2014

Numerical methods and stochastic simulation algorithms for reaction-drift-diffusion systems

https://hdl.handle.net/2144/15259

BOSTON UNIVERSITY

GRADUATE SCHOOL OF ARTS AND SCIENCES

Dissertation

NUMERICAL METHODS AND STOCHASTIC SIMULATION ALGORITHMS FOR REACTION-DRIFT-DIFFUSION SYSTEMS

by

AVA J. MAURO

M.A., University at Albany, SUNY, 2009
B.S., University of Pittsburgh, 2007

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

2014

Approved by

First Reader

Samuel A. Isaacson, Ph.D.
Associate Professor of Mathematics

Second Reader

C. Eugene Wayne, Ph.D.
Professor of Mathematics

Third Reader

Konstantinos Spiliopoulos, Ph.D.
Assistant Professor of Mathematics

Dedicated to my family, immediate and extended,

for their love and support

Acknowledgments

First and foremost, I would like to thank Sam Isaacson for being a great advisor. Sam has taught me so much not only about math itself, but also about the process of doing research. I greatly appreciated that he has always been available to help. Next I would like to thank the other members of my committee – Gene Wayne, Kostas Spiliopoulos, Nancy Kopell, and Tasso Kaper – for their advice and encouragement, and for taking the time to serve on my committee. I am grateful to Margaret Beck and Erin Munro, in addition to Gene, Nancy, and Sam, for everything that I learned in various courses with them. I would also like to acknowledge the support of Boston University and NSF grants DMS-0602204, DMS-0920886, and DMS-1255408.

Our collaborators, Paul Atzberger of UCSB and his former student Justin Shrake, worked on a version of the FPKMC algorithm that served as a precursor to the Dynamic Lattice FPKMC method developed in this dissertation. Paul and another of his former students, Jon Karl Sigurdsson, contributed to the application presented in Section 7.2. I appreciate all of their efforts (and the nice pictures).

I am very grateful to my fellow graduate students for being wonderful friends and engaging classmates. In particular, Ike Agbanusi (my mathematical brother) deserves the credit for introducing me to Sam's work. I am very glad that Ike was a year ahead of me in graduate school and always willing to share his experiences in everything from studying for the prelim to writing a dissertation. I would also like to thank my year mates - Brandon, Dan, Deniz, Hunter, Jon, Liz, Pak, and Rocío - for jumping into the unknown with me and providing an excellent support system ever since. The BU math department is a great community and I feel very fortunate to have been part of it. I am thankful to Allyson, Aly, Karen, Kathleen, and Theresa for keeping the department running smoothly and making all of our lives easier.

Mostly importantly, I am extremely grateful to my parents and my sister for being my motivation. I never could have made it this far without their continuous and unwavering love and support.

NUMERICAL METHODS AND STOCHASTIC SIMULATION ALGORITHMS FOR REACTION-DRIFT-DIFFUSION SYSTEMS

(Order No.)

AVA J. MAURO

Boston University, Graduate School of Arts and Sciences, 2014

Major Professor: Samuel A. Isaacson, Ph.D., Associate Professor of Mathematics

ABSTRACT

In recent years, there has been increased awareness that stochasticity in chemical reactions and diffusion of molecules can have significant effects on the outcomes of intracellular processes, particularly given the low copy numbers of many proteins and mRNAs present in a cell. For such molecular species, the number and locations of molecules can provide a more accurate and detailed description than local concentration. In addition to diffusion, drift in the movements of molecules can play a key role in the dynamics of intracellular processes, and can often be modeled as arising from potential fields. Examples of sources of drift include active transport, variations in chemical potential, material heterogeneities in the cytoplasm, and local interactions with subcellular structures.

This dissertation presents a new numerical method for simulating the stochastically varying numbers and locations of molecular species undergoing chemical reactions and drift-diffusion. The method combines elements of the First-Passage Kinetic Monte Carlo (FPKMC) method for reaction-diffusion systems and the Wang–

Peskin–Elston lattice discretization of the Fokker–Planck equation that describes drift-diffusion processes in which the drift arises from potential fields. In the FPKMC method, each molecule is enclosed within a “protective domain,” either by itself or with a small number of other molecules. To sample when a molecule leaves its protective domain or a reaction occurs, the original FPKMC method relies on analytic solutions of one- and two-body diffusion equations within the protective domains, and therefore cannot be used in situations with non-constant drift. To allow for such drift in our new method (hereafter Dynamic Lattice FPKMC or DL-FPKMC), each molecule undergoes a continuous-time random walk on a lattice within its protective domain, and the lattices change adaptively over time.

One of the most commonly used spatial models for stochastic reaction-diffusion systems is the Smoluchowski diffusion-limited reaction (SDLR) model. The DL-FPKMC method generates convergent realizations of an extension of the SDLR model that includes drift from potentials. We present detailed numerical results demonstrating the convergence and accuracy of our method for various types of potentials (smooth, discontinuous, and constant). We also present several illustrative applications of DL-FPKMC, including examples motivated by cell biology.

Contents

List of Tables

List of Figures

List of Abbreviations

act	actual
CDF	Cumulative Distribution Function
CN	Crank–Nicolson
cos	cosine function
DL-FPKMC	Dynamic Lattice First-Passage Kinetic Monte Carlo
DNA	Deoxyribonucleic Acid
ECDF	Empirical Cumulative Distribution Function
emp	empirical
exp	exponential function
FPKMC	First-Passage Kinetic Monte Carlo
GFRD	Green's Function Reaction Dynamics
KL	Kullback–Leibler divergence
KMC	Kinetic Monte Carlo
ln	natural logarithm function
low	lower
max	maximum
min	minimum
mRNA	Messanger Ribonucleic Acid
PD	Protective domain
PDE	Partial differential equation

PIDE	Partial integro-differential equation
Pr	Probability
RDME	Reaction-Diffusion Master Equation
rxn	reaction
SDE	Stochastic differential equation
SDLR	Smoluchowski Diffusion-Limited Reaction
sec	seconds
sin	sine function
SSA	Stochastic Simulation Algorithm
tanh	hyperbolic tangent function
TGA	Twizel–Gumel–Arigu
upp	upper
WPE	Wang–Peskin–Elston discretization

List of Symbols

$a_{i,j}$	Jump rate from the mesh point $\boldsymbol{x}_i$ to $\boldsymbol{x}_j$
D	Diffusion coefficient
δ	Delta function
$\boldsymbol{\eta}(\boldsymbol{x})$	Outward pointing normal at the point $\boldsymbol{x}$
$\mathbb{E}\left[\,\right]$	Expected value
$\mathbb{E}\left[T\right]$	Mean reaction time
$F^{\mathrm{rxn}}(x)$	Cumulative distribution function for reaction locations
h	Mesh width (lattice spacing)
h_{abs}	Mesh width at an absorbing boundary
h_p	Mesh width for a pair protective domain
$h_p^{\max}$	Maximum mesh width for pair protective domains
h_{ref}	Mesh width at a reflecting boundary
h_s	Mesh width for a single-molecule protective domain
$h_s^{\max}$	Maximum mesh width for single-molecule protective domains
$h^{\max}$	Maximum mesh width in the case that $h_s^{\max} = h_p^{\max}$
J	Flux
k_{B}	Boltzmann's constant
$KL(F\|\|G)$	Kullback–Leibler divergence of G from F
L	Domain length
λ_{Doi}	Reaction rate parameter in the Doi model

λ_{off}	Rate of exponential distribution for detachment times
m	Slope
$M^k(t)$	Number of molecules of the k^{th} chemical species at time t
Ω	Spatial domain
p_i	Discrete probability
ρ	Probability density
$\boldsymbol{Q}_l^k(t)$	Position vector of the l^{th} molecule of the k^{th} species at time t
$\boldsymbol{Q}^{\mathrm{rxn}}$	Location of a reaction
r_{pair}	Threshold distance for placing two molecules in a pair protective domain
r_{PD}	Distance from the initial location of a molecule within a symmetric protective domain to the nearest point on the boundary of the protective domain
r_{R}	Reaction radius
s	Arc length
$S(t)$	Survival probability at time t
t	Time
$\mathcal{T}$	Absolute temperature
T	Random variable for the time at which two molecules react
T_n	Random variable for the time at which the n^{th} reaction occurs
U	Protective domain
$\mathcal{U}(a,b)$	Uniform random distribution on the interval (a,b)
$V(\boldsymbol{x})$	Potential function at the point $\boldsymbol{x}$
$\boldsymbol{W}(t)$	Wiener process which describes Brownian motion

Chapter 1

Introduction

Parts of this chapter were previously published in [37] © 2013 Elsevier Inc.

A fundamental challenge in mathematical cell biology is to understand how to accurately model and simulate the dynamics of cellular processes. In recent years, there has been increased awareness that stochasticity in chemical reactions [4, 10, 18, 36, 49, 57] and diffusion of molecules [31, 33, 41] can have significant effects on the outcomes of intracellular processes, particularly given the low copy numbers of many proteins and mRNAs present in a cell. For such molecular species, the number and locations of molecules can provide a more accurate and detailed description than local concentration. In addition to diffusion, drift in the movements of molecules [6, 31, 44, 53] can also play a key role in the dynamics of intracellular processes, and can be modeled as arising from potential fields. Examples of sources of such drift include active transport, variations in chemical potential, material heterogeneities in the cytoplasm and nucleoplasm, and local interactions with cellular structures.

In this dissertation, we present a new numerical method [37] which allows for the explicit simulation of the stochastically varying numbers and locations of molecular species undergoing chemical reactions and drift-diffusion. The method combines elements of the First-Passage Kinetic Monte Carlo (FPKMC) method for reaction-diffusion systems [17, 42, 43, 58] and the Wang–Peskin–Elston lattice discretization of the Fokker–Planck equation that describes drift-diffusion processes in which the drift arises from potential fields [31, 63]. In our new method [37] (hereafter Dynamic Lattice FPKMC or DL-FPKMC), each molecule undergoes a continuous-time random walk on its own lattice, and the lattices change adaptively over time.

In Section 1.1, we will give an overview of and compare three of the most commonly used spatial models for stochastic reaction-diffusion systems: the spatially-continuous Smoluchowski Diffusion-Limited Reaction (SDLR) [54, 32] and Doi [59, 15, 16] models, and the lattice-based Reaction-Diffusion Master Equation (RDME) model [21, 38]. Then, in Section 1.2, we will present an extension of the SDLR model to incorporate drift due to potential fields, in addition to reactions and diffusion. Section 1.3 will mention methods for simulating the SDLR, Doi, and RDME models. Section 1.4 will include: (i) an overview of the recently developed First-Passage Kinetic Monte Carlo (FPKMC) method [17, 42, 43, 58], which can be used to generate exact realizations of the stochastic processes described by the SDLR model; (ii) an introduction to our new method which uses dynamically changing lattices within the FPKMC method to generate convergent realizations of the extended SDLR model with drift due to potentials, and (iii) a summary of this dissertation.

1.1 Commonly Used Stochastic Reaction-Diffusion Models

1.1.1 The Smoluchowski Diffusion-Limited Reaction (SDLR) Model

One of the most commonly used spatial models for stochastic reaction-diffusion systems is the Smoluchowski diffusion-limited reaction (SDLR) model. The SDLR model [54] is a spatially-continuous model in which molecules move by Brownian motion.

Bimolecular, or second-order, reactions between two molecules are modeled using boundary conditions. When the separation between two bimolecular reactants reaches a specified distance, called the reaction radius, there are two possible reaction mechanisms that are commonly used. Either the reaction occurs instantaneously (with probability one) [3] and is modeled using a pure-absorption Dirichlet boundary condition, or the reaction occurs with some probability less than one based on an intrinsic reaction rate [19, 58, 32] and is modeled using a partial-absorption Robin boundary condition [32]. Bimolecular reactants are not allowed to approach closer than their reaction radius.

Unimolecular, or first-order, reactions involving a single molecule represent internal processes, such as decay or splitting, and are modeled as occurring with specified probabilities per unit time, i.e., with exponentially distributed times based on a specified reaction-rate constant.

We will assume that reaction products are placed at the locations specified in Table 1.1. In the cases of two reaction products, for either unimolecular or bimolecular reactions, the angular orientation of the product separation vector about the center of mass is chosen randomly, as in [3].

The state of a system over time is given by the collection of stochastic processes for the numbers of molecules and the locations of each molecule of each chemical species. Alternatively, the SDLR model can be described by a system of partial integro-differential equations (PIDEs) for the probability densities of having a given number of molecules of each chemical species at a specified set of locations, similar to the stochastic reaction-diffusion PIDE models in [27] and [15, 16]. Due to the high-dimensionality of the system of PIDEs for the probability densities, numerical methods for solving the SDLR model typically use Monte Carlo approaches that

Table 1.1: Placement of reaction products.

	One Reaction Product	*Two Reaction Products*
Unimolecular Reaction	The product is placed at the same location as the reactant.	The products are placed a specified distance apart, with their center of mass at the location of the reactant.
Bimolecular Reaction	The product is placed at the center of mass of the two reactants.	The products are placed a specified distance apart, with their center of mass at the same location as the center of mass of the reactants.

approximate the underlying stochastic processes.

1.1.2 The Doi Model

The Doi model [59, 15, 16] is similar to the SDLR model: it is a spatially-continuous model in which molecules move by Brownian motion and unimolecular reactions occur with specified probabilities per unit time. However, in the Doi model, bimolecular reactants may approach arbitrarily close to each other, and a bimolecular reaction is modeled as occurring with a fixed probability per unit time, λ_{Doi}, when the separation between the two reactants is *less than* the specified reaction radius. In the case that the location of one of the two reactants is fixed, it has been shown that as the reaction parameter λ_{Doi} is taken to infinity, the solution of the Doi model converges to the solution of the pure-absorption version of the SDLR model at an $O(\lambda_{\mathrm{Doi}}^{-1/2})$ rate [1].

1.1.3 The Reaction-Diffusion Master Equation (RDME) Model

In contrast to the spatially-continuous SDLR and Doi models, the Reaction-Diffusion Master Equation model [21, 38] is a spatially-discrete lattice-based model. In the RDME model, a mesh is used to partition space into a collection of voxels. The diffusion of molecules is approximated by a continuous-time random walk between voxels. In the standard RDME model, bimolecular reactions can only occur between molecules within the same voxel, and are modeled as occurring with fixed probabilities per unit time. Unimolecular reactions are modeled in the same manner as in the SDLR and Doi models.

In applications, the RDME is frequently used as a more macroscopic approximation to the SDLR model [48, 19, 25]. However, the error in this approximation cannot be made arbitrarily small. It has been shown that for systems with bimolecular reactions in two or more dimensions, the RDME does not converge to the SDLR model [28, 30, 25]. As the lattice spacing in the RDME is taken to zero, the time required for two molecules to react diverges to infinity, since two molecules must be in the same voxel to react. Therefore, the standard RDME does not converge to any underlying spatially-continuous model of systems with bimolecular reactions. A Convergent RDME (CRDME) model has recently been developed, which modifies the RDME to create a model that converges to the Doi model as the lattice spacing is taken to zero [29].

Relationships between the models discussed in this section are summarized in Fig. 1.1.

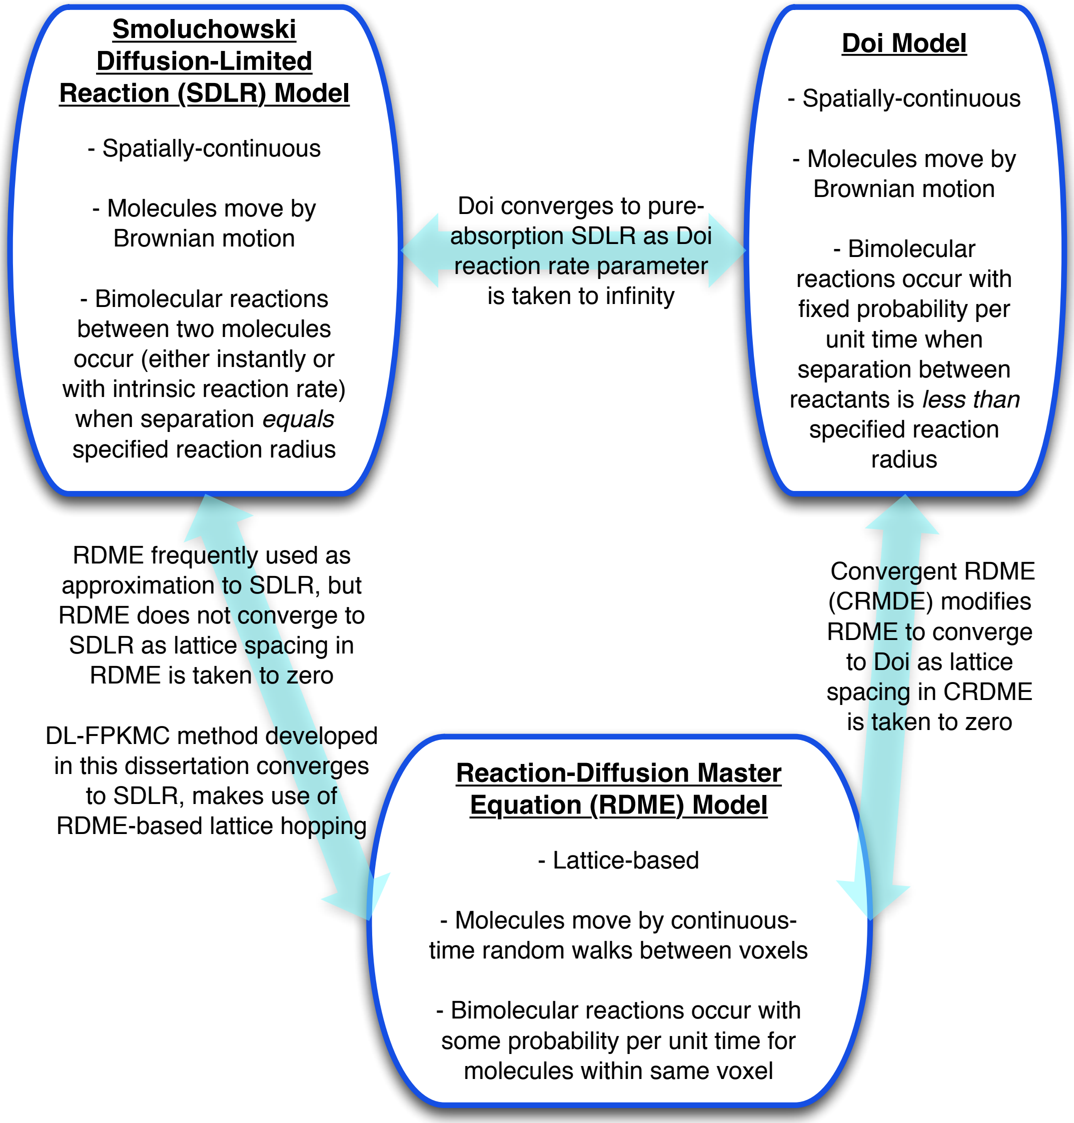

Figure 1.1: Relationships between three common stochastic reaction-diffusion models. Each of these models can be extended to include drift due to potential fields. See text of Sections 1.1 and 1.3 for more detailed discussion of the models and methods summarized in this figure.

1.2 Incorporating Drift into the SDLR Model

In this section a modification of the SDLR model incorporating drift due to potential fields is presented.

In a system with K chemical species, we denote by $\{M^k(t) : t \geq 0\}$ the stochastic process for the number of molecules of the k^{th} species at time $t \geq 0$, for $k = 1, \ldots, K$. The position vector of the l^{th} molecule of the k^{th} species at time $t \geq 0$ is given by the vector stochastic process $\{\boldsymbol{Q}_l^k(t) : t \geq 0\}$, where $\boldsymbol{Q}_l^k(t) \in \mathbb{R}^n$, for $l = 1, \ldots, M^k(t)$.

When incorporating drift due to potentials, the underlying reaction process remains unchanged (using boundary conditions as described above), but molecules move by a drift-diffusion process instead of pure Brownian motion. The l^{th} molecule of species S^k with position $\boldsymbol{Q}_l^k(t)$ undergoes diffusion with diffusion coefficient D^k and experiences drift due to a potential $V^k\big(\boldsymbol{Q}_l^k(t)\big)$. In this case, $\boldsymbol{Q}_l^k(t)$ satisfies the stochastic differential equation (SDE)

$$d\boldsymbol{Q}_l^k(t) = \frac{-D^k}{k_{\mathrm{B}}\mathcal{T}} \nabla V^k\big(\boldsymbol{Q}_l^k(t)\big)\, dt + \sqrt{2D^k}\, d\boldsymbol{W}_l^k(t), \tag{1.1}$$

where k_{B} is Boltzmann's constant, $\mathcal{T}$ is absolute temperature, ∇ denotes the gradient in the coordinates of $\boldsymbol{Q}_l^k(t)$, and $\boldsymbol{W}_l^k(t)$ denotes the standard n-dimensional Wiener process which describes Brownian motion. The $-\nabla V(\boldsymbol{Q})$ in the drift term represents the force arising from the potential field, and the $d\boldsymbol{W}(t)$ in the diffusion term is a fluctuating random force [20, 61]. Note that when V is constant, the drift-diffusion process of Eq. 1.1 reduces to pure Brownian motion. By the Einstein relation (also known as the Einstein–Smoluchowski relation), $D = k_{\mathrm{B}}\mathcal{T}/\zeta$ where ζ is the friction coefficient [8, 61, 63]. The friction coefficient of a spherical particle in a fluid, for example, is given by Stokes' Law: $\zeta = 6\pi\eta r$ where η is the viscosity of the fluid and

r is the radius of the particle [8]. Using the Einstein relation, Eq. 1.1 may also be written as

$$\zeta\, d\boldsymbol{Q}_l^k(t) = -\nabla V^k\big(\boldsymbol{Q}_l^k(t)\big)\, dt + \sqrt{2\zeta k_{\mathrm{B}}\mathcal{T}}\, d\boldsymbol{W}_l^k(t).$$

1.3 Simulation Methods

In this section, we mention methods for simulating the stochastic reaction-diffusion models described in Section 1.1, with the possibility of including drift.

Exact realizations of the RDME and CRDME models can be generated using the event-driven Stochastic Simulation Algorithm (SSA) [23], which is a well-established Kinetic Monte Carlo (KMC) method [11]. Ref. [31] extends the RDME model to include drift due to potentials, and presents simulations of this extension using the SSA.

The recently proposed First-Passage Kinetic Monte Carlo (FPKMC) method [17, 42, 43, 58] can be used to generate exact realizations of the stochastic processes described by the SDLR model for reaction-diffusion system. This method relies on analytic solutions of the diffusion equation, and therefore cannot be used in situations where non-constant drift is also present.

Note, the newer version of the method called Green's Function Reaction Dynamics (GFRD) is equivalent to FPKMC [58]. The earlier GFRD method of [62] was approximate. In this manuscript, we will consider exact GFRD methods to be included among the methods that we refer to by the term FPKMC.

An exact FPKMC method has been introduced incorporating spatially and temporally varying transition or annihilation rates for single particles [51], which could be used to simulate transitions from diffusive to ballistic modes in models of intracellular transport [35]. The newer GFRD method [58] has been modified to allow for advection due to a *spatially-uniform, constant* velocity field along a one-dimensional track [55].

Although the extension of the SDLR model to include drift due to arbitrary potentials (Section 1.2) is conceptually straightforward, there are no well-established standard methods for generating realizations of the stochastic processes described by this extended model. To address the need for such methods, we have developed a new numerical method [37], which will be the subject of this dissertation. Our method combines elements of the FPKMC method [17, 42, 43, 58] and the Wang–Peskin–Elston lattice discretization of the Fokker–Planck equation that describes drift-diffusion processes in which the drift arises due to potential fields [31, 63].

Lastly, Brownian Dynamics is a time-step based method that can be used to simulate the SDLR or Doi models. Brownian Dynamics can be modified to included drift due to a potential. Unlike the Wang–Peskin–Elston lattice discretization, Brownian Dynamics methods do not preserve detailed balance (zero net flux at equilibrium).

1.4 A DL-FPKMC Method for the SDLR Model with Drift

In this section we present an overview introducing the DL-FPKMC method. In the body of the paper, we will develop the details of the method and analyze convergence.

The recently proposed First-Passage Kinetic Monte Carlo (FPKMC) method [17, 42, 43, 58] can be used to generate exact realizations of the stochastic processes described by the SDLR model without drift. In the FPKMC method, each molecule is enclosed within a “protective domain,” either by itself or possibly with a small number of other molecules. The protective domains serve to define regions within which the movement of each molecule is independent of molecules in other protective domains. A significant change in the state of the system occurs only when a molecule leaves its protective domain (first-passage event) or a reaction occurs. The protective domains are generally chosen to be spherical or rectangular regions. As such, FPKMC can be interpreted as an extension to include chemical reactions in the Walk on Spheres [40] and Walk on Rectangles [13] methods for simulating Brownian motion in complex geometries.

In the rather special case of pure Brownian motion in simple protective domains (spheres or rectangular regions), the first-passage time distributions for a molecule to leave a protective domain [43, 42, 17, 58] or for two molecules to reach a threshold radius for reaction [58] can be computed analytically by solving the diffusion equation. By making use of these expressions, the FPKMC method can generate exact realizations of the stochastic processes described by the SDLR model. However, for many situations in cell biology, pure Brownian motion does not provide the most realistic description of the movement of molecules as a consequence of active transport, chemical gradients, interactions with cellular structures, etc. In such cases, significant drift terms are inherent to the particle dynamics and can be modeled as arising from a fixed potential field. The DL-FPKMC method we develop extends the FPKMC to allow for such drift. Analytical expressions for the first-passage time distributions from protective domains are no longer possible with the addition of spatially varying drift. In DL-FPKMC, we therefore approximate the drift-diffusion process each molecule undergoes within a protective domain by a continuous-time random walk on a discretized mesh. Exact sample paths of the molecules' random walks are generated using the SSA [23]. The transition rates or jump rates for these walks are determined from the Wang–Peskin–Elston finite-difference discretization of the Fokker-Planck equation [63]. When new protective domains are created during the course of a simulation, meshes within each protective domain are constructed dynamically. For this reason, our method can be interpreted as a dynamic-lattice master equation model. Unlike the standard RDME, it has the benefit of converging to the SDLR model as the lattice spacing is reduced. In the course of developing the DL-FPKMC method, we have derived two variations of the Wang–Peskin–Elston discretization, one for non-uniform lattice spacings and the other for partial-absorption Robin boundaries.

We present detailed numerical results demonstrating both the convergence and accuracy of the DL-FPKMC method as the mesh spacing in the discretization is decreased. In particular, we apply our algorithm to the general bimolecular reaction $A + B \to \varnothing$ where the molecules of species A and B undergo drift-diffusion subject to various types of potential functions (smooth, discontinuous, and constant). Our results indicate that the method is approximately second-order accurate for smooth potentials and approximately first-order accurate for discontinuous potentials.

