

UNIVERSITY OF CALIFORNIA,
IRVINE

Rare Events in Spatial Stochastic Simulations of T Cell Receptor Triggering

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Physics

by

Robert Taylor

Dissertation Committee:
Professor Jun Allard, Chair
Professor Elizabeth Read
Professor Jin Yu
Professor Steve Gross

2022

DEDICATION

To my family, who helped me stay grounded.

TABLE OF CONTENTS

	Page
LIST OF FIGURES	v
LIST OF TABLES	vi
ACKNOWLEDGMENTS	vii
VITA	viii
ABSTRACT OF THE DISSERTATION	x
1 Introduction	1
1.1 T Cell Receptor Triggering	2
1.2 Spatial Stochastic Processes and How to Computationally Study Them	3
1.3 Introduction to Weighted Ensemble Sampling	6
1.4 Summary of Research	7
2 Surface Dynamics of CD45 Evacuation	8
2.1 Introduction	8
2.2 Results	11
2.2.1 A Rare-Event Reaction-Diffusion Simulation	11
2.2.2 Oligomerization	15
2.2.3 Close Contacts	18
2.3 Discussion	21

2.4	Methods	24
2.4.1	Model and Dynamics	24
2.4.2	Weighted Ensemble	28
2.4.3	Implementation	31
2.4.4	Analysis	31
2.4.5	Method validation	32
2.5	Supplemental material	39
3	Quantifying Uncertainty in Weighted Ensemble Simulations and Optimizing Weighted-Ensemble Specific Parameters	43
3.1	Introduction	43
3.1.1	Error Quantification	43
3.1.2	Investigating quantitative differences between Weighted Ensemble meta-parameters.	45
3.2	Measuring and Understanding Uncertainty in WE	46
3.2.1	The Search for Key Metrics in Weighted Ensemble	47
3.2.2	The Sawtooth Effect	51
3.2.3	Metaparameter Dependence	53
3.3	Discussion	54
4	Discussion	58
4.0.1	Impact of oligomerization and close contacts on receptor triggering	59
4.0.2	Uncertainty quantification and metaparameter selection in Weighted Ensemble rare event simulation	62
	Bibliography	64
	Appendix A Analytic derivation of $p_{\text{entry}} = 0$ FPT distribution	72

LIST OF FIGURES

	Page
1.1 Artistic depiction of large surface molecule evacuation	4
2.1 Features of surface dynamics	10
2.2 Weighted Ensemble Algorithm	34
2.3 Diffusion-Only Evacuation and Computation Times	35
2.4 Impact of Dimerization on TCR Evacuation	36
2.5 Higher-order Oligomers Impact on Evacuation Times	37
2.6 Surface Dynamics of Close Contacts	38
2.7 Weighted Ensemble Re-weighting Algorithms	39
2.8 Domain Impact on Cell-Surface Evacuation	40
2.9 2D Reaction-Diffusion and Rate Constants	41
3.1 Methods of measuring output and uncertainty from WE simulations.	48
3.2 Creation of a modified Kolmogorov-Smirnov statistic	49
3.3 Candidate measures of flux measurement of WE simulations and simulation accuracy	50
3.4 Emergent properties of WE bin weights	52
3.5 Impact of m_{targ} on WE traces	53
3.6 Influence of m_{targ} on MFPT estimates	55
3.7 WE Metaparameter Accuracy Heatmap	56

LIST OF TABLES

	Page
2.1 Surface Dynamics Variables, Parameters, and Hyperparameters	42

ACKNOWLEDGMENTS

I would like to thank thank Omer Dushek (Oxford), Jay Newby (U Alberta), Brian Chu (UC Irvine), and Dhiman Ray (UC Irvine) for valuable discussion. This work was supported by: NSF grant DMS 1715455 to ELR; NSF CAREER grant DMS 1454739 to JA; NSF grant DMS 1763272; and a grant from the Simons Foundation (594598, QN).

