, 86, 6].

The vast majority of stochastic models has thus far been staged in a well-stirred setting. Assuming that molecules have time to explore the whole reaction volume before participating in any reactive encounter eliminates the need to explicitly model the movement of molecules in the domain. Even so, compared to the macroscopic RREs the model is harder to analyze. The probability density function (PDF) of the system is governed by the chemical master equation (CME). The CME is linear and simple in its structure, but presents a formidable challenge for direct solution methods due to the large number

of degrees of freedom that grows exponentially with the number of chemical species in the model. The Stochastic Simulation Algorithm (SSA) [44] can be used to generate independent realizations of the stochastic process, and quantities of interest can be estimated in a Monte Carlo fashion. Though simple in principle, stochastic simulation with SSA can also be computationally challenging due to the multiscale properties frequently displayed by models of reaction kinetics in molecular cell biology. This has motivated extensive method development, including the work presented here in papers I– III.

The well-stirred assumption facilitates the computational analysis of a system compared to a spatially resolved case. However, there are many biological phenomena for which a well-stirred assumption is too restrictive. Localization or compartmentalization of components in reaction networks are known to be a strategy employed by the cell for regulation. For examples where both the spatial distribution of molecules as well as stochastic fluctuations play a role see [118, 38, 108]. Recent advances in experimental techniques can provide data not only on the molecular copy number of a mRNA or a protein, but also on their location in the cell with individual molecule resolution [88, 31]. As these techniques develop further, methodology to accurately and efficiently simulate stochastic models with spatiotemporal resolution will continue to grow in importance.

The CME is readily generalized to a spatially heterogenous setting, the reaction-diffusion master equation (RDME), by discretizing the domain and using the well-stirred assumption locally within the voxels in the mesh. Molecular transport is then simulated as jumps between adjacent grid cells using methods derived from the Gillespie algorithm [8, 30, 61, 105]. Public domain software for mesoscopic reaction-diffusion modeling and simulation are described in [2, 22, 51, 115]. Unstructured meshes offer a greater geometrical flexibility and have the potential to resolve curved surfaces in a better way than uniform Cartesian meshes do. Simulation of the RDME on unstructured, tetrahedral meshes is considered here in paper IV.

Almost all work in the area of spatial stochastic simulation has focused on reaction-diffusion models. However, diffusion is not the only mechanism with which intracellular cargo can be transported in the living, eukaryote cell. In paper V, we propose a method to extend the mesoscopic simulation framework to include intracellular cargo transport on microtubules. This type of transport is active in the sense that is requires an energy input, and can be much faster than diffusion.

If diffusion is fast and/or the mesh is fine, the majority of the simulated events will be due to movement of molecules. This renders spatially resolved simulation much more computationally demanding than their well-stirred counterparts. By aggregating diffusive transfers simulations can be sped up [27, 72, 94]. In paper VI we propose a space-time adaptive method that dynamically switches between the mesoscopic and macroscopic modeling level to reduce the cost of simulating diffusion.

Another complication in the spatially resolved case is the fact that as the mesh becomes very fine, the RDME modeling framework breaks down and simulation results become inaccurate if the model contains bimolecular reactions [59, 36, 37]. This goes against any intuition derived from numerical solution of partial differential equations. If a fine mesh is needed to resolve features of the model, this is an indication that we are approaching a regime where a more detailed model is necessary. A more fine scaled alternative to the RDME framework is obtained by explicitly accounting for each molecule's position continuously in space. Diffusion is described by Brownian motion and the time until chemical reactions occur between nearby molecules is obtained by solving well-defined partial differential equations (PDEs) in methods based on Smoluchowski dynamics [120, 119, 4, 20, 108, 119, 56]. Such microscale methods can provide insight into a model's behaviour where the macroscopic or mesoscopic modeling framework fails [36, 37, 108, 119], but to a higher simulation cost. For models where microscopic resolution is needed in parts of the domain or for certain species, we propose a hybrid method in paper VII that couples mesoscopic simulation using URDME [22] with a microscopic method [56].

Multiscale aspects of the simulation of kinetic reaction-transport models permeate the work presented in this thesis. We make contributions to stochastic simulation of both well-stirred models (papers I–III) and of spatially heterogenous systems (papers IV–VII). A unifying theme for papers I, VI and VII is the aim to improve simulation efficiency for systems with multiscale properties. The approach is to combine the macroscale, mesoscale and microscale modeling frameworks in different ways. An alternative viewpoint is that the proposed methods improve on a coarser level model by locally increasing modeling resolution at a moderately increased cost. In paper V it is rather the desire to model active transport on the mesoscale without explicitly accounting for all the microtubule fibers that calls for multiscale modeling. Starting out from a microscale view of the cargo's movement, we formulate a mesoscopic model which we then connect to equations on the macroscopic scale. In this way we arrive at a convenient way of seamlessly incorporating active transport in software for simulation of stochastic reaction-diffusion processes [22].

The rest of the thesis is organized as follows. In Chapter 2 we review the mathematical model underlying the CME and the RDME along with computational methods to study them numerically. There, we also introduce the use of unstructured meshes to sample realizations of the stochastic process. The extension of the RDME framework to incorporate active transport is presented in Chapter 3. Finally, in Chapter 4 we discuss the multilevel approaches to stochastic simulation proposed in papers I, VI and VII.

2. Mesoscopic modeling and simulation of reaction-diffusion processes

In this chapter we will introduce the mesoscopic modeling framework that lies at the heart of all of the papers that make up this thesis. Starting out from the classical treatment of well-stirred mesoscopic models and some of the most frequently used methods to simulate them, we will progress to spatially heterogenous models. The state of the art in methodology and software for simulating such systems will be reviewed and we will discuss how realizations can be generated using unstructured, triangular and tetrahedral meshes. We will also comment briefly on recent work on the validity of the RDME framework.

2.1 Mesoscale description of well-stirred reaction networks

On the mesoscopic scale, the dynamics of a biochemical reaction network is modeled as a continuous time discrete space Markov processes. The state of the system is described by the random vector $\mathbf{X}(t)$ whose components $\{X_1,\ldots,X_N\}$ describe the total copy number of each species $j = 1\ldots N$. The stochastic process $\{\mathbf{X}_t\}_{t\geq 0}$ is a collection of such random variables indexed by time and describes the time evolution of the system. The species can take part in M different chemical reactions in the well-stirred container Ω. When a reaction r occurs after a time τ_r it changes the current state according to the *stoichiometry vector* $\mathbf{n}_r$. Introduce a realization of the process at time t, $\mathbf{x}(t) \in \mathbb{Z}_+^N$ having components x_j. Reaction r then changes the state according to $\mathbf{x}' = \mathbf{x}(t+\tau_r) = \mathbf{x}(t) + \mathbf{n}_r$. The *propensity function*, $\omega_r(\mathbf{x})$, describes the rate of the reaction and does not depend on time. It is interpreted as the probability that the reaction occurs in an infinitesimal time interval $[t, t+dt]$

$$\omega_r(\mathbf{x}) = \lim_{dt\to 0} \frac{P(\mathbf{x}+\mathbf{n}_r, t+dt) - P(\mathbf{x},t)}{dt}. \tag{2.1}$$

Depending on the particular description of the chemical reaction, ω_r will take different forms.

As an example, consider a small system with species A, B and C and a reaction pair describing the reversible binding between A and B forming the complex C.

$$(1)\ A+B \xrightarrow{k_1} C,\ (2)\ C \xrightarrow{k_2} A+B. \tag{2.2}$$

Assuming mass action kinetics the propensity function for reaction one is $\omega_1 = k_1 ab$ and that for reaction two is $\omega_2 = k_2 c$ where lower case letters denote molecular copy numbers of the respective chemical species. The stoichiometry vectors are in this case $\mathbf{n}_1 = [-1,-1,1]$ and $\mathbf{n}_2 = -\mathbf{n}_1$.

The Markov process has the important memoryless property, i.e. given the state of the system at time t, the state at any future time depends only on the current state and not on anything that happened at a time s before t, $s < t$. As a consequence, the time to the occurrence of reaction r is an exponentially distributed random variable $\tau_r \sim Exp(1/\omega_r(\mathbf{x}))$.

As a direct consequence of the Markov property the probability density function (PDF) of the system is governed by the forward Kolmogorov equation, or the chemical master equation (CME). For brevity of notation, in the following we will write $p(\mathbf{x},t) = p(\mathbf{x},t|\mathbf{x}_0,t_0)$ for the probability that the system can be found in state $\mathbf{x}$ at time t conditioned on the initial condition $\mathbf{x}_0$ at time t_0. The CME then takes the form

$$\frac{\mathrm{d}}{\mathrm{dt}} p(\mathbf{x},t) = \mathscr{M}\, p(\mathbf{x},t) \equiv \sum_{r=1}^{M} \omega_r(\mathbf{x}-\mathbf{n}_r)p(\mathbf{x}-\mathbf{n}_r,t) - \omega_r(\mathbf{x})p(\mathbf{x},t). \tag{2.3}$$

The first term in the sum describes flux of probability into state $\mathbf{x}$, and the second term the outflow of probability. The sum is only taken over *feasible* reactions, i.e. those reactions that result in a non-negative state, $\mathbf{x}+\mathbf{n}_r \geq 0$. For a derivation of the master equation, consult e.g. [113, Chap. V] or [45].

An ordinary differential equation for the first moment of X_j is obtained by summing over all states, $E[X_j(t)] = \sum_{\mathbf{x}\in\mathbb{Z}_+^N} x_j p(\mathbf{x},t)$,

$$\frac{\mathrm{d}}{\mathrm{dt}} E[X_j] = \sum_{r=1}^{M} n_{rj} E[\omega_r(X)], \quad j = 1,\ldots,N. \tag{2.4}$$

Let the volume of the reaction container be $|\Omega|$ and introduce the concentration of species X_j, $\langle X_j \rangle = E[X_j/|\Omega|]$. If we assume that $E[\omega_r(X)] \approx \omega_r(E[X])$ we recognize the much used macroscopic reaction rate equations, RREs, for the concentrations of the species,

$$\frac{\mathrm{d}}{\mathrm{dt}} \langle X_j \rangle = \sum_{r=1}^{M} \mathbf{n}_{rj} \tilde{\omega}_r(\langle X \rangle), \quad j = 1,\ldots,N, \tag{2.5}$$

where $\tilde{\omega}_r$ is an appropriate scaling of the mesoscopic rate function ω_r. This set of equations, however, is not an exact consequence of (2.3) or (2.4) unless all propensity functions $\omega_r(\mathbf{x})$ are linear in $\mathbf{x}$. However, it provides a good approximation when the copy numbers are large and the solution converges to the exact expected value in the thermodynamic limit [71].

The CME can be solved analytically only for special, simple models e.g. the classical birth-death process, or for systems with only monomolecular reactions (a subset of reactions with linear propensity functions ω_r) [64]. In other cases, one has to rely on simplifying approximations or numerical methods to study the model.

Frequently used approximations of the CME include the linear noise approximation (LNA), for examples in the systems biology context see e.g. [49, 29]. The LNA is the second order approximation in the Ω–expansion [113, Chap. X], the first order approximation is simply the reaction rate equations (2.5). By expanding (2.4) and closing the equations by truncating at higher order moments more accurate deterministic equations can be formulated [34, 46]. The use of the Fokker-Planck approximation is explored in [100]. The RREs, LNA and the two moment approximation (2MA) [34] are compared for several examples in [40].

2.2 Numerical simulation methods for well-stirred systems

Numerical methods to study the behavior of reaction networks as described by the Markov model above fall into two main classes: direct solution of the CME (2.3) with matrix based methods and methods based on stochastic sampling and Monte Carlo.

2.2.1 Direct solution metods

In order to solve the CME directly, the infinite state space has to be restricted by imposing upper bounds on the number of the different species. In principle, the resulting truncated master equation is simple to solve using an appropriate time-stepping scheme. However, even if upper bounds on the variables can be imposed, the size of the system will still be prohibitive, and techniques have thus been developed to reduce the state space further. By approximating the CME by the Fokker-Planck equation (FPE), Sjöberg et. al. make use of a continuous state space model to reduce the number of degrees of freedom [102, 32, 101]. Using space-time adaptivity, Ferm et al. further improve on the method [41].

By using the observation that in a short time interval, most of the probability mass will remain localized in a region in state space close to the starting support, Munsky and Khammash proposed the finite state projection (FSP)

method [85]. There, the local time steps in the time-stepping scheme are chosen in such a way that the probability of leaving the (locally restricted) computational domain is kept below some critical tolerance. Burrage et al. use FSP together with Krylov subspace methods for computation of the matrix exponential [10, 99] and this is further improved on by MacNamara et. al. [78]. MacNamara et al. also extend this scheme to study Markov processes with explicit time dependence [77].

By aggregating states, Ferm and Lötstedt reduced the state space in [39]. Hegland has proposed the use of sparse grids [52] and Krylov metods [99]. Jahnke et. al. use wavelet compression to reduce the state space [63, 65], and Engblom and Deufelhardt et. al. independently developed discrete spectral methods for the CME [35, 25]. Based on adaptive uniformization Mateescu et. al. constructed an efficient solver in [81].

Despite this extensive work on solving the CME for increasingly large systems, the approach is still limited to models with relatively few species, or with special properties that keep the state space small.

2.2.2 Stochastic simulation and Monte Carlo

Realizations of the stochastic process can be generated in a simple way. Given the current state of the system, the times until each reaction fires, $\tau_r,\ r = 1,\ldots,M$ are sampled from the exponential distribution $\mathrm{Exp}(\omega_r(\mathbf{x})^{-1})$. Then, the system is updated one step by executing the reaction corresponding to the shortest sampled reaction time, $\min_r \tau_r$ by updating the state according to $\mathbf{n}_r$. By using the observation that the minimum of independent identically distributed exponential random variables is itself an exponential random variable a more efficient formulation of the algorithm is Gillespie's direct method [44], commonly known as the stochastic simulation algorithm (SSA). In other contexts than chemical kinetics, the algorithm is known as kinetic Monte Carlo (kMC) and was introduced independently in [117, 1]. The direct method is outlined in Algorithm 1.