The FPKMC method was originally presented as an efficient way to simulate reaction-diffusion systems at low particle densities "without all the hops" by using larger "superhops" [43]. While DL-FPKMC uses more hops than FPKMC due to the random walk approximation of molecular motion, we demonstrate that DL-FPKMC maintains efficiency at low particle densities by requiring far fewer hops than fixed lattice methods with comparable resolution. By discretizing each individual protective region, DL-FPKMC allows fine meshes to be used in localized regions when needed for accuracy considerations. Examples where fine meshes may be necessary include resolving bimolecular reactions, boundary conditions, or rapidly varying potential fields. For protective domains in which such features are not present, coarser meshes can be used to improve efficiency.

The dissertation is organized as follows. In Section 1, we introduced the SDLR, Doi, and RDME models, an extension of the SLDR model to include drift, and the FPKMC and DL-FPKMC methods. Chapter 2 will give an a more detailed overview of the implementation and steps of the FPKMC or DL-FPKMC algorithm to generate realizations of the stochastic processes described by the SDLR model. Chapter 3 presents our numerical method for using dynamic lattices to incorporate drift into the FPKMC algorithm, and in Chapters 4 and 5 we demonstrate the convergence and accuracy of this DL-FPKMC method. In Chapter 6 we provide a running time analysis of DL-FPKMC, in which we demonstrate $O(N)$ scaling with the number of molecules in the system and compare DL-FPKMC to a fixed lattice method. Chapter 7 presents applications illustrating the use of DL-FPKMC. In 7.1 we compare the effects of drift due to several potentials on reaction time and location statistics. In 7.2 we investigate a simplified model of a coupled protein-polymer system, in which two reacting molecules undergo drift-diffusion along a polymer, and may also unbind from the polymer and diffuse in three dimensions. We study the interaction between polymer geometry, binding potentials along the polymer, and unbinding rate, and find that the interaction between these factors can significantly influence the reaction process.

Chapter 2

First-Passage Kinetic Monte Carlo Methods

Parts of this chapter were previously published in [37] © 2013 Elsevier Inc.

This chapter gives an overview of First-Passage Kinetic Monte Carlo methods, including DL-FPKMC and earlier methods, and then presents the steps for implementing FPKMC or DL-FPKMC.

2.1 Overview of FPKMC Approaches

Recently-developed First-Passage Kinetic Monte Carlo methods can be used to generate exact realizations of the SDLR model in the absence of drift [43, 42, 17, 58, 51].[1] In these FPKMC algorithms, a spherical or rectangular region called a 'protective domain' is drawn around every molecule in the system. Each protective domain generally contains only one or two molecules, but may contain more molecules if necessary (e.g. due to a cluster of closely packed molecules). The movements of molecules in separate protective domains are treated independently. The first-passage time for a single molecule in its own protective domain, meaning the time when the molecule

---

[1] The original FPKMC method [43, 42, 17] was exact for cube-shaped particles, but relied on time-stepping when spherical bimolecular reactants approached each other.

will first hit the boundary of the protective domain, can be sampled exactly using the corresponding analytical solution to the diffusion equation. The molecule that exits its protective domain first is updated to its exit position, and a new protective domain is defined. For two reactants in one protective domain, a candidate time and location for either their reaction or for one to exit can be sampled analytically, by converting to separation and center-of-mass coordinates and solving the corresponding diffusion equations [58].

When drift due to a potential is present in addition to diffusion, the probability densities for the locations of one or two molecules within their protective domains are no longer described by the diffusion equation, but rather by a Fokker-Planck equation. The DL-FPKMC method will address the issue that the Fokker-Planck equation cannot be solved analytically for general potentials. Let $\Omega \subset \mathbb{R}^n$ denote the overall domain with boundary $\partial\Omega$, and let $U \subset \Omega$ label a protective domain with boundary ∂U. We are interested in the time a molecule first leaves U, leading to a zero Dirichlet boundary condition on ∂U. Let $V(\boldsymbol{x})$ denote the strength of the potential at $\boldsymbol{x} \in \Omega$. The probability density, $\rho(\boldsymbol{x}, t)$, for a single molecule with diffusion coefficient D to be at location $\boldsymbol{x}$ within its protective domain U at time $t > 0$ evolves according to the equations

$$\begin{aligned} \frac{\partial \rho(\boldsymbol{x}, t)}{\partial t} = D \nabla \cdot \left(\rho(\boldsymbol{x}, t) \frac{\nabla V(\boldsymbol{x})}{k_{\mathrm{B}} \mathcal{T}} + \nabla \rho(\boldsymbol{x}, t) \right), \qquad & \text{on } U, \\ \rho(\boldsymbol{x}, t) = 0, \qquad & \text{on } \partial U \backslash (\partial U \cap \partial \Omega), \\ \rho(\boldsymbol{x}, 0) = \delta(\boldsymbol{x} - \boldsymbol{x}_0), \qquad & \end{aligned} \tag{2.1}$$

where $\boldsymbol{x}_0$ is the initial position of the molecule within the protective domain. If ∂U intersects $\partial\Omega$, the boundary conditions on $\partial U \cap \partial\Omega$ will agree with the boundary conditions on $\partial\Omega$. These boundary conditions may be reflecting, absorbing, partially absorbing, or periodic.

For two bimolecular reactants within one protective domain $U \subset \Omega \subset \mathbb{R}^n$, the joint probability density $\rho(\boldsymbol{x}, \boldsymbol{y}, t)$ for one molecule to be at location $\boldsymbol{x} \in U$ and the other molecule to be location $\boldsymbol{y} \in U$ at time $t > 0$ can be described by a Fokker-Planck equation in $\mathbb{R}^{2n}$. Define the domain in $\mathbb{R}^{2n}$ to be $W = \{(\boldsymbol{x}, \boldsymbol{y}) : \boldsymbol{x} \in U,\ \boldsymbol{y} \in U,\ \text{and } ||\boldsymbol{x} - \boldsymbol{y}|| > r_{\mathrm{R}}\}$. D_1 and D_2 will denote the respective diffusion coefficients of the two molecules. We define $\mathcal{D}$ to be a $2n \times 2n$ diagonal matrix, with the first n elements of the diagonal equal to D_1 and the next n elements equal to D_2. Let $V_1(\boldsymbol{x})$ and $V_2(\boldsymbol{y})$ be the potential fields that impart drift to the $\boldsymbol{x}$ molecule and the $\boldsymbol{y}$ molecule respectively, and define $V(\boldsymbol{x}, \boldsymbol{y}) = V_1(\boldsymbol{x}) + V_2(\boldsymbol{y})$. The behavior on the boundary ∂W may be different on each of the following three components: the reactive boundary $\partial W_{\mathrm{rxn}} = \partial W \cap \{||\boldsymbol{x} - \boldsymbol{y}|| = r_{\mathrm{R}}\}$, the (possibly empty) intersection with the overall domain boundary $\partial W_{\mathrm{outer}} = (\partial W \cap \{(\boldsymbol{x}, \boldsymbol{y}) : \boldsymbol{x} \in \partial\Omega \text{ or } \boldsymbol{y} \in \partial\Omega\}) \backslash \partial W_{\mathrm{rxn}}$, and the remaining component $\partial W \backslash (\partial W_{\mathrm{rxn}} \cup \partial W_{\mathrm{outer}})$. Then, $\rho(\boldsymbol{x}, \boldsymbol{y}, t)$ satisfies

$$\begin{aligned}
\frac{\partial \rho(\boldsymbol{x}, \boldsymbol{y}, t)}{\partial t} &= \nabla \cdot \mathcal{D} \left(\rho(\boldsymbol{x}, \boldsymbol{y}, t) \frac{\nabla V(\boldsymbol{x}, \boldsymbol{y})}{k_{\mathrm{B}} \mathcal{T}} + \nabla \rho(\boldsymbol{x}, \boldsymbol{y}, t) \right), \quad && \text{on } W, \qquad (2.2) \\
\rho(\boldsymbol{x}, \boldsymbol{y}, t) &= 0, && \text{on } \partial W_{\mathrm{rxn}}, \\
\rho(\boldsymbol{x}, \boldsymbol{y}, t) &= 0, && \text{on } \partial W \backslash (\partial W_{\mathrm{rxn}} \cup \partial W_{\mathrm{outer}}), \\
\rho(\boldsymbol{x}, \boldsymbol{y}, 0) &= \delta(\boldsymbol{x} - \boldsymbol{x}_0, \boldsymbol{y} - \boldsymbol{y}_0),
\end{aligned}$$

where $\boldsymbol{x}_0$ and $\boldsymbol{y}_0$ are the initial positions of the molecules within the protective domain, and the gradient and divergence operators are in the $(\boldsymbol{x}, \boldsymbol{y})$ coordinates. Similar to the single particle case, if ∂U intersects $\partial\Omega$, the boundary conditions on

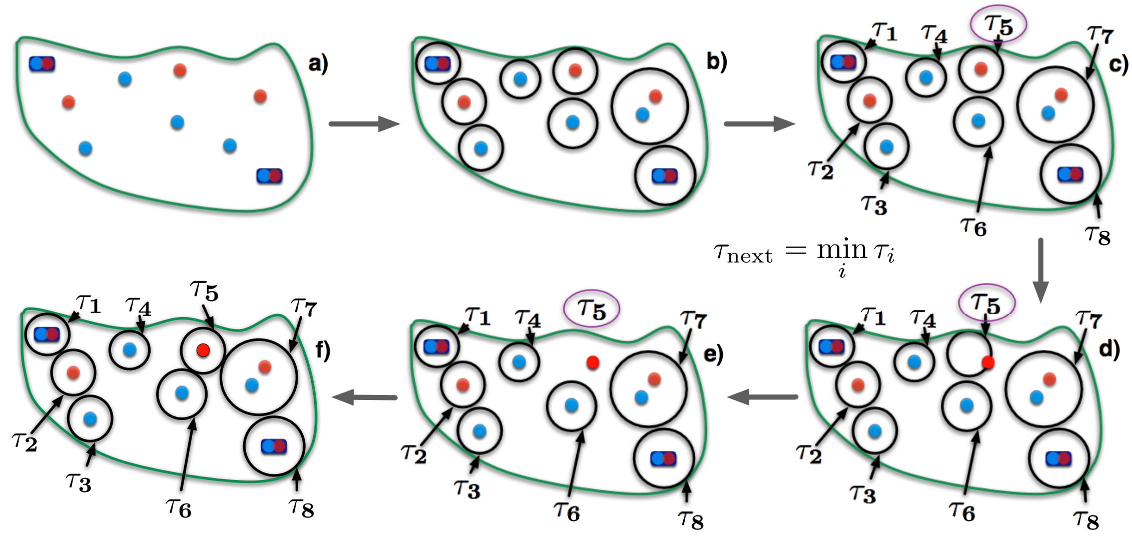

Figure 2.1: Schematic of the FPKMC or DL-FPKMC algorithm as described in Subsection 2.2. Here, τ_i denotes the next event time for the molecule or pair of molecules in the i^{th} protective domain, and $\tau_{\text{next}} = \tau_5$ is the time of a global event.

$\partial W_{\text{outer}}$ will agree with the boundary conditions on $\partial\Omega$. The boundary condition $\rho(\boldsymbol{x}, \boldsymbol{y}, t) = 0$ when $||\boldsymbol{x} - \boldsymbol{y}|| = r_{\text{R}}$ models the pure-absorption reaction mechanism. This reactive boundary condition could be modified to use a Robin partial-absorption mechanism if desired. In the case that the potential V is a constant function, Eqs. (2.1) and (2.2) reduce to diffusion equations.

Our method for using discretizations of these Fokker-Planck equations to sample times and locations of first-passage and reaction events is described in Chapter 3.

2.2 Main Steps of the FPKMC or DL-FPKMC Algorithm

In this subsection we describe the role of protective domains and the processing of events in our implementation of the FPKMC or DL-FPKMC algorithm. We then list the main steps of the algorithm. Our implementation is based on the FPKMC algorithm developed in [43, 42, 17], with some modifications. The content of this section applies to both the FPKMC and DL-FPKMC algorithms. The only theoretical difference between the two methods is the approach for sampling event times within a protective domain. In FPKMC, which can only be used in the case of constant $V(\boldsymbol{x})$, event times are sampled from *exact* solutions of the diffusion equation. In DL-FPKMC, which allows for arbitrary $V(\boldsymbol{x})$, event times are *approximated* by generating sample paths of continuous-time random walks on meshes within protective domains (see Chapter 3).

2.2.1 Protective Domains

To apply the algorithm, every molecule in the system is placed in a protective domain. In one dimension the protective domains are intervals and in higher dimensions the protective domains are usually rectangular or spherical regions. In general, the boundaries of protective domains are absorbing. The boundary of a protective domain can contain a portion of the boundary of the overall spatial domain, in which case the protective domain boundary conditions will depend on the overall domain boundary conditions. We allow protective domains to contain either one or two molecules. Protective domains containing only one molecule are referred to as ‘single protective domains’, and those containing two molecules are referred to as ‘pair protective domains’. Molecules in separate protective domains behave independently. Each molecule undergoes drift-diffusion within its protective domain, and may undergo unimolecular reactions. Two molecules in the same protective domain may additionally participate in bimolecular reactions. To maintain independence when bimolecular reactants are in different protective domains, we require a separation of at least one reaction radius. For non-reacting molecular species we allow for overlap to prevent the size of protective domains from going to zero. Additional details on constructing and updating protective domains are provided in Appendix A.

2.2.2 Events, Times, and Updating

Each event that may occur will have a type, time, and location. The two major event types are first passage from a protective domain and reaction. First passage from a protective domain occurs when a molecule first reaches an absorbing boundary of its protective domain. In DL-FPKMC with general $V(\boldsymbol{x})$ (resp. FPKMC with constant $V(\boldsymbol{x})$), times and locations for first-passage events from single protective domains are sampled from probability densities determined from approximate (resp. exact) solutions of Eq. (2.1). Similarly, for pair protective domains, solutions of Eq. (2.2) are used to sample times and locations for first-passage events or for bimolecular reactants to first reach a separation of one reaction radius. The time for a unimolecular reaction to occur is sampled from an exponential distribution with a specified reaction rate, and a corresponding reaction location is sampled from a "no-passage" probability density for the molecule involved (see below).

To facilitate the discussion of these first-passage and reaction events, we use specific names for three times. The 'global time' will refer to when the most recent event has occurred, irrespective of its particular type or which molecules were involved. An 'individual time' and a 'next event time' will be associated with each particular molecule. 'Individual time' will refer to when the molecule was last updated, and 'next event time' will refer to the sampled time at which the molecule might next undergo an event. Individual times are less than or equal to the global time, and next event times are greater than the global time. All times are stored as absolute times.

Usually, the individual time and location of a molecule will only be updated when the molecule undergoes a major event (first-passage or reaction). In this case, the time and location of the molecule will be updated to the time and location of the event. However, a molecule can also be updated to any specified time prior to its next event time, by sampling a new position for the molecule within the protective domain from the conditional probability density for the molecule to be at a position within the domain, at the specified time, and not yet have undergone a first-passage or reaction event. This procedure is called a 'no-passage' update.

2.2.3 Overall Algorithm

The FPKMC or DL-FPKMC algorithm is carried out according to the following steps:

1. Protective domains are defined around each molecule or pair of molecules, as shown in Fig. 2.1b.
2. The next individual event for each molecule or pair of molecules is determined by sampling an event type, time, and location. In Fig. 2.1c, each next event time is labeled by a τ_i.
3. To determine global events, the individual events are sorted in a priority queue ordered from the shortest event time to the longest event time. For example, τ_5 denotes the shortest event time in Fig. 2.1c.
4. The next global event is determined from the priority queue using the next individual event with the shortest time. The global time and the individual time(s) of the participating molecule(s) are updated to the event time. In the case of a first-passage event for a molecule to leave its protective domain, the molecule's location is updated to the sampled first-passage location, as shown in Fig. 2.1d. If this molecule is in a pair protective domain with another

molecule, the other molecule is no-passage updated to the new global time. In the case of a reaction event, the reaction products are placed at or about the reaction location, as specified in Table 1.1.

5. Molecules in protective domains that are close to or overlap the newly updated molecules are no-passage updated to the new global time.
6. New protective domains are constructed only for those molecules that have undergone an update to reach the current global time, as shown in Fig. 2.1f. New events are sampled for these updated molecules, and the event times are sorted into the priority queue. All other molecules and events remain unchanged.
7. Steps 4 through 6 are then repeated.

Note that Step 5 is used to keep the sizes of the protective domains from becoming too small [17], in which case the *effective* time steps used in the FPKMC or DL-FPKMC methods could become very short.

We remark that information about the state of any molecule in the system can be obtained for any particular time during the course of a simulation. For instance, if one would like to sample the locations of all molecules at a specified time, this can be done by taking the state of the system at the largest global time before or equal to the specified time and then no-passage updating each molecule to the specified time.

2.3 Protective Domain Changes during One Simulation

During simulations, updates are made to the protective domains sequentially as events occur changing the state of the system. To illustrate this process, we consider

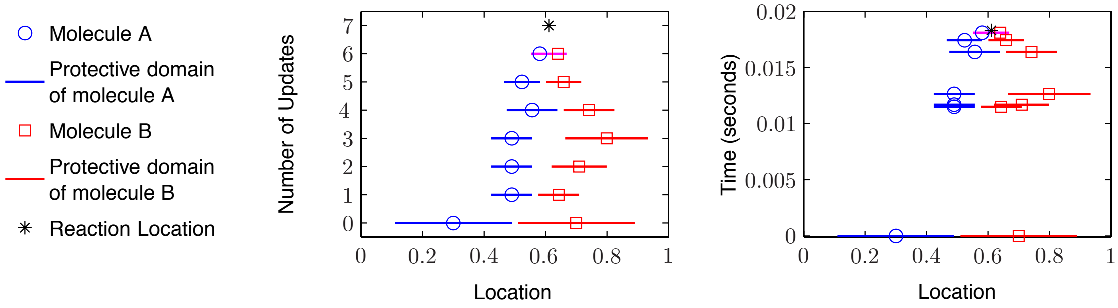

Figure 2.2: One simulation of the reaction $A + B \to \varnothing$ in one dimension, with one molecule each of A and B present initially and $V(x) = 0$. Both panels show the same run of the simulation. In the left panel, the vertical axis is the number of times the simulation cycled through Steps 4 to 6 of the algorithm. In the right panel, the vertical axis is the time of the most recent event.

the simulation in one dimension of the reaction $A+B \to \varnothing$ starting with one molecule each of A and B and using our DL-FPKMC algorithm. One simulation is shown in Fig. 2.2.

In the left panel of Fig. 2.2, the vertical axis is the number of times that the simulation cycled through Steps 4 to 6 of the algorithm; we call this number N_{update}. In the right panel, the vertical axis is the time of the most recent event. At $N_{\text{update}} = 1$ in this particular run of the simulation, molecule A is first-passage updated to the right endpoint of its initial protective domain. This location is close to the left endpoint of molecule B's protective domain, so molecule B is no-passage updated and new protective domains are defined around each molecule. From $N_{\text{update}} = 1$ to $N_{\text{update}} = 3$, molecule B is first-passage updated but does not come close to the protective domain of molecule A, so molecule A is not updated. At $N_{\text{update}} = 6$, the distance between molecules A and B is less than a specified pair threshold, so they are placed in a pair protective domain. At $N_{\text{update}} = 7$, the distance between the molecules reaches the reaction radius and the reaction occurs.

Chapter 3

Propagation of Molecules within Protective Domains in DL-FPKMC

Parts of this chapter were previously published in [37] © 2013 Elsevier Inc.

In this chapter we introduce a lattice discretization of the Fokker-Planck equation [63] which is used within each protective domain. The discretization is chosen to have the form of a master equation, so that the discretization weights can be interpreted as transition rates for continuous-time random walks by the molecules within each protective domain. We then describe how the Stochastic Simulation Algorithm (SSA) [23, 11] is used to generate exact realizations of these random walks within each protective domain, giving the next event times and locations needed by the DL-FPKMC algorithm (see Subsection 2.2.3). It should be stressed that our DL-FPKMC method modifies FPKMC by using this lattice method to propagate molecules *within* their protective domains.

We allow the simulation domain to have reflecting, absorbing, partially-absorbing, or periodic boundaries. Reflecting boundaries are modeled using zero-flux boundary conditions, absorbing boundaries are modeled by zero Dirichlet boundary conditions, and partially-absorbing boundaries by Robin boundary conditions. Protective domains are proper subdomains of the overall domain.

3.1 Lattice Discretization of the Fokker-Planck Equation

In the case of pure diffusion, the probability distributions for first-passage times, first-passage locations, and no-passage locations can all be determined from analytic solutions of the diffusion equation [43, 42, 17, 58]. In contrast, once drift is included, such analytic approaches are no longer possible in general. Instead one must consider probability densities that satisfy Fokker-Planck equations, such as Eqs. (2.1) and (2.2) in the case that the drift arises from a spatially varying potential energy function $V(x)$.

To sample event times and locations in DL-FPKMC, we introduce approximations by treating the movement of each molecule within its protective domain as a discrete-space continuous-time Markov chain, more specifically a continuous-time random walk on discrete mesh points. Jump rates between neighboring mesh points are obtained using the Wang–Peskin–Elston [63] (WPE) spatial discretization of the Fokker-Planck equation

$$\frac{\partial \rho(\boldsymbol{x},t)}{\partial t} = D\nabla\cdot\left(\rho(\boldsymbol{x},t)\nabla V(\boldsymbol{x}) + \nabla\rho(\boldsymbol{x},t)\right). \tag{3.1}$$

Here, we have absorbed the factor of $(k_B T)^{-1}$ into the potential function V. The solution $\rho(\boldsymbol{x},t)$ of the Fokker-Planck equation (3.1), gives the probability density of being at location $\boldsymbol{x}$ at time t. The WPE discretization was derived to have the form of a spatially discrete Markov chain, so that the discretization weights determine the jump rates (i.e. probabilities per unit time) for molecules to hop from one mesh point to another.

Following the approach of [63], we present the discretization for the one-dimensional case. In higher dimensions, the same discretization holds in each coordinate. We first consider the case of a uniform mesh of width h. Let $p_i(t)$ be the probability that a molecule is located at mesh point x_i at time t in the discrete Markov chain model. The mesh point x_i is considered to represent the interval $(x_i - \frac{h}{2}, x_i + \frac{h}{2})$ in the sense that

$$p_i(t) \approx \int_{x_i-\frac{h}{2}}^{x_i+\frac{h}{2}} \rho(x,t)\,dx \approx \rho(x_i,t)h. \tag{3.2}$$

The jump rate for a molecule to hop from the mesh point x_i to a neighboring mesh point x_j, in the case of a uniform mesh of width h, is then given by [63]:

$$a_{ij} = \begin{cases} \frac{D}{h^2}\frac{V(x_j)-V(x_i)}{\exp[V(x_j)-V(x_i)]-1} & \text{for } V(x_i) \neq V(x_j) \\ \\ \frac{D}{h^2} & \text{otherwise.} \end{cases} \tag{3.3}$$

The following alternative form of the jumps rates, which agrees with (3.3) to second order, was later derived in [31][1]:

$$a_{ij} = \frac{D}{h^2}\frac{2}{\exp[V(x_j) - Vx(_i)] + 1} \tag{3.4}$$

Then the time evolution of $p_i(t)$ is described by the master equation

$$\frac{d\,p_i(t)}{dt} = a_{i-1,i}\,p_{i-1}(t) - (a_{i,i-1} + a_{i,i+1})\,p_i(t) + a_{i+1,i}\,p_{i+1}(t). \tag{3.5}$$

If $x_{i\pm1}$ is an absorbing boundary, then $p_{i\pm1}(t) = 0$ in Eq. (3.5). We shall extend Eqs. (3.3) and (3.5) for non-uniform meshes in Subsection 3.3 and in Appendix B.

The discretization given by Eq. (3.3) or Eq. (3.4) has the following properties:

- Converges at second-order for smooth potentials, and can handle discontinuous

---

[1] We use the form of the jump rates in Eq. (3.3) in our initial simulations in one-dimensional spatial domains, and then switch to the form in Eq. (3.4) for later simulations in two and three-dimensional spatial domains.

potentials [63].

- Satisfies a discrete version of detailed balance (zero net flux at equilibrium), which helps reduce artificial drift due to numerical discretization errors [63].
- Is consistent with the standard second-order-accurate discretization of the Laplacian operator, in that a_{ij} converges to D/h^2 as $V(x_j) - V(x_i)$ approaches zero.
- Can incorporate a spatially dependent diffusion coefficient $D(x)$. For example, in the case that $D(x)$ is continuous, the constant D in Eq. (3.3) can be replaced by $[D(x_i) + D(x_j)]/2$ [63].

3.2 Generating Sample Paths

To make use of the WPE discretization, a mesh is defined within each protective domain so that every molecule is located at a mesh point. Rather than numerically solve the master equation (3.5), and then sample this solution to determine next event times, we generate realizations of the jump process described by Eq. (3.5). Each molecule then undergoes a continuous-time random walk on the mesh, with the transition rate from a mesh point x_i to a neighboring point x_j given by a_{ij} of Eq. (3.3). Exact sample paths of the molecules' random walks are generated using the event-driven Stochastic Simulation Algorithm (SSA), specifically Gillespie's "direct method" version of the SSA [23]. In this algorithm, the times of the hops are sampled from exponential distributions. There is no fixed time step. By varying the mesh width, the resolution of this process can be adjusted depending on the desired trade-off between computational efficiency and accuracy. Our specific approach for choosing the mesh width and the locations of mesh points is described in more detail in Subsection 3.4.

Any portion of a protective domain boundary on the interior of the overall domain is absorbing, as are any portions that coincide with an absorbing boundary of the overall domain. As mentioned above, boundary conditions on the intersection of a protective domain boundary with the overall domain boundary agree with the overall domain boundary conditions. Since protective domains are proper subdomains of the overall domain, each protective domain will have at least part of its boundary on the interior of the overall domain. Thus, part of the boundary of each protective domain has absorbing Dirichlet boundaries.

For a newly constructed protective domain containing a single molecule, we determine the molecule's next event time by using the SSA to sample an exact random-walk path for the molecule to hop on the mesh points until it first reaches an absorbing boundary of the protective domain, or is absorbed at a partially-absorbing boundary. The time and location where the molecule exits are called the first-passage time and the first-passage location. For pair protective domains with two molecules, we perform random walks for each molecule until either: (i) one molecule exits the protective domain; or (ii) a reaction occurs when the distance between the two molecules is equal to the reaction radius r_{R}. A no-passage location at any specified time before the next event time can be obtained by finding the last time in the sample path less than or equal to the specified time and taking the location of the molecule at that time.

As will be described in more detail in Section 3.4, the mesh width for pair protective domains in one-dimension is always chosen to exactly divide r_{R}, so that the reaction occurs when the two molecules are exactly one reaction radius apart. In higher dimensions, the mesh width changes adaptively when two reactants are near each other, so that the reaction still occurs at a separation of exactly one reaction radius.