VITA

Robert Taylor

EDUCATION

Doctor of Philosophy in Physics University of California, Irvine	2022 <i>Irvine, CA</i>
Master of Science in Physics Case Western Reserve University	2015 <i>Cleveland, OH</i>
Bachelor of Science in Mathematics and Physics Case Western Reserve University	2015 <i>Cleveland, OH</i>

RESEARCH EXPERIENCE

Graduate Research Assistant University of California, Irvine	2017–2022 <i>Irvine, California</i>
--	---

TEACHING EXPERIENCE

Teaching Assistant University of California, Irvine	2016–2017, 2021 <i>Irvine, California</i>
---	---

ARTICLES

Simulation of receptor triggering by kinetic segregation shows role of oligomers and close-contacts https://doi.org/10.1101/2021.09.29.462451	2021
Locally-correlated kinetics of post-replication DNA methylation reveals processivity and region-specificity in DNA methylation maintenance https://doi.org/10.1101/2021.09.28.462223	2021

CONFERENCE PUBLICATIONS

- Weighted Ensemble Simulation of Kinetic Segregation Shows Role of Oligomers and Close Contacts** **Feb 2022**
Biophysical Society Annual Meeting (Talk)
- Immune Cell Triggering by Spatial Segregation Studied using Stochastic Rare Event Simulation** **Feb 2020**
Biophysical Society Annual Meeting (Poster)
- Immune Cell Triggering by Spatial Segregation Studied using Stochastic Rare Event Simulation** **Oct 2019**
2nd Annual Symposium on Multiscale Cell Fate Research
- Weighted Ensemble Methods Applied to Cell Surface Evacuation** **Oct 2018**
Inaugural Symposium on Multiscale Cell Fate (Poster)

SOFTWARE

LibsmolWE <http://https://github.com/dydtaylor/LibsmolWE>
C code combining biochemical simulator Smoldyn with rare event sampling method Weighted Ensemble

ABSTRACT OF THE DISSERTATION

Rare Events in Spatial Stochastic Simulations of T Cell Receptor Triggering

By

Robert Taylor

Doctor of Philosophy in Physics

University of California, Irvine, 2022

Professor Jun Allard, Chair

On the surface of T Cells, large transmembrane molecules such as CD45 are distributed across the surface. However, during engagement with antigen-presenting cells, these proteins are noticeably absent from the region surrounding the T-Cell receptors (TCR) involved. Computer simulations can aid in studying these processes. However, simulations of spatiotemporal dynamics of protein reorganizations on cell surfaces can be computationally expensive when they involve rare events – that is, when processes of interest are slow relative to other relevant timescales in the system. To better understand the process of TCR triggering, we build a biochemical simulator that implements the rare event sampling method Weighted Ensemble together with particle-based reaction diffusion simulation. We investigate how two simple processes, oligomerization and formation of close-contacts, can impact these otherwise rare events. While working with Weighted Ensemble, we further investigated the properties of the outputs from rare event simulations. We investigated how best to quantify error from Weighted Ensemble simulations, developed metrics toward prediction of accuracy, and tested Weighted Ensemble specific metaparameters’ impact on simulation output. This work finds that the processes of oligomerization and close-contact formation can both have drastic impacts on molecular evacuation rates, but even the reduced evacuation times are still unrealistically long, suggesting that a yet-to-be-described mechanism drives evacuation. Furthermore, we reveal key characteristics of Weighted Ensemble simulation traces and suggest future methods of exploring these characteristics.

Chapter 1

Introduction

The world of cell biology is a noisy, crowded, chaotic place. Typical cell surface densities of membrane-bound proteins on the surface of the cell can range from 500-3000 molecules / μm^2 [65]. In such crowded environments, however, some of the most interesting phenomena involve instances of “uncrowding”, or rapid depletion of certain protein species from specific regions. For instance, when T Cell receptors (TCRs) are triggered in the process of receptor-antigen binding, large surface molecules such as CD45 rapidly deplete the region surrounding the TCR. While many specific details of the triggering are not known, it is known that this evacuation of CD45 is necessary for the triggering to occur. A basic model for this process involves measuring evacuation times, or the mean first passage time (MFPT), for a large number of freely-diffusing molecules, but this lowest order model shows discrepancies between the likelihood of seeing an evacuation event and the time-scales observed for TCR triggering. For instance, if one assumes there are n independent molecules uniformly distributed in the region of interest (ROI) surrounding the TCR, the probability of simply observing an evacuated region of interest surrounding the TCR can be approximated $p = (1 - A_{\text{ROI}}/A_{\text{domain}})^n \sim 10^{-10}$ for some of the lower density estimates for CD45, with $A_{\text{ROI}}, A_{\text{domain}}$ giving the area of the ROI and area of the 2D diffusion domain, respectively (see Model parameters for specifics regarding the quantities A_{ROI} and the density). While the