Algorithm 1 SSA (Gillespie's direct method)

1. Initialize: Compute all propensities $\omega_r(\mathbf{x}), r = 1,\ldots,M$, $t = 0$.

while $t < T$ **do**

2. Compute the sum of the propensities, $\lambda = \sum_{r=1}^{M} \omega_r(\mathbf{x})$.
3. Draw two random numbers u_1 and u_2 from the uniform distribution.
4. Sample the next reaction time (by inversion), $\tau = -\frac{\log(u_1)}{\lambda}$.
5. Sample the next reaction event (by inversion), i.e. find r such that $\sum_{i=1}^{r-1} \omega_i(\mathbf{x}) < \lambda u_2 \leq \sum_{i=1}^{r} \omega_i(\mathbf{x})$.
6. Update the state vector, $\mathbf{x} = \mathbf{x} + \mathbf{n}_r$.
7. $t = t + \tau$.

end while

The SSA has been used extensively in computational systems biology, and many improvements to its implementation have been suggested. In the direct method the step to determine the next reaction to fire given the time until the next event is of linear complexity. For systems with many reaction channels, implementations that reduce this complexity have been developed, for overviews see [98, 97]. Best known in the systems biology context is the logarithmically scaling next reaction method (NRM) by Gibson and Bruck [43]. This algorithm has found much use recently as a basis for algorithms tailored for the spatially dependent case [30]. An extensive overview of many varieties of the implementation of the SSA can be found in [82].

Frequently, models of biochemical reaction networks possess a property that has been coined *stochastic stiffness*. A few reactions in the system are occurring on much faster timescales than the others. In an exact stochastic simulation every event is simulated and the majority of the computational time will be spent sampling the fastest reactions. At the same time, the most interesting stochastic dynamics are likely to occur on the slower time scales in the system, for species with low copy numbers. Three main categories of methods targeted at dealing with stiffness have emerged: tailored implementations of the exact method that explicitly account for stiffness, approximate stochastic methods and hybrid methods. In the first category, notable contributions are the Optimized Direct Method (ODM) [17] by Cao et. al. and the sorting direct method (SDM) by McCollum et al. [84].

The most popular approximate, stochastic method is the tau-leaping method [48] which relies on the existence of a macroscopic time where the propensities change slowly, cf. [47]. The basic idea is to let several (fast) reactions fire in the same time interval. The number of events that occur for a certain reaction in a time interval Δt is approximately a Poisson distributed random variable, $N_r \sim \mathrm{Po}(\omega_r(\mathbf{x})\Delta t)$, and can be conveniently sampled by standard methods for generating Poisson variates. By aggregating events in a time interval Δt many fast reactions can be "leaped over", improving simulation times. Many versions and improvements of the original method have been proposed [111, 19, 87, 12, 14, 15, 3]. Several versions of the tau–leaping algorithm, as well as exact simulation methods that take stochastic stiffness into account are implemented in the StochKit software [16] package.

Hybrid methods achieve their improved simulation speed by treating parts of the system at a less accurate modeling level. Typically, hybrid methods can improve on simulation times with a large factor, often greater than that of tau-leaping, but they rely on a good partitioning of the system, something that can be challenging in practice unless the modeller possesses some prior knowledge about the system under study. The chemical Langevin equation is combined with SSA in [50, 95]. In the nested SSA method [28] an inner SSA is used for the fast scale reactions and an outer for the slow scale. The method developed in paper I combines the stochastic and deterministic modeling levels and will be discussed further in Chapter 4. Based on the quasi

steady-state assumption the QSSA algorithm relies on Michaelis-Menten kinetics to achieve speed-up [89]. A similar approach is taken in the slow-scale SSA [13, 11]. The validity of the Michaelis-Menten approximation is studied in [96].

There has been an extensive development of methods of the type summarized here, especially for tau-leaping methods. The review above is by no means complete. More extensive summaries on methods to simulate well-stirred kinetics on the mesoscopic level can be found in e.g. [54, 75].

2.2.3 Monte Carlo simulation of reaction networks using CellMC

Whenever stochastic simulation as described in the previous section is to be applied to a model it will be necessary to generate multiple realizations in order to study the system in a Monte Carlo fashion. From ensembles of trajectories the PDF solving (2.3) or other metrics of the system's performance can be estimated to within a certain confidence level. Generating a large set of independent realizations is inherently a task-parallel problem. The SSA algorithm, being very simple in its structure, is a natural target for implementation on specialized multicore architectures. The use of Field Programmable Gate-Arrays (FPGA) was explored in [116] and in [76, 69] general purpose graphics processing units (GPUs) were used to generate large ensembles in parallel.

In paper II, we describe the development of a software package, CellMC, that uses SIMD vector instructions and threading to generate ensembles efficiently on multicore x86 architectures or the Cell Broadband engine, Cell/BE. The software is freely available from http://www.cellmc.org. More details regarding the implementation can be found in the master's thesis report by Emmet Caulfield [18]. CellMC is implemented as a code-generating program that will translate a model expressed in the Systems Biology Markup Language (SBML) [58] into highly optimized code for the platform on which it is executed. That source code is then internally compiled into executables using the GNU compiler *gcc*. At the time of publication, CellMC was the only software to utilize vector instructions on the PC to harness the full capability of modern x86 processors and to be able to utilize Cell/BE for stochastic simulation with SSA.

Table 2.1 shows simulation performance of CellMC in both absolute simulations times and million reactions per second (Mrps). Ensembles were generated on different platforms (PC x86 and Cell/BE) for two model problems (abbreviated dd and hsr) commonly used as benchmarks in methods development, see e.g. [48, 76]. Runs on the PC were conducted using single precision arithmetic to be comparable to results for the Cell/BE platforms.

At the time of writing, the Cell/BE is no longer pursued as an architecture for general purpose high performance computing. Even so, the techniques and implementation developed in CellMC are generally applicable and provide

Machine	Type	CPU/SPUs	Time		Mrps	
			dd	hsr	dd	hsr
espirit	PC	2	22.03	49:49	35.89	20.77
os	PC	4	8.58	23:58	91.9	44.61
skara	PS3	6	6.97	15:46	113.24	66.01
grad	PC	8	3.65	9:13	218.23	112.27
cell2	QS22	16	2.53	6:03	312.83	174.43
Cluster	4×PS3	24	1.80	4:10	439.60	261.14

Table 2.1: *CellMC performance for a range of different machines (ordered in increasing order of performance). Note how a single PS3 game console outperforms a quad-core PC by almost a factor of two. The timings are for* 10^3 *hsr trajectories run to 500s and* 3×10^4 *dd trajectories run to 10s. Timings are given in seconds (dd) and minutes and seconds (min:sec, hsr).*

good speed-up also on x86 architectures. When talking about the performance of a new implementation on e.g. a GPU card over that of a single CPU or even the host workstation, it is worth remembering that using vector instructions, 4 realizations can be generated simultaneously per core using single precision, see paper II.

2.2.4 Variance reduction based on quasi-Monte Carlo

In many cases Monte Carlo techniques are the only feasible way of numerically studying the dynamics of reaction networks as described by the mesoscopic model. Monte Carlo methods suffer from a slow convergence rate, but have the advantage of being simple and relatively insensitive to the dimensionality of the system. In the previous section we illustrated how efficient, parallel implementations could aid in the collection of large amounts of realizations from which statistics about the system can be computed. The approximate and hybrid methods discussed in Section 2.2.2 have the common aim to reduce the time with which a single realization can be generated. If the goal is to produce large ensembles to estimate some unknown quantity, another approach is to try to reduce the variance in the estimates [92, 24]. In such a way one can achieve speed-up by reducing the number of trajectories that needs to be sampled in order to achieve some specified error tolerance.

Monte Carlo methods converge as $L^{-1/2}$ where L is the ensemble size. The width of a 95% confidence interval for a quantity estimated by repeated sampling is

$$\varepsilon = \frac{1.96\bar{s}}{\sqrt{L}} \tag{2.6}$$

where $\bar{s}$ is the empirical standard deviation. Variance reduction techniques then aim at reducing $\bar{s}$ and thus the uncertainty ε. In paper III, we explore the possibility to use quasi-Monte Carlo (QMC) techniques to reduce the cost of estimating the PDF of the system using SSA. Quasi-random sequences, such as the Sobol sequence, can be used in the same way as pseudo-random numbers but they have a low-discrepancy property; they fill space more evenly than regular pseudo-random number generators. Figure 2.1 illustrates the intuitive difference by generating random numbers in a unit square. By chance, the pseudo-random numbers form clusters and regions where the point-density is lower (right) while the Sobol sequence fills the square more evenly (left). When used for numerical integration, intuitively this means that for a fixed number of points S, the support for the integral is sampled more uniformly and the probability of missing important regions of the integrand is lower than when using classical random number generators. QMC methods have the potential to improve on the convergence rate in Monte Carlo methods. For sufficiently smooth integrands, a convergence rate of almost L^{-1} can be obtained.

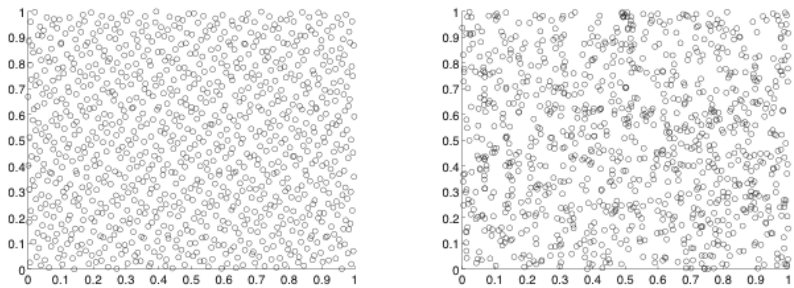

Figure 2.1: A Sobol sequence (left) samples the unit square with lower discrepancy than a pseudo-random sequence (right).

To use quasi-random sequences in the SSA we converted the continuous-time Markov process to a discrete-time chain subordinate to a Poisson processes via the method of uniformization [66]. In practice, this is done by introducing a null-reaction in the model so that the distribution for the next event-time is the same in every iteration of the algorithm, $\tau \sim \mathrm{Po}(\lambda_{max})$ where λ_{max} is an upper bound on the sum of the propensities over the whole time-course of the simulation. We could then apply previously published QMC methods for discrete chains [73] to estimate the PDF of the system.

In Figure 2.2 we visually compare the difference in variance between the marginal PDF of species A and B from the model problem used in paper III when estimated using the classical SSA (a) and the QMC extension (b) for a fixed number of trajectories (10^5). As can be seen, the PDF produced with the QMC methods is much smoother. For this example, the convergence rate

is improved using quasi-random sequences (c). The execution time to achieve a certain error tolerance can be more than an order of magnitude lower for the QMC method (d). As an example, to achieve the same resolution with SSA as that in Figure 2.2b obtained with 10^5 realizations, 10^7 trajectories would have to be sampled, giving a speed-up of 33 compared to our reference implementation of SSA (for timings for the reference implementation see paper II).

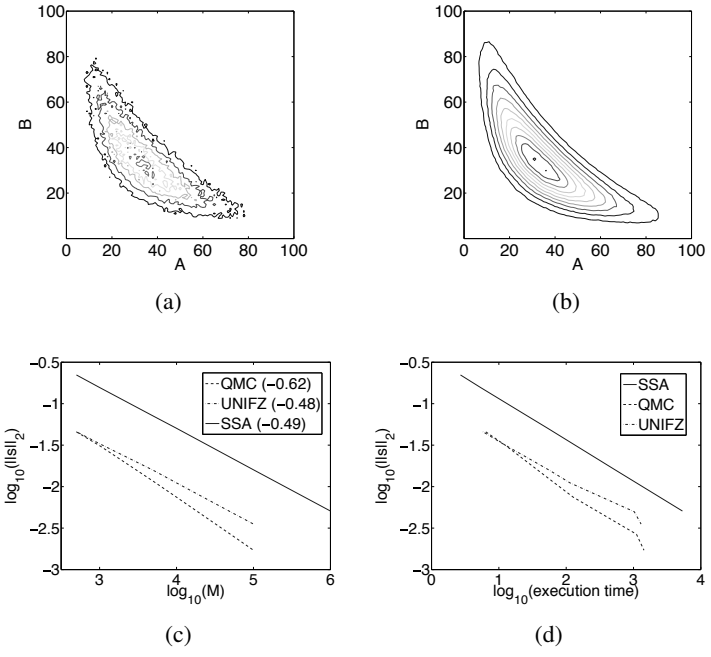

Figure 2.2: The marginal densities of A and B computed with 10^5 trajectories by SSA (a) and QMC SSA (b). The convergence rate (c) and execution time (d) are improved in the QMC method.

2.3 Introducing spatial dependence

The assumption of a well-stirred reaction container can be rather restrictive. Many processes in a living cell involves spatial localization or separation of proteins to function properly. Especially in euakaryotes, different parts of a pathway will occur in different cellular compartments. The most obvious example is gene expression, where transcription occurs in the nucleus and translation in the cytoplasm, often located at the membrane of the endoplasmic reticulum.

The mesoscopic description can be extended to the spatially hetreogenous case. To introduce molecular transport by diffusion we discretize space and

divide the domain Ω into K non-overlapping voxels, or subvolumes, $\mathscr{V}_i$. For a system with N species the system can now be described by the $K \times N$ matrix of random variables $\mathbf{X}(t)$. Again, denote a realization of the system with $\mathbf{x} \in \mathbb{Z}_+^{KN}$. The entries of the row $\mathbf{x}_{i\cdot}$ in the matrix are the copy numbers of species $j = 1, \ldots, N$ in voxel i. Chemical reactions occur within the voxels by treating them as in the spatially homogenous case. In other words, we assume that inside any voxel in the mesh, the assumption of a well-stirred system holds so that molecules are effectively uniformly distributed in the subvolume at any time. For example, a bimolecular reaction r involving species X_j and X'_j in voxel i can in this setting be written

$$X_{ij} + X_{ij'} \xrightarrow{\omega_{ir}(\mathbf{x}_{i\cdot})} X_{ij''}. \tag{2.7}$$

The propensity function depends only on the copy number of the species in voxel i so that reactions are completely local. In general, the propensity can take different forms in different voxels. This will be the case if the rate constants vary in space, or more commonly, if a reaction is only active in parts of the domain. If no molecular motion occurs, the master equation for the full system will be

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}p(\mathbf{x},t) = \mathscr{M}\,p(\mathbf{x},t) \equiv \\ \sum_{i=1}^{K}\sum_{r=1}^{M} \omega_{ir}(\mathbf{x}_{i\cdot} - \mu_{ir})p(\mathbf{x}_{1\cdot}, \ldots, \mathbf{x}_{i\cdot} - \mu_{ir}, \ldots, \mathbf{x}_{K\cdot}, t) \\ - \sum_{i=1}^{K}\sum_{r=1}^{M} \omega_{ir}(\mathbf{x}_{i\cdot})p(\mathbf{x},t),\end{aligned} \tag{2.8}$$

where μ_{ir} is the $1 \times N$ stoichiometry vector for reaction r in voxel i. For the case of one subvolume, $K = 1$, (2.8) reduces to the CME (2.3).