3.3 Discretization for Non-uniform Mesh Cells in One Dimension

In our simulations in one-dimensional spatial domains, non-uniform mesh cells are used when needed to conform to a boundary or to move molecules onto a uniform mesh where the mesh width exactly divides the reaction radius, as will be described in Section 3.4. In higher dimensions, we avoid the need for non-uniform mesh cells by allowing the mesh to change dynamically within each protective domain (also described in Section 3.4), in addition to changing when a molecule's protective domain is updated.

Let x_0 be the initial location of a molecule on a non-uniform mesh, with x_1 and x_2 denoting the locations of the neighboring mesh points in either direction. Note that we may have either $x_1 < x_0 < x_2$ or $x_2 < x_0 < x_1$ (see Fig. 3.1, top row). Let $h_1 = |x_0 - x_1|$ and $h_2 = |x_0 - x_2|$. The jump rates from x_0 to x_j for $j = 1, 2$ are given by

$$a_{0j} = \begin{cases} \frac{2D}{h_j(h_1+h_2)} \frac{V(x_j)-V(x_0)}{\exp[V(x_j)-V(x_0)]-1} & \text{for } V(x_i) \neq V(x_j) \\ \\ \frac{2D}{h_j(h_1+h_2)} & \text{otherwise.} \end{cases} \tag{3.6}$$

The non-uniform discretization in Eq. (3.6) is derived in Appendix B by a finite-volume approach using the fluxes from the WPE discretization of the Fokker-Planck equation [63]. In the case of constant $V(x)$, the non-uniform discretization (3.6) reduces to the non-uniform spatial discretization of the Laplacian at a Dirichlet boundary given by equation (20) of [22]. For solving the Poisson equation with

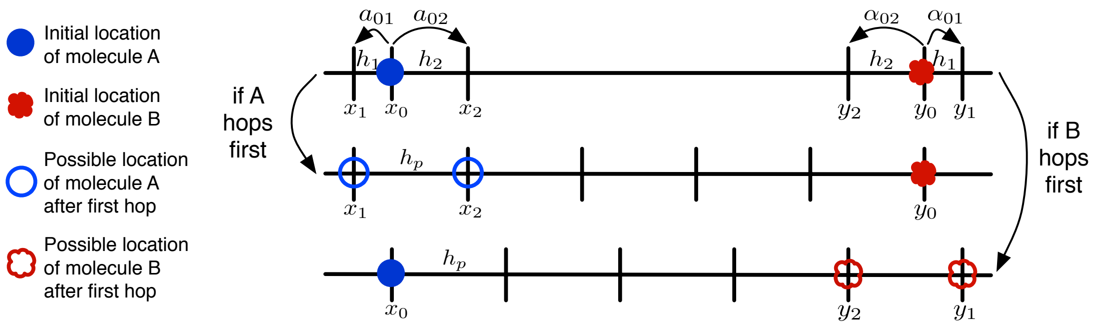

Figure 3.1: The top row shows the non-uniform sub-meshes that are defined when two molecules are first placed in a pair protective domain, for simulations in one-dimensional domains. The second (resp. third) row shows the uniform mesh that is chosen if molecule A (resp. B) hops first. The user-specified uniform mesh width, h_p, is chosen to exactly divide r_{R}. Let d be the initial distance between the two molecules at the time that the pair protective domain is defined. Then, h_2 is defined to be the remainder of the quotient d/h_p, and $h_1 = h_p - h_2$. By choosing h_1 and h_2 in this way, the distances $|y_0 - x_1|$, $|y_0 - x_2|$, $|y_1 - x_0|$, and $|y_2 - x_0|$ are all exactly divisible by h_p. The rates a_{01} and a_{02} are given by Eq. (3.6). The SSA is used to simulate a single hop of one of the molecules to a new point on its sub-mesh, after which the new distance between the two molecules will be one of the four distances divisible by h_p listed above. Then, a new mesh of uniform width h_p can be defined so that both molecules lie exactly on mesh points, and the generation of a sample path using the SSA proceeds until one molecule exits the pair protective domain or the distance between the two molecules reaches r_{R}.

Dirichlet boundary conditions, using a uniform interior mesh and non-uniform mesh cells at the boundaries, this discretization is second-order accurate [22]. To our knowledge, Eq. (3.6) gives a new discretization of the Fokker-Planck equation for non-uniform meshes.

3.4 Choosing the Mesh within Protective Domains

In this subsection we describe how a mesh is defined within each protective domain, given user-specified maximum mesh widths. The convergence tests in Chapters 4 and 5 will demonstrate that DL-FPKMC converges to the underlying SDLR model as the mesh widths are decreased.

In one-dimension protective domains are intervals, and in higher dimensions protective domains will be chosen to be rectangular regions, so that Cartesian grids can be used with the protective domains. For single-molecule protective domains with absorbing boundaries, it will generally be possible to define the protective domain and the mesh in such a way that a single, fixed, uniform mesh can be used within the protective domain. For protective domains intersecting an overall domain boundary and for pair protective domains, non-uniform mesh cells or adaptively changing meshes are used to conform with the reactive boundaries or protective domain boundaries. Note, in our implementation we have used non-uniform mesh cells in one-dimension and adaptively changing meshes in two- and three-dimensions, but either approach could be used in any dimension.

Single-molecule protective domains with absorbing boundaries are chosen to be symmetric about the location of the molecule. A maximum mesh width for all single protective domains, $h_s^{\max}$, is specified by the user. In one-dimension, let r_{PD} be the distance from the molecule to either endpoint of the protective domain. In

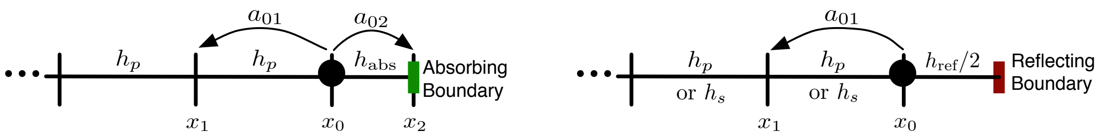

Figure 3.2: Non-uniform mesh cell at an absorbing Dirichlet boundary of a pair protective domain (left panel), or a reflecting boundary of a pair or single protective domain (right panel). Arrows are shown only where the jump rates differ from those given by Eq. (3.3) for uniform mesh cells. The point labeled x_0 is the mesh point closest to the boundary (not necessarily the initial location of a molecule when the protective domain is first defined). When using a non-uniform mesh cell at an absorbing boundary, h_{abs} is defined to be the distance from the absorbing boundary to nearest mesh point of the uniform mesh, x_0. Note that $h_{\text{abs}} \leq h_p$, because otherwise the uniform mesh would have additional mesh points. Then, the jump rates a_{01} and a_{02} in the absorbing boundary case are given by the non-uniform rates in Eq. (3.6) with $h_1 = h_p$ and $h_2 = h_{\text{abs}}$. At a reflecting boundary, h_{ref} is defined to be twice the distance from the boundary to the nearest mesh point, x_0. Then, $h_{\text{ref}}/2 \leq h_p$ or h_s. In this case, a molecule cannot jump from x_0 toward the boundary due to the reflecting zero-flux boundary condition. The jump rate a_{01}, going away from the reflecting boundary, is given by Eq. (3.6) with $h_1 = h_p$ or h_s, and $h_2 = h_{\text{ref}}$.

higher dimensions, let r_{PD} be the perpendicular distance from the molecule to any boundary of the protective domain. Then, for each individual protective domain, the mesh width h_s is calculated according to the formula

$$h_s = \frac{r_{\text{PD}}}{\text{ceil}(r_{\text{PD}}/h_s^{\max})}\,. \tag{3.7}$$

In this way, h_s is always chosen to be the largest value less than or equal to $h_s^{\max}$ that exactly divides r_{PD}. Generally h_s will satisfy $h_s^{\max}/2 < h_s \leq h_s^{\max}$, unless $r_{\text{PD}} \leq h_s^{\max}/2$, in which case $h_s = r_{\text{PD}} \leq h_s^{\max}/2$. In practice, h_s will almost always be strictly less than $h_s^{\max}$, since it is unlikely that $h_s^{\max}$ will exactly divide r_{PD}. After calculating h_s, a uniform mesh with spacing h_s is constructed so that the molecule and the boundaries of the protective domain lie exactly on mesh points. Having the boundaries lie on mesh points allows enforcement of the absorbing Dirichlet boundary conditions at the endpoints without modification of the jumps rates in Eq. (3.3).

In one-dimension, for single-molecule protective domains with one absorbing endpoint and one reflecting endpoint, the mesh width h_s is chosen to be the largest value less than or equal to $h_s^{\max}$ that exactly divides the distance from the molecule to the absorbing endpoint. A mesh is defined so that the molecule and the absorbing endpoint lie exactly on mesh points. The mesh is uniform with the exception of one non-uniform cell used immediately adjacent to the reflecting boundary, as shown in Fig. 3.2 (right panel).

In pair protective domains in one-dimension, the mesh width h_p is a user-specified value chosen to exactly divide r_{R}. Each time that two molecules are placed in a new pair protective domain, the initial distance between the molecules will not necessarily be divisible by h_p. Rather than perturbing the molecules, non-uniform mesh cells are used to move one of the molecules, as shown in Fig. 3.1, so that a uniform mesh of width h_p can be defined with both molecules lying exactly on mesh points. Since this uniform mesh is chosen based on the locations of the two molecules, the endpoints of the protective domain may not conform with the mesh. In this case, one non-uniform mesh cell is used at each endpoint, which may be absorbing or reflecting (see Fig. 3.2).

In two or three dimensions, for pair protective domains or any protective domains with reflecting or partially-absorbing boundaries, we allow the mesh to change dynamically with each protective domain. When molecules are initially placed in such protective domains, each molecule will hop on it's own mesh of spacing $h_p^{\max}$ for pair protective domains or $h_s^{\max}$ for single protective domains, where $h_p^{\max}$ and $h_s^{\max}$ are specified by the user. This continues until a single hop could take a molecule out of the protective domain or until the distance between two molecules is small enough the a single hop could take the molecules to within less than the reaction radius r_{R}.

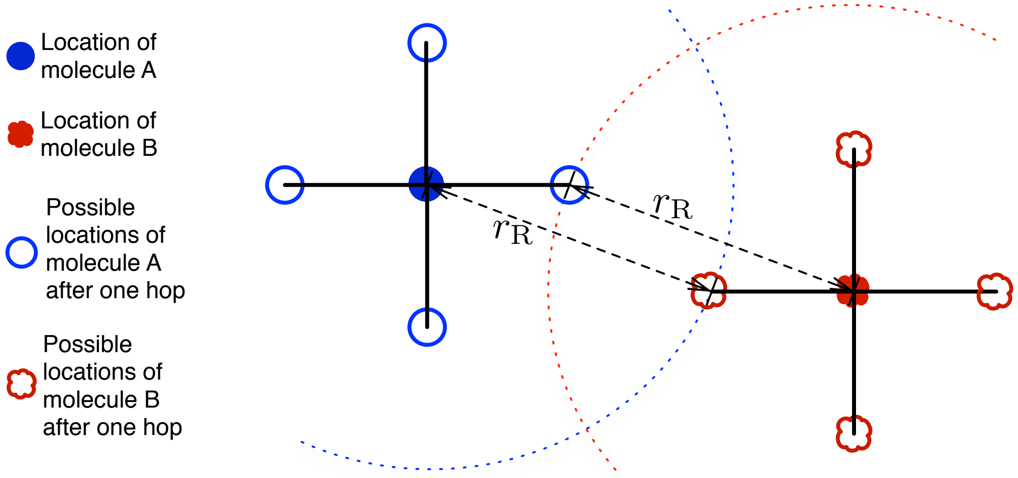

Figure 3.3: Local meshes in pair protective domain when two molecules are near each other, in the case that the pure-absorption reaction mechanism is being used. When two molecules are initially placed in a pair protective domain, each molecule will hop on it's own mesh of spacing $h_p^{\max}$, where $h_p^{\max}$ is a value specified by the user. This continues until the distance between the two molecules is small enough that a single hop could take the molecules to within less than the reaction radius r_{R}. Then, based on the locations of the two molecules and the value of r_{R}, a local mesh of spacing $\leq h_p^{\max}$ is defined for each molecule as shown in the figure. If one molecule hops to the point on its mesh that is nearest to the other molecule, the new distance between the two molecules will exactly equal r_{R}. If a molecule hops to another point on its mesh the molecules' separation will be greater than r_{R}; then, depending on how large the separation is, each molecule will proceed to hop on its own mesh of spacing $h_p^{\max}$, or new local meshes of spacing $\leq h_p^{\max}$ will be defined for each molecule.

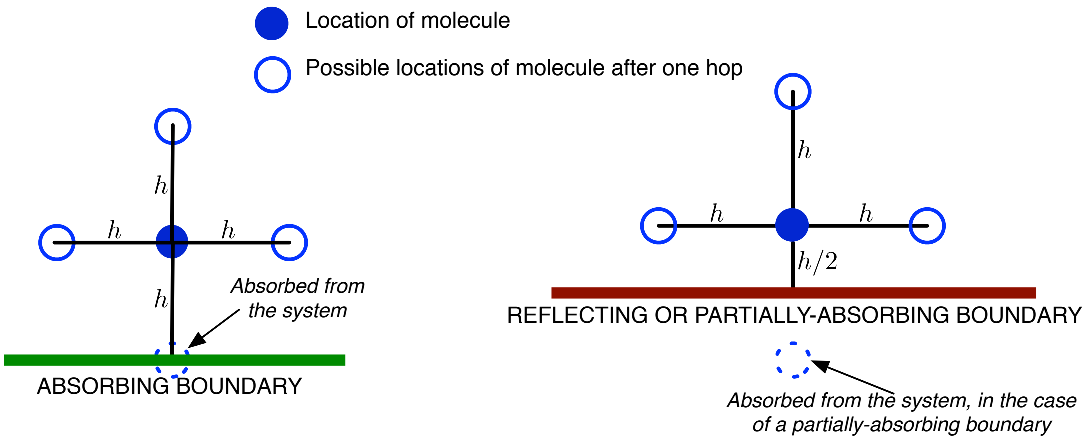

Figure 3.4: Local meshes at an absorbing domain boundary (left panel), or a reflecting or partially-absorbing domain boundary (right panel). The mesh width $h = h_s$ or h_p and locations of mesh points are chosen based on the location of the molecule and the location of the boundary. In the case of an absorbing boundary, h is set equal to the perpendicular distance from the boundary to the molecule. For a reflecting or partially-absorbing boundary, h is defined to be twice the perpendicular distance from the boundary to the molecule.

In the latter case, when using the pure-absorption reaction mechanism, the mesh for the next hop is chosen so that after a single hop by either molecule, the distance between the two molecules may be exactly r_{R} but not smaller (see Fig. 3.3). In the case that the perpendicular distance from a molecule to the nearest absorbing boundary of the protective domain is less than h_s^{max} or h_p^{max}, then the mesh for the next hop is chosen so that a single hop may take the molecule exactly to the boundary. (see Fig. 3.4, left panel). In the case of reflecting or partially-absorbing boundary (either for the reactive boundary, or a domain boundary) the mesh will be defined so that the reflecting or partially absorbing boundary is centered between mesh points (see Fig. 3.4, right panel).

Next, we consider that case that a molecule is near a curved boundary or a corner of a piecewise linear/planar boundary. Since the WPE jump rates give a coordinate-wise discretization, a different mesh width can be used in each coordinate and the

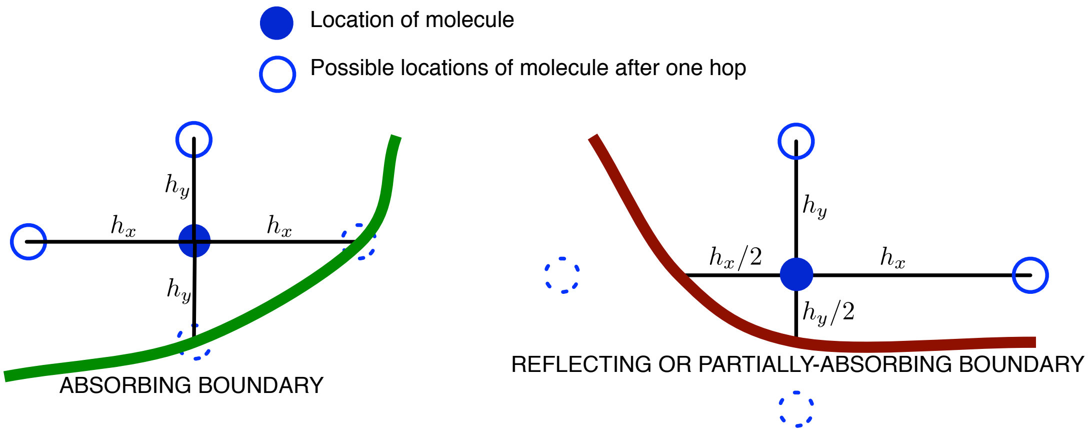

Figure 3.5: Local meshes at a curved domain boundaries: absorbing boundary (left panel), or reflecting or partially-absorbing boundary (right panel). The mesh widths in each coordinate and locations of mesh points are chosen based on the location of the molecule and the location of the boundary, as shown in the figure. The local mesh near a corner of a piecewise linear/plane boundary would be defined in a similar manner.

jump rates along each coordinate are independent of the mesh widths in the other coordinates. As shown in Fig. 3.5, we choose the mesh in each coordinate to conform with the boundary (either in a cell-centered or edge-centered manner depending on whether the boundary is absorbing, partially-absorbing, or reflecting). In the case that two molecules are very near to each other and to a boundary, the local mesh width(s) will be chosen to be the minimum of the values obtained using the approach for pairs (as in Fig. 3.3) and the approach for boundaries (as in Fig. 3.4 or 3.5).

Since h_p in one-dimension is always chosen to exactly divide r_{R}, it necessarily follows that $h_p \leq r_{\mathrm{R}}$. For single molecule protective domains we allow $h_s^{\max}$ to be larger than r_{R}. In higher dimensions, we allow both $h_s^{\max}$ and $h_p^{\max}$ to be larger than r_{R}. In the convergence studies of the next chapters, we set $h_s^{\max} = k h_p^{(\max)}$ where $k \geq 1$, and hold the ratio of $h_s^{\max}$ to $h_p^{(\max)}$ constant as both are reduced to study convergence. As discussed above, the actual mesh widths h_s used in single protective domains are almost always strictly less than $h_s^{\max}$, and either non-uniform or dynamic mesh widths are used in both single and pair protective domains. For these reasons, we keep track of the mean of the mesh widths that are actually used in each simulation. In calculating this mean, each mesh width is weighted by the number of times that it is actually used in a sample path. Then, for all simulations performed with fixed $h_s^{\max}$ and $h_p^{(\max)}$, we calculate an overall mean mesh width by taking the arithmetic mean of the means for each simulation.

Chapter 4

Convergence of DL-FPKMC in One Dimension

This chapter was previously published in [37] © 2013 Elsevier Inc.

In this chapter we perform convergence studies of DL-FPKMC for the annihilation reaction $A+B \to \varnothing$, where the molecular species A and B undergo drift-diffusion subject to various potentials on the interval $[0,1]$. Our results demonstrate both the convergence and accuracy of our method as the mesh widths, $h_s^{\max}$ and h_p, in the discretization are decreased. We denote by $M^A(t)$ and $M^B(t)$ the number of molecules of A and B, respectively, at time t. In the first set of convergence studies (Subsection 4.1), only two molecules are simulated, $M^A(0) = M^B(0) = 1$. Each simulation runs until the two molecules have reacted. A large number (10^7) of simulations are performed in order to minimize the statistical error, so that the error due to the spatial discretization and the rate of convergence can be studied. In the next set of convergence studies (Subsection 4.2), multiple molecules each of A and B are simulated, $M^A(0) = M^B(0) = 10$ or $M^A(0) = M^B(0) = 50$, and each simulation runs until all the molecules have reacted. We will denote the i^{th} molecule of species A by A_i, and the location of A_i at time t by $Q_i^A(t)$. B_j and $Q_j^B(t)$ are defined analogously. In the case $M^A(0) = M^B(0) = 1$, we will drop the subscripts i and j.

4.0.1 Potential Functions and Parameters

The convergence studies are performed using three different potential functions: (i) zero potential, $V_{\text{zero}}(x) = 0$ (which results in pure diffusion); (ii) a cosine potential with two energy wells, $V_{\cos}(x) = \cos(4\pi x)$; and (iii) a step potential with one step,

$$V_{\text{step}}(x) = \begin{cases} 2 & \text{if } x < \frac{1}{2} \\ 0 & \text{if } x \geq \frac{1}{2} \end{cases}.$$

The step potential is used to demonstrate that our DL-FPKMC algorithm with the WPE discretization of the Fokker-Planck equation can handle discontinuous potentials. Note that adding a constant to any potential would not change the results, since the Fokker-Planck equation depends on the derivative of the potential but not the potential itself. In particular, any constant potential would produce the same results as $V(x) = 0$.

In all the convergence studies in this chapter, the length L of the overall domain is 1 unit, the boundaries of the overall domain are reflecting, and the diffusion coefficient D is 1 unit^2/sec for both A and B. The values used for the reaction radius r_{R} will be specified in each subsection. We will use the notation $\mathcal{U}(a, b)$ for the uniform random distribution on the interval (a, b). The initial locations $Q_i^A(0)$ and $Q_j^B(0)$ are drawn from $\mathcal{U}(a, b)$, where $(a, b) \subseteqq (0, L)$ will be specified in each subsection. If $|Q_{i*}^A(0) - Q_{j*}^B(0)| \leq r_{\text{R}}$ for some i^* and j^*, then A_{i*} and B_{j*} will react immediately, at $t = 0$. For $t > 0$, a reaction occurs if $|Q_{i*}^A(t) - Q_{j*}^B(t)| = r_{\text{R}}$.

Note, for most of the convergence studies we run, these initial conditions correspond to non-equilibrium spatial distributions. In the absence of reactions, the equilibrium probability density for a molecule to be at location $x \in (0, L)$ is given by the Gibbs-Boltzmann distribution

$$\frac{e^{-V(x)}}{\int_0^L e^{-V(x)}\,dx}. \tag{4.1}$$

For non-constant $V(x)$, the systems with uniformly distributed initial locations on $(0, L)$ are not initially in spatial equilibrium, since the molecule positions are not distributed according to Eq. (4.1). Furthermore, the systems with initial locations restricted to subintervals of the overall domain are not in spatial equilibrium for any of the potentials used here.

4.0.2 Comparison of DL-FPKMC simulation results in the two-molecule case to analytic and numerical solutions

In general, to describe a stochastic reaction-drift-diffusion system of many molecules by the probability density of having a given number of molecules at specified positions, a large coupled system of partial integro-differential equations is required [27]. In the special case of only two molecules, $M^A(0) = M^B(0) = 1$, with both molecules having the same diffusion coefficient D, the SDLR model for the reaction-drift-diffusion system $A + B \to \varnothing$ in 1D can be described by a single 2D PDE: a Fokker-Planck equation (or a diffusion equation when V is constant). Let $\rho(x, y, t)$ denote the joint probability density for finding molecule A at location x and molecule B at location y at time $t > 0$. We consider the case where the molecules move in an interval domain of length L with zero-flux boundary conditions at both endpoints. The domain for the corresponding 2D PDE is then $\Omega = \{(x, y) : 0 < x, y < L \text{ and } |x - y| > r_{\mathrm{R}}\}$, and $\rho(x, y, t)$ evolves according to the equations

$$
\begin{aligned}
&\frac{\partial \rho(x,y,t)}{\partial t} = D\nabla \cdot \left(\rho(x,y,t) \left(\frac{dV(x)}{dx} + \frac{dV(y)}{dy} \right) + \nabla \rho(x,y,t) \right), && \text{on } \Omega, \\
&\rho(x,y,t) = 0, && \text{on } \partial\Omega \cap \{|x-y| = r_{\mathrm{R}}\}, \\
&\rho(x,y,t)\frac{\partial (V(x)+V(y))}{\partial \boldsymbol{\eta}} + \frac{\partial \rho(x,y,t)}{\partial \boldsymbol{\eta}} = 0, && \text{on } \partial\Omega \backslash \{|x-y| = r_{\mathrm{R}}\},
\end{aligned} \tag{4.2}
$$

where $\boldsymbol{\eta} = \boldsymbol{\eta}(x,y)$ denotes the outward pointing normal at the point (x,y), and the gradient and divergence operators are in the (x,y) coordinates. The 2D domain for Eq. (4.2) is illustrated in Appendix C, left panel of Fig. C.1. When the initial locations of the two molecules in the DL-FPKMC simulations are drawn from $\mathcal{U}(0,L)$, the corresponding initial condition for the 2D Fokker-Planck or diffusion equation is a constant, $\rho(x,y,0) = 1/L^2$. Note, in the following we define ρ on $0 \leq x, y \leq L$ by also defining $\rho(x,y,t) = 0$ for $|x-y| < r_{\mathrm{R}}$ and $t > 0$. The 2D Fokker-Planck and diffusion equations (4.2) can both be solved numerically by finite difference methods, and the diffusion equation can be solved analytically using an eigenfunction expansion. These numerical and analytic solutions are discussed in Appendix C, and provide a baseline against which we compare the results of two-molecule DL-FPKMC simulations in Subsection 4.1.

Let T denote the random variable for the time at which the two molecules react. Using the solution $\rho(x,y,t)$ of Eq. (4.2), we can calculate the survival probability,

$$
S(t) = \Pr\left[T > t\right] = \int_0^L \int_0^L \rho(x,y,t)\, dx\, dy, \tag{4.3}
$$

and the mean reaction time,

$$
\mathbb{E}\left[T\right] = -\int_0^\infty t\, S'(t)\, dt = \int_0^\infty S(t)\, dt. \tag{4.4}
$$

Note that $1 - S(t)$ is the reaction time cumulative distribution function (CDF), and $-S'(t)$ is the corresponding density function.

4.0.3 Statistical Error and Discretization Error

In what follows, we will use the term 'statistical error' to refer to the difference between the empirical value of a statistic (e.g., mean reaction time) estimated from the DL-FPKMC simulations and the upper, or lower, bound of the 99% confidence interval for the statistic. By 'discretization error,' we will mean the difference between the empirical value from the DL-FPKMC simulations and the actual value. In the two-molecule case, actual values are known exactly from the analytic solution $\rho(x, y, t)$ of Eq. (4.2) when V is constant, and are estimated from the numerical PDE solver described in Appendix C when V is not constant. Note that the measured discretization error is the sum of two unknown quantities: sampling error, and the true discretization error due to the spatially discretized nature of the method.