mean time to reach an evacuated state does not have a relationship to this probability that can be immediately stated, the exponential decay of this probability with the number of evacuating molecules reveals how in this very basic model, evacuation is a rare event. Indeed, preliminary estimations for this MFPT yielded values on the order of $\sim 10^{10}$ s [74], drastically longer than the actual timescales physiologically observed for this triggering, ~ 1 s. Creating a theoretical solution for these MFPTs including only diffusion is complicated as-is, and extending these theoretical solutions to a wider array of physiological scenarios and interactions that CD45 might undergo on the cell surface is a gargantuan task.

Rather than developing a new model for cell-surface based evacuations, it became evident that much progress could be made by instead combining biochemical simulators with rare event sampling methods, using them to better understand what sort of impact various biochemical and mechanical properties might have on this evacuation time. Thus, to bridge the gap between these theoretical diffusion-only based values and the observed physiological values, we chose to combine biochemical simulation with rare-event sampling methods.

T Cell Receptor Triggering

In the immune system, a key step involved in T cells sensing pathogens is a process called T cell receptor triggering. In this process, a T Cell receptor (TCR) binds to an antigen presenting cell (APC) and initiates an immunological response. While many specific details of the triggering are not known, there are three prevailing models used to describe it: aggregation, conformational changes, and segregation. Aggregation and conformational changes both describe changes in the TCR complex after binding to the APC, and both remain outside the scope of this work. Kinetic segregation, the focus of this study, describes molecular movement prior to TCR triggering [93].

The cell surface is a dynamic environment, and TCR triggering is associated with reorganization

of cell surface proteins. Kinetic segregation refers to a model where large transmembrane proteins (e.g. CD45) that are too close to the TCR block receptor-antigen binding. In this model, only after all of these large molecules move outside (evacuate) a contact region surrounding the TCR can binding occur (see Fig 1.1A). CD45 has been shown to have an inhibitory effect on TCR triggering [29, 80]. CD45 is found in high density close to TCRs before a TCR is triggered as well as much in lower concentrations close to triggered TCRs [20, 29]. This behavior is consistent with the kinetic segregation model, as well as with the known geometry of the rigid extracellular domain of CD45 (~ 21 nm) [21, 20] and the proximity of triggered T cells to APCs, (~ 13 nm) [25, 20].

Simple, passive diffusion, i.e. Brownian motion, of these large surface molecules is not sufficient to explain these evacuations. 2D diffusion-only simulations find complete evacuation timescales to be many orders of magnitude larger than the second-long timescale of TCR triggering [74]. However, CD45 diffusion across the membrane is not simple Brownian motion, as it might interact chemically with other surface molecules, dimerize [92], and experiences diffusional barriers such as actin filaments [10, 41]. Thus, we pose the question: what role do these more complicated reaction-diffusion features have on evacuation times, and are they sufficient to explain the gap in timescales?

Spatial Stochastic Processes and How to Computationally Study Them

Cells are highly organized and complex objects with millions of molecules. A complete description of the cell, while desirable, is computationally expensive and not feasible at present. Relevant length scales can vary from 10^{-10} to 10^{-5} m [59], and time scales as low as 10^{-12} s [59] up to several hours or longer. Depending on the cellular process of interest and their relevant length and timescales, there are multiple modeling methods to select between. On an atomic and molecular




Figure 1.1: (A) Artistic rendition of kinetic segregation on the T Cell surface. When in the orange region of the membrane surrounding the TCR (circular sector on lower membrane), large surface molecules block receptor antigen binding. Once they leave this region, the antigen can get close enough for receptor-antigen binding to occur. Rendition taken from [74]. (B) 2D representation of CD45 evacuation used for the base model. The TCR is presumed to be at the center of a square domain, side length L . The circular region of radius R is the region of the domain that must be completely devoid of any CD45 molecules (black heptagons) before binding can occur. In the simplest model, these CD45 molecules are independent Brownian particles that only interact with the domain boundaries.