The diffusive motion of a molecule from one voxel to an adjacent one is modeled by the linear jump event

$$X_{ij} \xrightarrow{d_{ijk}} X_{kj} \tag{2.9}$$

with propensity function $a(\mathbf{x}_{i\cdot}) = d_{ijk}x_{ij}$. It takes a species j in voxel i to one of the immediate neighbors $\mathscr{V}_k$. The transition can be written $\mathbf{x}'_{\cdot j} = \mathbf{x}_{\cdot j} + \nu_{ijk}$. The $K \times 1$ stoichiometry vector ν_{ijk} acts on the column $\mathbf{x}_{\cdot j}$ and has all components zero except for $\nu_{ijk}(i) = -1$ and $\nu_{ijk}(k) = 1$.

For the case of molecular transport only, the master equation then takes the form

$$\frac{\mathrm{d}}{\mathrm{d}t}p(\mathbf{x},t) = \mathscr{D}p(\mathbf{x},t) \equiv$$
$$\sum_{j=1}^{N}\sum_{i=1}^{K}\sum_{k=1}^{K} a(\mathbf{x}_{\cdot j} - \nu_{ijk})p(\mathbf{x}_{\cdot 1},\ldots,\mathbf{x}_{\cdot j} - \nu_{ijk},\ldots,\mathbf{x}_{\cdot N},t)$$
$$-\sum_{j=1}^{N}\sum_{i=1}^{K}\sum_{k=1}^{K} a(\mathbf{x}_{\cdot j})p(\mathbf{x},t), \tag{2.10}$$

For a system with both reactions and diffusion, the full RDME is then simply

$$\frac{\mathrm{d}}{\mathrm{d}t}p(\mathbf{x},t) = \mathscr{M}p(\mathbf{x},t) + \mathscr{D}p(\mathbf{x},t). \tag{2.11}$$

Next, we will discuss how realizations of the process can be sampled using unstructured, tetrahedral meshes.

2.4 Simulating the RDME on unstructured meshes

In the formulation of the RDME (2.11) in the previous section, we have yet to specify the values of the rate constants d_{ijk} in the expression of the diffusion propensity function (2.9). Those values will depend on the nature of the spatial discretization, i.e. the shapes and sizes of the voxels. In this section, we will focus on the case where the nature of the molecular motion is diffusion; in the next chapter we will discuss an extension to include active transport. To simplify notation, we assume that the diffusion constant is the same for every species everywhere in space and drop one index in d_{ik}.

In the Markov model, the value of the propensity function for an event corresponds to the inverse of the expected time for that event to occur. In other words, the value $1/(d_{ik}x_{ij})$ is the expected time for the first species of type X_j to leave voxel $\mathscr{V}_i$ and become well-mixed in voxel $\mathscr{V}_k$. A reasonable requirement on the rate constants d_{ik}, irrespective of the mesh, is that they are chosen in such a way so that the concentration process converges to the diffusion equation in the thermodynamic limit. On the macroscale, normal diffusion is described by the heat equation

$$\frac{\partial u(\zeta,t)}{\partial t} = \sigma\Delta u(\zeta,t). \tag{2.12}$$

This means that d_{ik} has to correspond to the elements in the matrix D that arise from a convergent discretization of (2.12)

$$\hat{u}_t = D_{(h)}\hat{u} \tag{2.13}$$

where h is a discretization length parameter as in the previous section and

$$\lim_{h\to 0} \hat{u} = u. \tag{2.14}$$

In addition, all coefficients d_{ik} need to be non-negative for the stochastic process to be well defined. Also, in the standard formulation of the RDME, interactions should be local, occurring only between nearest neighbors in the mesh. In other words, if boundaries are reflective, a matrix D that is suitable for a stochastic simulations has the properties

$$D_{ik} \geq 0, \quad D_{ii} \leq 0, \quad \sum_{k=1}^{K} D_{ik} = 0. \tag{2.15}$$

Figure 2.3 shows parts of a Cartesian mesh (a) and a triangular mesh (b) in 2D. For the unstructured mesh, molecules are assumed to be well-mixed in the dual elements indicated by the thin solid lines in the figure. The same interpretation holds for a Cartesian mesh, where the dual is the staggered grid with respect to the primal mesh (dashed lines).

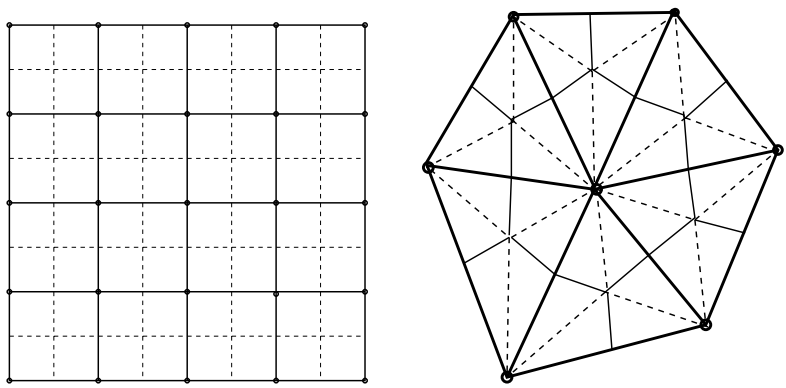

(a) (b)

Figure 2.3: Parts of a Cartesian mesh (a) and an unstructured triangular mesh (b). The dual of the triangular mesh is obtained by connecting the midpoints of the edges and the centroids of the triangles (thin solid lines). In every such subvolume molecules are assumed to be well-stirred, i.e. have a uniform distribution.

For a uniform Cartesian discretization the jump rates in the RDME are naturally defined by the corresponding coefficients in a second order central finite difference discretization, $d_{ik} = \sigma/h^2$, where h is the length of the voxels [42, 113]. For an unstructured mesh made up from tetrahedra, the most natural choice is to obtain the rates from either a first order finite element (FE) discretization with linear Lagrange elements or from a two-point flux finite volume scheme for (2.12). For a Cartesian mesh, both those schemes

are equivalent to the finite element scheme, i.e. would result in the same rate constants and (2.15) is fulfilled.

In paper IV we choose to base the jump coefficients on a FE discretization. There are two main reasons for this choice: the convergence of the finite volume scheme is sensitive to the quality of the mesh and moreover, there are many software packages available for solutions of partial differential equations based on FEM. We have made use of this to to implement a flexible software package for spatial stochastic simulation, URDME [22].

With a FE discretization of (2.12) with linear basis and test functions, one obtains the following system of equations for the values of u at the vertices in the primal mesh

$$Mu_t = \sigma Su. \tag{2.16}$$

M is commonly termed the mass matrix while S is known as the stiffness matrix. Consult e.g. [9] for an introduction to the finite element method. M and S have the same sparsity pattern, and in order to simplify time integration a technique known as mass lumping is frequently employed. In our case, for the stochastic simulation, this is essential for the molecules to jump only to neighboring grid cells. Let $\hat{M}$ be the diagonal, lumped mass matrix M with $\hat{M}_{ii} = \sum_{k=1}^{K} M_{ik}$. Then, for the dual elements described in Fig. 2.3, $\hat{M}_{ii} = |\mathcal{V}_i|$, i.e. the volume of voxel i [110]. Together, this defines the jump coefficients as the entries in the matrix

$$D \equiv \sigma \hat{M}^{-1} S. \tag{2.17}$$

Figure 2.4 is adopted from paper IV and shows an example of the use of unstructured meshes to simulate a reaction network in a non-trivial geometry. The computational domain is representative for a yeast cell. Two large internal structures are present, the nucleus and a large vacuole. The finite element software Comsol Multiphysics 3.4 [21] was used to define the geometry and mesh. The stochastic simulation was based on the next subvolume method (NSM) [30] as implemented in an early version of the URDME software [23]. For details about the model see [30].

In a finite volume discretization, the jump coefficients d_{ik} will always take positive values irrespective of the shapes of the voxels, but the scheme does not converge for general meshes. The conditions on the mesh for which (2.15) holds for a FE discretization are discussed in paper IV. For FEM, the convergence properties are better. However, for a general mesh such as those readily available from third-party, state of the art mesh generators there is no guarantee that the off-diagonals in D will have the same sign. In a stochastic simulation, this is unacceptable, and in that case we need to truncate the matrix D to remove the elements with wrong sign. This will render also the FEM dis-

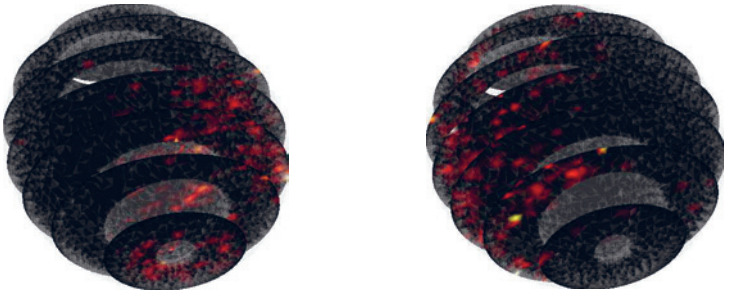

(a) t=700s (b) t=700s

Figure 2.4: Snapshots of a model proposed by Elf. et. al. [30] simulated using an unstructured mesh. The left panel shows a species A and the right panel species B, see [30] and paper IV for details. Light and dark colors represent high and low concentrations, respectively. Due to the combination of slow diffusion and feedback in the model, A and B form localized clusters in different parts of the domain.

cretization inconsistent with (2.12). Even so, the truncated matrix describes a discrete diffusion, it is a local operator and has the properties:

1. The conditions (2.15) are satisfied.
2. The discrete maximum principle is satisfied [70].
3. The non-negativity of the species is preserved [70].
4. Mass is conserved.

The potential inconsistency does not necessarily mean that the error it causes in a computed quantity will be large compared to other sources of error. What it means is that at some mesh resolution, the discretization error in the diffusion part of the reaction-diffusion equation does no longer decrease as the mesh size is reduced. However, what is more important is the level of the error when that happens, and how this is related to other errors in the simulation. aFurther work is needed to fully establish what errors this incurs in the stochastic simulations for realistic models, but early work indicates that they are small compared to the errors caused by the inconsistencies due to small subvolumes [68].

Another approach to incorporate diffusion in complex geometries is proposed by Isaacson and Peskin [61, 60]. There, the domain is discretized by a uniform Cartesian mesh and the curved boundary is treated by an embedded boundary method. Also in their approach, the jump coefficients are derived from a numerical discretization of the macroscale diffusion equation (2.12) or as pointed out in [61], equivalently from discretizing the Fokker-Planck

equation for Brownian motion. The PDF for the position of a single molecule undergoing Brownian motion is governed by the Fokker–Planck equation

$$\frac{\partial p(\zeta,t)}{\partial t} = \sigma \Delta p(\zeta,t) \tag{2.18}$$

In the limit of vanishing voxel size, the strategy described above to obtain jump constants for use within the mesoscopic RDME framework is thus consistent with both the concentrations of species in the macroscopic limit and the probability of a single diffusing Brownian particle to be found inside a given voxel in the mesh (in steady-state this probability is simply given by $|\mathscr{V}_i|/|\Omega|$ by (2.15)). Compared to our approach in paper IV, the method in [61] has the advantage to rely on structured Cartesian meshes in interior regions having the potential for higher interior accuracy and possible optimizations for an implementation of the stochastic simulation algorithm which are not possible in the case of general meshes. A method using tetrahedral unstructured meshes on the other hand, is simpler to apply in the general case and provides a natural way of handling also species diffusing on the actual boundary.

2.5 Computational approaches to spatial stochastic simulation

The introduction of molecular motion into the model has not changed the basic mathematical properties of the master equation. However, even for few species and a very coarse discretization, the dimensionality of the model will be so high that direct solution methods based on matrix methods are infeasible. The total number of events in the system is high for a fine discretization, and simply applying Gillespie's direct method to simulate the system would be very inefficient. Elf and Ehrenberg proposed the next subvolume method (NSM), an adaptation of NRM tailored for reaction-diffusion systems. The complexity of NSM is then logarithmic in the number of voxels in the mesh. Note however, that this does not mean that the total simulation time for a system grows logarithmically as the mesh becomes finer. Table 2.2 shows the execution time for an implementations of NSM [30] as implemented in the URDME [22] software package, when applied to the simulation of Min oscillations in the rod-shaped bacterium *E. Coli* [38].

As will be made clear in the next section, the number of diffusion events scales as the inverse of the mesh resolution squared as seen in the final column of the table. For the finest mesh resolution, 99.98% of the events are due to diffusion in this example. This stiffness is characteristic for space dependent systems with diffusion. This can be compared to solution of partial differential equations, where the stiffness of diffusion equations are known to render explicit time stepping schemes unsuitable. For a first order explicit

$h_{\max}$	# voxels	t [s]	# events/s $\times 10^6$	# diffusion events $\times 10^8$
2×10^{-7}	1555	56	3.0	1.7
1×10^{-7}	10759	333	1.8	6.0
5×10^{-8}	80231	2602	0.8	22.8

Table 2.2: *Simulation times for a spatial dependent system simulated with NSM as implemented in URDME. The execution time (third column) grows rapidly with increased mesh resolution. $h_{\max}$ is the maximum local mesh size allowed. The system was simulated to a final time 200s.*

time marching scheme, the time step must be chosen such that $\Delta t \leq \frac{Ch^2}{\sigma}$ for stability, where C is a constant and σ is the diffusion constant. In Table 2.2 the increased cost of simulation when the mesh size is halved comes from two sources. When h is reduced by a factor of two, the number of diffusion events increases by a factor of four. Also, since the number of voxels grows the cost of generating a single event increases (column four in the table).

As illustrated by the example above, reaction-diffusion systems tend to be diffusion dominated, at least for well-resolved geometries. Unlike the well-stirred case where a few reaction channels are causing stochastic stiffness, here it is rather the whole diffusion operator that causes the stiffness and there will most likely be a continuum of average inter-event times, even if they will form groups due to species having different diffusion constants. Rossinelli et al. use tau-leaping and combines tau-leaping and deterministic diffusion to speed up simulations [94]. Binomial tau-leaping for spatial stochastic simulation was proposed in [80]. Additional implementations of spatial tau-leaping are introduced in [62] and compared to the implementations in [80, 94]. When explicit tau-leaping is applied in a spatial context, it is necessary to generate one Poisson variate in every direction (edges) in the mesh for every voxel in order for the scheme to be conservative. Lampoudi attempts to reduce the number of variates for 2D systems in a multinomial simulation algorithm [72]. The performance of spatial tau-leaping, the Gillespie Multiparticle Method (GMP) [93] and NSM are compared in [67] and the GMP has been implemented for GPUs in [114]. Another approach to reduce the cost of simulating diffusion is taken in the Diffusive Finite State Projection (DFSP) algorithm [27], in which the FSP method [85] is used locally in space to aggregate diffusive transfers in a small time interval. For systems with scale separation, we propose a space-time adaptive hybrid algorithm in paper VI. There, deterministic diffusion, explicit tau-leaping and exact SSA are combined and the appropriate method is dynamically chosen in each voxel in space based on easily computable criteria. By adaptive mesh refinement it is possible to locally increase the resolution and in that way save computational time over a uniformly fine mesh [7].