Since we perform a large number of simulations (10^7) at each mesh width in the two-molecule case, the statistical error is quite small, generally between 0.04% and 0.19% for reaction time statistics. Although we perform fewer simulations (4×10^4) when using multiple molecules each of species A and B, the statistical error is still reasonably small, generally between 0.4% and 1%. Our results show that as the mesh width is decreased, the discretization error rapidly decreases to below the statistical error. This demonstrates that the DL-FPKMC algorithm converges and accurately resolves the underlying reaction-drift-diffusion processes.

4.1 Results of Two-Molecule Convergence Studies

In this subsection we consider the reaction $A + B \to \varnothing$ for a system with $M^A(0) = M^B(0) = 1$. Here, the initial locations $Q^A(0)$ and $Q^B(0)$ are drawn randomly from $\mathcal{U}(0,1)$; the reaction radius r_R is 0.02 units; and the pair threshold r_pair is equal to $2r_\mathrm{R}$ (i.e. the molecules are placed in a pair protective domain when $|Q^A(t) - Q^B(t)| \leq r_\mathrm{pair}$). We study the convergence of several statistics as the mesh widths h_s^max and h_p are reduced.

For $V = 0$, the errors in the DL-FPKMC simulation results are generally calculated relative to the exact analytic solution, as determined from the eigenfunction expansion in Appendix C. For V_cos and V_step, the errors are relative to the numerical solution from the PDE solver described in Appendix C. In the PDE solver, we use either the Crank–Nicolson method or the Twizel–Gumel–Arigu method for the time-stepping as explained in the appendix.

As the mesh widths in the DL-FPKMC simulations are decreased, the empirical values of the statistics studied approach the actual values for all three potentials. In each of Figures 4.1, 4.2, and 4.4–4.6, the first panel will show results for $V = 0$, the second panel for V_cos, and the third panel for V_step.

4.1.1 Mean Reaction Times

Let $\mathbb{E}_\mathrm{emp}[T]$ be the empirical mean reaction time calculated from the DL-FPKMC simulations. Define $\mathbb{E}_\mathrm{upp}[T]$ and $\mathbb{E}_\mathrm{low}[T]$ to be, respectively, the upper and lower bounds of the 99% confidence interval for the empirical mean reaction time. We denote by $\mathbb{E}_\mathrm{act}[T]$ the exact analytic mean reaction time in the case that $V = 0$, or the mean reaction time determined from the numerical PDE solution in the $V \neq 0$

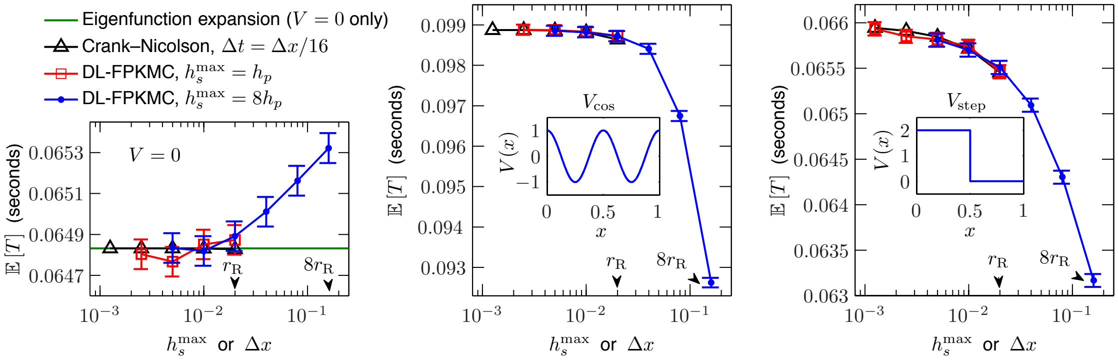

Figure 4.1: Convergence of the mean reaction time, $\mathbb{E}[T]$, for the two-molecule $A+B \to \varnothing$ reaction, as the mesh width is decreased. Note the expanded scales of the vertical axes. Each DL-FPKMC data point shows $\mathbb{E}_{\text{emp}}[T]$ from 10^7 simulations, with 99% confidence intervals. The resulting values of $\mathbb{E}[T]$ from solving Eq. (4.2) using our Crank–Nicolson numerical PDE solver are shown here in the $V = 0$ case to demonstrate the solver's accuracy (see Appendix C for more detail). The mean reaction time calculated from the numerical PDE solution was resolved to an absolute error tolerance of less than 10^{-5} for $V_{\cos}$ and less than 10^{-4} for V_{step}. The PDE solver error tolerances are smaller than the DL-FPKMC statistical errors, allowing the PDE solver estimate for $\mathbb{E}[T]$, determined with the finest value of Δx, to be used in the absence of an analytic value when V is non-constant.

cases (see Eqs. (4.4) and (C.3)). We calculate the relative error by

$$\frac{|\mathbb{E}_{\text{act}}[T] - \mathbb{E}_{\text{emp}}[T]|}{\mathbb{E}_{\text{act}}[T]} \pm \frac{|\mathbb{E}_{\text{bd}}[T] - \mathbb{E}_{\text{emp}}[T]|}{\mathbb{E}_{\text{act}}[T]} \qquad \text{where } \mathbb{E}_{\text{bd}}[T] = \mathbb{E}_{\text{upp}}[T] \text{ or } \mathbb{E}_{\text{low}}[T]. \tag{4.5}$$

The 99% confidence interval for $\mathbb{E}_{\text{emp}}[T]$ is symmetric about the mean, so $\mathbb{E}_{\text{upp}}[T] - \mathbb{E}_{\text{emp}}[T] = \mathbb{E}_{\text{emp}}[T] - \mathbb{E}_{\text{low}}[T]$. If $\mathbb{E}_{\text{act}}[T]$ is contained in the interval $(\mathbb{E}_{\text{low}}[T], \mathbb{E}_{\text{upp}}[T])$, then the discretization error $|\mathbb{E}_{\text{act}}[T] - \mathbb{E}_{\text{emp}}[T]|$ is less than the statistical error $|\mathbb{E}_{\text{bd}}[T] - \mathbb{E}_{\text{emp}}[T]|$.

Figure 4.1 shows $\mathbb{E}_{\text{emp}}[T]$ plotted against $h_s^{\max}$ as the mesh widths are varied, and in Fig. 4.2 the relative errors in $\mathbb{E}_{\text{emp}}[T]$ calculated by Eq. (4.5) are plotted against the mean mesh width. The insets in Fig. 4.1 show the respective potentials. As the mesh widths are reduced, the discretization errors decrease to less than the corresponding statistical errors. Note that the statistical errors are very small since

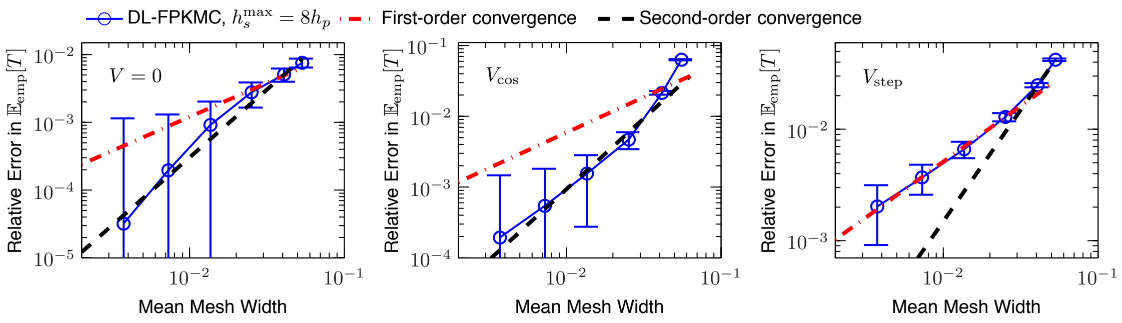

Figure 4.2: Relative errors (4.5) in the empirical mean reaction time, $\mathbb{E}_{\text{emp}}[T]$, for the two-molecule $A + B \to \varnothing$ reaction. In the $V = 0$ case, the errors in $\mathbb{E}_{\text{emp}}[T]$ from the DL-FPKMC simulations are calculated relative to the exact analytic $\mathbb{E}_{\text{act}}[T]$ given by Eq. (C.3). For $V \neq 0$, the errors in $\mathbb{E}_{\text{emp}}[T]$ are calculated relative to the $\mathbb{E}_{\text{act}}[T]$ determined from the numerical solution of Eq. (4.2) using the Crank–Nicolson method with the finest value of Δx (see Appendix C). Note that the vertical axes have different scales in each panel. Here $h_s^{\max} = 8h_p$, and the ratio is held constant as both are decreased by successive halving. Each data point is based on 10^7 simulations with fixed values of $h_s^{\max}$ and h_p. The error bars are determined by Eq. (4.5) using the 99% confidence intervals for the empirical mean reaction times. Note, the error bars are symmetric (and similar in size for the different mesh widths), but appear asymmetric (and larger for finer meshes) due to the log scale. Relative errors are not plotted for the simulations with $h_s^{\max} = h_p \leq r_{\text{R}}$, because the discretization errors are small compared to the statistical errors (cf. Fig. 4.1, red lines with square markers).

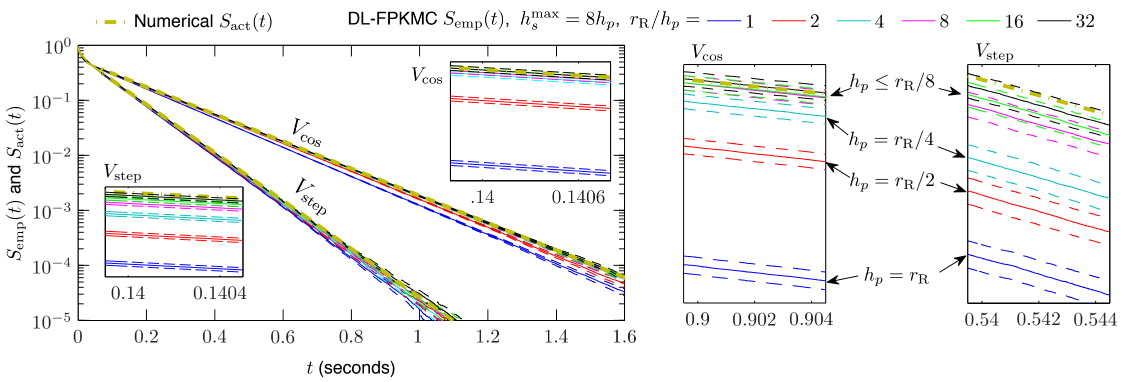

Figure 4.3: Convergence of empirical survival probabilities, $S_{\mathrm{emp}}(t)$, for the two-molecule $A + B \to \varnothing$ reaction with $V_{\cos}$ and V_{step}. Each empirical survival probability function is based on 10^7 simulations. The dashed lines show 99% confidence bounds. (Note that $r_{\mathrm{R}} = 0.02$ units is constant; r_{R}/h_p changes as h_p and $h_s^{\max}$ are successively halved.) $S_{\mathrm{emp}}(t)$ from the $V = 0$ case is not plotted because it is difficult to visually distinguish between the results for different mesh widths. In the case $V = 0$ and $h_p = r_{\mathrm{R}}/32$, the analytic $S_{\mathrm{act}}(t)$ is contained within the 99% confidence bounds of $S_{\mathrm{emp}}(t)$ at *every* time point t_i at which $S_{\mathrm{act}}(t)$ was evaluated ($> 6 \times 10^4$ time points); for $h_p = r_{\mathrm{R}}/8$ and $h_p = r_{\mathrm{R}}/16$, $S_{\mathrm{act}}(t)$ is contained within the confidence bounds of $S_{\mathrm{emp}}(t)$ at more than 96% of the time points. This demonstrates that the discretization error is small compared to the statistical error.

10^7 simulations were performed at each mesh width.

We estimate the rate of convergence to be approximately second-order for $V_{\cos}$ and approximately first-order for V_{step}. This is consistent with the convergence rates of the WPE discretization of the Fokker-Planck equation for smooth and discontinuous potentials. In the $V = 0$ case, it is difficult to draw a conclusion about the rate of convergence since the discretization errors are small relative to the statistical errors; however, for the same reason, we can conclude that the results are very accurate in this case.

4.1.2 Survival Probabilities

For each mesh width, the empirical survival probability $S_{\mathrm{emp}}(t)$ and the associated 99% confidence bounds are calculated using the MATLAB function '`ecdf`'. The confidence bounds determined by '`ecdf`' are symmetric to within numerical precision at all but a few points, which are at the tails of the distributions. $S_{\text{act}}(t)$ will denote the analytic survival probability for $V = 0$, and the numerical survival probability determined from the PDE solver for $V \neq 0$ (see Eqs. (4.3), (C.2) and Appendix C for more information). Figure 4.3 plots the empirical survival probabilities $S_{\text{emp}}(t)$ with 99% confidence bounds, and shows convergence of $S_{\text{emp}}(t)$ to $S_{\text{act}}(t)$ as the mesh widths are reduced.

To quantify the magnitude of the error, we calculate the distances between $S_{\text{emp}}(t)$ and $S_{\text{act}}(t)$ using the following methods. The discrete L^1, L^2, and L^∞ norms of a function $u(t)$ are given by

$$||u(t)||_{L^1} = \sum_i |u(t_i)|\Delta t_i, \quad ||u(t)||_{L^2} = \left\{ \sum_i u(t_i)^2 \Delta t_i \right\}^{\frac{1}{2}}_{,} \quad ||u(t)||_{L^\infty} = \max_i |u(t_i)|. \tag{4.6}$$

The relative error for $S_{\text{emp}}(t)$ we report in each norm is then given by

$$\frac{||S_{\text{act}}(t) - S_{\text{emp}}(t)||}{||S_{\text{act}}(t)||}, \tag{4.7}$$

where the norms are evaluated on the interval $t \in [0,\ S^{-1}(10^{-6})]$ with $\Delta t_i = t_{i+1} - t_i$. The time points, t_i, used in evaluating the norms correspond to those at which the numerical PDE solutions were calculated (see Appendix C for more information). In the $V = 0$ case, the analytic expression for the survival probability was evaluated at those same t_i's. In all cases, the empirical DL-FPKMC survival probabilities were linearly interpolated to obtain values at every t_i.

The L^∞ error is equivalent to the Kolmogorov distance between distributions [12], which is a statistical distance. Another statistical distance is the Kullback–Leibler (KL) divergence [34], which is a measure of the information lost when a distribution

$G(t)$ is used to approximate a (true) distribution $F(t)$. For continuous probability distributions, let $g(t)$ and $f(t)$ be the corresponding densities. Then, the KL divergence of $G(t)$ from $F(t)$ is given by

$$KL(F||G) = \int f(t) \ln\left(\frac{f(t)}{g(t)}\right) dt. \tag{4.8}$$

For discrete distributions

$$KL(F||G) = \sum_i F_i \ln\left(\frac{F_i}{G_i}\right). \tag{4.9}$$

Since empirical densities obtained from Monte Carlo simulations are noisy, we approximate $KL(S_{\text{act}}||S_{\text{emp}})$ of Eq. (4.8) by binning the data into a finite number of time intervals and then applying Eq. (4.9). We use either 10 or 20 bins: 9 or 19 evenly-sized bins on the interval $t \in [0, S^{-1}(0.01)]$, and one bin for $t > S^{-1}(0.01)$. This approximation of the KL divergence becomes noisier as more bins are used. Note, while the KL divergence can be interpreted as a measure of the difference between two distributions, it is does not define a metric.

Figure 4.4 shows the relative errors measured in norm (4.7) and the KL divergence (4.9), plotted against the mean mesh width. The magnitude of the relative errors and the estimated convergence rates for $S_{\text{emp}}(t)$ are similar to those for $\mathbb{E}_{\text{emp}}[T]$. For each measure we see that the error in the survival probabilities decreases as the mesh widths are reduced. For finer mesh sizes, the discretization error becomes small compared to the statistical error (see Fig. 4.3). As such, the observed rate of convergence decreases and the graphs become less regular in Fig. 4.4.

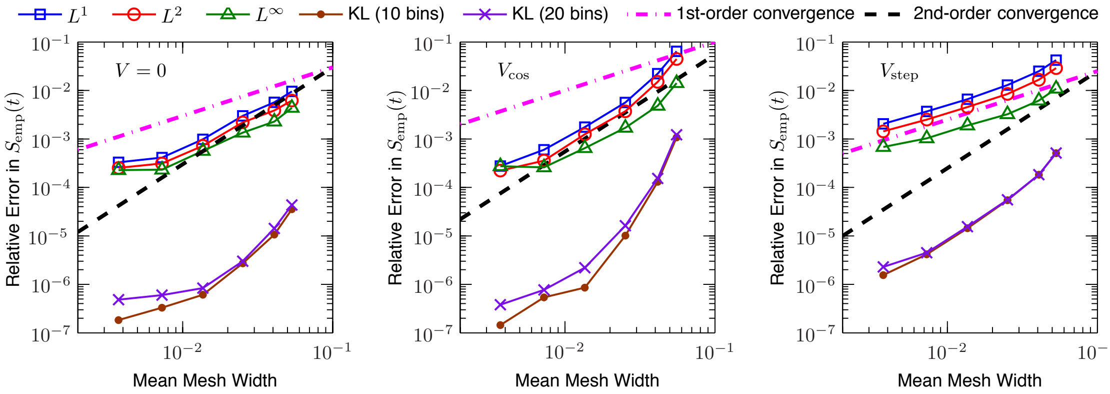

Figure 4.4: Errors in the empirical survival probability $S_{\text{emp}}(t)$, relative to $S_{\text{act}}(t)$, for the two-molecule $A + B \to \varnothing$ reaction: relative errors measured in norm (4.7) and KL divergence (4.9). Here $h_s^{\max} = 8h_p$ as both are decreased. Each empirical survival probability function is based on 10^7 simulations. Note that the slight plateauing of the errors for finer meshes is due to sampling error, as illustrated in Fig. 4.3 by the size of the confidence bounds relative to the smaller distance between $S_{\text{emp}}(t)$ and $S_{\text{act}}(t)$.

4.1.3 Discrete Joint Spatial Probabilities

Recall that the joint density, $\rho(x, y, t)$, is the solution of Eq. (4.2). We now study the convergence of the joint probability that $(Q^A(t), Q^B(t))$ is contained in one of $N_x \times N_x$ subregions of the domain Ω of Eq. (4.2). For $1 \le i, j \le N_x$, define the discrete joint probability function for the locations of the two molecules at time t by

$$\begin{aligned} p_{\text{act}}(i,j,t) &= \Pr\left[\frac{i-1}{N_x}L < Q^A(t) < \frac{i}{N_x}L,\ \frac{j-1}{N_x}L < Q^B(t) < \frac{j}{N_x}L\right] \qquad (4.10)\\ &= \int_{(j-1)L/N_x}^{jL/N_x} \int_{(i-1)L/N_x}^{iL/N_x} \rho(x,y,t)\,dx\,dy. \end{aligned}$$

We study the convergence of the empirical $p_{\text{emp}}(i,j,t)$ from the DL-FPKMC simulations at two fixed times, $t = S_{\text{act}}^{-1}(0.75)$ and $t = S_{\text{act}}^{-1}(0.5)$, where $p_{\text{emp}}(i,j,t)$ is determined by binning the locations of the molecules at time t into the $N_x \times N_x$ subregions.

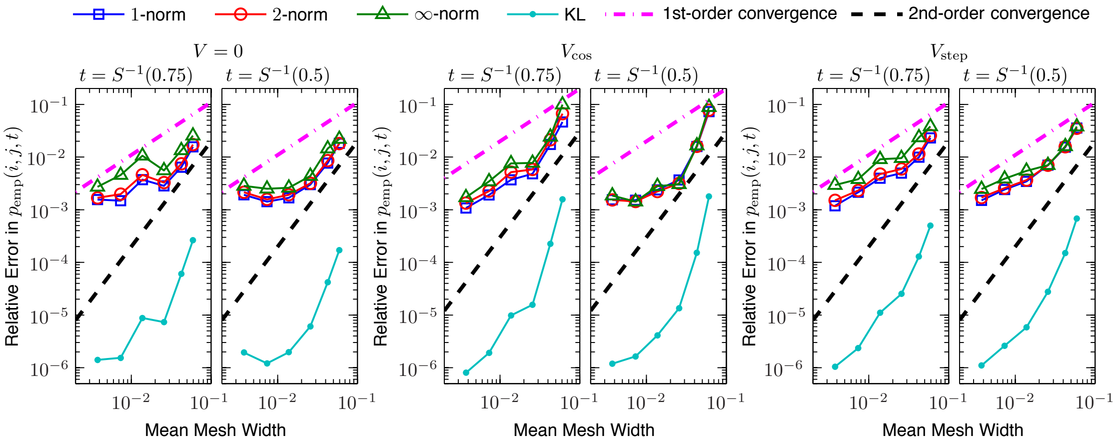

Figure 4.5: Relative errors and KL divergence of the empirical joint spatial probabilities $p_{\text{emp}}(i, j, t)$ at $t = S_{\text{act}}^{-1}(0.75)$ and $t = S_{\text{act}}^{-1}(0.5)$ for the two-molecule $A + B \to \varnothing$ reaction. For all graphs we divide the domain into 25 equally-sized bins ($1 \leq i, j \leq 5$). The errors are calculated relative to $p_{\text{act}}(i, j, t)$, as determined from the PDE solver with Twizell–Gumel–Arigu time-stepping described in Appendix C. Here $h_s^{\text{max}} = 8h_p$ as both are decreased. Each data point is based on 10^7 simulations. At $t = S_{\text{act}}^{-1}(0.75)$ (resp. $t = S_{\text{act}}^{-1}(0.5)$) the two molecules have not yet reacted in approximately 75% (resp. 50%) of the simulations. We also recorded $p_{\text{emp}}(i, j, t)$ at $t = S_{\text{act}}^{-1}(0.25)$, however the convergence results were more strongly affected by noise, since the molecules had already reacted by that time in most of the simulations.

We calculate the relative error in $p_{\text{emp}}(i,j,t)$ at a fixed time t by

$$\frac{||p_{\text{act}}(i,j,t) - p_{\text{emp}}(i,j,t)||}{||p_{\text{act}}(i,j,t)||}, \tag{4.11}$$

where the norms in the numerator and denominator are given by

$$||q(i,j,t)||_1 = \sum_{i=1}^{N_x}\sum_{j=1}^{N_x} |q(i,j,t)|, \qquad ||q(i,j,t)||_2 = \left\{\sum_{i=1}^{N_x}\sum_{j=1}^{N_x} q(i,j,t)^2\right\}^{\frac{1}{2}}, \tag{4.12}$$
$$\text{or} \quad ||q(i,j,t)||_\infty = \max_{1\le i\le N_x}\max_{1\le j\le N_x} |q(i,j,t)|,$$

with $q = p_{\text{act}} - p_{\text{emp}}$ for the numerator and $q = p_{\text{act}}$ for the denominator.

We also calculate the KL divergence (4.9) of $p_{\text{emp}}(i,j,t)$ from $p_{\text{act}}(i,j,t)$. Since the KL divergence is only defined for probabilities that sum to one, we include an extra "bin" representing that probability that the two molecules have already reacted, which is given by $1 - S(t) = 1 - \sum_{i=1}^{N_x}\sum_{j=1}^{N_x} p(i,j,t)$. The relative errors and KL divergence are shown in Fig. 4.5. We see similar convergence behavior as in the previous subsection, with the errors clearly decreasing for coarser mesh widths, but plateauing as sampling error becomes dominant at smaller mesh widths.

4.1.4 Reaction Location Distributions

For a reaction between A and B, the location of the reaction will be taken to be $Q^{\text{rxn}} = (Q^A(t) + Q^B(t))/2$ where t is the time of the reaction. Then $F^{\text{rxn}}(x) = \Pr[Q^{\text{rxn}} \le x]$ is the cumulative distribution function for the reaction locations. In studying convergence of the empirical $F^{\text{rxn}}_{\text{emp}}(x)$ as the mesh widths are reduced, we consider only those reactions that occur after $t = 0$. (When $|Q^A(0) - Q^B(0)| \le r_{\text{R}}$, the reaction occurs immediately at $t = 0$ regardless of the mesh width.) For the

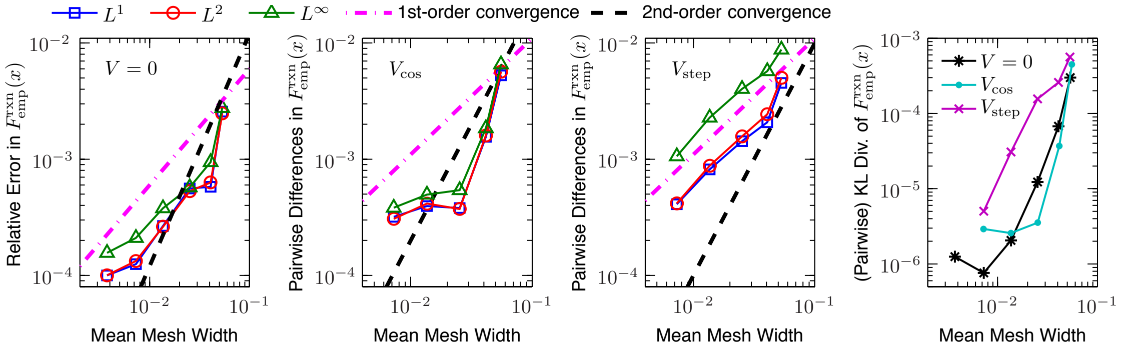

Figure 4.6: Errors in $F^{\mathrm{rxn}}_{\mathrm{emp}}(x)$, the empirical reaction location CDF, for the two-molecule $A + B \to \varnothing$ reaction: errors relative to the analytic $F^{\mathrm{rxn}}_{\mathrm{act}}(x)$ for $V = 0$, and pairwise relative differences between $F^{\mathrm{rxn}}_{\mathrm{emp}}(x)$ from successive mesh widths for V_{cos} and V_{step}. The pairwise relative differences are plotted against the mean mesh width for the coarser of the two meshes. Here $h^{\max}_s = 8h_p$ as both are decreased. Each empirical CDF is based on the simulations in which the two molecules react after time $t = 0$, out of 10^7 total simulations. The plateaus in the errors or differences at the finer mesh widths are due to sampling error. For example, in the case of V_{cos}, the 99% confidence bounds for $F^{\mathrm{rxn}}_{\mathrm{emp}}(x)$ from the two finest mesh widths overlap at more than 96% of the 10^4 points at which they were evaluated.

parameters used here approximately 3.96% of the reactions occur at $t = 0$.

For $V = 0$, the exact distribution $F^{\mathrm{rxn}}_{\mathrm{act}}(x)$ can be calculated analytically from the solution of Eq. (4.2). See Eq. (C.4) of Appendix C for the analytic result. In this case, we study convergence by comparing $F^{\mathrm{rxn}}_{\mathrm{emp}}(x)$ from the DL-FPKMC simulations to the analytic $F^{\mathrm{rxn}}_{\mathrm{act}}(x)$. For the $V \neq 0$ cases, we examine the successive pairwise differences between $F^{\mathrm{rxn}}_{\mathrm{emp}}(x)$ from the DL-FPKMC simulations as the mesh widths are decreased. In all cases, the distributions $F^{\mathrm{rxn}}(x)$ are evaluated at 10^4 evenly spaced points.