level, dynamics can be investigated in detail with molecular dynamics simulations. On larger length and timescales, deterministic reaction kinetics can describe species evolution without spatial dynamics, but that is valid when both the number of particles is large and chemical species are well-mixed. To bridge the gap between these levels, many simulations function on a mesoscale level: molecules are treated as individual particles with motion described by Langevin dynamics, or Brownian dynamics if inertia can be neglected [59]. Reactions are taken to occur when the appropriate species come into physical contact, based on a framework developed by Smoluchowski [95] to relate how macroscopic chemical rate constants translate into a binding radius for diffusing spherical molecules.

For a kinetic segregation model involving complete evacuation, detailed spatial organization of individual molecules is desired, which lends itself towards choosing a particle based stochastic simulator that operates on the mesoscale. One such simulator is Smoldyn, initially published by Steve Andrews and Dennis Bray in 2004 [14]. It is a particle-based reaction-diffusion simulator that has a long track record of use in published research [35, 67, 23].

Smoldyn as a biochemical simulator treats space as a continuous medium and uses discrete time steps in its calculations [14]. While using a finite time step for calculating evacuation times causes well known inaccuracies in the specific measurement of evacuation times, trends upon changing parameters such as number of evacuating particles and chemical reactions persist as long as the time step remains the same between subsequent simulations. Additionally Smoldyn is written with much flexibility in the user's ability to define chemical species, chemical reactions, and spatial features such as barriers. As a continuous space, discrete time simulator, Smoldyn's accuracy is comparable to the enhanced Green's Function Reaction Dynamics [90], but scales much more efficiently with an increase in particles. I selected Smoldyn over other simulation methods for its relative computational efficiency as well as the flexibility required to implement more complicated reaction-diffusion features.

Introduction to Weighted Ensemble Sampling

Previous studies done on the evacuation of Brownian molecules in a 2D domain have found this evacuation to be a relatively rare event many orders of magnitude larger than typical time scales of the system [74]. A useful computational tool for studying rare events is the Weighted Ensemble sampling algorithm. This method, originally proposed by Huber and Kim [51], has been shown to be statistically exact for many stochastic processes [103]. Weighted Ensemble has been applied to studying rare events in a variety of situations, most notably in molecular dynamics simulations [104]. Weighted Ensemble obtains information of long timescale processes through multiple short timescale trajectories, hereafter denoted as replicas (figure 2.2). A group of simulation replicas are initialized and attributed a probabilistic weight. These replicas are allowed to evolve into a steady state distribution based of the ensemble space, which is then organized into “bins” based on location inside the state space. After the definition of these bins, the simulations are allowed to evolve for a fixed amount of time with periodic duplications and deletions of certain simulations. These duplications and deletions, dubbed splitting and merging, respectively, are done in such a way that preserves the total probabilistic weight of the system: splitting involves separating a simulation into 2 identical simulations each with half as much weight, and merging involves giving the weight from a deleted simulation to a simulation inside the same bin as the deleted simulation. While the total weight and its distribution between bins might change as the simulation evolves, the number of simulations inside each bin is manipulated so that computation power is evenly split between bins.

Finding the first time of evacuation requires fine temporal resolution as well as fine spatial resolution, but the timescales of evacuation make these events extremely rare. It is instead more useful to consider an equivalent problem: given a steady state ensemble of system states, what is the average probability flux into evacuated states per unit time? The equivalence of these two problems, commonly referred to as the Hill relation [49, 12], relates the mean first passage time T to the average

probability flux $\overline{\Phi}$ as follows

$$T = \frac{1}{\overline{\Phi}}. \tag{1.1}$$

While Weighted Ensemble is statistically exact, the measurements of any one Weighted Ensemble simulation have a highly correlated nature. Because of this, convergence and error analysis for any single Weighted Ensemble simulation, is not straightforward. Furthermore, key metaparameters of Weighted Ensemble simulations are typically chosen ad-hoc. As stated by Bogetti et al., "choosing WE parameters is something of an art" [13]. In the proposal presented, we lay foundations for better analysis and WE metaparameter selection.