2.6 Validity of the RDME

The mesh-size parameter h is an important parameter in a simulation using the RDME and in one important respect the mesoscopic model evades intuition based on the numerical solution of PDEs. There, provided that the numerical scheme is convergent, a smaller mesh size will lead to a more accurate solution of the equation. For the RDME modeling framework on the other hand, in the limit of vanishing voxel sizes the model breaks down in the sense that it does not become a good approximation to the desired microscopic model [59]. From physical intuition, the voxel sizes cannot be taken to be too small nor too large. The average time it takes for a molecule to diffuse the length of a voxel should be small compared to the time scale with which the molecule reacts so that spatial correlations are short-lived compared to the time between reactions. When in this regime, simulations based on the RDME were found experimentally to agree well with molecular dynamics simulations in [5].

The mesh size must be taken small enough as to resolve the geometrical features of the domain, otherwise the discretization error in the diffusion and reaction operators will be large and simulations will be inaccurate. To the best of our knowledge, no satisfactory *a priori* error estimates for the effects of a large mesh size on the results obtained from simulation of the stochastic model is available however. For error in expectation, in the thermodynamic limit, one can of course apply standard theory from numerical solution of PDEs to obtain such error-estimates and this would provide bounds in a weak sense.

In a Cartesian mesh the expected time it takes for a species A with diffusion constant σ_A to transition to and become well-mixed in a neighboring grid-cell is $\tau_D = h^2/\sigma_A$. If A participates in a bimolecular reaction, the expected time for it to react with a species B in voxel containing one molecule of each kind is $\tau_r = k_a/h^3$ where k_a is the reaction constant and has units of volume over time. Comparing the two times, requiring that $\tau_D \ll \tau_r$, the mesh size should fulfill $h \gg k_a/max(\sigma_A, \sigma_B)$. Expanding the RDME to second order in the reaction radius Isaacson proved in [59] that for a chemically reacting system with the reaction $A + B \xrightarrow{k_a} \emptyset$, when $h \to 0$ the second order term in the expansion diverges as $1/h$ compared to the corresponding term in an expansion of the equations of the Smoluchowski model. Intuitively, contributions of the bimolecular reactions become successively smaller (and completely lost in the limit). Isaacson suggests that in order to alleviate this problem it is necessary to "appropriately renormalize the bimolecular reaction rate and/or extend the reaction operator to couple in neighboring voxels". However, he also notes that for parameters frequently occurring in real biochemical systems it is possible to find a mesh resolution that gives less than 1% error in distribution between the correspondingly truncated models. For highly diffusion limited parameter combinations however, the error incurred by too fine a mesh can be substantial [36, 37].

In [36] Erban and Chapman derives a modified expression for the propensity function of a bimolecular reaction in order to obtain the correct steady-state distribution in the case of a box-geometry and uniform Cartesian mesh. In their approach, each reaction event is local to the voxel and they find a smallest mesh resolution h for which it is possible to correct the model in this way. If h is sufficiently small compared to the side length of the cubic domain, this critical value is approximated by $h_{crit} \approx 0.25272 k_a/(\sigma_A + \sigma_B)$ [36]. Note that this value is close to the intuitive lower bound on h above. Fange et. al. [37] both derive corrected propensities (in 2D and 3D) and allow for reactions between molecules in neighboring grid-cells. In this way, they obtain mesh-independent steady-state levels of molecules for spatial resolutions all the way down to the reaction-radii of the molecules and good agreement in numerical experiments to methods based on the Smoluchowski framework for the model proposed in [108]. Both approaches [36, 37] are compared for a few model problems in [68] and found to work reasonably well. Neither of the methods have been applied to unstructured meshes but the method [37] seems more straightforward to extend to the case of a general mesh. As an alternative approach to handle this problem, in paper VII we demonstrate how a mesoscopic-microscopic multiscale algorithm can be used to simulate the same model problem [108] without the need for a very fine mesh.

3. Beyond diffusive transport

Diffusion is the dominating mechanism of molecular transport in prokaryotes such as *E. Coli*, and it was to models of such systems the reaction-diffusion master equation and the next subvolume method were first applied in a systems biology context [30, 38]. However, diffusion is not the only way molecules are transported in the eukaryote cell. Molecular cargo can be transported actively on cytoskeleton structures made up of microtubule and actin polymers. Molecular motor proteins bind to the cargo and to the cytoskeleton and move the cargo in a certain direction depending on the molecular nature of the motor. This transport is usually much faster than diffusion but requires an energy input. Vesicles, organelles, mRNA and proteins involved in signaling are examples of molecular cargo that is transported actively in the living cell, for reviews see [57, 79, 112].

In paper V we propose a multiscale method for introducing active transport of cellular cargo on microtubules. In order to make the method computationally feasible, we connect the micro- and mesoscopic model of cargo transport to a discretization of a model on the phenomenological, macroscopic level. This is done in order to be able to apply the same machinery as described in the previous section to simulate the resulting stochastic process on unstructured (and structured) meshes using previously available algorithms and software. Transport of convective nature can be incorporated into the model in a completely analogous way to that of diffusion by allowing for a more general transport operator $\mathscr{T}$ in the master equation,

$$\frac{\mathrm{d}}{\mathrm{dt}}p(\mathbf{x},t) = (\mathscr{M}+\mathscr{T})p(\mathbf{x},t), \quad \mathscr{T} = \mathscr{D}+\mathscr{A}, \tag{3.1}$$

where $\mathscr{A}$ will describe molecular transfers due to active transport modeled as convection. To arrive at jump coefficients that will dictate the rates of jumps between voxels, we start out with the macroscopic linear transport equation, written here in conservative form

$$\frac{\partial u(\zeta,t)}{\partial t} = \nabla\cdot(\mathbf{v}(\zeta,t)u(\zeta,t)) \tag{3.2}$$

where ζ are coordinates in a Cartesian coordinate system. The velocity field $\mathbf{v}(\zeta,t)$ then models the rate and directivity with which the molecules are transported. Following the same logic as in Section 2.4, we require that the discrete

approximation of (3.2) is local, limiting the connection of voxels to their immediate neighbors. As for diffusion, we also require the jump constants to take non-negative values. The natural choice for a discretization fulfilling those requirements is a first order, upwind finite volume scheme. In the finite volume method, the equation (3.2) is written on integral form using Gauss' theorem. Let the unknown quantities be the average concentrations in the voxels

$$\bar{u}_i = \frac{1}{|\mathscr{V}_i|}\int_{\mathscr{V}_i} u(\zeta,t)d\Omega. \tag{3.3}$$

Let $\partial\mathscr{V}_i$ be the polygonal boundary of voxel $\mathscr{V}_i$ and $\mathbf{n}$ its unit surface normal. Integrating (3.2) and using Gauss' theorem we obtain

$$\int_{\mathscr{V}_i}\frac{\partial u}{\partial t}\,d\Omega + \int_{\mathscr{V}_i}\nabla\cdot(\mathbf{v}u)\,d\Omega = \frac{\partial}{\partial t}\int_{\mathscr{V}_i} u\,d\Omega + \int_{\partial\mathscr{V}_i}\mathbf{n}\cdot\mathbf{v}u\,dS = 0. \tag{3.4}$$

Using the average value (3.3) we have

$$|\mathscr{V}_i|\frac{\partial\bar{u}_i}{\partial t} = -\int_{\partial\mathscr{V}_i}\mathbf{n}\cdot\mathbf{v}u\,dS. \tag{3.5}$$

By approximating the integral in the right hand side of (3.5) we have a finite volume discretziation of (3.2). The simplest choice that leads to a scheme with local connections and all positive off diagonal elements in the corresponding jump-matrix is obtained by evaluating the velocity field $\mathbf{v}_{ik}$ at the center of the face e_{ik} connecting voxels $\mathscr{V}_i$ and $\mathscr{V}_k$ and the surface normal $\mathbf{n}_{ik}$ directed from $\mathscr{V}_i$ into $\mathscr{V}_k$ and letting

$$q_{ik} = \frac{1}{|\mathscr{V}_i|}\int_{e_{ik}}\mathbf{n}\cdot\mathbf{v}u\,dS \approx \frac{|e_{ik}|}{|\mathscr{V}_i|}f(\mathbf{v}_{ik}\cdot\mathbf{n}_{ik},u_i,u_k), \tag{3.6}$$

where

$$f(\mathbf{v}_{ik}\cdot\mathbf{n}_{ik},u_i,u_k) = \begin{cases} \mathbf{v}_{ik}\cdot\mathbf{n}_{ik}u_i, & \mathbf{v}_{ik}\cdot\mathbf{n}_{ik}\geq 0 \\ \mathbf{v}_{ik}\cdot\mathbf{n}_{ik}u_k, & \mathbf{v}_{ik}\cdot\mathbf{n}_{ik}<0 \end{cases} \tag{3.7}$$

The matrix A discretizing the right hand side in (3.2) has off-diagonal elements q_{ik} and diagonal entries $q_{ii} = -\sum_{k\neq i}q_{ik}$. An event of convective nature taking a molecule of type j in $\mathscr{V}_i$ to $\mathscr{V}_k$ can then be simulated as in (2.9) with propensity $a(\mathbf{x}_{i\cdot}) = q_{ik}x_{ij}$ allowing for a seamless integration with software to generate realization of the RDME.

This procedure to define the stochastic jump process, together with that for diffusion which was outlined in Section 2.4 and in detail in paper IV can be used to simulate a general transport process. Here, in the same manner as for the diffusive case the correct jump coefficients on an unstructured mesh

are conveniently obtained by discretizing a phenomenological equation, in this case the linear transport equation. The key in any biological modeling using this framework is then the modeling of the macroscopic velocity field $\mathbf{v}(\zeta,t)$. In paper V, we actually start out in the other end, by considering the microscopic properties of cargo transport by molecular motor proteins on a microtubule filament network. By aggregating the effect of many fibers in the voxels, we derived a mesoscopic propensity without first invoking the macroscopic level. In principle, the mesoscopic propensity could be obtained directly if it was possible to, from an image of the fiber network, count the number of fibers entering and exiting through every face in a dual mesh overlaying the image, as well as measuring the total length of the fibers in each voxel.

If we assume that on the scale of the mesh resolution, the fiber network is "well-behaved", i.e. that fibers do not cross each other and change direction appreciably, we can relate the mesoscopic formulation of our model to the framework outlined above with the velocity field $\mathbf{v} = v\mathbf{t}(\zeta)$ where $\mathbf{t}(\zeta)$ is a smoothly varying vector field approximating the fiber network and v is the average speed of the motor protein (typically kinesin or dynein). However, as we point out in the paper, the procedure to formulate the mesoscopic model outlined here is general. It can be followed if another macroscopic velocity field could be derived by incorporating more microscale details. Also, the velocity field can of course model transport of convective nature arising from completely different sources than motor-driven cargo transport.

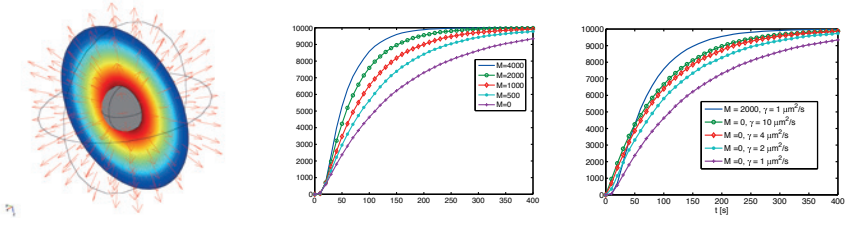

(a) (b) (c)

Figure 3.1: Response curves showing how the response species R accumulates at the cell membrane with varying filament numbers (M) (b) and for different diffusion constants (c). The macroscopic velocity field and density modeling the fiber network is shown in (a).

Figure 3.1a shows the geometry and the macroscopic velocity field used to model the motor-driven transport in paper V. Microtubules are thought to originate from an area at the nucleus (pictured in grey) and extend radially outwards towards the plasma membrane giving rise to a radially directed velocity field on the macroscale. We consider a simple example with two species that can either diffuse or be transported on the microtubule network after binding to a fiber. An abstract signal species originates at the plasma membrane, is

transported down to the nucleus where it is converted into an abstract response species which is then in turn transported back to the plasma membrane.

In Figures 3.1b and 3.1c we show the differences in response of the system for varying fiber densities (or equivalently, rate constant for binding to fibers) and different diffusion constants. This illustrates the different behavior between the two transport mechanisms and the potential for regulation of the timing of delivery by modulating the amount of time the motor proteins spends bound to the fiber network. This is a mechanism of regulation that has been proposed as one means of regulating the intricate transport machinery that helps maintaing the highly polarized state of healthy neurons. Little is known about the details of such regulation though microtubule associated proteins (MAPs) such a Tau seem likely to play an important role by being able to differentially regulate the affinity of kinesin and dynein to the fibers [26].

4. Multilevel modeling and simulation of reaction-diffusion models

Papers I, VI and VII share a common theme. They all deal with computational methods that aim at reducing the cost of numerical study of the biochemical network by formulating different parts of the system on different modeling levels. Many realistic reaction networks in molecular biology will have multiscale properties. Some species might be abundant and modeled sufficiently accurately on the macroscopic level, while in the same model, other parts may need a mesoscopic or even microscopic modeling resolution. In those cases, resolving the whole system on a modeling level dictated by the finest scales in the model will lead to unnecessarily costly simulations.

4.1 An hybrid method for well-stirred systems

For a general, high-dimensional well-stirred model, solving for the PDF of the full system via the CME is intractable. The high dimensionality of the problem makes the state space too large to study using matrix based direct methods. Stiffness in the model can render a stochastic simulation based on SSA slow, making also Monte Carlo study of the model challenging. Figure 4.1 shows an example of the frequencies of firings of the different reactions when simulating a model of the heat shock response in *E. Coli* [17]. This model has been used frequently in the literature as a benchmark for new, more efficient stochastic simulation methods as well as optimized implementations of the SSA.