For $V = 0$, the relative errors in each norm (4.6) are calculated by $||F^{\mathrm{rxn}}_{\mathrm{act}}(x) - F^{\mathrm{rxn}}_{\mathrm{emp}}(x)|| \,/\, ||F^{\mathrm{rxn}}_{\mathrm{act}}(x)||$. For V_{cos} and V_{step}, the pairwise relative differences are calculated by the preceding formula with $F^{\mathrm{rxn}}_{\mathrm{emp}}(x)$ from the finer of the two meshes replacing $F^{\mathrm{rxn}}_{\mathrm{act}}(x)$. The corresponding density is approximated by binning the reaction locations into 20 evenly-sized bins, so that the KL divergence (4.9) can be calculated (comparing to the analytic values for $V = 0$, and comparing the empirical values from successive mesh widths for $V_{\cos}$ and V_{step}). The results are shown in Fig. 4.6. Note that the relative errors and pairwise relative differences are all less than 1%. Although the results for the spatial statistics are more affected by noise (Figs. 4.5, 4.6), the convergence rates appear similar to those for the time statistics (cf. Figs. 4.2, 4.4).

4.2 Results of Multiple-Molecule Convergence Studies

In the convergence studies for multiple molecules undergoing the reaction $A+B \to \varnothing$, we start the simulations with either 20 molecules ($M^A(0) = M^B(0) = 10$) or 100 molecules ($M^A(0) = M^B(0) = 50$). In the 20-molecule simulations: $r_{\text{R}} = 0.02$ units; $r_{\text{pair}} = 2r_{\text{R}}$; $Q_i^A(0) \sim \mathcal{U}(0.1, 0.4)$ and $Q_j^B(0) \sim \mathcal{U}(0.6, 0.9)$ for $1 \leq i,\ j \leq 10$. In the 100-molecule simulations: $r_{\text{R}} = 0.001$ units; $r_{\text{pair}} = 4r_{\text{R}}$; $Q_i^A(0)$ and $Q_j^B(0) \sim \mathcal{U}(0, 1)$ for $1 \leq i,\ j \leq 50$. Two molecules A_{i*} and B_{j*} are placed in a pair protective domain if they are closer to each other than to any other molecules of the opposite type, i.e. if

$$|Q_{i*}^A(t) - Q_{j*}^B(t)| = \min_i |Q_i^A(t) - Q_{j*}^B(t)| = \min_j |Q_{i*}^A(t) - Q_j^B(t)|,$$

and if the distance between them satisfies

$$\begin{aligned} &|Q_{i*}^A(t) - Q_{j*}^B(t)| \leq r_{\text{pair}} \\ &\qquad\qquad\text{and} \\ &|Q_{i*}^A(t) - Q_{j*}^B(t)| \leq r_{\text{R}} + \min(\min_{i \neq i*}(|Q_{i*}^A(t) - Q_i^A(t)|,\ \min_{j \neq j*} |Q_{j*}^B(t) - Q_j^B(t)|). \end{aligned} \tag{4.13}$$

The last condition in Eq. (4.13) was added to prevent the length of the protective domain of a molecule, say A_i, from approaching zero when A_i is close to another

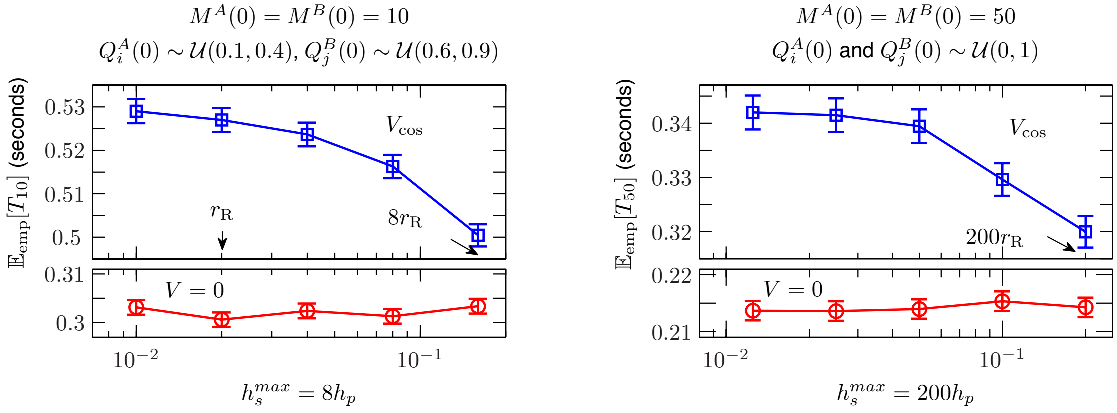

Figure 4.7: Mean time for all molecules to react via the reaction $A + B \to \varnothing$. Error bars show 99% confidence intervals, based on 4×10^4 simulations per data point. *Left panel*: $r_{\text{R}} = 0.02$ units. *Right panel*: $r_{\text{R}} = 0.001$ units.

molecule of the same type, say A_{i*}, where A_{i*} would otherwise have been placed in a pair with B_{j*}.

Each simulation runs until all of the molecules have reacted. Let T_n denote the random variable for the time at which the n^{th} reaction occurs. Since we are working with the irreversible reaction $A + B \to \varnothing$, we have that $M^A(t) = M^A(0) - n$ for $t \in [T_n, T_{n+1})$, and similarly for $M^B(t)$. Figure 4.7 shows the empirical mean time for all molecules to react; this is $\mathbb{E}_{\text{emp}}[T_{10}]$ when $M^A(0) = M^B(0) = 10$ and $\mathbb{E}_{\text{emp}}[T_{50}]$ when $M^A(0) = M^B(0) = 50$. With these values of $M^A(0)$ and $M^B(0)$, the SDLR model will correspond to a large coupled system of partial integro-differential equations for the probability densities of having a specified number of molecules at specified locations. It is no longer feasible to solve these equations numerically to obtain high-accuracy solutions for assessing the empirical convergence of our DL-FPKMC method. As such, we now estimate the accuracy of reaction time statistics by comparing DL-FPKMC results at coarser mesh widths to results obtained with the finest mesh width.

When $V(x) = \cos(4\pi x)$ (Fig. 4.7, top panels), we see convergence as the mesh

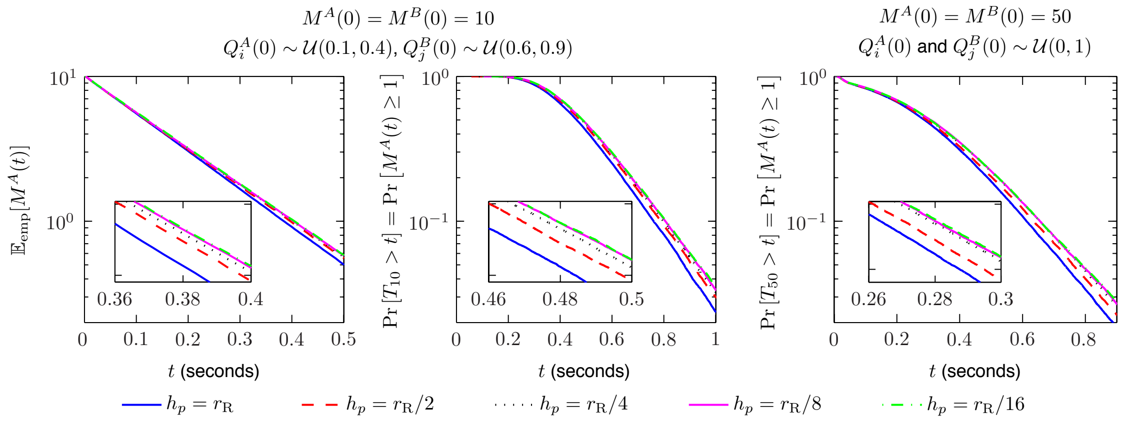

Figure 4.8: Mean number of molecules of A remaining at time t, $\mathbb{E}_{\text{emp}}[M^A(t)]$ (left panel), and probability that at least one molecule of A remains at time t, $\Pr\left[M^A(t) \geq 1\right]$ (center and right panels). $V = \cos(4\pi x)$. Each graph is based on 4×10^4 simulations. $\mathbb{E}_{\text{emp}}[M^A(t)]$ is not plotted in the case $M^A(0) = M^B(0) = 50$, because the results for the different mesh widths are essentially indistinguishable until $\mathbb{E}_{\text{emp}}[M^A(t)] \lesssim 4$. *Left and center*: $r_{\text{R}} = 0.02$ units, $h_s^{max} = 8h_p$. *Right*: $r_{\text{R}} = 0.001$ units, $h_s^{max} = 200h_p$.

width is decreased. The percent difference between $\mathbb{E}_{\text{emp}}[T_{10}]$ for the coarsest mesh width compared to the finest mesh width is approximately 5.4%; this percent difference for $\mathbb{E}_{\text{emp}}[T_{50}]$ is approximately 7%. These differences are comparable in size to the explicitly calculated discretization error of approximately 6.4% at the coarsest mesh width for $\mathbb{E}_{\text{emp}}[T]$ in the two-molecule case (previous subsection).

When $V = 0$ (Fig. 4.7, bottom panels), the confidence intervals for all mesh widths overlap, indicating that the discretization error is less than the statistical error even for the coarsest mesh width. The statistical errors when $V = 0$ are between 0.49% and 0.81%. If the unknown discretization error here is comparable in size to the known discretization error in $\mathbb{E}_{\text{emp}}[T]$ in the two-molecule case, which was approximately 0.87% at the coarsest mesh width, that would provide an explanation for why can not observe convergence when $V = 0$.

$\mathbb{E}_{\text{emp}}[M^A(t)]$ is the mean number of A molecules remaining at time t, and $\Pr\left[M^A(t) \geq 1\right]$ is the probability that at least one molecule of A remains at time t.

Figure 4.8 shows convergence of $\mathbb{E}_{\mathrm{emp}}[M^A(t)]$ and $\Pr\left[M^A(t) \geq 1\right]$ as the mesh width is decreased, in the case $V(x) = \cos(4\pi x)$. Results are not plotted for the $V = 0$ case, because the confidence bounds for different mesh widths overlap; this indicates that the results are resolved to within the statistical error, even for coarser mesh widths.

Chapter 5

Convergence of DL-FPKMC in Two and Three Dimensions

In this chapter we perform convergence studies of DL-FPKMC simulating the two-molecule annihilation reaction $A + B \to \varnothing$, where the two molecules undergo drift-diffusion in a two-dimensional square domain (Section 5.1) or three-dimensional free space (Section 5.2).

5.1 Results of Two-Dimensional Convergence Studies

In this section we consider a square domain with sides of length $L = 0.2$ μm. The initial locations of the two molecules, $\boldsymbol{Q}^A(0)$ and $\boldsymbol{Q}^B(0)$, are drawn from a uniform random distribution over the square. The reaction radius r_{R} is 5 nm, and the pair threshold r_{pair} is set to $8r_{\mathrm{R}}$. As in the previous chapter, two molecules may be placed in a pair protective domain if $|\boldsymbol{Q}^A(t) - \boldsymbol{Q}^B(t)| \leq r_{\text{pair}}$. For the diffusion coefficients, $D_A = 1\,\mu\mathrm{m}^2 s^{-1}$ and $D_B = 2\,\mu\mathrm{m}^2 s^{-1}$. Here, we set $h_s^{\max} = h_p^{\max}$, and study convergence as both are reduced from $8r_{\mathrm{R}}$ to $r_{\mathrm{R}}/8$. For each fixed value of $h_s^{\max} = h_p^{\max}$, 10^5 simulations are performed and each simulation runs until the two molecules have reacted. As in Section 4.1, T denotes the random variable for the time at which the two molecules react. Two potential functions are considered:

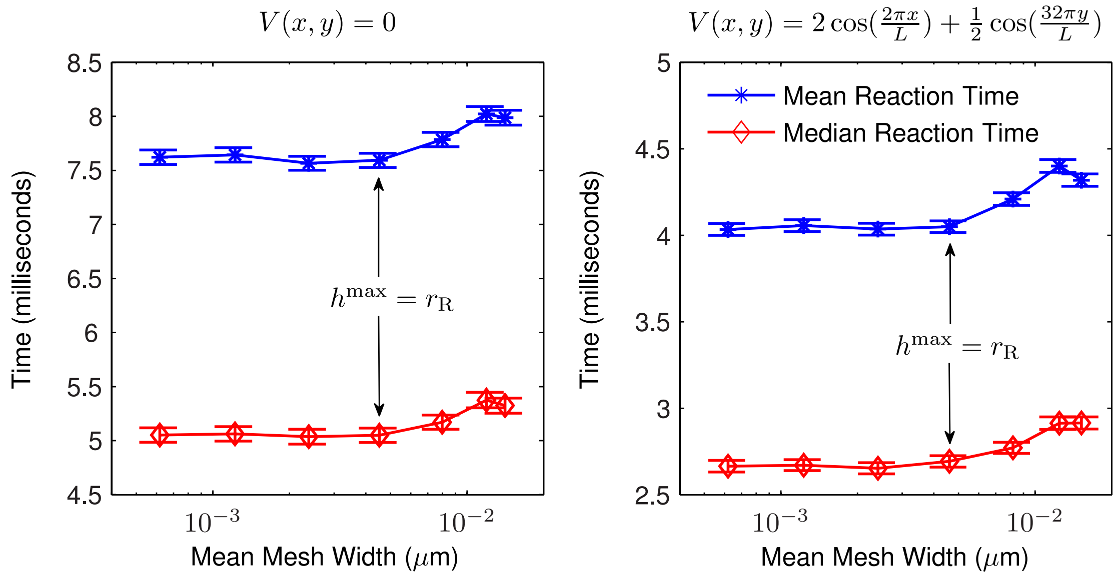

$V(x,y) = 0$

$V(x,y) = 2\cos(\frac{2\pi x}{L}) + \frac{1}{2}\cos(\frac{32\pi y}{L})$

Figure 5.1: Empirical mean and median reaction times for the two-molecule $A + B \to \varnothing$ reaction in a two-dimensional square domain with sides of length $L = 0.2$ μm. The statistical error bars show 99% confidence intervals. The confidence intervals for the median reaction time are calculated using the V_{PB_2} estimator for the variance of the median [46]. The confidence intervals for the mean reaction time are determined from the usual Standard Error of the Mean. Each data point is based on 10^5 simulations with $h_s^{\max} = h_p^{\max}$, $r_{\mathrm{R}} = 5\,\mathrm{nm}$, $D_A = 1\,\mu\mathrm{m}^2 s^{-1}$, $D_B = 2\,\mu\mathrm{m}^2 s^{-1}$, and $r_{\mathrm{pair}} = 8\,r_{\mathrm{R}}$. $Q^A(0)$ and $Q^B(0)$ are uniformly randomly distributed over the square.

$V(\boldsymbol{x}) = V(x,y) = 0$ and $V(\boldsymbol{x}) = V(x,y) = 2\cos\left(\frac{2\pi x}{L}\right) + \frac{1}{2}\cos\left(\frac{32\pi y}{L}\right)$.

Figure 5.1 shows the convergence of the empirical mean and median reaction times as the mesh widths are reduced. When $h_s^{\max} = h_p^{\max}$ has been reduced to $\lesssim r_{\mathrm{R}}$, the discretization error is less than the statistical error. The empirical survival probabilities $S_{\mathrm{emp}}(t)$ are plotted in Fig. 5.2. Again, the confidence bounds overlap for $h_s^{\max} = h_p^{\max} \lesssim r_{\mathrm{R}}$, indicating that the results have converged to within statistical error.

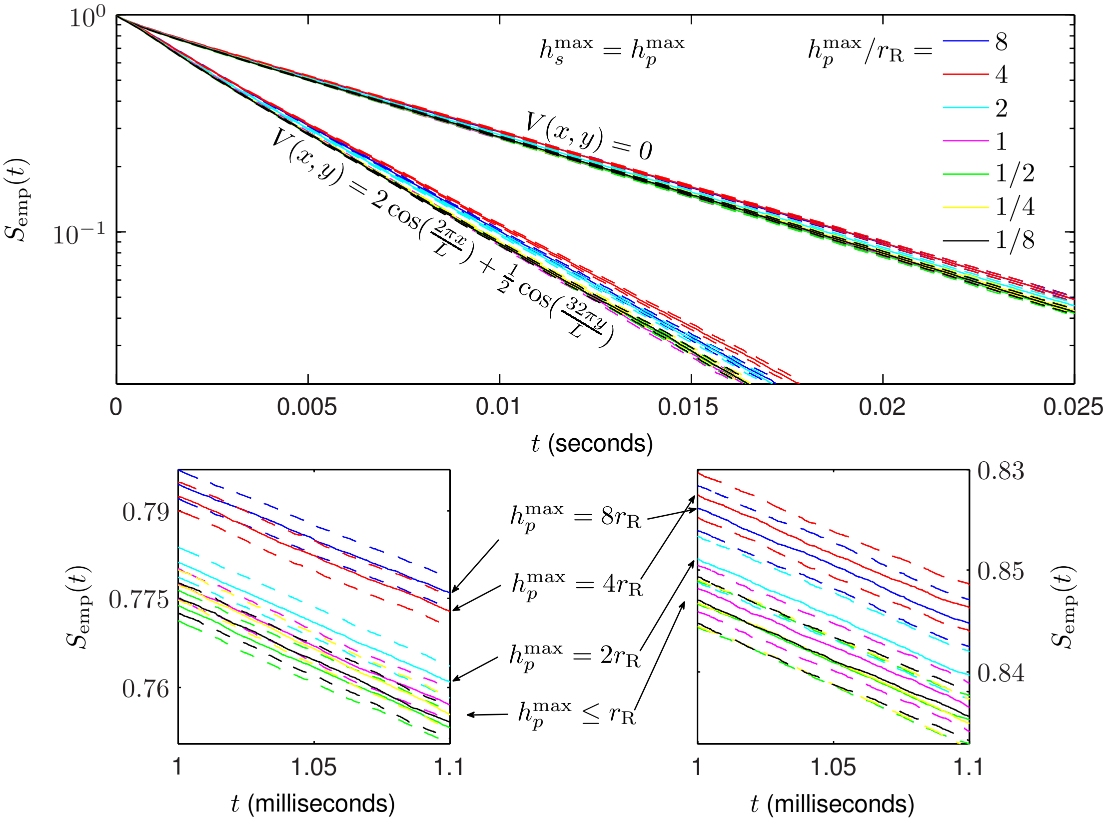

Figure 5.2: Empirical survival probabilities $S_{\mathrm{emp}}(t)$ for the two-molecule $A + B \to \varnothing$ reaction in a two-dimensional square domain with sides of length $L = 0.2\ \mu$m. The dashed lines show 95% confidence intervals. Each empirical survival probability function is based on 10^5 simulations with $h_s^{\max} = h_p^{\max}$. The parameters values are the same as in Fig. 5.1.

5.2 Results of Three-Dimensional Convergence Studies

In this section we study convergence of the DL-FPKMC method simulating the two-molecule $A + B \to \varnothing$ reaction in three-dimensional free space. We set $r_{\mathrm{R}} = 10\,\mathrm{nm}$, $D_A = 4\,\mu\mathrm{m}^2 s^{-1}$, $D_B = 6\,\mu\mathrm{m}^2 s^{-1}$, and $r_{\mathrm{pair}} = 4\,r_{\mathrm{R}}$. For the initial locations, $\boldsymbol{Q}^A(0) = (0,0,0)$ and $\boldsymbol{Q}^B(0)$ is uniformly randomly distributed over the sphere of radius $2r_{\mathrm{R}}$ centered at the origin. Thus, the initial separation between the two molecules is $||\boldsymbol{Q}^A(0) - \boldsymbol{Q}^B(0)|| = 2\,r_{\mathrm{R}}$. The potential function will be $V(\boldsymbol{x}) = V(x,y,z) = 0$ or $V(\boldsymbol{x}) = V(x,y,z) = -2\cos(\frac{\pi x}{4 r_{\mathrm{R}}})$.

In the case that $V(\boldsymbol{x}) = 0$, the diffusion equation for this two-body problem in free space can be solved analytically by switching to separation coordinates, $\boldsymbol{Q}^A(t) - \boldsymbol{Q}^B(t)$ [30]. In free-space there is some probability that the two molecules will never react, but T will still be used to denote the random variable for the time at which they react. The analytic solution of the diffusion equation in the separation coordinates can be used to derive the following formulas [30] for the reaction time CDF

$$F_{\mathrm{act}}(t) = \Pr[T \leq t] = \frac{r_{\mathrm{R}}}{||\boldsymbol{Q}^A(0) - \boldsymbol{Q}^B(0)||}\,\mathrm{erfc}\left(\frac{||\boldsymbol{Q}^A(0) - \boldsymbol{Q}^B(0)|| - r_{\mathrm{R}}}{\sqrt{4t(D_A + D_B)}}\right) \tag{5.1}$$

and for the probability that the two molecules will ever react

$$\lim_{t\to\infty} \Pr[T \leq t] = \frac{r_{\mathrm{R}}}{||\boldsymbol{Q}^A(0) - \boldsymbol{Q}^B(0)||}. \tag{5.2}$$

For the parameters chosen in these convergence studies, the probability that the two molecules will ever react is 50%.

We set $h_s^{\max} = h_p^{\max}$, and study convergence as both are reduced from $4\,r_{\mathrm{R}}$ to $r_{\mathrm{R}}/16$. For each fixed value of $h_s^{\max} = h_p^{\max} := h^{\max}$, 10^6 simulations are performed. We run each simulation until $t = 0.0015$ seconds for $h^{\max} \geq r_{\mathrm{R}}/4$, $t = 8 \times 10^{-5}$

seconds for $h^{\max} = r_{\mathrm{R}}/8$, and $t = 1.2 \times 10^{-5}$ seconds for $h^{\max} = r_{\mathrm{R}}/16$. When $V(\boldsymbol{x}) = 0$, the probability that the two molecules have reacted at $t = 0.0015$ is $\sim 95.4\%$ of the probability that they will ever react. At $t = 8 \times 10^{-5}$ the probability is $\sim 80.3\%$, and at $t = 1.2 \times 10^{-5}$ it is $\sim 51.9\%$.

Figure 5.3 shows the empirical reaction time CDFs, $F_{\mathrm{emp}}(t)$, and the upper and lower 95% confidence bounds, $F_{\mathrm{upp}}(t)$ and $F_{\mathrm{low}}(t)$, as determined using MATLAB's '`ecdf`' function. In the $V(\boldsymbol{x}) = 0$ case, the analytic CDF (5.1) is also plotted for comparison. For each empirical CDF, the relative errors shown in Fig. 5.4 are calculated by

$$\frac{||F_{\mathrm{act}}(t) - F_{\mathrm{emp}}(t)||}{||F_{\mathrm{act}}(t)||} \pm \frac{||F_{\mathrm{upp}}(t) - F_{\mathrm{low}}(t)||/2}{||F_{\mathrm{act}}(t)||} \tag{5.3}$$

in each of the L^1, L^2, and L^∞ norms (4.6). The norms are evaluated on the interval $t \in [0, 1.2 \times 10^{-5}]$. In the $V = 0$ case, $F_{\mathrm{act}}(t)$ is the known analytic CDF (5.1). For $V(x, y, z) = -2\cos(\pi x/4r_{\mathrm{R}})$, no analytic or high-accuracy numerical solution is available, so errors are calculated relative to the DL-FPKMC simulation results with the finest mesh width used ($h^{\max} = r_{\mathrm{R}}/16$). Similar to the one-dimensional convergence results of the previous chapter, the rate of convergence appears to be approximately second order.

The Kullback–Leibler divergence (4.8) of $F_{\mathrm{emp}}(t)$ from $F_{\mathrm{act}}(t)$, also shown in Fig. 5.4, is approximated by binning the data into a finite number of time intervals and then calculating the discrete KL divergence defined in Eq. (4.9). We use 5×2^n $(n = 0, \ldots, 6)$ evenly-sized bins on the interval $t \in [0, 1.2 \times 10^{-5}]$, and one bin for $t > 1.2 \times 10^{-5}$.