Summary of Research

The work for my PhD thesis is thus split into three parts. First, I developed freely available software combining Weighted Ensemble sampling with the biochemical simulator Smoldyn. Next, I used this tool to investigate CD45 evacuation times for a variety of physical scenarios including dimerization and limited reentry. Finally, I analyzed two key metaparameters for Weighted Ensemble, the frequency at which re-weighting occurs and the number of replicas maintained in each bin, investigated their impact on simulation convergence and error, and investigated the impact of various methods of measuring error.

Chapter 2

Surface Dynamics of CD45 Evacuation

Introduction

The physical and biochemical mechanisms of T cell receptor triggering are not fully understood, though a number of different models have been proposed, reviewed in [93, 87]. These models synthesize a variety of experimental observations and help further understanding of how T cells can achieve exquisite sensitivity to small amounts of antigen, while also discriminating antigen from self.

As previously stated, T cell receptor triggering is associated with reorganization of cell surface proteins, shown schematically in Fig. 2.1A. Local depletion of surface proteins with large ectodomains, such as the phosphatases CD45 and CD148, in the receptor's vicinity has been demonstrated to be an important step [64, 94]. The kinetic segregation model is supported by various lines of evidence [29, 80, 20, 56, 25]. However, the kinetic segregation model cannot by itself fully explain T cell receptor triggering: other mechanisms likely contribute (e.g., see [93]) and CD45 plays somewhat contradictory roles [29].

How and when the local evacuation of large ectodomain phosphatases from the receptor vicinity occurs remains unclear. It could happen pre-contact-formation (e.g., does local evacuation of CD45 clear the way for receptor-antigen binding?) or post-close-contact formation (e.g., a scenario where receptor-antigen binding is first enabled by close-contact due to a microvillus [57], or active membrane protrusion). The question of what drives this evacuation can be cast in three different lights: In the absence of receptor ligation, what mechanisms ensure evacuation does not happen accidentally? In the presence of ligation, what ensures evacuation happens in sub-second timescales necessary to initiate and sustain a signal? And, finally, if the process of evacuation tips the balance from inhibitory to stimulatory signaling in T cells, could modulating the evacuation process itself be an avenue accessible to engineered therapeutics?

Various mechanisms for this evacuation process have been proposed. These include: simple Brownian motion of CD45 in the plasma membrane [74]; oligomerization reactions between the molecules [91, 98] shown schematically in Fig. 2.1B; and mechanical compression by a nearby surface, such as that of the cell presenting the antigen. The compression region can be conceptually categorized as either a close-contact of $\sim 100\text{nm}$ [40, 22], shown schematically in Fig. 2.1C, where there is no net lateral pressure on the molecules within the close contact, or something more similar to the wedge or tent shape resulting from a force on a single receptor pulling the membrane at a point [3]. The latter case has been studied theoretically [3, 76]. Beyond the scope explored in this work, there are many more possible mechanisms, including: spatial heterogeneity due to lipid composition or interaction with the cytoskeleton [42, 86, 77]; modulation of the configurational state of the individual molecules themselves can modulate their organization, e.g., by electrostatic interaction with the lipid membrane [27, 26]; or crowding out by CD3 [66].

Computer simulations of the spatiotemporal events involved in receptor triggering and immune synapse formation can provide a means to investigate phenomena that lie beyond the spatiotemporal resolution of measurement techniques [63]. In this paper, we investigate the evacuation process through computer simulations tracking reaction-diffusion dynamics of protein molecules




Figure 2.1: Surface dynamics for a large, membrane-bound surface molecule like CD45. (A) T Cell Receptor (TCR) binding to an antigen presented by another cell. (Left) If CD45 is uniformly distributed around the receptor, signaling is inhibited. (Right) The local depletion of CD45 from the region of interest (ROI) near the receptor, approximated here as a circle with radius R_{ROI} , is a key step in T Cell receptor signaling [93, 87]. (B) Intermolecular interactions between CD45, e.g., oligomerization of CD45 into dimers. Diffusion coefficients for monomers and dimers are D_1 and D_2 , respectively [91, 98]. (C) Close contacts created by, e.g., microvilli, could lead to biased movement of CD45 due to compressional resistance of the molecule [40]. We represent these close contacts by probabilistically limiting entry into the ROI.

on the cell surface. To simulate wide ranges of parameters, including different molecular phenomena, with a physiological and near-physiological numbers of molecules, we made recourse to the Weighted Ensemble algorithm [51], an enhanced sampling simulation method.