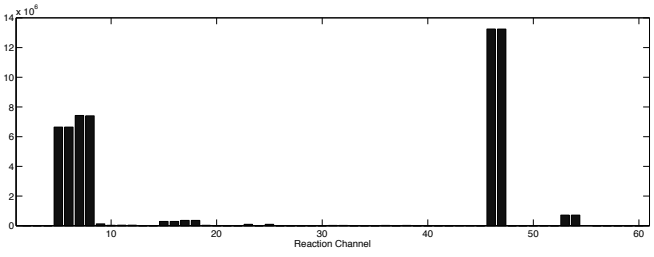

Figure 4.1: The frequencies of the reaction events in a model of the heat shock response in *E. Coli*. Six of the reactions account for almost all events occurring up to the end-time of the simulation. Figure adopted from [54]

As can be seen, only a few of the reactions account for almost all the events that occurred until the end time of the simulation. The multiscale nature of the model is clearly visible here, and if one is most interested in the statistics of a species that is involved in reactions occurring on one of the slower timescales, an attractive approach is to try to reduce the size and the stiffness of the model by dimension reduction. In the method described in paper I, we propose to reduce the dimensionality of the model by splitting the species in the model in two parts: one mesoscopic and one macroscopic part. The idea is that species with a small relative variance can be modeled on the macroscopic level, while the other species need a mesoscopic description since the fluctuations in their copy numbers are substantial compared to their mean values. This separation in variance was illustrated in paper I for a model of a mitogen-activated protein kinase (MAPK) cascade [74] where the relative standard deviation of the species varied in the range 0.02–1.08.

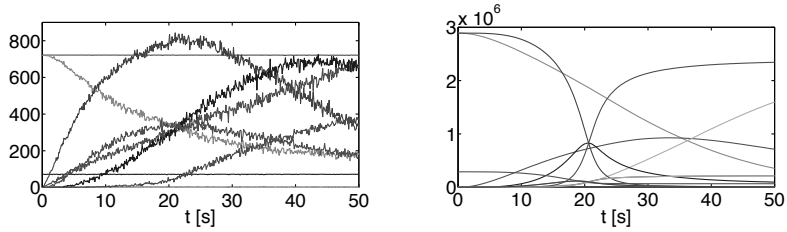

(a) (b)

Figure 4.2: In the model problem considered in [53], the separation in copy numbers is large. The species in (a) all stay below 1000 molecules for most of the time, while other species are present in the order of 10^6 molecules. Figure adopted from [53].

The relative variances will often correlate strongly with high or low copy numbers. A big spread in copy numbers is frequently occurring in realistic models. Consider for example a generic model of a signaling network, where the input to the system is a signal emanating at the cell surface, propagating through the cytosol via different proteins or even second messengers, eventually resulting in the translocation of a transcription factor to the nucleus. There, the transcription factor regulates a gene whose transcription might result in the production of a protein that acts as a negative or positive regulator. In such a case, copy numbers can range from one (the gene) to several thousands (some of the proteins) to tens or hundreds of thousands (second messengers). Figure 4.2 shows how the copy number varies between species in another model of a MAPK cascade used as model problem in [53]. The model consists of 22 species and 30 reactions. As can be seen to the left in the figure, eight of them have copy numbers below 1000 (a) while some species are present in the orders of 10^6 molecules (b). This particular model is very stiff, and was chosen as a test problem in [53] to accentuate the need for hybrid

simulations. Systems that are not amenable to good splitting and thus not to simulation with the hybrid method are discussed in paper I and in [53].

While not the only possible reason, see [11], large discrepancies in copy numbers amongst species is a frequent cause for stiffness in a model. When using mass action rate expressions, the propensity of a reaction is at least directly proportional to the copy number of the involved species, so unless there is a big difference in rate constants between two given reactions, a reaction involving high copy number species will be much faster than a reaction with all reactants present in low copy numbers.

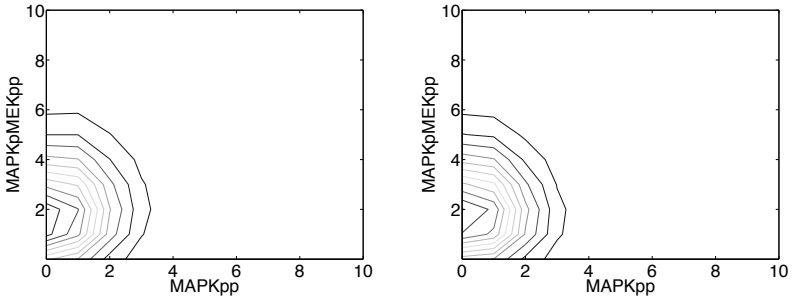

Figure 4.3: Isolines of the marginal probability distribution for MAPKpp and MAPKpMEKpp with SSA (a) and the hybrid method (b). The PDF obtained with the hybrid method approximates the PDF of the full system well.

Figure 4.3 shows marginal probability density functions for two of the species in Table 7 in paper I. For this model, by only treating a few high variance species as mesoscopic variables, we reduce simulation times by up to an order of magnitude compared to our reference implementation of SSA. For the extremely stiff MAPK model used as a test problem in [53] and shown in Figure 4.2, the speed-up was close to 300 when only eight of the species in the left panel was simulated as mesoscopic variables.

In summary, the hybrid method described in paper I improves simulation times by reducing the dimension of the system, thus leading to a smaller reaction network to simulate on the mesoscopic modeling level. If in addition the splitting also reduces the stiffness of the model, the speed-up over stochastic simulation of the full system can be substantial.

4.2 A space-time adaptive method for spatial stochastic simulation

In the spatially dependent case, as illustrated in Table 2.2, stiffness additionally arises from the molecules' diffusion. Diffusion events typically outnum-

ber reactions events, and matters become worse with a fine mesh resolution because of the $\mathscr{O}(1/h^2)$ scaling of the propensity function. For models where some species are diffusing fast or are abundant, or if the mesh is fine, it may become challenging to propagate a system on the time scale of interest using exact simulation on the mesoscopic level; most of the CPU time will be spent on propagating molecules in space in between reactive encounters. By simulating high copy number species on the macroscopic level, species with intermediate copy numbers by an approximate stochastic scheme and low copy number species by exact, mesoscopic stochastic simulation, we reduce the work involved in updating the diffusion part of the system in an adaptive method in paper VI. Importantly, the method is capable of dynamically changing the modeling resolution locally in space based on a few easily computable criteria. This space-time adaptivity is an essential feature of the algorithm, since one cannot expect an initial partitioning to be valid for long times if the system displays e.g. transient local bursts in protein or molecule levels.

To be able to update the degrees of freedom with different methods for diffusion, the reaction-diffusion operator is split into three parts,

$$\frac{\partial p(\mathbf{x},t)}{\partial t} = \mathscr{D}_m p(\mathbf{x},t) + \mathscr{D}_\tau p(\mathbf{x},t) + (\mathscr{M}+\mathscr{D})p(\mathbf{x},t). \tag{4.1}$$

$\mathscr{D}_m$ describes macroscopic diffusion, $\mathscr{D}_\tau$ diffusion by the approximate tau-leaping method [48, 14, 91], $\mathscr{M}$ is all chemical reactions in the system and $\mathscr{D}$ diffusion with the exact method. To simulate the hybrid system, (4.1) is evolved using an operator splitting approach [104]. Expanding the formal solution of (4.1) in a Taylor series at $t+\Delta t$ gives

$$\begin{aligned} p(\mathbf{x},t+\Delta t) &= e^{\Delta t(\mathscr{D}_m+\mathscr{D}_\tau+(\mathscr{M}+\mathscr{D}))}p(\mathbf{x},t) \\ &= e^{\frac{1}{2}\Delta t\mathscr{D}_m}e^{\frac{1}{2}\Delta t\mathscr{D}_\tau}e^{\Delta t(\mathscr{M}+\mathscr{D})}e^{\frac{1}{2}\Delta t\mathscr{D}_\tau}e^{\frac{1}{2}\Delta t\mathscr{D}_m}p(\mathbf{x},t)+\mathscr{O}(\Delta t^3). \end{aligned} \tag{4.2}$$

Truncating the $\mathscr{O}(\Delta t^3)$ terms gives a second order approximation of the PDF. The macroscopic part is then integrated in time using the forward Euler method, $\mathscr{D}_\tau$ with the explicit tau-leaping method [48]. The reactions $\mathscr{M}$ and the exact part of the diffusion operator $\mathscr{D}$ is updated using NSM.

To determine which species in which parts of the domain to simulate with which method, we first determine a time step such that all degrees of freedom can be simulated with the tau-leaping method. The error estimate is adapted from [91]. The local splitting time-step in step n of the algorithm is chosen such that

$$\Delta t^n \leq \left(\min_{j=1\ldots N} \frac{2\varepsilon_1}{\max_{i=1\ldots K} |(D^2\mathbf{x}^n_{\cdot j})_i / x^n_{ij}|} \right)^{1/2}, \tag{4.3}$$

where ε_1 is a desired bound on the relative error per time step in the discretization, thus making the global error proportional to this tolerance. D is the matrix with jump coefficients in (2.17). If, in addition, the following holds for a user specified tolerance ε_2

$$\sqrt{(|D|(\mathbf{x}^n_{\cdot j})_i}/\bar{x}_{ij} \leq \varepsilon_2 \tag{4.4}$$

where $|D|$ is the matrix with entries being the absolute values of the entries in D, then that degree of freedom (species j in $\mathscr{V}_i$) can be simulated on the macroscopic level. Note that the time step has been chosen such that *all* species can be updated with tau-leaping in that time step. This is not what we want because it will lead to small time steps. One of the major issues with many tau-leaping schemes is the fact that copy numbers can become negative. To determine which species should be simulated with SSA we estimate the probability that a species in a voxel will end up negative after Δt. If we accept that species X_{ij} end up negative with probability ε_3 we can use the following relation to determine whether it should be updated with SSA or tau-leaping

$$x^n_{ij} + \mu_i - \kappa(\varepsilon_3)\sqrt{v_i} < -1. \tag{4.5}$$

This expression is a consequence of the fact that the change in the copy number after Δt^n is Skellam-distributed (the difference between two Poisson random variables). μ_i is the difference between the parameters λ^i_+ and λ^i_- in the Poisson distributions describing the outflow and inflow of molecules to voxel $\mathscr{V}_i$ and v_i is the variance of $\lambda^i_+ + \lambda^i_-$. μ and v both depend on the time-step Δt^n. Using a normal approximation for this distribution we get (4.5). With $\varepsilon_1 = 0.01$, $\kappa = 2.33$, and with $\varepsilon_1 = 0.025$, $\kappa = 1.96$. When a certain number of degrees of freedom have been assigned the exact simulation method, we could most likely allow for a larger time step in (4.3). Thus in practice, (4.3), (4.4) and (4.5) is evaluated in an iterative manner for a few steps until the partitioning of the system changes little between evaluations.

The method is summarized in Algorithm 2, where the order of the steps is that used in [55].

Figure 4.4b shows a snapshot of the partitioning chosen by the algorithm at a given instance in time for the model problem considered in paper VI. A high copy number species that was initially present outside the cell has started spreading into the cytoplasm. Outside the cell, it is simulated on the macroscopic level (indicated in red). As is shown in Fig. 4.4a the mesh is finer near the plasma membrane, leading the algorithm to use tau-leaping in those regions (colored green). Inside the cytosol, and especially in the nucleus, the concentration in some of the voxels (colored blue) is not yet high enough for neither the macroscopic nor tau-leap solver to be applicable and the exact stochastic simulation algorithm is used.

Algorithm 2 The adaptive hybrid algorithm

Initialize: Choose error tolerances ε_1, ε_2, ε_3, a final time T_f and an initial state $\mathbf{x} = \mathbf{x_0}$.

while $t < T_f$ **do**

1. Compute the time step Δt and partition the degrees of freedom (dofs) according to the error estimates in paper VI. Place dofs that can be advanced in time by macroscopic diffusion in the set $\mathscr{J}_m$, dofs that can be tau-leaped (diffusion only) in the set $\mathscr{J}_\tau$ and the remaining dofs in $\mathscr{J}_s$. Set $\mathbf{x_1} = \mathbf{x}(t)$.
2. Advance $\mathbf{x_1}$ $\Delta t/2$ in time by updating dofs in $\mathscr{J}_m$ by macroscopic diffusion to get $\mathbf{x_2}$.
3. Advance $\mathbf{x_2}$ $\Delta t/2$ in time by diffusing dofs in $\mathscr{J}_s$ and updating all reactions by NSM to get $\mathbf{x_3}$.
4. Advance $\mathbf{x_3}$ Δt in time by updating dofs in $\mathscr{J}_\tau$ by explicit tau-leaping for diffusion to get $\mathbf{x_4}$.
5. Advance $\mathbf{x_4}$ $\Delta t/2$ in time by diffusing dofs in $\mathscr{J}_s$ and updating all reactions by NSM to get $\mathbf{x_5}$.
6. Advance $\mathbf{x_5}$ $\Delta t/2$ in time by updating dofs in $\mathscr{J}_m$ by macroscopic diffusion to get $\mathbf{x}(t+\Delta t)$.
7. $t = t + \Delta t$.

end while

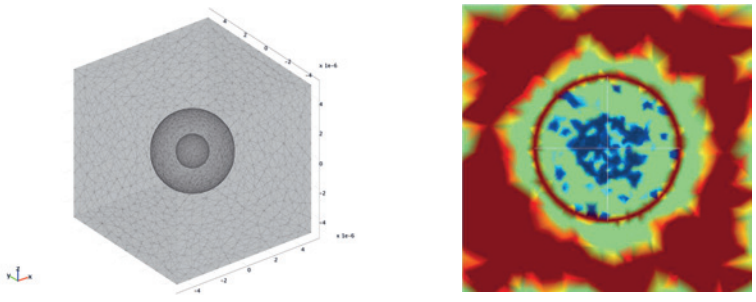

(a) (b)

Figure 4.4: In paper VI we consider a simple model of a reaction network inside a yeast cell contaning a nucleus. The cell is immersed in an external medium (a). The unstructured mesh is coarser outside the spherical domain and densest inside the nucleus and near the nuclear membrane. The hybrid algorithm dynamically partitions the species in every voxel into three different categories based on the method to use for diffusion.

Depending on the model's properties, we report speedups in paper VI and in [55] of up to two orders of magnitude compared to our base line implementation of NSM. We also give an example of a scenario where an exact simulation of the full system performs better due to the overhead of partitioning the system.

4.3 Coupling the microscopic and mesoscopic modeling levels

A better modeling accuracy as compared to the mesoscopic scale can be achieved by microscale methods based on Brownian motion for diffusion and Smoluchowski reaction dynamics [103]. This model underlies the spatial stochastic software GFRD [119, 120] and Smoldyn [4]. There, the position of individual particles are tracked explicitly and the spatial resolution is better than in a typical mesoscopic simulation based on the RDME. However, the increased accuracy comes with a price and microscale simulations are much more expensive than their mesocopic counterparts. In many cases, the RDME in its classical form can be expected to provide an accurate description of the reaction-diffusion kinetics, but there are cases where it fails to capture important features of a corresponding microscopic simulation [36, 37, 108, 119].