As discussed in Section 3.4 the actual mesh widths that are used, h_s and h_p in single and pair protective domains respectively, change adaptively and are frequently smaller than $h_s^{\max}$ and $h_p^{\max}$. Figure 5.5 shows histograms of the mesh widths used as

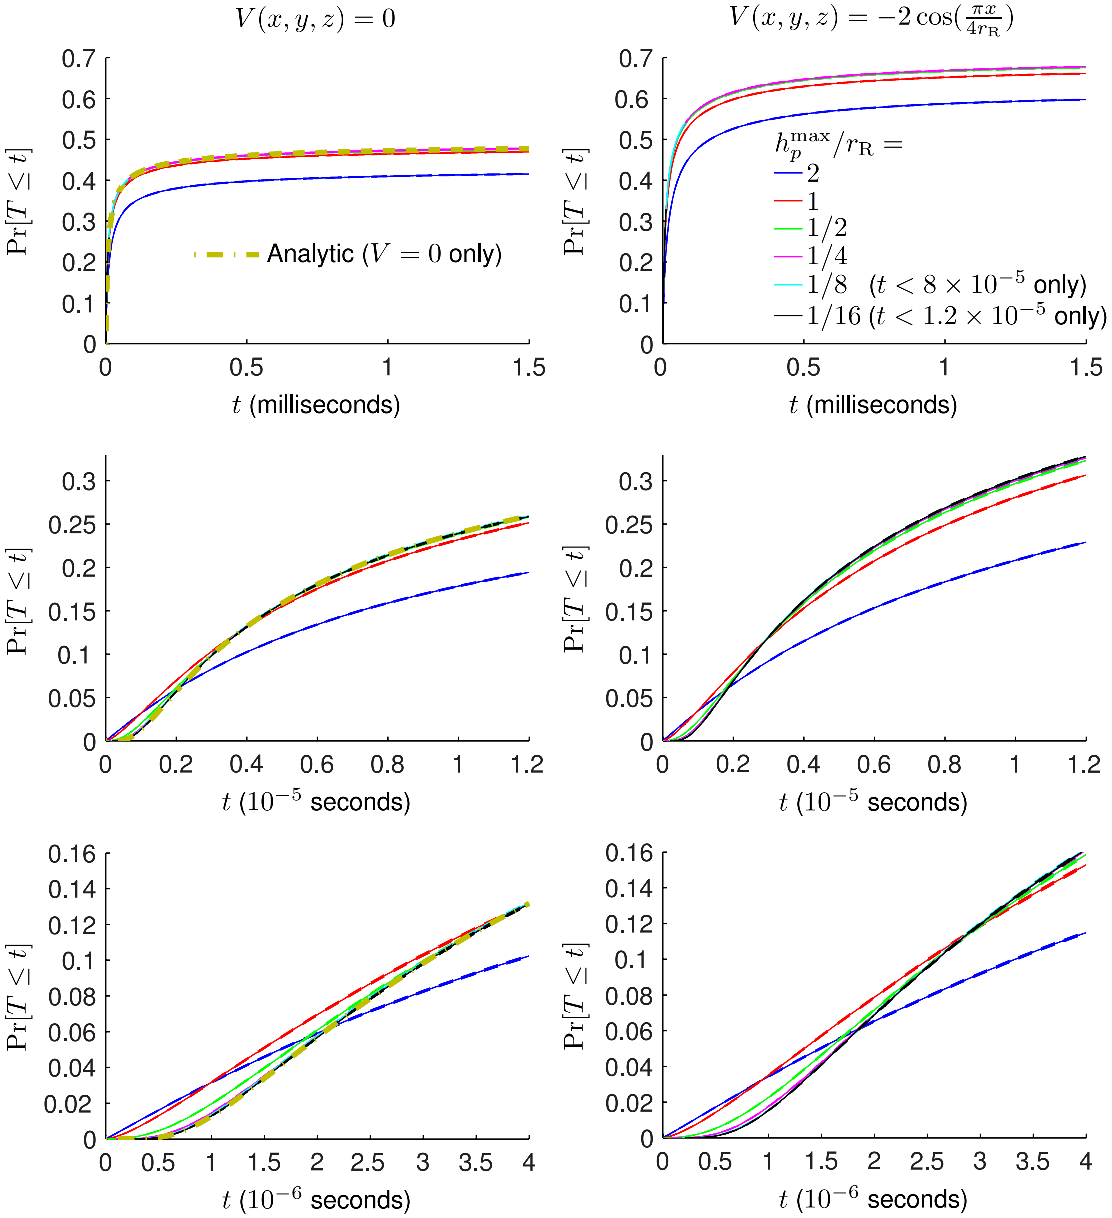

Figure 5.3: Reaction time CDFs, $F(t) = \Pr[T \leq t]$, for the two-molecule $A + B \to \varnothing$ reaction in three-dimensional free space. Each row shows the same results, but plotted on different time scales (zooming in from top row to bottom row). The dashed lines show 95% confidence intervals. Each empirical CDF is based on 10^6 simulations with $h_s^{\max} = h_p^{\max}$, $r_R = 10\,\text{nm}$, $||\boldsymbol{Q}^A(0) - \boldsymbol{Q}^B(0)|| = 2\,r_R$, $D_A = 4\,\mu\text{m}^2 s^{-1}$, $D_B = 6\,\mu\text{m}^2 s^{-1}$, $r_{\text{pair}} = 4\,r_R$. $\boldsymbol{Q}^A(0) = (0,0,0)$ and $\boldsymbol{Q}^B(0)$ is uniformly randomly distributed over the sphere of radius $2\,r_R$ centered at the origin. The results for $h^{\max} = 4\,r_R$ are not plotted, because they are similar to the results for $h^{\max} = 2\,r_R$.

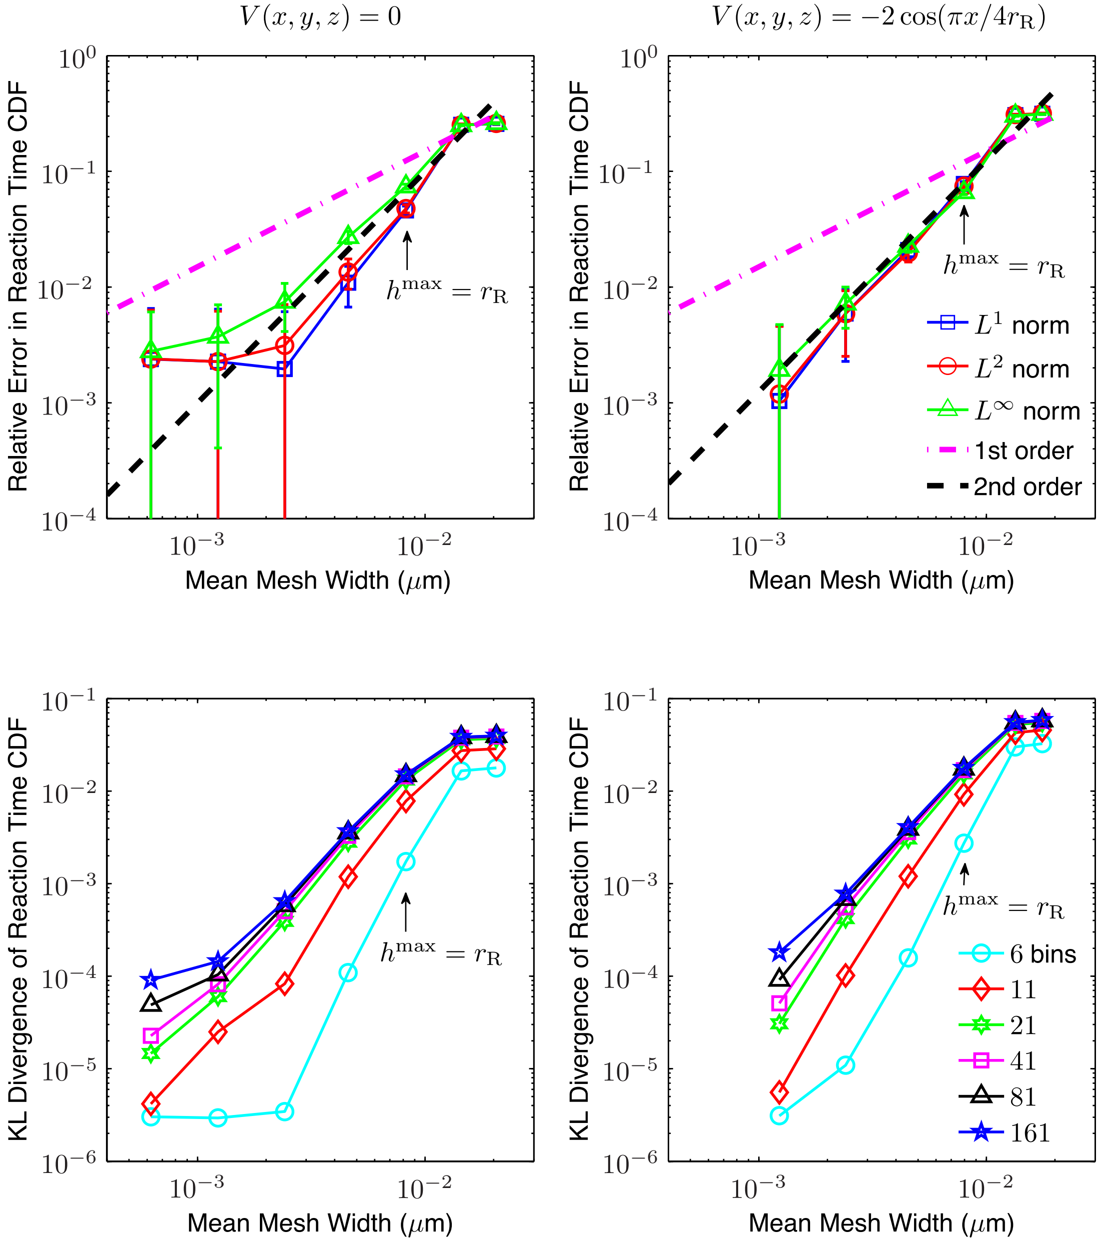

Figure 5.4: Relative errors in the empirical reaction time CDFs, $\Pr[T \leq t]$, on the interval $t \in [0, 1.2 \times 10^{-5}]$ for the two-molecule $A + B \to \varnothing$ reaction in three-dimensional free space (same parameters values as Fig 5.3). The maximum mesh width, $h^{\max} = h_s^{\max} = h_p^{\max}$, is varied from $4r_{\mathrm{R}}$ to $r_{\mathrm{R}}/16$. In the case of $V(\boldsymbol{x}) = 0$, the errors are calculated relative to the analytic CDF. For $V(\boldsymbol{x}) = V(x, y, z) = -2\cos(\pi x/2r_{\mathrm{R}})$, the errors are relative to the simulation results obtained using the finest mesh width, $h^{\max} = r_{\mathrm{R}}/16$. The statistical error bars are determined from the formula in Eq. 5.3 using the 95% confidence bounds for the empirical CDFs.

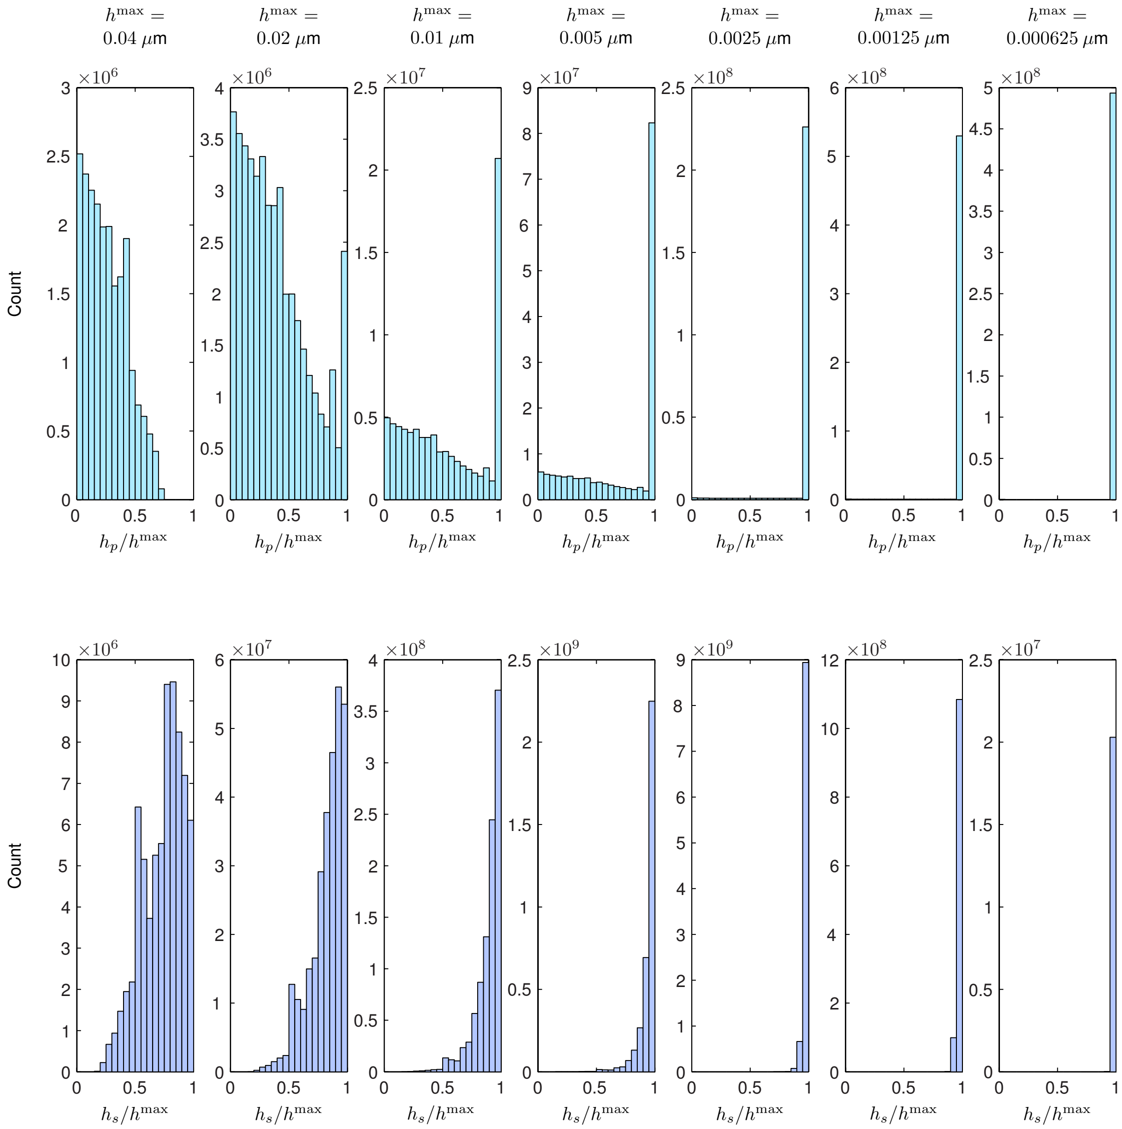

$h^{\max} = 0.04\ \mu$m

$h^{\max} = 0.02\ \mu$m

$h^{\max} = 0.01\ \mu$m

$h^{\max} = 0.005\ \mu$m

$h^{\max} = 0.0025\ \mu$m

$h^{\max} = 0.00125\ \mu$m

$h^{\max} = 0.000625\ \mu$m

Figure 5.5: Histograms of the mesh widths used (h_s and h_p) as the maximum mesh width ($h^{\max} = h_s^{\max} = h_p^{\max}$) is decreased from $4\,r_{\mathrm{R}}$ to $r_{\mathrm{R}}/16$, for the two-molecule $A + B \to \varnothing$ reaction in three-dimensional free space with $V(\boldsymbol{x}) = 0$. The parameters values are the same as in Figs. 5.3 and 5.4. Each histogram is based on 10^6 simulations. Each simulation was run until $t = 0.0015$ seconds for $h^{\max} \geq r_{\mathrm{R}}/4$, $t = 8 \times 10^{-5}$ seconds for $h^{\max} = r_{\mathrm{R}}/8$, and $t = 1.2 \times -5$ seconds for $h^{\max} = r_{\mathrm{R}}/16$.

the maximum mesh widths are decreased. When $h_s^{\max}$ and $h_p^{\max}$ are large, h_s and h_p are frequently much smaller than the maximum values. In pair protective domains, this occurs when the separation between the two molecules is small compared to $h_p^{\max}$. From the formula in Eq. 3.7, we saw that h_s will usually satisfy $h_s^{\max}/2 < h_s \leq h_s^{\max}$, but can be smaller if the length of the protective domain is less than $h_s^{\max}$.

Chapter 6

Running Time Analysis

This chapter was previously published in [37] © 2013 Elsevier Inc.

In this chapter we demonstrate that the running time of the DL-FPKMC algorithm when simulating the $A + B \to \varnothing$ reaction scales linearly with the number of molecules in the system. We also compare the computational performance of DL-FPKMC to a second method in which all molecules hop on a fixed, global, uniform lattice. Both methods were implemented in MATLAB, and an attempt was made to take advantage of reasonable and standard optimizations. That said, we make no claim that our implementation of either method provides optimal computational performance. All DL-FPKMC and fixed lattice simulations were performed in MATLAB on a Sun Fire X4600 M2 x64 server. The server was configured with four AMD Opteron Model 8220 processors (2.8 GHz dual-core) and 16 GB of RAM.

In both methods, bimolecular reactions occur when the distance between two reactants is exactly equal to r_{R}. To enforce this condition in the fixed lattice method, the lattice spacing h will be chosen to equal r_{R}. We set $V(x) = 0$, so the spatial hopping rates in the fixed lattice method are simply D/h^2. The results of the convergence studies in Section 4 indicate that in DL-FPKMC taking $h_p \approx r_{\mathrm{R}}$ and $h_s^{\max} \approx L/50$ is sufficient to resolve the reaction and diffusion processes to within statistical error for $n \leq 10^7$ simulations. For biologically relevant parameter values, e.g. $L = 10\mu\text{m}$ and $r_\text{R} = 1\text{nm}$, this will allow $h_s^{\max}$ in DL-FPKMC to be a couple orders of magnitude larger than h in the fixed lattice method without compromising accuracy.

To our knowledge, the only reaction system that has previously been simulated *in one dimension* using any FPKMC-type method is the $A + A \to \varnothing$ [43, 42] reaction, in which any two molecules annihilate as soon as they collide. In contrast, for the reaction system $A+B \to \varnothing$, molecules of the same type do not react. As mentioned in Subsection 2.2, we allow the protective domains of non-reacting molecules to overlap to prevent the size of protective domains from going to zero when two non-reacting molecules approach each other. We expect that this issue could also be addressed by allowing protective domains to contain more than two molecules. In order to simulate the same underlying process with both DL-FPKMC and the fixed lattice method, non-reacting molecules in the fixed lattice simulations are allowed to occupy the same lattice site and cross each other.

6.1 Parameters

In comparing the DL-FPKMC and fixed lattice methods, the following parameters values are used: $L = 10\mu\text{m}$, $r_\text{R} = 1\text{nm}$, $D = 10\mu\text{m}^2/\text{sec}$, and $V(x) = 0$. The overall simulation domain, the interval $(0, L)$, has reflecting boundaries. The initial number of molecules of A and B are equal, $M^A(0) = M^B(0)$, and each simulation runs until all molecules have reacted. The initial locations of molecules are uniformly distributed over the interval $(0, L)$. In both DL-FPKMC and the fixed lattice method, if the initial distance between a molecule of A and a molecule of B is less than or equal to r_R, then they react immediately. All later reactions occur when the distance

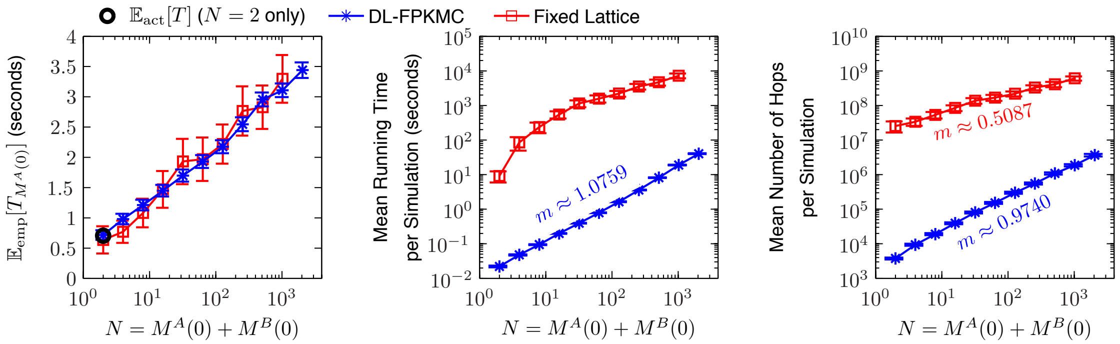

Figure 6.1: Comparison of DL-FPKMC and fixed lattice methods for the reaction $A + B \to \varnothing$ as the number of molecules present initially, $N = M^A(0) + M^B(0)$, is increased. $M^A(0) = M^B(0)$, $Q_i^A(0)$ and $Q_j^B(0) \sim \mathcal{U}(0, L)$ for $1 \le i, j \le M^A(0)$, $L = 10\mu\text{m}$, $r_{\text{R}} = 1\text{nm}$, $D = 10\mu\text{m}^2/\text{sec}$, and $V(x) = 0$. Each DL-FPKMC data point is based on 10^3 simulations, and each fixed lattice data point is based on 10^2 simulations. Error bars indicate 99% confidence intervals.

between an A molecule and a B molecule equals r_{R}. In the fixed lattice simulations, $h = r_{\text{R}} = 1\text{nm}$. In DL-FPKMC, $h_s^{\max} = L/50 = 200\text{nm}$, $h_p = r_{\text{R}} = 1\text{nm}$, and $r_{\text{pair}} = 50\text{nm}$. To check that using a coarser value for $h_s^{\max}$ than h_p is still sufficient to obtain accurate results, we ran 10^6 simulations with $M^A(0) = M^B(0) = 1$. In this case the exact mean reaction time, $\mathbb{E}_{\text{act}}[T]$, is known analytically, see Eq. (C.3). Using the preceding parameters, the resulting $\mathbb{E}_{\text{emp}}[T]$ from the DL-FPKMC simulations agrees with $\mathbb{E}_{\text{act}}[T]$ to within statistical error, which is approximately 0.34%.

6.2 Results of Running Time Analysis

Figure 6.1 compares the simulation results and computational performance of DL-FPKMC to those of the fixed lattice method as the total number of molecules in the system, $N = M^A(0) + M^B(0)$, is varied. As shown in the left panel, the resulting values of $\mathbb{E}_{\text{emp}}[T_{M^A(0)}]$, the mean time for all molecules to react within the simulations, from the two methods agree to within statistical error. The running time (center panel) for DL-FPKMC is two to three orders of magnitude faster than the running time for the fixed lattice. DL-FPKMC also requires fewer hops than the fixed lattice by two to three orders of magnitude (right panel). The slope of the line for DL-FPKMC on the log-log plot of running time (center panel) indicates that the method is approximately $O(N)$, i.e., the running time scales linearly with the number of molecules, $N = M^A(0) + M^B(0)$. Although the fixed lattice method has better scaling with N, it would not appear to become more efficient than DL-FPKMC until N is substantially larger than 10^5. At that large a number of molecules, it is common to transition to more macroscopic stochastic reaction-diffusion models.

As the number of molecules in the system is increased, the proportion of running time spent on different steps of the DL-FPKMC algorithm changes. For the parameter values used in this section, the changes in proportions of running time as N is increased from 2 to 2048 are as follows (approximated using the MATLAB '`profile`' function):

- Generating sample paths: decreases from 67% to 38%,
- Identifying neighboring molecules and defining new protective domains for updated molecules based on locations of neighbors: increases from 17% to 52%,
- Determining which molecules (possibly none) to no-passage update after each first-passage update: decreases from 4% to 3%,
- Sorting the event queue: decreases from 7% to 3%.

These percentages could vary substantially for different parameter values and mesh widths, but we expect the overall trends would hold.

As was described in Subsection 3.2, the DL-FPKMC algorithm generates a sample path for each molecule within its protective domain until a first-passage or reaction event occurs. Whenever a molecule is no-passage updated before its next event time, part of its sample path goes unused. The mean number of hops per simulation shown in Fig. 6.1 (right panel) is based on the total number of hops in all sample paths that are generated, not just the hops that are actually used. Particularly long paths tend to be generated for protective domains touching an overall domain boundary, since molecules can only exit through one endpoint of such protective domains. If we cap the length of single-molecule protective domains that are near the boundaries, then the total number of hops decreases, but the number of times that the protective domains are updated increases and the overall running time also increases. Some guidelines for and difficulties in optimizing the partitioning of space among protective domains in FPKMC are discussed in [17], but we have not yet attempted to address this difficult optimization problem.

6.3 Expectations for Future Studies

It is expected that the implementation of DL-FPKMC in higher dimensions can be done in such a way as to maintain the $O(N)$ scaling for the following reasons:

- FPKMC in two and three dimensions can be implemented to have $O(N)$ scaling [17]. One might expect identifying neighboring molecules to be more costly in higher dimensions, however, in practice the near-neighbor list (NNL) method allows for constant per event costs, leading to $O(N)$ scaling for the overall method [17];

- The only algorithmic difference between DL-FPKMC and FPKMC is that DL-FPKMC generates sample paths within protective domains using continuous-time random walks, whereas FPKMC samples from analytic solutions of the diffusion equation;

- Since the fraction of running time spent generating sample paths in DL-FPKMC decreases as N increases, the scaling of overall running time as N increases is expected to depend mainly on other steps of the algorithm. These steps can all be implemented in the same way in DL-FPKMC as in FPKMC.

See Table 5 of [14] for the computational cost of several reaction-diffusion methods, including methods for simulating the SDLR model and the RDME. The original, approximate GFRD [62] and Smoldyn [3] exhibit $O(N)$ scaling with the total number of molecules for diffusive movements and $O(\sum_{N_R} \prod_{S \in R} N_S)$ scaling for reactive distances, where N_R is the number of reaction channels and N_S is the number of molecules of a given species [14]. We expect FPKMC and DL-FPKMC will have similar scaling with number of reaction channels.

Although $O(N)$ scaling would eventually cause the computational cost to become too high for very large N, many systems of biological interest could still be simulated at low cost. FPKMC has been used to simulate systems with 10^8 particles in one dimension [43] and 216×10^6 particles in three dimensions [42]. For many relevant biological systems, the number of molecules would less than 1000, or even in the single digits in many cases. For example, the chemotaxis system in bacteria shows "sensitivity to concentrations as low as 3 ligands per cell volume" [56]. In mammalian cells, the number of molecules of a particular mRNA was on the order of tens or hundreds of molecules per cell [47]. Ref. [7] measured the copy numbers of approximately 7300 proteins in a common human tissue culture cell line (U2OS), and found that proteins involved in signaling, cell communication, regulation of cellular processes, catalysis of post-translational modifications, and lipid metabolism tend to be present at copy numbers of less than 500 per cell.

Chapter 7

Applications

Parts of this chapter were previously published in [37] © 2013 Elsevier Inc.

In Section 7.1 we investigate the effects of drift due to several potentials on reaction time and location statistics. In Section 7.2 we present a simplified model of a protein-polymer system, in which two reacting molecules undergo drift-diffusion along a polymer, and may also unbind from the polymer and diffuse in three dimensions. We study the interaction between polymer geometry, binding potentials along the polymer, and unbinding rates.

7.1 Comparison of Potentials

To demonstrate the contrasting effects that can be produced by different drifts, we consider the reaction $A + B \to \varnothing$ where the molecules diffuse within various potential energy landscapes in a one-dimensional domain. We consider the following three cases: (i) zero potential, $V(x) = 0$; (ii) a one-well potential, $V(x) = \cos(2\pi x)$; and (iii) a two-well potential, $V(x) = \cos(4\pi x)$. We use a domain of length $L = 1$ unit with reflecting boundaries, and diffusion coefficient $D = 1$ unit^2/sec for both molecular species A and B.

In the absence of reactions, the equilibrium probability density for a molecule to

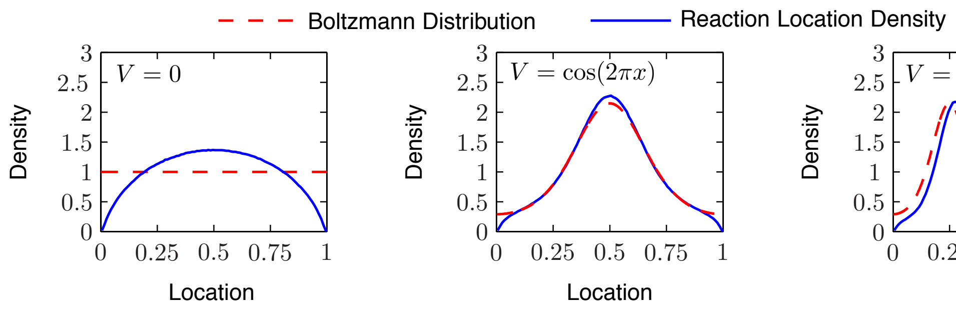

Figure 7.1: Reaction locations from $A + B \to \varnothing$ DL-FPKMC simulations with $M^A(0) = M^B(0) = 1$, and the Gibbs-Boltzmann distribution for each potential. The reaction location is $(Q^A(t) + Q^B(t))/2$, where t is the time of the reaction. Each graph of reaction locations is based on 10^7 simulations. $r_{\mathrm{R}} = 0.02$. $h_s^{\max} = h_p = r_{\mathrm{R}}/8$. $Q^A(0)$ and $Q^B(0) \sim \mathcal{U}(0, 1)$. The plotted densities were determined by binning the reaction locations into 100 bins.

be at location $x \in (0, L)$ is given by the Gibbs-Boltzmann distribution in Eq. (4.1). We compare the Gibbs-Boltzmann distributions for each of the three potentials to the reaction locations from the $A+B \to \varnothing$ DL-FPKMC simulations in the particular case of two molecules, $M^A(0) = M^B(0) = 1$. The results are shown in Fig. 7.1. The potentials serve to spatially "confine" molecules, in the sense that molecules are most likely to be found in locations where the potential energy is lowest. Consequently, the reactions are most likely to occur in such low energy locations.