We confirm mathematical estimates that, at physiological concentrations, spontaneous evacuation is extremely rare [74]. We find that dimerization decreases the timescale of evacuation for even weak bi-molecular interaction by several orders of magnitude. The formation of higher-order oligomers reduces evacuation to a sub-second process, opening the possibility that an engineered

oligomer of CD45 could significantly modulate receptor triggering. We find that formation of close contacts also decreases the timescale of evacuation. However, for large regions of close-contact, such as those induced by large microvilli, our model predicts evacuation times that are still too long by several orders of magnitude, using current estimates for the molecular size and compressibility of CD45. This suggests that the change in molecular motion driven by close contact alone is not sufficient to drive receptor triggering.

Results

A Rare-Event Reaction-Diffusion Simulation

Molecules in our model are represented as individual particles moving on a two-dimensional surface. Their mean density is ρ , and we assume there is a region of interest (ROI), for example near a single T cell receptor, that we approximate as a disk with radius R_{ROI} . We make the simplifying assumption that the receptor motion is negligible relative to the motion of the individual CD45 molecules. The ROI lies in the center of a 2D square domain with edge length L . The quantity we wish to compute is the mean time until the ROI is empty—which we refer to as the evacuation time or mean first-passage time (MFPT)—under various assumptions about the dynamics of the molecules. In the base case, we assume motion is purely diffusive with coefficient D (which could be thermal or include active, random forces [44]). Note that previous work [74] has shown that diffusion alone is too slow to be consistent with experimental data, giving MFPTs of $\sim 10^{10}$ s whereas triggering can occur in reality within seconds [54, 50]. That is, the simultaneous evacuation of all molecules from the ROI by simple diffusion, given physiological surface density and ROI size, is a rare event. We are interested in the rare event limit because, first, its quantification helps reject this null hypothesis, and, second, it provides a necessary starting point for hypothesizing what may accelerate the biological process of T cell activation out of the rare event limit.

Summary of key parameters in the base model

We parameterize our model based on experiment-derived estimates. The key parameters of the base model are the size of the evacuation region, the diffusion coefficient of CD45, and the surface density of CD45. The size of the evacuation region (that is, the radius of the ROI) provides a characteristic lengthscale for the model system. Assuming $R_{\text{ROI}} = 100 \text{ nm}$ ([17, 40, 3]), and assuming the diffusion coefficient to be $D = 0.01 \mu\text{m}^2/\text{s}$ ([18]), these provide a characteristic timescale for the system of $R_{\text{ROI}}^2/D \approx 1\text{s}$. We define scaled units for molecule density, ρ , as molecules per R_{ROI}^2 . In these scaled units, we estimate the physiological surface density of CD45 to be $\rho \approx 9$ molecules per R_{ROI}^2 [3]. Parameters are listed in Table 2.1; further details of parameter values are discussed in Methods.

Combining Weighted Ensemble and Smoldyn allows simulation of long-timescale stochastic spatial phenomena

To simulate evacuation in the rare-event limit, we combine the particle-based reaction-diffusion simulator Smoldyn [7] with a Weighted Ensemble algorithm [51, 12, 103, 32]. The algorithm has been shown to be statistically exact for many stochastic processes [103] and has been applied to many systems, especially in molecular dynamics [104]. In this algorithm, shown schematically in Fig. 2.2, many Smoldyn simulations are run in parallel, each assigned a weight w_i , with $\sum_i w_i = 1$. An order parameter is used (here, number of particles in the ROI, n_{in}) to bin the state space (Fig. 2.2A). Periodic redistribution of weights and trajectories among the bins occurs after each time interval τ . More simulations are run in probabilistically less-likely regions of state space, but assigned a smaller weight. This focuses computational power on rare events while still maintaining an algorithmically exact statistical ensemble [103, 104]. Every τ time units, we record the probability flux into the evacuated state (i.e., the summed weights of trajectories that reached the bin with $n_{\text{in}} = 0$ during that time interval), as shown schematically by red arrow and circle in