The model in Table 4.1 was introduced in [108] where it was demonstrated that when the reactivation rate in reactions (13,14) is high and diffusion of the proteins slow, microscopic resolution is needed to accurately simulate the system. A mesoscopic model could in principle also be used if the mesh resolution is fine enough, but in those regimes the model are likely to break down [59, 36]. In [37] the mesoscopic RDME model was then modified to correct for errors induced by small subvolumes and to allow for reactions between molecules in neighboring voxels. In that way the mesoscopic model could reproduce the results obtained with GFRD, provided that the mesh resolution was taken to be very fine (approaching the reaction radii of the molecules).

Another approach to overcome the limitations of a pure mesoscopic simulation is taken in paper VII. There, instead of refining the mesh to approach a microscopic modeling resolution in a mesoscopic simulation, we propose to refine the modeling level locally in space and species. For those species and in those regions of space where a microscopic modeling resolution is required we switch to a flexible implementation of the GFRD algorithm based on operator splitting and numerical computation of PDFs [56], further extended in paper VII for flexible simulations on surfaces embedded in 3D. Given a partitioning of the degrees of freedom the hybrid method advances one time step $\Delta\tau$ from τ^n to τ^{n+1} in the following way:

(1,2)	$KK + K \underset{k_2}{\overset{k_1}{\rightleftharpoons}} KK - K$	(3)	$KK - K \xrightarrow{k_3} KK^* + K_p$
(4,5)	$KK + K_p \underset{k_5}{\overset{k_4}{\rightleftharpoons}} KK - K_p$	(6)	$KK - K_p \xrightarrow{k_6} KK^* + K_{pp}$
(7,8)	$P + K_{pp} \underset{k_2}{\overset{k_1}{\rightleftharpoons}} P - K_{pp}$	(9)	$P - K_{pp} \xrightarrow{k_3} P^* + K_p$
(10,11)	$P + K_p \underset{k_5}{\overset{k_4}{\rightleftharpoons}} P - K_p$	(12)	$P - K_p \xrightarrow{k_6} P^* + K$
(13)	$KK^* \xrightarrow{k_7} KK$	(14)	$P^* \xrightarrow{k_7} P$

Table 4.1: *MAPK model from [108]. When diffusion is slow and the reactivation reactions (13,14) are fast, mesoscale simulations fail to accurately represent the system unless the spatial resolution is very high.*

1. Freeze all microscopic state variables and update the mesoscopic part of the system with NSM during the time step $\Delta\tau$. The chemical reactions that are executed by the mesoscale solver only involve mesoscale reactants.
2. Freeze the mesoscopic state variables and update the microscopic part of the system $\Delta\tau$ in time using the method described in [56]. Mixed reactions involving both mesoscopic and microscopic variables are executed by the microsolver.
3. Update the state at τ^{n+1} by determining if any meso- or microvariables have been created or destroyed in the time step and repeat from Step 1 until the final simulation time has been met.

Figure 4.5 shows the model in Table 4.1 simulated with the hybrid algorithm in paper VII. As can be seen, the method reproduces the results of a pure microscale simulation (a), while reducing the number of particles simulated by the more expensive microscale method nearly an order of magnitude. Importantly, there is no need to use an extremely fine mesh resolution to simulate the mesoscopic part of the hybrid system.

In the example above, it is the dynamics of a fast association-dissociation reaction pair that renders the mesoscale method insufficient. Another reason why one would expect fine mesh resolutions to be necessary to capture the model dynamics is if the geometry has fine scaled structures, or if there are membranes embedded in the 3D bulk that has high curvature. Furthermore, accurate mesoscopic propensities for phenomena such as reversible binding to a membrane is not currently known for general meshes. In the case of interaction between a molecule and a surface it is thus motivated to use the microscale solver, for which such interactions are clearly defined. Figure 4.6a shows a full microscale simulation of the model problem proposed in Section 6.3 in paper VII. Molecules of type C are created in the cytosol and will diffuse and bind reversibly to the nuclear membrane as part of their search for

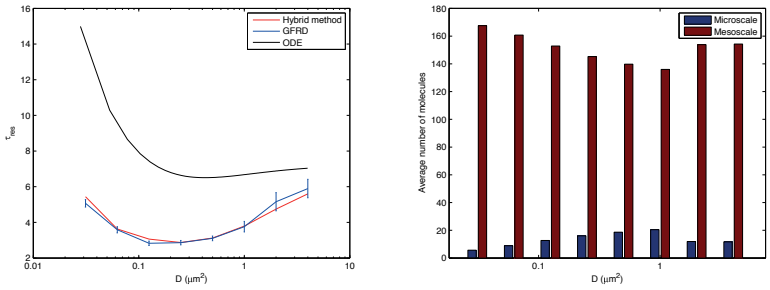

(a) (b)

Figure 4.5: The average response time τ_{res}, defined as the time to reach on average 50% of the steady state value of the doubly phosphorylated substrate K_{pp}, as a function of the diffusion constant D (a). The response is faster for the microscopic model than in the corresponding macroscopic RRE description (ODE) . The average number of molecules simulated on the microscale in the beginning of a time step compared to the average number of molecules simulated on the mesoscale (b).

a nuclear pore (a species localized to the surface of the inner sphere). When they find a pore, they are translocated into the nucleus by a simple reaction.

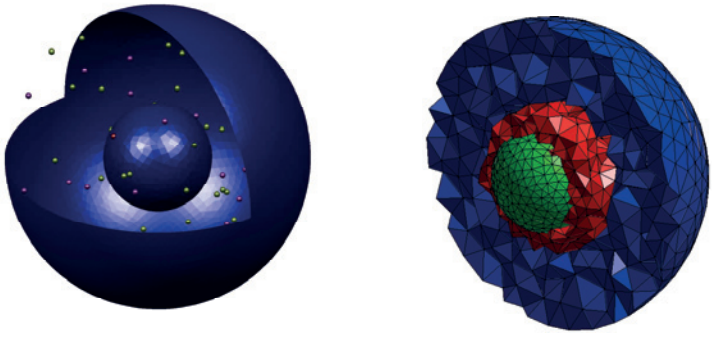

(a) (b)

Figure 4.6: A snapshot of a model of translocation through nuclear pores simulated with the microscale method (a) and the partitioning of space used by the hybrid method (b).

Using the hybrid algorithm, we resolve the binding to the membrane and the membrane bound species and reactions on the microscopic level while treating most parts of the 3D-bulk using the mesoscopic modeling level. The partitioning in space used by the hybrid solver is shown to the right in Figure 4.6. To resolve the partial adsorption of C molecules to the nuclear membrane (colored green) on the microscopic level, we let C be a microscopic degree

of freedom in a layer of elements around the nucleus (red). In the rest of the domain (blue), the system can be updated by the mesoscopic model.

5. Conclusions

In general, methods based on a single scale seem to be insufficient to simulate larger, complex models of biochemical reaction networks. Stochastic, microscale methods have a high spatial resolution, and arguably capture the reaction-diffusion dynamics accurately enough compared to other modeling assumptions. Unfortunately, microscopic simulations quickly become very expensive if many molecules are present. In that case, mesoscopic or even macroscopic methods may provide a sufficiently accurate description and are much more computationally tractable. For many models however, some components will be present in few copies or participate in reactions for which the parameters call for a microscale description. In the same time other species will be present in intermediate to high numbers. Resolving the whole model on the microscale will render it hard to simulate, and basing the simulation on a purely mesoscopic or macroscopic description will result in failure to capture important aspects of the dynamics.

This multiscale nature of the networks calls for specialized computational methods capable of dealing with modeling across scales. Three such methods have been proposed in this thesis in papers I, VI and VII. Common for all of the methods is the need to partition the system in different parts. Choosing a good partitioning is essential for the outcome of the simulation, both in terms of accuracy and efficiency. The hybrid method in VI goes the furthest in terms of automation. There, it is enough to choose three relative error tolerances and the hybrid algorithm will then adaptively partition the degrees of freedom according to what method it will use for simulating diffusion. For the hybrid methods in I and VII it is still necessary for the modeler to input some knowledge or guesses for a partitioning, based on intuition or prior knowledge about the system. Hybrid methods are a promising approach to handle systems with multiscale properties, but for a hybrid method to become a generally accepted tool used by practitioners it needs to behave like a black-box solver. Hence, future work in this area should include development of automatic and adaptive partitioning strategies. The hybrid methods proposed for the spatially dependent case is both based on the URDME software framework and thus shares a common platform. As a consequence, they could easily be integrated into one method.

Another reason for the need of different modeling levels might be dictated by the geometrical features of the domain. Narrow regions or interior boundaries with curvature might require a fine mesh locally. Unstructured meshes

facilitates the resolution of such regions. We have developed methodology for spatial stochastic simulation on tetrahedral and triangular meshes and implemented a software framework, URDME, that uses such meshes for stochastic simulation with NSM. A fine mesh resolution can potentially introduce errors in mesoscopic simulation and more work is needed to understand the balance of this effect and the discretization errors introduced by a too coarse mesh resolution. As many important processes occur on or near membranes, this is an important area of further study.

Transport of cellular cargo driven by motor proteins is known to be an important mechanism of transport in eukaryote cells. Mesoscale simulations have so far only focused on reaction-diffusion mechanisms, but with the extension of the model that we propose in paper V it is is possible to model, albeit in a coarse way, cargo transport on microtubule networks. How inclusion of this type of transport changes the behavior of e.g. models of signaling in the cell with respect to noise should be studied further. Another interesting area of further study would be to examine the validity of the proposed mesoscopic model by comparing it to a microscale description of the process. That would also lead to potential refinements of the mesoscopic method.

For the well-stirred case, we have developed two new methods as part of this thesis. In paper III we use the method of uniformization in conjunction with a quasi-Monte Carlo method in order to reduce the ensemble size needed for a good estimation of the PDF of the system. The method will work best for reasonably small, non-stiff models. Stiffness is frequently occurring in many realistic models, and this was one of the main motivations for the hybrid method in paper I. That method also aims at reducing the size of the system to be simulated on the mesoscopic level, and needs the PDF or the mean of the propensity functions to be estimated. As the subsystem obtained after a successful splitting in the hybrid model will be small and non-stiff, integration of the two methods would be beneficial.

6. Summary in Swedish

Matematisk modellering är ett viktigt verktyg för att undersöka de mekanismer och designprinciper som ligger till grund för olika reglersystem i levande celler. Genom att anta en uppsättning reaktioner mellan olika makromolekyler, till exempel olika proteiner, kan relativt enkla modeller formuleras i form av ordinära differentialekvationer (ODE). Dessa deterministiska *makrosopiska* modeller kan sedan simuleras med hjälp av datorer, och på så sätt kan modellens egenskaper studeras för ett stort antal olika parametrar och scenarier.

I celler kan antalet molekyler av nyckelproteiner vara mycket litet, och även små variationer kan därför göra en stor skillnad för systemets beteende. En viktig aspekt av ett cellulärt reglersystems funktion är därför hur fluktuationer i antalet molekyler påverkar dess pålitlighet [107, 33, 106, 90]. Stokastiska modeller som fångar upp den till synes slumpmässiga variationen i molekylantal ger i dessa fall en bättre beskrivning än de deterministiska ekvationerna. I stället för koncentrationen av ett protein är man nu intresserad av sannolikheten att finna ett visst antal molekyler vid en viss tid. Dessa *mesoskopiska* modeller har visat sig användbara när de makroskopiska modellerna brister i sin beskrivning av dynamiken [83, 6, 33, 109, 86, 6].

De stokastiska mesoskopiska modellerna är även de enkla att formulera baserat på samma antaganden som de deterministiska modellerna. Tyvärr är de betydligt svårare att studera numeriskt. En förklaring till detta är den höga dimensionaliteten – komplexiteten för att räkna ut sannolikheterna växer exponentiellt med antalet molekylslag om matrisbaserade metoder används. Om man istället studerar systemet genom att generera statistiskt korrekta realisationer av den stokastiska processen [44] och använder sig av Monte Carlometoder blir dimensionaliteten inte ett lika stort problem. Svårigheter uppstår då istället för system med tidsskaleseparation – väntetiderna för de olika reaktionerna varierar ofta över storleksordningar.

I artikel I tar vi oss an detta problem genom att simulera en del av modellen på makroskalan och en annan, mindre del på mesoskalan. Molekylslagen delas upp baserat på deras relativa varians, ett mått som ofta sammanfaller med antalet molekyler. På så sätt reduceras dimensionaliteten i den stokastiska delen av modellen. Ofta kommer molekylslag som placeras i den deterministiska delen att vara delaktiga i de snabbaste reaktionerna. På så sätt kan även problemet med tidsskaleseparation minskas i och med att dessa reaktioner kommer

att hanteras av den deterministiska ODE-lösaren. Resultatet blir alltså en reducerad stokastisk modell som är enklare att simulera.

När stokastisk simulering används för att studera modellen behöver man allt som oftast generera ett stort antal oberoende realisationer av processen. Denna statistiska ensemble kan sedan användas för att uppskatta olika mått på systemets prestanda och egenskaper. I artikel II presenterar vi en mjukvara, CellMC, kapabel att använda sig av vektorinstruktioner för att effektivt generera ensembler på både x86-arkitekturer (vanliga PC-datorer) och Cell Broadband Engine (Cell/BE). I artikel III utnyttjar vi istället kvasi-Monte Carlo-metodik för att reducera storleken på ensemblen givet ett krav på felet.

I artiklarna I, II och III antas den volym i vilken proteinerna reagerar vara väl omrörd. Molekyler transporteras i mediet genom diffusion vilket modellerar den sammanlagda effekten av många kollisioner med vattenmolekyler. Genom att säga att systemet är väl omrört antar vi att det finns tillräckligt med tid för varje molekyl att utforska hela området innan den deltar i en kemisk reaktion. Mer realistiska modeller behöver ta hänsyn till skillnader i rummet. Genom att diskretisera domänen med hjälp av ett beräkningsnät kan den mesoskopiska modellen utökas till det rumsberoende fallet. Lokalt i varje beräkningscell antar vi att systemet är väl omrört. Reaktioner sker på samma sätt som tidigare, men nu blir varje beräkningscell en egen, mindre reaktionsvolym. Diffusion modelleras sedan som hopp mellan beräkningsceller i nätet. Mjukvara för att simulera rumsberoende modeller med hjälp av strukturerade nät uppbyggda av kuber har utvecklats tidigare [51, 2]. I artikel IV demonstrerar vi hur processerna kan simuleras på ett ostrukturerat nät som utgörs av tetraedrer i 3D och trianglar i 2D. Sådana nät har stora fördelar för generella geometrier.