As would be expected, the mean reaction time for the two-molecule $A + B \to \varnothing$ reaction is faster with the one-well potential (~ 0.03481 sec) than with no potential (~ 0.06483 sec), while slower with the two-well potential (~ 0.09887 sec). Figure 7.2 (left and center panels) compares the survival probabilities $S_{\mathrm{emp}}(t)$ for the three different potentials. The semi-log graphs of $S_{\mathrm{emp}}(t)$ for all three potentials appear linear except at short times, indicating that the reaction time distributions could be well approximated by exponential distributions for larger times (Fig. 7.2, left panel). For $V = 0$, this is as expected since the survival probability distribution is known analytically and is given by the eigenfunction expansion (C.2) in Appendix C. At

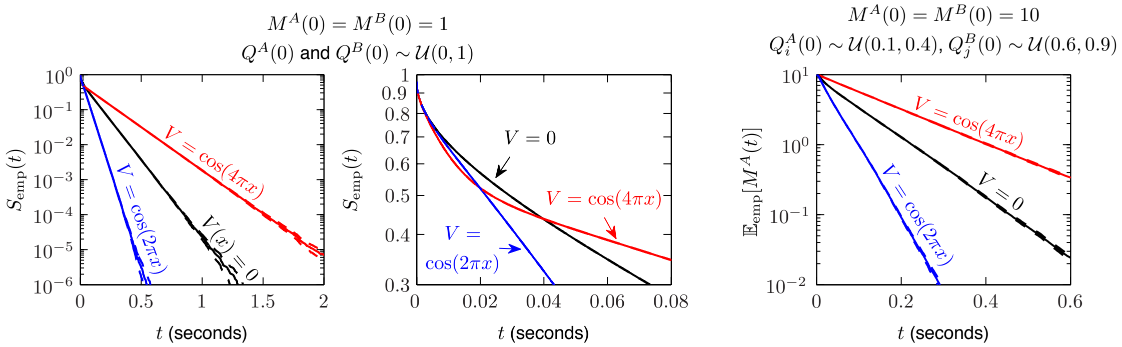

Figure 7.2: *Left and center panels*: Empirical survival probability $S_{\mathrm{emp}}(t)$ for the two-molecule $A + B \to \varnothing$ reaction. Both panels show the same survival probabilities, however the axes have different scales. The dashed lines show 99% confidence bounds based on 10^7 simulations. $r_{\mathrm{R}} = 0.02$. $h_s^{\max} = h_p = r_{\mathrm{R}}/8$. *Right panel*: Empirical mean number of molecules of A remaining at time t, for the reaction $A+B \to \varnothing$ with $M^A(0) = M^B(0) = 10$. The dashed lines show 99% confidence bounds based on 4×10^4 simulations. $r_{\mathrm{R}} = 0.02$. $h_p = r_{\mathrm{R}}/16$. $h_s^{\max} = 8h_p$.

short times (Fig. 7.2, center panel), the graphs of $S_{\mathrm{emp}}(t)$ for the different potentials are not linear and behave differently from each other. Initially, reactions occur more quickly with the two-well potential than with either the one-well potential or no potential. However, after $S_{\mathrm{emp}}(t)$ has decreased to below 40−50%, reactions occur more slowly with the two-well potential than with the other potentials. As would be expected, if the initial locations of the two molecules are in the same well of a potential, then they tend to react more quickly; whereas, if the two molecules start in different energy wells, then the time until they react tends to be longer.

Figure 7.2 (right panel) shows $\mathbb{E}_{\mathrm{emp}}[M^A(t)]$, the mean number of molecules of A remaining at time t, when $M^A(0) = M^B(0) = 10$, $Q_i^A(0) \sim \mathcal{U}(0.1, 0.4)$ and $Q_j^B(0) \sim \mathcal{U}(0.6, 0.9)$. In this case, $\mathbb{E}_{\mathrm{emp}}[M^A(t)]$ could be described by exponential distributions for all three potentials. In the case of $M^A(0) = M^B(0) = 50$ with $Q_i^A(0)$ and $Q_j^B(0) \sim \mathcal{U}(0,1)$, $\mathbb{E}_{\mathrm{emp}}[M^A(t)]$ is not plotted but behaves very similarly to $S_{\mathrm{emp}}(t)$ in the two-molecule case with $Q^A(0)$ and $Q^B(0) \sim \mathcal{U}(0,1)$ (Fig. 7.2, left

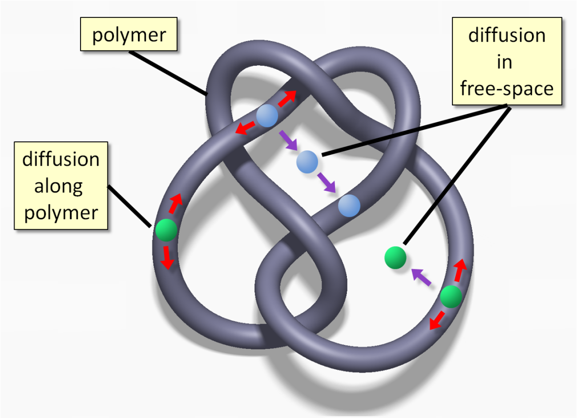

Figure 7.3: Protein-polymer system in which two reacting molecules undergo drift-diffusion along a polymer, and may also unbind from the polymer, diffuse in three dimensions, and rebind to the polymer. The geometry of the polymer affects the times and locations of rebinding. *Figure courtesy of P. Atzberger.*

and center panels).

7.2 Role of Biopolymer Geometry and Binding Potentials in Protein Diffusive Search

The application presented in this subsection is joint work with collaborators Paul Atzberger and Jon Karl Sigurdsson of the University of California at Santa Barbara.

7.2.1 Model of the Biopolymer Drift-Diffusion Process and Three-Dimensional Excursions

Many proteins can diffuse in one dimension by moving along biopolymer filaments, such as actin, microtubules, or DNA [2]. Position-dependent forces along the polymer (e.g., heterogeneity in binding affinity) can be modeled by a potential. In the case of DNA regulatory proteins, single molecule experiments and theoretical work have examined the combination of one-dimensional diffusion along the DNA with three-dimensional diffusion of a protein searching for a DNA binding site [64, 50, 52, 24,

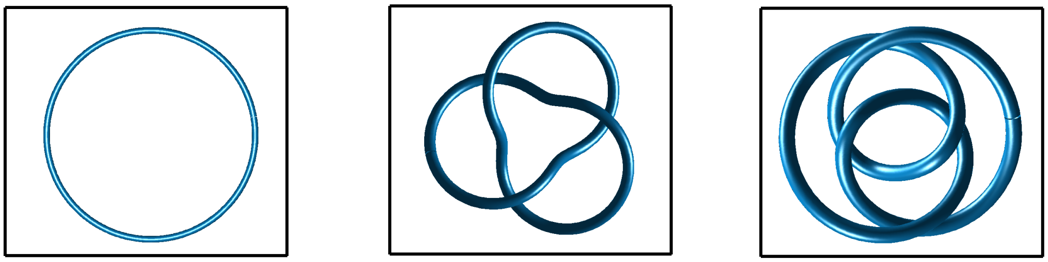

Figure 7.4: Biopolymer conformations: circular loop, trefoil knot, and figure-eight knot. *Figure courtesy of J. K. Sigurdsson.*

53, 39, 26, 9].

We study a model of two proteins undergoing a search process until they encounter each other ($A + B \to \varnothing$) on a biopolymer. In this process the two proteins undergo drift-diffusion along polymers with various geometric conformations (circle, trefoil knot, figure eight) and may also unbind from the polymer, diffuse is three dimensions, and rebind. For example, this process could model the formation of a regulatory complex at a non-specific DNA binding site.

The drift-diffusion process along the biopolymer is simulated using the FPKMC method in one-dimension with periodic boundary conditions. A reaction between the two molecules can only occur when both are on the biopolymer. Detachment times are sampled from an exponential distribution with rate λ_{off}. Rebinding times and locations, which depend on the geometric conformation of the biopolymer, are sampled from results of a diffusion PDE solver provided by our collaborators Jon Karl Sigurdsson and Paul Atzberger. Details of the diffusion solver are given in [37]

We remark that reflecting boundary conditions or potential barriers in our DL-FPKMC method could be used to model obstructions on the biopolymer. Either potential sinks or absorbing Dirichlet boundary conditions could be used to model irreversible binding sites on the polymer.

The studied geometric configurations for the biopolymer are shown in Fig. 7.4.

Table 7.1: Parameters for protein-biopolymer application.

3D Parameters	Value	1D Parameters	Value
3D diffusion coeff.	$2.183823\,\mu\mathrm{m}^2\mathrm{sec}^{-1}$	1D diffusion coeff.	$0.01\,\mu\mathrm{m}^2\mathrm{sec}^{-1}$
Domain Size	$500\,\mathrm{nm} \times 500\,\mathrm{nm} \times$	λ_{off}	0.02 to $200\,\mathrm{sec}^{-1}$
	$500\,\mathrm{nm}$	r_{R}	$20\,\mathrm{nm}$
Polymer Length	$1000\,\mathrm{nm}$	$h_s^{\max}$	$10\,\mathrm{nm}$
Δx in PDE solver	$12.5\,\mathrm{nm}$	h_p	$5\,\mathrm{nm}$
Δt in PDE solver	4292.9	r_{pair}	$4\,r_{\mathrm{R}} = 80\,\mathrm{nm}$

The polymer conformations are parameterized in the form $X(s) = c * \Big(x(s), y(s), z(s)\Big)$ for $0 \leq s \leq 1$, where c is chosen so that the arc length of each of polymer is 1000nm:

Circle: $x(s) = \cos(2\pi s)$, $y(s) = \sin(2\pi s)$, $z(s) = 0$;

Trefoil: $x(s) = (2 + \cos(6\pi s)) \cos(4\pi s)$, $y(s) = (2 + \cos(6\pi s)) \sin(4\pi s)$,
$z(s) = \sin(6\pi s)$; and

Figure-eight: $x(s) = (2 + \cos(4\pi s)) \cos(6\pi s)$, $y(s) = (2 + \cos(4\pi s)) \sin(6\pi s)$,
$z(s) = \sin(8\pi s)$.

We assume one molecule each of protein species A and B are initially present on the biopolymer. The initial positions are drawn from a uniform distribution over the length of the biopolymer. Parameters values are listed in Table 7.1.

7.2.2 Simulation Results: Diffusion-Excursion Search with Different Biopolymer Geometries

For each biopolymer conformation, the reaction location distributions and the mean reaction times are investigated as the detachment rate, λ_{off}, is varied (Fig. 7.5). These simulations are for the case that $V(s) = 0$ along the polymer. As λ_{off} increases, the effects of the trefoil and figure-eight biopolymer geometries on the reaction locations become more pronounced (Fig. 7.5, left and center panels). For the unknotted circle

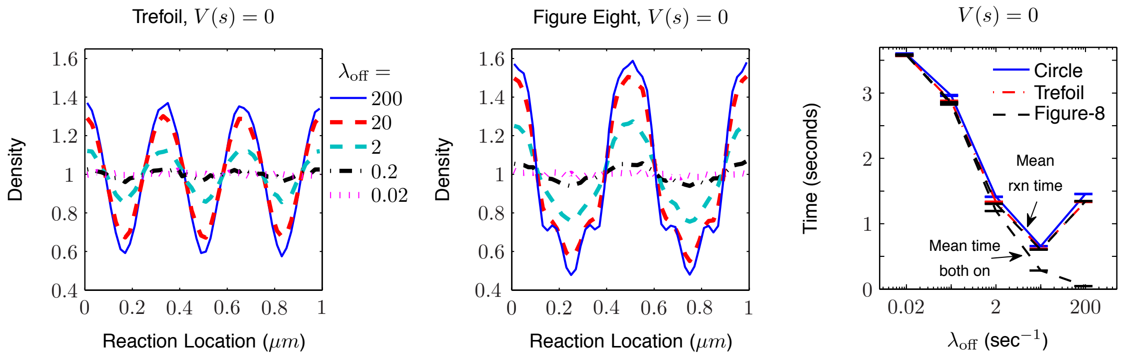

Figure 7.5: *Left and center panels:* Reaction location densities for $A + B \to \varnothing$ with $M^A(0) = M^B(0) = 1$ and $V(s) = 0$. The reaction location is the midpoint between the locations of A and B at the time of the reaction. Each graph is based on 10^6 simulations. The plotted densities were determined by binning the reaction locations into 50 bins. Note that no potentials are imposed in this figure. The non-uniformity in the reaction location density is a result of the polymer geometry. *Right panel:* Mean reaction time and mean time that both molecules are on the biopolymer. Each data point is based on 10^6 simulations. Error bars indicate 99% confidence intervals. The mean time that both molecules are on the biopolymer is shown for the figure-eight conformation only, but the results for the circle and trefoil conformations are similar.

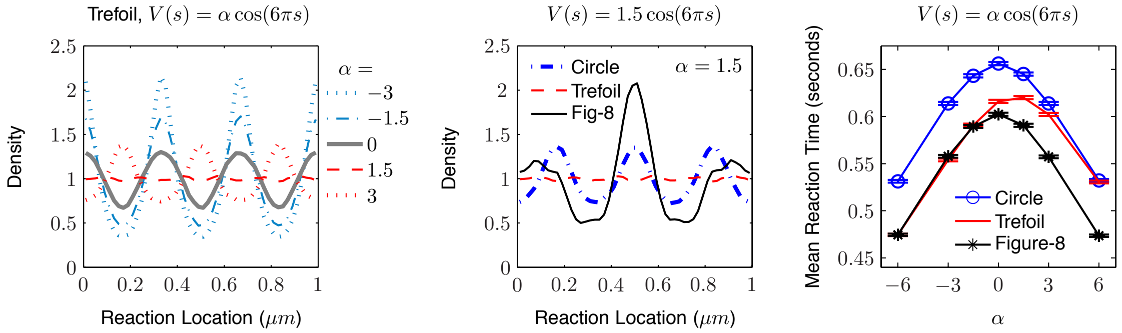

Figure 7.6: *Left and center panels:* Reaction location densities for $A + B \to \varnothing$, with $M^A(0) = M^B(0) = 1$, $V(s) = \alpha \cos(6\pi s)$, and $\lambda_{\text{off}} = 20$ sec^{-1}. Each graph is based on 10^6 simulations. The plotted densities were determined by binning the reaction locations into 50 bins. *Right panel:* Mean reaction times. Each data point is based on 10^6 simulations. Error bars indicate 99% confidence intervals.

the reaction locations are approximately uniform over the biopolymer regardless of the choice of λ_{off}. The mean reaction times attain a minimum at approximately $\lambda_{\text{off}} = 20$ and are slightly faster with the trefoil and figure-eight biopolymer conformations than with the circle conformation (Fig. 7.5, right panel). The mean total time that both molecules are on the biopolymer decreases monotonically as λ_{off} is increased (Fig. 7.5, right panel). Since the proteins spend more time diffusing in the three-dimensional space (off of the polymer) when λ_{off} is large, the polymer geometries affect the re-absorption locations more strongly, which in turn have a greater effect on the reaction locations (Fig. 7.5).

A natural question is how diffusion in a potential energy landscape may enhance or counteract the effects of biopolymer geometry on the reaction locations. From Subsection 7.1, we expect the density of reaction locations to decrease in areas where the potential is large and to increase in areas where the potential is small. To test this idea, we investigated the trefoil knot conformation with potentials of the form $V(s) = \alpha \cos(6\pi s)$ for several values of $\alpha \in [-6, 6]$. The resulting reaction location densities as α is varied are shown in Fig. 7.6 (left panel). As expected, the potential enhances the effects of the trefoil knot geometry when $\alpha < 0$, and counteracts the effects when $\alpha > 0$. When $\alpha = 1.5$, the effects of the potential and the trefoil geometry essentially cancel each other out. For comparison, we also ran simulations with $\alpha = 1.5$ for each of the other two biopolymer geometries (Fig. 7.6, center panel).

When the potential $V(s) = \alpha \cos(6\pi s)$ is used, the mean reaction time is fastest when the amplitude $|\alpha|$ is large (see Fig. 7.6, right panel). As discussed in Subsection 7.1, molecules that are in the same energy well of a potential tend to react more quickly. However, if the proteins can only move on the biopolymer, potentials with large amplitudes as used here would represent large energy barriers, greatly slowing the time for the two molecules to find each other. For example, with $\alpha = -3$ and $\lambda_{\text{off}} = 0$, the mean reaction time is approximately 80 ± 9 seconds, based on 10^3 simulations (compare to Fig. 7.6, right panel). The three-dimensional diffusion excursion provides an alternative path for proteins to circumvent the energy barriers present on the biopolymer. This could be an important mechanism in protein-protein interactions associated with biopolymers. As shown in Fig. 7.6 (right panel), for the circle and figure-eight conformations, the maximum mean reaction time occurs when $\alpha = 0$, and the mean reaction times with positive and negative α of the same magnitude are similar. In contrast, for the trefoil conformation, the maximum mean reaction time occurs when $\alpha = 1.5$, which is the same value of α that resulted in the most uniform distribution of reaction locations (as was shown in Fig. 7.6, left panel). This demonstrates that the effect of the potential is not independent of the polymer geometry.

Appendix A

Constructing Protective Domains (PDs)

This appendix describes our method for constructing protective domains (PDs). It should be noted that this is just one of many possible approaches, and it is still an open question how best to optimize the partitioning of space among PDs. For the reader interested in this question we refer and defer to [17]. Regardless of the method used for constructing the PDs, our lattice approach of Chapter 3 can be used to propagate molecules within their respective PDs.

The PDs should be defined in a such a way that the distance between any two bimolecular reactants in separate PDs will remain strictly greater than their reaction radius, r_{R}, for as long as they remain in their respective PDs. This is necessary to ensure that the movements of molecules in separate PDs are independent. We allow the PDs of non-reacting molecules to overlap.

The following is the approach for defining PDs that resulted in the scaling demonstrated in Section 6. Steps 3 through 5 are specific to one dimension, but could be extended to higher dimensions by following a similar procedure in each coordinate. When constructing PDs for all molecules initially, or when updating the PDs of more than one molecule, we begin with the molecule closest to the left endpoint of the overall domain and then proceed to the right. An alternative approach would be to begin with a seed molecule and proceed outward.

1. For each molecule in need of a protective domain, identify the nearest potential reaction partners. In one dimension, this can by done by keeping a list of all molecules ordered by location. In higher dimensions, the near-neighbor list (NNL) method described in [17] can be used.

2. Determine which molecules will be placed in pair PDs. Two potential reaction partners are placed in a pair if:

    - the two molecules are closer to each other than to any other potential reaction partners, and
    - the distance between the two molecules is less than a pair threshold, r_{pair}, which is a parameter chosen by the user.

    For the reaction system $A + B \to \varnothing$, we also enforce the condition in Eq. (4.13) when determining if two molecules will be placed in a pair. No other PDs are allowed to overlap with a pair PD.

3. For each molecule, identify it's "limiting neighbor," which we define in the following way:

    - For each molecule that will be in a pair PD, the limiting neighbor is the next nearest molecule of any type outside the pair.
    - For a molecule that will be in a single PD, the limiting neighbor is the nearest molecule that is either a potential reaction partner or in a pair with another molecule.
    - Let d_{nbr} be the distance from a molecule to it's limiting neighbor. If a PD for the limiting neighbor has already been defined, let d_{nbrPD} be the

distance from the molecule in question to the nearest endpoint of it's limiting neighbor's PD.

- Note, a molecule is not necessarily the limiting neighbor of it's limiting neighbor.

4. Determine the size of each pair PD and then the size of each single PD as follows, ignoring overall domain boundaries until Step 5:

    - Pair PDs will be symmetric about the midpoint of the two molecules' locations. Single PDs will be symmetric about the location of the molecule.
    - Let r_{PD} for a pair PD be the distance from either molecule to the nearest endpoint of the PD. For a single PD, r_{PD} will denote the distance from the molecule to either endpoint of the PD.
    - Define Condition 1 to be that the limiting neighbor is a potential reaction partner, and Condition 2 to be that a PD for the limiting neighbor has not yet been defined.

      For a single PD, calculate r_{PD} by

      $$r_{\mathrm{PD}} = \begin{cases} (d_{\mathrm{nbr}} - r_{\mathrm{R}})/2 & \text{if Conditions 1 and 2 hold} \\ d_{\mathrm{nbrPD}} - r_{\mathrm{R}} & \text{if only Condition 1 holds} \\ d_{\mathrm{nbr}}/2 & \text{if only Condition 2 holds} \\ d_{\mathrm{nbrPD}} & \text{otherwise.} \end{cases} \tag{A.1}$$

      For a pair PD, calculate the quantity in Eq. (A.1) for each molecule and then set r_{PD} for the pair to be the minimum of the two quantities.
    - We recommend capping the size of pair PDs, so that two molecules will not remain in a pair if they have moved sufficiently far away from each

other. For example, r_{PD} for a pair PD could be set to the minimum of the value calculated above and half the initial distance between the two molecules in the pair.[1]

5. If a PD as defined in Step 4 extends beyond an endpoint of the overall domain, truncate the PD so that one endpoint of the PD will coincide with the overall domain endpoint. Such PDs will no longer be symmetric. An alternative approach would be to treat the overall domain boundary as a neighbor in Steps 1 through 4, and only allow a molecule's PD to touch the domain boundary if the distance to the boundary is less than r_{pair}.

---

[1] The size of single PDs may be capped also. Generally speaking, the size of the PDs should be made as large as possible. However, making a particular molecule's PD as large as possible will result in less space for the PDs of neighbor molecules. Introducing a cap on the size of all PDs may result in a more equitable partitioning of space among them. We have not yet determined what value for the cap, if any, results in the most efficient performance of the overall algorithm. An optimal cap would mostly likely take into account the potential field (in DL-FPKMC), and would therefore vary spatially. See [17] for discussion of the case where different molecules have different diffusion coefficients in FPKMC.

Appendix B

Derivation of Jump Rates for Non-Uniform Mesh Cells

Parts of this appendix were previously published in [37] © 2013 Elsevier Inc.

To derive the non-uniform rates in Eq. (3.6), we use the fluxes from the WPE discretization [63]. Note, as described in Sections 3.3 and 3.4, we only use a non-uniform mesh for mesh cells bordering a boundary, or to move molecules onto a uniform mesh in a newly formed pair protective domain in one-dimension.

Following the derivation approach of [63], we consider the one-dimensional Fokker-Planck equation

$$\frac{\partial \rho(x,t)}{\partial t} = D\frac{\partial}{\partial x}\left(\rho(x,t)\frac{dV(x)}{dx} + \frac{\partial \rho(x,t)}{\partial x}\right). \tag{B.1}$$

As illustrated in Fig. 3.1 (top row, left), let $x_1 < x_0 < x_2$ be the locations of mesh points with non-uniform spacing $h_j = |x_0 - x_j|$. The jump rates a_{0j} from x_0 to x_j, $j = 1, 2$, are derived in this appendix. As discussed in Section 3.1, the solution $\rho(x,t)$ of the Fokker-Planck equation (B.1), gives the probability density of being at location x at time t. Let $\rho^{\text{eq}}(x)$ denote the equilibrium value of $\rho(x,t)$. Define $p_i(t)$ to be the probability of being at the mesh point x_i at time t in the discrete master equation model. We consider the point x_1 to represent the interval $(x_1 - \frac{h_1}{2}, x_1 + \frac{h_1}{2})$ in the sense that

$$p_1(t) \approx \int_{x_1-\frac{h_1}{2}}^{x_1+\frac{h_1}{2}} \rho(x,t)\,dx \approx \rho(x_1,t)h_1. \tag{B.2}$$

Similarly, x_0 represents $(x_0 - \frac{h_1}{2}, x_0 + \frac{h_2}{2})$ and x_2 represents $(x_2 - \frac{h_2}{2}, x_2 + \frac{h_2}{2})$, so

$$p_0(t) \approx \rho(x_0,t)\frac{h_1+h_2}{2} \qquad \text{and} \qquad p_2(t) \approx \rho(x_2,t)h_2. \tag{B.3}$$

Let $J(x,t)$ denote the flux

$$J(x,t) = -D\left(\rho(x,t)\frac{\partial V(x)}{\partial x} + \frac{\partial \rho(x,t)}{\partial x}\right).$$

For convenience, we define

$$A_{ik}(t) = \begin{cases} \frac{V(x_k)-V(x_i)}{\exp[V(x_k)-V(x_i)]-1} & \text{for } V(x_k) \neq V(x_i) \\ \\ 1 & \text{otherwise.} \end{cases}$$

Based on the WPE discretization, we approximate the unidirectional outward flux from x_0 to x_j by

$$J_{0j}(t) = \frac{D}{h_j}A_{0j}(t)\rho(x_0,t),$$

and the unidirectional inward flux from x_j to x_0 by

$$J_{j0}(t) = \frac{D}{h_j}A_{j0}(t)\rho(x_j,t).$$

Then, the net flux from x_0 to x_j is $J_{0j}(t) - J_{j0}(t)$. In the case that x_j, for $j = 1$ or 2, lies on an absorbing Dirichlet boundary, $\rho(x_j,t) = 0$, and so we would have

$J_{j0} = 0$ throughout the following calculation. The Fokker-Planck PDE (3.1) at x_0 is approximated by

$$\begin{aligned}
\frac{\partial \rho(x_0,t)}{\partial t} &= -\frac{\partial}{\partial x} J(x_0,t) \approx \frac{2}{h_1+h_2} \int_{x_0-\frac{h_1}{2}}^{x_0+\frac{h_2}{2}} -\frac{\partial}{\partial x} J(x,t)\,dx \\
&= \frac{2}{h_1+h_2}\left(-J(x_0+\frac{h_2}{2},t) + J(x_0-\frac{h_1}{2},t)\right) \qquad \text{(B.4)} \\
&\approx \frac{2}{h_1+h_2}\Big((J_{20}-J_{02}) - (J_{01}-J_{10})\Big) \\
&= \frac{2}{h_1+h_2}\left(\frac{D}{h_2}\Big(A_{20}\,\rho(x_2,t) - A_{02}\,\rho(x_0,t)\Big) - \frac{D}{h_1}\Big(A_{01}\,\rho(x_0,t) - A_{10}\,\rho(x_1,t)\Big)\right). \qquad \text{(B.5)}
\end{aligned}$$

For a reflecting boundary at $x_0 - \frac{h_1}{2}$ or $x_0 + \frac{h_2}{2}$, the corresponding flux term in Eq. (B.4) is zero (since we assume reflecting boundaries are at the edges of mesh cells). The corresponding terms involving A_{0j} and A_{j0} then drop out of Eq. (B.5).

Multiplying Eq. (B.5) through by $\frac{h_1+h_2}{2}$ and making the substitutions in Eqs. (B.2) and (B.3) yields the master equation

$$\begin{aligned}
\frac{d\,p_0(t)}{dt} = &\left(\frac{D}{h_2^2}A_{20}\,p_2(t) - \frac{2D}{h_2(h_1+h_2)}A_{02}\,p_0(t)\right) - \\
&\left(\frac{2D}{h_1(h_1+h_2)}A_{01}\,p_0(t) - \frac{D}{h_1^2}A_{10}\,p_1(t)\right).
\end{aligned}$$

Thus, we obtain the non-uniform jump rates (3.6)

$$a_{0j} = \frac{2D}{h_j(h_1+h_2)}A_{0j} \qquad j = 1, 2.$$

If $h_1 = h_2 = h$, the WPE discretization is recovered. When one of x_1 or x_2 corresponds to a (non-uniform) Dirichlet boundary point, and all interior mesh cells are uniform, the spatial discretization in Eq. (B.5) in the case that $V(x) = 0$ is known to be second-order accurate for the Poisson equation [22].

Whenever x_j is *not* a Dirichlet boundary point, the jump rate a_{j0} in the opposite direction agrees with the uniform rate in Eq. (3.3) with $h = h_j$. Hence, the system still satisfies a discrete detailed balanced condition at equilibrium (similar to the uniform WPE discretization):

$$a_{0j}\, \rho^{\mathrm{eq}}(x_0) \frac{h_1 + h_2}{2} = a_{j0} \rho^{\mathrm{eq}}(x_j) h_j,$$

since $\rho^{\mathrm{eq}}(x)$ is given by the Gibbs-Boltzmann distribution $\rho^{\mathrm{eq}}(x) \propto \exp[-V(x)]$ and

$$\frac{\rho^{\mathrm{eq}}(x_j)}{\rho^{\mathrm{eq}}(x_0)} = \frac{\exp[-V(x_j)]}{\exp[-V(x_0)]} = \frac{1}{\exp[V(x_j) - V(x_0)]}.$$

Appendix C

Analytic and Numerical Solutions for the Two-Molecule Annihilation Reaction $A + B \to \varnothing$

In this Appendix we discuss the numerical solution of the Fokker-Planck equation and the analytic solution of the diffusion equation to which we compared the DL-FPKMC simulation results in Section 4.1. The 1D reaction-drift-diffusion system of two molecules undergoing the reaction $A + B \to \varnothing$ can be described by Eq. (4.2) on the 2D domain in the left panel of Fig. C.1. The solutions of the Fokker-Planck equation on the two disjoint triangular components of this domain are independent of each other. By symmetry, solving the Fokker-Planck equation on the domain in the left panel of Fig. C.1 can be reduced to solving the same equation on the single triangular domain in the center panel. We have written a PDE solver to solve the Fokker-Planck equation on the triangular domain. The PDE solver uses the rates in Eq. (3.3) from the WPE discretization [63] of the Fokker-Planck equation. This discretization is second-order accurate for smooth potentials and first-order accurate for discontinuous potentials. The mesh in the PDE solver is uniform.

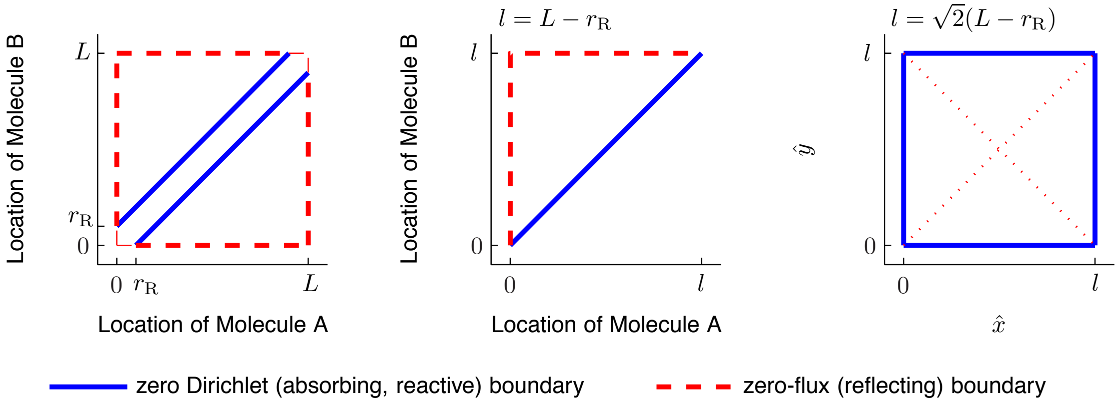

Figure C.1: *Left*: 2D domain equivalent to 1D simulation domain of length L in which the two molecules are located. The zero Dirichlet boundaries on the diagonals correspond to the reaction occurring between the two molecules when they are one reaction radius apart. The zero-flux boundaries on the outer edges correspond to the reflecting endpoints of the 1D simulation domain. *Center*: PDE solver domain. *Right*: Eigenfunction expansion domain.

As discussed in Chapter 4 after Eq. (4.2), when the initial locations of the two molecules in the DL-FPKMC simulations are drawn from $\mathcal{U}(0, L)$, the corresponding initial condition for the 2D Fokker-Planck or diffusion equation on the domain in the left panel of Fig. C.1 is a constant, $\rho_0 = \rho(x, y, 0) = 1/L^2$. In order to recover the same survival probabilities and reaction time statistics, an initial condition of $\rho_0 = 2/L^2$ is used when numerically solving the Fokker-Planck equation on the domain in the center panel of Fig. C.1.

The solution to the diffusion equation on the triangular domain in Fig. C.1 (center) can be recovered by solving on a square domain formed by adjoining four copies of the triangular domain at their reflecting (zero Neumann) edges, as shown in the right panel of Fig. C.1. This domain transformation was suggested by P. Atzberger [5]. On the square domain, an eigenfunction expansion for the solution of the diffusion equation can be determined analytically. The solution on the square domain in the right panel has the property that the normal derivative along the diagonals is zero;

therefore, the solution restricted to any of the four triangular subdomains will agree with the solution of the boundary value problem on the triangular domain in the center panel.

The analytic solution provides a check for the DL-FPKMC simulations in the $V = 0$ case, as well as a check for the PDE solver in that case. An initial condition of $\rho_0 = 1/(2L^2)$ on the square domain (right panel) results in the same survival probabilities and reaction time statistics as an initial condition of $\rho_0 = 2/L^2$ on the triangular domain (center panel). The solution $\rho(x, y, t)$ to the diffusion equation with constant initial condition ρ_0 and diffusion coefficient D on a square domain with sides of length l is given by [45]:

$$\begin{aligned}\rho(x,y,t) = \frac{16\rho_0}{\pi^2} &\left[\sum_{n=0}^{\infty} \frac{1}{2n+1} \sin\left(\frac{(2n+1)\pi x}{l}\right) \exp\left(\frac{-\pi^2(2n+1)^2 Dt}{l^2}\right)\right] \\ \times &\left[\sum_{m=0}^{\infty} \frac{1}{2m+1} \sin\left(\frac{(2m+1)\pi y}{l}\right) \exp\left(\frac{-\pi^2(2m+1)^2 Dt}{l^2}\right)\right].\end{aligned} \tag{C.1}$$

The survival probability $S(t)$, the probability that the two particles have not yet reacted by time t, is

$$S(t) = \int_0^l \int_0^l \rho(x,y,t)\,dx\,dy = \frac{64\,\rho_0\,l^2}{\pi^4} \left[\sum_{n=0}^{\infty} \left(\frac{1}{2n+1}\right)^2 \exp\left(\frac{-\pi^2(2n+1)^2 Dt}{l^2}\right)\right]^2. \tag{C.2}$$

The mean reaction time is

$$\begin{aligned}\mathbb{E}\left[T\right] &= -\int_0^{\infty} t\,S'(t)\,dt = \int_0^{\infty} S(t)\,dt \\ &= \frac{64\,\rho_0\,l^4}{D\,\pi^6} \sum_{n=0}^{\infty}\sum_{m=0}^{\infty} \left(\frac{1}{2n+1}\right)^2 \left(\frac{1}{2m+1}\right)^2 \frac{1}{(2n+1)^2+(2m+1)^2}\end{aligned}$$

$$= \frac{16\,\rho_0\,l^4}{D\,\pi^5} \sum_{n=0}^{\infty} \left(\frac{1}{2n+1}\right)^4 \left(\frac{\pi}{2} - \frac{\tanh(\frac{\pi}{2}(2n+1))}{(2n+1)}\right). \tag{C.3}$$

Evaluating $\mathbb{E}\left[T\right]$ at the parameter values corresponding to the two-molecule DL-FPKMC simulations in Section 4.1 gives a mean reaction time of 0.064831881311 seconds.

Next, we derive an analytic expression for the distribution of reaction locations along the reactive boundary (the diagonal boundaries in the left or center panel of Fig. C.1). As defined in Section 4.1.4, $Q^{\mathrm{rxn}} = (Q^A(t) + Q^B(t))/2$ is considered to be the location of a reaction between molecule A and molecule B at time t. Each side of the square domain in the right panel of Fig. C.1 is a copy of the reactive boundary. Let $\boldsymbol{\eta} = \boldsymbol{\eta}(x, y)$ be the outward pointing normal at the point (x, y). The value of the density of reaction locations at any point on the reactive boundary can be obtained by integrating the outward probability flux over time

$$\int_0^{\infty} -D \frac{\partial}{\partial \boldsymbol{\eta}} \rho(x, y, t)\, dt,$$

where ρ is given by Eq. (C.1) Since the outward flux along each side is equivalent, the reaction location distribution can be calculated using any one of the four sides of the square domain in Fig. C.1, right panel. Using the boundary where $y = 0$, the reaction location distribution is

$$F^{\mathrm{rxn}}(x) = \Pr\left[Q^{\mathrm{rxn}} \le x\right] = 4 \int_0^x \int_0^{\infty} \left[-D \frac{\partial}{\partial \boldsymbol{\eta}} \rho(\xi, y, t)\right]_{y=0} dt\, d\xi \tag{C.4}$$

$$= 4\,D \int_0^x \int_0^{\infty} \left[\frac{\partial}{\partial y} \rho(\xi, y, t)\right]_{y=0} dt\, d\xi$$

$$= \frac{64\,\rho_0\,l^2}{\pi^4} \sum_{n=0}^{\infty} \sum_{m=0}^{\infty} \left(\frac{1}{2n+1}\right)^2 \frac{1}{(2n+1)^2 + (2m+1)^2} \left(1 - \cos\left(\frac{(2n+1)\pi x}{l}\right)\right)$$

$$= \frac{16\,\rho_0\,l^2}{\pi^3} \sum_{n=0}^{\infty} \left(\frac{1}{2n+1}\right)^3 \tanh\left(\frac{\pi(2n+1)}{2}\right) \left(1 - \cos\left(\frac{(2n+1)\pi x}{l}\right)\right).$$

The factor of 4 in the first line is included to normalize the distribution, assuming the initial condition ρ_0 on the square domain (right panel of Fig. C.1) has been chosen to equal $1/(2L^2)$ as described above.

For each of the potential functions, we ran the PDE solver using the Crank–Nicolson method in time and the WPE discretization in space, with spatial step sizes ranging from $\Delta x = r_\mathrm{R} = 0.02$ to $\Delta x = r_\mathrm{R}/16$ and time steps $\Delta t = \Delta x/16$. For $V = 0$, we can check the results of the PDE solver against the analytic solution; in this case, the numerical mean reaction times determined using the Crank–Nicolson method converge at approximately second-order ($m \approx 2.0025$) to the analytic mean reaction time. For $V \neq 0$, there is no analytic solution to which the numerical solutions can be compared; however, the decrease in the pairwise differences between the mean reaction times determined from the numerical PDE solutions at successive step sizes indicates convergence ($m \approx 1.9968$ for V_cos and $m \approx 0.9844$ for V_step). In the DL-FPKMC convergence results in Section 4.1, the empirical mean reaction times are compared to the analytic mean reaction time in the case $V = 0$, and to the numerical mean reaction times determined using the Crank–Nicolson PDE solver with the finest spatial step size, $\Delta x = r_\mathrm{R}/16$, in the cases of V_cos and V_step.

The survival probabilities calculated using the Crank–Nicolson method show small numerical oscillations during the first few time steps. This is due to the incompatibility of the initial condition with the zero Dirichlet boundary condition. In order to numerically resolve the survival probabilities more accurately at short times, we re-ran the PDE solver using the Twizell–Gumel–Arigu (TGA) method [60] from $t = 0$ to $t = 0.07$ seconds. The TGA method is a second order, L_0 stable time

Table C.1: Improved errors in numerical survival probabilities.

	Crank–Nicolson $\Delta x = r_\mathrm{R}/16$, $\Delta t = \Delta x/16$	TGA $\Delta x = r_\mathrm{R}/16$, $\Delta t = (\Delta x)^2$
L^1 Error	4.1117e-6	4.7909e-8
L^2 Error	8.2513e-5	1.6171e-6
L^∞ Error	5.8634e-3	2.2463e-4

discretization. We also used a finer time step, $\Delta t = (\Delta x)^2$ where $\Delta x = r_\mathrm{R}/16$, to further improve the accuracy at short times when the survival probabilities change most rapidly. We then determined the numerical survival probabilities using the results from the TGA method for $t = 0$ to $t = 0.07$ seconds and using the results from Crank–Nicolson method for $t > 0.07$ seconds. At $t = 0.07$, the absolute difference in survival probabilities between the two methods is less than 10^{-8} for V_cos and less than 10^{-7} for V_step.

To check that using the TGA method with a finer time step improved the accuracy of the numerical survival probabilities over the Crank–Nicolson method, we compared the numerical survival probabilities for $V = 0$ to the analytic survival probability on the interval from $t = 0$ to $t = 0.07$. Table C.1 shows the absolute errors between the numerical survival probability calculated using each method and the analytic survival probability.

Bibliography

[1] I.C. Agbanusi and S.A. Isaacson. A comparison of bimolecular reaction models for stochastic reaction-diffusion systems. *Bulletin of Mathematical Biology*, 76(4):922–946, 2014. http://dx.doi.org/10.1007/s11538-013-9833-6.

[2] B. Alberts, A. Johnson, J. Lewis, M. Raff, K. Roberts, and P. Walter. *Molecular Biology of the Cell*. Garland Science, New York, 5th edition, 2007.

[3] S. S. Andrews and D. Bray. Stochastic simulation of chemical reactions with spatial resolution and single molecule detail. *Physical Biology*, 1:137–151, 2004.

[4] A. M. Arias and P. Hayward. Filtering transcriptional noise during development: concepts and mechanisms. *Nature Reviews Genetics*, 7(1):34–44, 2006.

[5] P. J. Atzberger. Personal communication, 2012.

[6] P. J. Atzberger and C. S. Peskin. A Brownian dynamics model of kinesin in three dimensions incorporating the force-extension profile of the coiled-coil cargo tether. *Bulletin of Mathematical Biology*, 68(1):131–160, 2006.

[7] M. Beck, A. Schmidt, J. Malmstroem, M. Claassen, A. Ori, A. Szymborska, F. Herzog, O. Rinner, J. Ellenberg, and R. Aebersold. The quantitative proteome of a human cell line. *Molecular Systems Biology*, 7:1–8, November 2011.

[8] H. C. Berg. *Random Walks in Biology*. Princeton University Press, Princeton, NJ, 1993.

[9] O. G. Berg, R. B. Winter, and P. H. Von Hippel. Diffusion-driven mechanisms of protein translocation on nucleic acids. 1. models and theory. *Biochemistry*, 20(24):6929–6948, November 1981.

[10] W. J. Blake, M. Kaern, C. R. Cantor, and J. J. Collins. Noise in eukaryotic gene expression. *Nature*, 422:633–637, April 2003.

[11] A. B. Bortz, M. H. Kalos, and J. L. Lebowitz. A new algorithm for Monte Carlo simulation of Ising spin systems. *Journal of Computational Physics*, 17(1):10–18, 1975.

[12] Y. Cao and L. Petzold. Accuracy limitations and the measurement of errors in the stochastic simulation of chemically reacting systems. *Journal of Computational Physics*, 212(1):6–24, February 2006.

[13] M. Deaconu and A. Lejay. A random walk on rectangles algorithm. *Methodology and Computing in Applied Probability*, 8(1):135–151, Mar 2006.

[14] M. Dobrzynski, J. V. Rodriguez, J. A. Kaandorp, and J. G. Blom. Computational methods for diffusion-influenced biochemical reactions. *Bioinformatics*, 23(15):1969–1977, May 2007.

[15] M. Doi. Second quantization representation for classical many-particle system. *Journal of Physics A: Mathematical and General*, 9(9):1465–1477, 1976.

[16] M. Doi. Stochastic theory of diffusion-controlled reaction. *Journal of Physics A: Mathematical and General*, 9(9):1479–1495, 1976.

[17] A. Donev, V. V. Bulatov, T. Oppelstrup, G. H. Gilmer, B. Sadigh, and M. H. Kalos. A first-passage kinetic Monte Carlo algorithm for complex diffusion–reaction systems. *Journal of Computational Physics*, 229(9):3214–3236, January 2010.

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

[19] D. Fange, O. G. Berg, P. Sjöberg, and J. Elf. Stochastic reaction-diffusion kinetics in the microscopic limit. *Proceedings of the National Academy of Sciences of the United States of America*, 107(46):19820–19825, November 2010.

[20] C. W. Gardiner. *Handbook of Stochastic Methods: For Physics, Chemistry, and the Natural Sciences*, volume 13 of *Springer Series in Synergetics*. Springer Verlag, New York, 2nd edition, 1996.

[21] C. W. Gardiner, K. J. McNeil, D. F. Walls, and I. S. Matheson. Correlations in stochastic theories of chemical reactions. *Journal of Statistical Physics*, 14:307–331, 1976.

[22] F. Gibou, R. P. Fedkiw, L. Cheng, and M. Kang. A second-order-accurate symmetric discretization of the Poisson equation on irregular domains. *Journal of Computational Physics*, 176:205–227, 2002.

[23] D. T. Gillespie. Exact stochastic simulation of coupled chemical-reactions. *Journal of Physical Chemistry*, 81(25):2340–2361, 1977.

[24] J. Gorman and E. C Greene. Visualizing one-dimensional diffusion of proteins along DNA. *Nature Structural & Molecular Biology*, 15(8):768–774, August 2008.

[25] S. Hellander, A. Hellander, and L. Petzold. Reaction-diffusion master equation in the microscopic limit. *Physical Review E*, 85(4):042901(1–5), April 2012.

[26] T. Hu, A. Y. Grosberg, and B. I. Shklovskii. How proteins search for their specific sites on DNA: The role of DNA conformation. *Biophysical Journal*, 90(8):2731–2744, April 2006.

[27] S. A. Isaacson. Relationship between the reaction-diffusion master equation and particle tracking models. *Journal of Physics A: Mathematical and Theoretical*, 41(6):065003 (15pp), 2008.

[28] S. A. Isaacson. The reaction-diffusion master equation as an asymptotic approximation of diffusion to a small target. *SIAM Journal on Applied Mathematics*, 70(1):77–111, 2009.

[29] S. A. Isaacson. A convergent reaction-diffusion master equation. *Journal of Chemical Physics*, 139(5):054101, 2013.

[30] S. A. Isaacson and D. Isaacson. Reaction-diffusion master equation, diffusion-limited reactions, and singular potentials. *Physical Review E*, 80(6):066106 (9pp), 2009.

[31] 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 of the United States of America*, 108(9):3815–3820, March 2011.

[32] J. Keizer. Nonequilibrium statistical thermodynamics and the effect of diffusion on chemical reaction rates. *Journal of Physical Chemistry*, 86:5052–5067, 1982.

[33] B. N. Kholodenko, J. F. Hancock, and W. Kolch. Signalling ballet in space and time. *Nature Reviews Molecular Cell Biology*, 11:414–426, June 2010.

[34] S. Kullback and R. A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22(1):79–86, 1951.

[35] C. Loverdo, O. Benichou, M. Moreau, and R. Voituriez. Enhanced reaction kinetics in biological cells. *Nature Physics*, 4:134–137, 2008.

[36] H. Maamar, A. Raj, and D. Dubnau. Noise in gene expression determines cell fate in Bacillus Subtilis. *Science*, 317:526–529, July 2007.

[37] A. J. Mauro, J. K. Sigurdsson, J. Shrake, P. J. Atzberger, and S. A. Isaacson. A first-passage kinetic Monte Carlo method for reaction-drift-diffusion processes. *Journal of Computational Physics*, 259:536–567, 2014. DOI: `http://dx.doi.org/10.1016/j.jcp.2013.12.023`, Preprint available at `http://arxiv.org/abs/1302.0793` [math.NA] and at `http://math.bu.edu/people/avamauro` (2013) 38 pp.

[38] D. A. McQuarrie. Stochastic approach to chemical kinetics. *Journal of Applied Probability*, 4:413–478, 1967.

[39] L. Mirny, M. Slutsky, Z. Wunderlich, A. Tafvizi, J. Leith, and A. Kosmrlj. How a protein searches for its site on DNA: the mechanism of facilitated diffusion. *Journal of Physics A: Mathematical and Theoretical*, 42(43):434013, Jan 2009.

[40] M. E. Muller. Some continuous Monte Carlo methods for the Dirichlet problem. *Annals of Mathematical Statistics*, 27:569–589, 1956.

[41] J. Muñoz-García, Z. Neufeld, B. N. Kholodenko, and H. M. Sauro. Positional information generated by spatially distributed signaling cascades. *PLoS Computational Biology*, 5(3):e1000330, March 2009.

[42] T. Oppelstrup, V. V. Bulatov, A. Donev, M. H. Kalos, G. H. Gilmer, and B. Sadigh. First-passage kinetic Monte Carlo method. *Physical Review E*, 80(6):066701, December 2009.

[43] T. Opplestrup, V. V. Bulatov, G. H. Gilmer, M. H. Kalos, and B. Sadigh. First-passage Monte Carlo algorithm: Diffusion without all the hops. *Physical Review Letters*, 97(23):230602, Dec 2006.

[44] C. S. Peskin and G. Oster. Coordinated hydrolysis explains the mechanical behavior of kinesin. *Biophysical Journal*, 68(4 Suppl):202S, 1995.

[45] A. D. Polyanin. *Handbook of Linear Partial Differential Equations for Engineers and Scientists*. Chapman & Hall/CRC, 2002.

[46] R. M. Price and D. G. Bonett. Estimating the variance of the sample median. *Journal of Statistical Computation and Simulation*, 68(3):295–305, 2001.

[47] A. Raj, C. S. Peskin, D. Tranchina, D. Y. Vargas, and S. Tyagi. Stochastic mRNA synthesis in mammalian cells. *PLoS Biology*, 4(10):e309, 2006.

[48] R. Ramaswamy and I. F. Sbalzarini. Exact on-lattice stochastic reaction-diffusion simulations using partial-propensity methods. *Journal of Chemical Physics*, 135:244103 (17pp), 2011.

[49] J. Raser and E. O'Shea. Control of stochasticity in eukaryotic gene expression. *Science*, 304:1811–1814, Jan 2004.

[50] A. D. Riggs, S. Bourgeois, and M. Cohn. The lac represser-operator interaction: Iii. kinetic studies. *Journal of Molecular Biology*, 53(3):401–417, November 1970.

[51] K. Schwarz and H. Rieger. Efficient kinetic Monte Carlo method for reaction-diffusion problems with spatially varying annihilation rates. *Journal of Computational Physics*, 237:396–410, 2013.

[52] N. Shimamoto. One-dimensional diffusion of proteins along dna. *Journal of Biological Chemistry*, 274(22):15293–15296, May 1999.

[53] M. Slutsky and L. A. Mirny. Kinetics of protein-DNA interaction: facilitated target location in sequence-dependent potential. *Biophysical Journal*, 87(6):4021–35, Dec 2004.

[54] M. V. Smoluchowski. Mathematical theory of the kinetics of the coagulation of colloidal solutions. *Zeitschrift für Physikalische Chemie*, 92:129–168, 1917.

[55] T. Sokolowski, L. Bossen, T. Miedema, and N. Becker. Green's function reaction dynamics–an exact and efficient way to simulate intracellular pattern formation. In T. E. Simos, G. Psihoyios, and Ch. Tsitouras, editors, *ICNAAM, Numerical Analysis and Applied Mathematics, International Conference 2010, Vol. II*, volume 1281 of *AIP Conference Proceedings*, pages 1342–1345, 2010.

[56] V. Sourjik and N. S. Wingreen. Responding to chemical gradients: bacterial chemotaxis. *Current Opinion in Cell Biology*, 24(2):262–268, April 2012.

[57] G. M. Süel, J. G. Garcia-Ojalvo, L. M. Liberman, and M. B. Elowitz. An excitable gene regulatory circuit induces transient cellular differentiation. *Nature*, 440:545–550, March 2005.

[58] K. Takahashi, S. Tanase-Nicola, and P. R. ten Wolde. Spatio-temporal correlations can drastically change the response of a MAPK pathway. *Proceedings of the National Academy of Sciences of the United States of America*, 107(6):2473–2478, February 2010.

[59] E. Teramoto and N. Shigesada. Theory of bimolecular reaction processes in liquids. *Progress of Theoretical Physics*, 37(1):29–51, 1967.

[60] E. H. Twizell, A. B. Gumel, and M. A. Arigu. Second-order, L_0-stable methods for the heat equation with time-dependent boundary conditions. *Advances in Computational Mathematics*, 6(3):333–352, 1996.

[61] N. G. Van Kampen. *Stochastic Processes in Physics and Chemistry.* North-Holland, Amsterdam, 2001.

[62] 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. *Physical Review Letters*, 94(12):128103, 2005.

[63] H. Wang, C. S. Peskin, and T. C. Elston. A robust numerical algorithm for studying biomolecular transport processes. *Journal of Theoretical Biology*, 221:491–511, 2003.

[64] Y. M. Wang, R. H. Austin, and E. C. Cox. Single molecule measurements of repressor protein 1D diffusion on DNA. *Physical Review Letters*, 97(4):048302, July 2006.

Curriculum Vitae

[redacted]