I eukaryota celler, som till exempel jäst, är inte diffusion den enda mekanism med vilken makromolekyler transporteras. I cellen finns ett intrikat nätverk av fibrer som utgör cytoskelettet. På dessa fibrer kan molekylära motorer, proteinkomplex med förmåga att "vandra" längs fibrerna, transportera organeller, vesiklar, mRNA och proteiner i olika riktningar beroende på motorns molekylära natur. Denna typ av transport har inte tidigare inkluderats i de mesoskopiska modellerna som vi beskrivit här. I artikel V föreslår vi ett sätt att göra detta som kan kombineras med existerande metodik för reaktions-diffusionssimuleringar.

Även för rumsberoende simuleringar uppstår problem med tidsskalorna. Förutom reaktionerna bidrar diffusionen med styvhet, och i många fall kommer antalet diffusionshändelser vida överstiga reaktionshändelserna. Detta är också ett fenomen som blir värre med en finare nätupplösning – om nätupplösningsparametern halveras så fyrdubblas antalet diffusionshändelser. I samma anda som i artikel I föreslår vi en hybridmetod i artikel VI. Reaktioner simuleras med den exakta stokastiska algoritmen medan diffusionen hanteras med tre olika metoder beroende på antalet molekyler av de olika molekylslagen. Den exakta stokastiska metoden [30] används när det

är ett fåtal molekyler i en beräkningscell, en approximativ stokastisk metod [48] för intermediärt antal och en makroskopisk, deterministisk metod när det är ett stort antal. På så sätt kan kostnaden för att simulera diffusionen reduceras. En viktig egenskap hos hybridalgoritmen är att den är adaptiv i både rummet och tiden.

Artikel VII fortsätter på samma tema, men där är det i stället den mesoskopiska modellen som står för den grövsta modelleringsnivån. En stokastisk *mikroskopisk* metod [56] kopplas ihop med den mesoskopiska metoden för att i vissa områden i domänen och för vissa molekylslag uppnå en högre noggrannhet. I den mikroskopiska metoden följs molekylernas positioner kontinuerligt i rummet, och den är därför mycket kostsammare att simulera än den mesoskopiska. Det är alltså fördelaktigt att använda den mesoskopiska beskrivningen för så stora delar av systemet som möjligt, utan att för den skull missa att fånga upp viktiga fenomen i den modell man studerar.

Multiskalaspekter genomsyrar denna avhandling. I tre olika hybridmetoder kopplas makroskopiska, mesoskopiska och mikroskopiska modelleringsnivåer för en så effektiv simulering som möjligt av en given modell. För att undvika att explicit ta hänsyn till alla fibrer kopplas en mikroskopisk, mesoskopisk och makroskopisk beskrivning av aktiv transport. På så sätt tar metoderna hänsyn till de skalseparationer som uppstår naturligt i systemen, något som är av största vikt för att kunna simulera mer och mer detaljerade modeller av cellbiologiska reglersystem.

7. Author's contributions

Paper I

The theoretical foundation of the method was developed by Lötstedt in earlier work. I proposed the use of quasi-Monte Carlo methodology, implemented the hybrid algorithm and designed and performed all numerical experiments. The paper was written in close collaboration with Lötstedt, and it built heavily on my master's thesis.

Paper II

I supervised the creation of the software and drafted the paper.

Paper III

I am the sole author of this paper.

Paper IV

The methodology was developed in close collaboration with the co-authors. I wrote parts of the manuscript and conducted the 3D numerical experiments.

Paper V

I designed the research and wrote the paper in close collaboration with Lötstedt. I implemented the method and designed and conducted all numerical experiments.

Paper VI

I developed the hybrid method in close collaboration with Lötstedt and Ferm. I designed, implemented and performed the 3D numerical experiments and wrote parts of the manuscript.

Paper VII

I implemented the hybrid method and designed the numerical experiments in close collaboration with S. Hellander. My contribution was mainly in the coupling between the mesoscopic and microscopic codes and the handling of the geometry and mesh. Lötstedt conducted the analysis of the approximation in the case of a sphere. I wrote parts of the manuscript.

8. Acknowledgment

Per, thank you for being a great advisor and collaborator and for always placing the needs of your students in the first room. Your passion for research rubs off easily and working with you has been both stimulating and rewarding.

Collaborators, past and present, thank you for all your hard work and for sharing your experience. Linda, thank you for hosting me and adopting me as a part of your research group during my extended visits to UCSB.

My family, mom and dad, thank you for always encouraging me. Your support throughout my time in graduate school has been invaluable.

Marie, thank you for all your patience and love. This thesis would have been much harder to complete without your company and support. Fabian, thank you for giving me a heartwarming smile in stressful times.

Comments from Sven-Erik Ekström, Elias Rudberg, Emil Kieri, Katharina Kormann, Magnus Gustafsson, Stefan Hellander, Sverker Holmgren and Marie Flinkfeldt improved the summary.

Financial support from the Swedish Graduate School in Mathematics and Computing (FMB), the Swedish Foundation for Strategic Research (SSF), the Swedish Royal Academy of Science (KVA), the Wallenberg foundation and Anna-Maria Lundin's stipendiefond have made this work and my travels possible.

Bibliography

[1] J. L. A.B. Bortz, M.H. Kalos. A new algorithm for Monte Carlo simulation of Ising spin systems. *J. Comput. Phys.*, 17(1):10–18, January 1975.

[2] M. Ander, P. Beltrao, B. D. Ventura, J. Ferkinghoff-Borg, M. Foglierini, C. Lemerle, I. Tomas-Oliveira, and L. Serrano. Smartcell, a framework to simulate cellular processes that combines stochastic approximation with diffusion and localisation: analysis of simple networks. *Systems Biology*, 1(1):129–138, 2004.

[3] D. F. Anderson. Incorporating postleap checks in tau-leaping. *J. Chem. Phys.*, 128(5):054103, 2008.

[4] S. S. Andrews, N. J. Addy, R. Brent, and A. P. Arkin. Detailed simulations of cell biology with Smoldyn 2.1. *PLoS Comput. Biol.*, 6(3):e1000705, 2010.

[5] F. Baras and M. M. Mansour. Reaction-diffusion master equation: A comparison with microscopic simulations. *Phys. Rev. E*, 54(6):6139–6148, Dec 1996.

[6] N. Barkai and S. Leibler. Circadian clocks limited by noise. *Nature*, 403:267–268, 2000.

[7] B. Bayati, P. Chatelin, and P. Koumoutsakos. Adaptive mesh refinement for stochastic reaction-diffusion processes. *J. Comput. Phys.*, 230(1):13–26, 2011.

[8] D. Bernstein. Simulating mesoscopic reaction-diffusion systems using the Gillespie algorithm. *Phys. Rev. E*, 71:041103, 2005.

[9] D. Braess. *Finite elements*. Cambridge University Press, Cambridge, 1992.

[10] K. Burrage, M. Hegland, S. MacNamara, and R. B. Sidje. A Krylov–based finite state projection algorithm for solving the chemical master equation arising in the discrete modelling of biological systems. In A. Langville and W. Stewardt, editors, *150th Markov Anniversary Meeting*, Boson Books, pages 21–38. Raleigh, NC, 2006.

[11] Y. Cao, D. Gillespie, and L. Petzold. Multiscale stochastic simulation algorithm with partial equilibrium assumption for chemically reacting systems. *J. Comput. Phys.*, 206:395–411, 2005.

[12] Y. Cao, D. T. Gillespie, and L. R. Petzold. Avoiding negative populations in explicit poisson tau-leaping. *J. Chem. Phys.*, 123(5):054104, 2005.

[13] Y. Cao, D. T. Gillespie, and L. R. Petzold. The slow-scale stochastic simulation algorithm. *J. Chem. Phys.*, 122, 014116, 2005.

[14] Y. Cao, D. T. Gillespie, and L. R. Petzold. Efficient step size selection for the tau-leaping simulation method. *J. Chem. Phys.*, 124(4):044109, 2006.

[15] Y. Cao, D. T. Gillespie, and L. R. Petzold. Adaptive explicit-implicit tau-leaping method with automatic tau selection. *J. Chem. Phys.*, 126(22):224101, 2007.

[16] Y. Cao, A. Hall, H. Li, and S. Lampoudi. *User's Guide for StochKit*. http://www.engr.ucsb.edu/∼cse, 2004.

[17] Y. Cao, H. Li, and L. Petzold. Efficient formulation of the stochastic simulation algorithm for chemically reacting systems. *J. Chem. Phys*, 121(9):4059–4067, 2004.

[18] E. Caulfield. CellMC: An XSLT-based SBML Model Compiler for Cell/BE and IA32. Master's thesis, Department of Information Technology, Uppsala University, 2009.

[19] A. Chatterjee, D. G. Vlachos, and M. A. Katsoulakis. Binomial distribution based tau-leap accelerated stochastic simulation. *J. Chem. Phys.*, 122(2):024112, 2005.

[20] J. S. Coggan, T. M. Bartol, E. Esquenazi, J. R. Stiles, S. Lamont, M. E. Martone, D. K. Berg, M. H. Ellisman, and T. J. Sejnowski. Evidence for ectopic neurotransmission at a neuronal synapse. *Science*, 309:446–451, 2005.

[21] COMSOL, Inc. http://www.comsol.com.

[22] J. Cullhed, B. Drawert, S. Engblom, and A. Hellander. Urdme 1.1: User's manual. Technical Report 003, Department of Information Technology, Division of Scientific Computing, Uppsala University, 2010.

[23] J. Cullhed, S. Engblom, and A. Hellander. The URDME manual version 1.0. Technical Report 2008-022, Department of Information Technology, Uppsala University, 2008.

[24] B. J. Daigle, M. K. Roh, D. T. Gillespie, and L. R. Petzold. Automated estimation of rare event probabilities in biochemical systems. *J. Chem. Phys.*, 134(4):044110, 2011.

[25] P. Deuflhard, W. Huisinga, T. Jahnke, and M. Wulkow. Adaptive discrete Galerkin methods applied to the chemical master equation. *SIAM J. Sci. Comput.*, 30(6):2990–2011, 2008.

[26] R. Dixit, J. L. Ross, Y. E. Goldman, and E. F. Holzbaur. Differential regulation of dynein and kinesin motor proteins by Tau. *Science*, 319(5866):1086–1089, 2008.

[27] 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. *J. Chem. Phys.*, 132(7):074101, 2010.

[28] W. E, D. Liu, and E. Vanden-Eijnden. Nested stochastic simulation algorithm for chemical kinetic systems with disparate rates. *J. Chem. Phys.*, 123, 194107, 2005.

[29] J. Elf and M. Ehrenberg. Fast evaluation of fluctuations in biochemical networks with the linear noise approximation. *Genome Res.*, 13(11):2475–2484, 2003.

[30] J. Elf and M. Ehrenberg. Spontaneous separation of bi-stable biochemical systems into spatial domains of opposite phases. *Syst. Biol.*, 1(2):230–236, 2004.

[31] J. Elf, G.-W. Li, and X. S. Xie. Probing transcription factor dynamics at the single-molecule level in a living cell. *Science*, 316(5828):1191–1194, 2007.

[32] J. Elf, P. Lötstedt, and P. Sjöberg. Problems of high dimension in molecular biology. In *Proc. 19th GAMM Seminar Leipzig*, pages 21–30, 2003.

[33] M. B. Elowitz, A. J. Levine, E. D. Siggia, and P. S. Swain. Stochastic gene expression in a single cell. *Science*, 297(5584):1183–1186, 2002.

[34] S. Engblom. Computing the moments of high dimensional solutions of the master equation. *Appl. Math. Comput.*, 180(2):498–515, 2006.

[35] S. Engblom. Galerkin spectral method applied to the chemical master equation. *Commun. Comput. Phys*, 5(5):871–896, 2009.

[36] R. Erban and J. Chapman. Stochastic modelling of reaction-diffusion processes: algorithms for bimolecular reactions. *Phys. Biol.*, 6:046001, 2009.

[37] D. Fange, O. G. Berg, P. Sjöberg, and J. Elf. Stochastic reaction-diffusion kinetics in the microscopic limit. *Proc. Natl. Acad. Sci. USA*, 107(46):19820–19825, 2010.

[38] D. Fange and J. Elf. Noise induced Min phenotypes in *E. coli*. *PLoS Comput. Biol.*, 2(6):e80, 2006.

[39] L. Ferm and P. Lötstedt. Adaptive solution of the master equation in low dimensions. *Appl. Numer. Math.*, 59(1):265–284, 2009.

[40] L. Ferm, P. Lötstedt, and A. Hellander. A hierarchy of approximations of the master equation scaled by a size parameter. *J. Sci. Comp.*, 34(2):127–151, 2008.

[41] L. Ferm, P. Lötstedt, and P. Sjöberg. Conservative solution of the Fokker–Planck equation for stochastic chemical reactions. *BIT*, 46(0):61–83, 2006.

[42] C. W. Gardiner. *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*. Springer-Verlag, 3rd edition, 2004.

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

[44] D. T. Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reacting systems. *J. Comput. Phys.*, 22:403–434, 1976.

[45] D. T. Gillespie. A rigorous derivation of the chemical master equation. *Physica A*, 188:404–425, Sept. 1992.

[46] D. T. Gillespie. *Markov Processes: An Introduction for Physical Scientists*. Academic Press, San Diego, 1992.

[47] D. T. Gillespie. The chemical Langevin equation. *J. Chem. Phys.*, 113(1):297–306, 2000.

[48] D. T. Gillespie. Approximate accelerated stochastic simulation of chemically reacting systems. *J. Chem. Phys.*, 115(4):1716–1733, 2001.

[49] A. Grönlund, P. Lötstedt, and J. Elf. Costs and constraints from time-delayed feed-back in small gene regulatory motifs. *Proc. Natl. Acad. Sci. USA*, 107:8171–8176, 2010.

[50] E. L. Haseltine and J. B. Rawlings. Approximate simulation of coupled fast and slow reactions for stochastic chemical kinetics. *J. Chem. Phys.*, 117(15):6959–6969, 2002.

[51] J. Hattne, D. Fange, and J. Elf. Stochastic reaction–diffusion simulation with MesoRD. *Bioinformatics*, 21(12):2923–2924, 2005.

[52] M. Hegland, C. Burden, L. Santoso, S. MacNamara, and H. Booth. A solver for the stochastic master equation applied to gene regulatory networks. *J. Comput. Appl. Math.*, 205(2):708–724, 2007.

[53] M. Hegland, A. Hellander, and P. Lötstedt. Sparse grids and hybrid methods for the chemical master equation. *BIT Num. Math.*, 48(2):265–283, 2008.

[54] A. Hellander. Numerical simulation of well stirred biochemical reaction networks governed by the master equation. *Licentiate thesis, Department of Information Technology, Uppsala University*, 2008.

[55] A. Hellander, L. Ferm, and P. Lötstedt. Adaptive stochastic hybrid simulation of biochemical reaction-diffusion models. In *Proc. 3rd International Conference on Foundations of Systems Biology in Engineering (FOSBE)*, 2009.

[56] S. Hellander and P. Lötstedt. Flexible single molecule simulation of reaction-diffusion processes. *J. Comput. Phys.*, 230:3948–3965, 2011.

[57] J. Howard. The movement of kinesin along microtubules. *Annu. Rev. Physiol.*, 58:703–729, 1996.

[58] M. Hucka, A. Finney, H. M. Sauro, H. Bolouri, J. C. Doyle, K. H., , the rest of the SBML Forum:, A. P. Arkin, B. J. Bornstein, D. Bray, A. Cornish-Bowden, A. A. Cuellar, S. Dronov, E. D. Gilles, M. Ginkel, V. Gor, I. I. Goryanin, W. J. Hedley, T. C. Hodgman, J.-H. Hofmeyr, P. J. Hunter, N. S. Juty, J. L. Kasberger, A. Kremling, U. Kummer, N. Le Novere, L. M. Loew, D. Lucio, P. Mendes, E. Minch, E. D. Mjolsness, Y. Nakayama, M. R. Nelson, P. F. Nielsen, T. Sakurada, J. C. Schaff, B. E. Shapiro, T. S. Shimizu, H. D. Spence, J. Stelling, K. Takahashi, M. Tomita, J. Wagner, and J. Wang. The systems biology markup language (SBML): a medium for representation and exchange of biochemical network models. *Bioinformatics*, 19(4):524–531, 2003.

[59] S. A. Isaacson. The reaction-diffusion master equation as an asymptotic approximation of diffusion to a small target. *SIAM J. Appl. Math.*, 70:77–111, 2009.

[60] S. A. Isaacson, D. M. McQueen, and C. S. Peskin. The influence of volume exclusion by chromatin on the time required to find specific DNA binding sites by diffusion. *Proceedings of the National Academy of Sciences*, 108(9):3815–3820, 2011.

[61] S. A. Isaacson and C. S. Peskin. Incorporating diffusion in complex geometries into stochastic chemical kinetics simulations. *SIAM J. Sci. Comput.*, 28:47–74, 2006.

[62] K. A. Iyengar, L. A. Harris, and P. Clancy. Accurate implementation of leaping in space: The spatial partitioned-leaping algorithm. *The Journal of Chemical Physics*, 132(9):094101, 2010.

[63] T. Jahnke. An adaptive wavelet method for the chemical master equation. *SIAM. J. Sci. Comput*, 31(6):4373–4394, 2010.

[64] T. Jahnke and W. Huisinga. Solving the chemical master equation for monomolecular reaction systems analytically. *Journal of Mathematical Biology*, 54:1–26, 2007. 10.1007/s00285-006-0034-x.

[65] T. Jahnke and T. Udrescu. Solving the chemical master equation by adaptive wavelet compression. *J. Comput. Phys.*, 229(16):5724–5741, 2010.

[66] A. Jensen. Markoff chains as an aid in the study of Markoff processes. *Skand. Aktuarietidskr.*, 36:87–91, 1953.

[67] M. Jeschke, R. Ewald, and A. M. Uhrmacher. Exploring the performance of spatial stochastic simulation algorithms. *Journal of Computational Physics*, 230(7):2562 – 2574, 2011.

[68] E. Kieri. Accuracy aspects of the reaction-diffusion master equation on unstructured meshes. Master's thesis, Uppsala University, 2011.

[69] G. Klingbeil, R. Erban, M. Giles, and P. K. Maini. STOCHSIMGPU: Parallel stochastic simulation for the systems biology toolbox 2 for matlab. *Bioinformatics*, doi: 10.1093/bioinformatics/btr068, 2011.

[70] S. Korotov, M. Křížek, and P. Neittaanmäki. Weakened acute type condition for tetrahedral triangulations and the discrete maximum principle. *Math. Comp.*, 70:107–119, 2000.

[71] T. G. Kurtz. Solutions of ordinary differential equations as limits of pure jump Markov processes. *J. Appl. Prob.*, 7:49–58, 1970.

[72] S. Lampoudi, D. Gillespie, and L. Petzold. The multinomial simulation algorithm for discrete stochastic simulation of reaction-diffusion systems. *J. Chem. Phys.*, 130(9):094104, 2009.

[73] P. L'Ecuyer, C. Lécot, and B. Tuffin. A randomized Quasi-Monte Carlo simulation method for Markov chains. In H. Niederreiter and D. Talay, editors, *Monte Carlo and Quasi Monte Carlo Methods 2004*, pages 331–342. Springer Berlin Heidelberg, 2007.

[74] A. Levchenko, J. Bruck, and P. W. Sternberg. Scaffold proteins may biphasically affect the levels of mitogen-activated scaffold proteins may biphasically affect the levels of mitogen-activated protein kinase signaling and reduce its threshold properties. *Proc. Nat. Acad. Sci. USA*, 97:5818–5823, 200.

[75] H. Li, Y. Cao, L. Petzold, and D. Gillespie. Algorithms and software for stochastic simulation of biochemical reacting systems. *Biotechnology Progress*, 24(1):56–61, 2008.

[76] H. Li and L. Petzold. Efficient parallelization of stochastic simulation algorithm for chemically reacting systems on the graphics processing unit. *Int. J. High. Perform. C.*, 24(7):107–116, 2009.

[77] S. MacNamara and K. Burrage. Stochastic modeling of na[i-umlaut]ve t cell homeostasis for competing clonotypes via the master equation. *Multiscale Model. Simul.*, 8(4):1325–1347, 2010.

[78] S. MacNamara, K. Burrage, and R. B. Sidje. Multiscale modeling of chemical kinetics via the master equation. *Multiscale Model. Simul.*, 6(4):1146–1168, 2008.

[79] R. Mallik and S. P. Gross. Molecular motors: Strategies to get along. *Current Biology*, 14:971–982, 2004.

[80] T. T. Marquez-Lago and K. Burrage. Binomial tau-leap spatial stochastic simulation algorithm for applications in chemical kinetics. *J. Chem. Phys.*, 127(10):104101, 2007.

[81] M. Mateescu, V. Wolf, F. Didier, and T. Henzinger. Fast adaptive uniformisation of the chemical master equation. *IET Systems Biology*, 4(6):441–452, 2010.

[82] S. Mauch and M. Stalzer. Efficient formulations for exact stochastic simulation of chemical systems. *IEEE/ACM Trans. Comput. Biol. Bioinf.*, 2009.

[83] H. H. McAdams and A. Arkin. It's a noisy business! Genetic regulation at the nanomolar scale. *Trends in Genetics*, 15(2):65–69, 1999.

[84] J. M. McCollum, G. D. Peterson, C. D. Cox, M. L. Simpson, and N. F. Samtova. The sorting direct method for stochastic simulation of biochemical systems with variying reaction execution behaviour. *Comput. Biol. Chem.*, 30:39–49, 2005.

[85] B. Munsky and M. Khammash. The finite state projection algorithm for the solution of the chemical master equation. *J. Chem. Phys.*, 124(4):044104, 2006.

[86] J. Paulsson, O. G. Berg, and M. Ehrenberg. Stochastic focusing: Fluctuation-enhanced sensitivity of intracellular regulation. *Proc. Nat. Acad. Sci. USA*, 97(13):7148–7153, 2000.

[87] M. F. Pettigrew and H. Resat. Multinomial tau-leaping method for stochastic kinetic simulations. *J. Chem. Phys.*, 126(8):084101, 2007.

[88] A. Raj, P. van den Bogaard, S. A. Rifkin, and A. van Oudenaarden. Imaging individual mRNA molecules using multiple singly labeled probes. *Nat. Meth.*, 5(10):877–879, 2008.

[89] C. V. Rao and A. P. Arkin. Stochastic chemical kinetics and the quasi-steady-state assumption: Application to the Gillespie algorithm. *J. Chem. Phys.*, 118(11):4999–5010, 2003.

[90] J. M. Raser and E. K. O'Shea. Noise in gene expression: Origins, consequences and control. *Science*, 309:2010–2013, 2005.

[91] M. Rathinam, L. R. Petzold, Y. Cao, and D. T. Gillespie. Consistency and stability of tau-leaping schemes for chemical reaction systems. *Multiscale Model. Simul.*, 4(3):867–895, 2005.

[92] M. Rathinam, P. W. Sheppard, and M. Khammash. Efficient computation of parameter sensitivities of discrete stochastic chemical reaction networks. *J. Chem. Phys.*, 132(3):034103, 2010.

[93] J. V. Rodriguez, J. A. Kaandorp, M. Dobrzyński, and J. G. Blom. Spatial stochastic modelling of the phosphoenolpyruvate-dependent phosphotransferase (pts) pathway in *Escherichia coli*. *Bioinformatics*, 22(15):1895–1901, 2006.

[94] D. Rossinelli, B. Bayati, and P. Koumoutsakos. Accelerated stochastic and hybrid method for spatial simulations of reaction-diffusion systems. *Chem. Phys. Lett.*, 451:136–140, 2008.

[95] H. Salis and Y. Kaznessis. Accurate hybrid stochastic simulation of a system of coupled chemical or biochemical reactions. *J. Chem. Phys.*, 122, 054103, 2005.

[96] K. R. Sanft, D. T. Gillespie, and L. R. Petzold. Legitimacy of the stochastic Michaelis-Menten approximation. *IET Syst. Biol.*, 5(1):58–69, 2011.

[97] T. P. Schulze. Kinetic Monte Carlo simulations with minimal searching. *Phys. Rev. E*, 65(3):036704, Feb 2002.

[98] T. P. Schulze. Efficient kinetic Monte Carlo simulation. *J. Comput. Phys.*, 227(4):2455–2462, 2008.

[99] R. B. Sidje. Expokit: a software package for computing matrix exponentials. *ACM Trans. Math. Softw.*, 24(1):130–156, 1998.

[100] P. Sjöberg. *Numerical Methods for Stochastic Modeling of Genes and Proteins*. PhD thesis, Uppsala University, 2007.

[101] P. Sjöberg. Partial approximation of the master equation by the Fokker–Planck equation. *Applied Parallel Computing. State of the Art in Scientific Computing, Lecture Notes in Scientific Computing, Springe Verlag*, pages 637–646, 2007.

[102] P. Sjöberg, P. Lötstedt, and J. Elf. Fokker-Planck approximation of the master equation in molecular biology,. *Comput. Vis. Sci.*, 12:37–50, 2009.

[103] M. v. Smoluchowski. Versuch einer mathematischen Theorie der Koagulationskinetik kolloider Lösungen. *Z. phys. Chemie*, 92:129–168, 1917.

[104] G. Strang. On the construction and comparison of difference schemes. *SIAM J. Numer. Anal.*, 5:506–517, 1968.

[105] A. B. Stundzia and C. L. Lumsden. Stochastic simulation of coupled reaction-diffusion processes. *J. Comput. Phys.*, 127:196–207, 1996.

[106] P. S. Swain. Efficient attenuation of stochasticity in gene expression through post-transcriptional control. *J. Mol. Biol.*, 344(4):965 – 976, 2004.

[107] P. S. Swain, M. B. Elowitz, and E. D. Siggia. Intrinsic and extrinsic contributions to stochasticity in gene expression. *Proceedings of the National Academy of Sciences of the United States of America*, 99(20):12795–12800, 2002.

[108] K. Takahashi, S. Tănase-Nicola, and P. R. ten Wolde. Spatio-temporal correlations can dratically change the response of a MAPK pathway. *Proc. Natl. Acad. Sci. USA.*, 107(6):2473–2478, 2010.

[109] M. Thattai and A. van Oudenaarden. Intrinsic noise in gene regulatory networks. *Proc. Nat. Acad. Sci. USA*, 98:8614–8619, 2001.

[110] V. Thomée and L. Wahlbin. On the existence of maximum principles in parabolic finite element equations. *Math. Comp.*, 77:11–19, 2008.

[111] T. Tian and K. Burrage. Binomial leap methods for simulating stochastic chemical kinetics. *J. Chem. Phys.*, 121(21):10356–10364, 2004.

[112] R. D. Vale. The molecular motor toolbox for intracellular transport. *Cell*, 112:467–480, 2003.

[113] N. G. van Kampen. *Stochastic Processes in Physics and Chemistry*. Elsevier, third edition, 2007.

[114] M. Vigelius, A. Lane, and B. Meyer. Accelerating reaction-diffusion simulations with general-purpose graphics processing units. *Bioinformatics*, 2010.

[115] S. Wils and E. D. Schutter. STEPS: modeling and simulating complex reaction-diffusion systems with Python. *Front. Neuroinf.*, 3:15, 2009.

[116] M. Yoshimi, Y. Nishikawa, Y. Osana, and A. Funahashi. Practical implementation of a network-based stochastic biochemical simulation system on an FPGA. In *FPL 2008 International Conference on Field Programmable Logic and Applications, IEEE*, 2008.

[117] W. Young and E. Elcock. Monte Carlo studies of vacancy migration in binary ordered alloys. *Proc. Phys. Soc.*, 89:735–747, 1966.

[118] J. S. van Zon and M. J. Morelli and S. Tănase-Nicola and P. R. ten Wolde. Diffusion of transcription factors can drastically enhance the noise in gene expression. *Biophys. J.*, 91:4350–4367, 2006.

[119] 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. *J. Chem. Phys.*, 123:234910, 2005.

[120] J. S. van Zon and P. R. ten Wolde. Simulating biochemical networks at the particle level and in time and space: Green's-function reaction dynamics. *Phys. Rev. Lett.*, 94:128103, 2005.

Acta Universitatis Upsaliensis

Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology 832

Editor: The Dean of the Faculty of Science and Technology

A doctoral dissertation from the Faculty of Science and Technology, Uppsala University, is usually a summary of a number of papers. A few copies of the complete dissertation are kept at major Swedish research libraries, while the summary alone is distributed internationally through the series Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology. (Prior to January, 2005, the series was published under the title "Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology".)

Distribution: publications.uu.se
urn:nbn:se:uu:diva-152098

ACTA
UNIVERSITATIS
UPSALIENSIS
UPPSALA
2011