Fig. 2.2B, and actual evacuated weights shown in Fig. 2.2C. Trajectories that reach evacuation are killed, and their weights are redistributed into the system. Once the system reaches steady state, the mean flux gives the reciprocal of the MFPT [12]. In practice, an estimate of the MFPT is obtained based on multiple independent simulations, as indicated by different color lines in Fig. 2.2C. For full details see Methods and Supplemental Fig. 2.7. We name this algorithm and combination WE-Smoldyn.

A summary of Weighted Ensemble and simulation hyperparameters is in Table 2.1.

WE-Smoldyn agrees with brute force stochastic simulation at low densities and approaches asymptotic calculation for high densities.

We first simulate evacuation for a density of particles ρ undergoing diffusion only, and compute the mean time to evacuation. In Fig. 2.3A, we demonstrate the computational ability to simulate evacuation and compute MFPTs for a range of surface density values, reaching $\rho = 10$, which is consistent with experimental estimates [3, 40] and which corresponds to a mean number inside the ROI of $\overline{n_{\text{in}}} = \pi\rho = 31$. The MFPT grows superexponentially with ρ , reaching $T \approx 10^{12}$ seconds at the highest simulated density. The MFPT has an uncertainty of less than one order of magnitude (error bars are standard error of the mean).

To validate our method, we compare with a brute-force simulation using Smoldyn, at low ρ (red). At high ρ , we compare to the asymptotic approximation from [74] (black dashed), which gives the following value for the MFPT

$$T = \frac{\kappa_{2D} R_{\text{ROI}}^2 e^{\overline{n_{\text{in}}}}}{D \overline{n_{\text{in}}}^2}, \quad \overline{n_{\text{in}}} \gg 1. \quad (2.1)$$

Here, $\kappa_{2D} \approx 0.7$ is a constant independent of all parameters, see [74]. In the rare event limit, agreement to the asymptotics is within two orders of magnitude, and we hypothesized that the dis-

agreement is due to a finite domain size in our simulation (whereas the asymptotic approximation is in the limit of infinite domain). We confirm that the MFPT, estimated from our simulations as a function of increasing domain size, approaches the value in Eq. 2.1 (see Supplemental Fig. 2.8). The domain-size effect can be intuitively understood to result from the difference in particle distributions for a small domain versus a large domain during an evacuation event. For evacuation to occur, the density of particles outside the ROI must increase, while the density of particles inside the ROI must decrease (to zero). Also, the ROI area makes up a greater fraction of the total system area for a small domain, as compared to a large domain. Therefore, more particles must be “squeezed” into a smaller area (that is, the density of particles outside the ROI must increase further from the equilibrium value), in order for evacuation to occur in the small domain case, as compared to the large domain.

The computational scaling in Fig. 2.3B suggests that simple timestepping would take 4×10^5 years of CPU time for 1000 evacuation events, whereas the Weighted Ensemble method took approximately 25 days per MFPT measurement (2.5 days per measurement, then repeated 10 times) on a single CPU core. Thus, the simulations we present throughout this paper would be unfeasible without recourse to an enhanced sampling algorithm like Weighted Ensemble.

The definition of evacuation time we use here and in [74] is instantaneous, in other words until the last molecule reaches the boundary of the ROI. This raises two notes: First, this is an approximation, since in the T Cell it is likely that a T cell receptor in an almost-evacuated ROI could still become triggered, just at a lower rate. Second, the timestepping algorithm we use here could lead to overestimates of the MFPT. To control for this second approximation, we confirm in Supplemental Fig. 2.8 that MFPTs are independent of simulation timestep Δt .

Oligomerization

Model for intermolecular interactions

There is some evidence that CD45 dimerizes [98, 91]. In this section, we explore the impact this would have on the evacuation process. To include dimerization, we add a reversible binding reaction to the model, with unbinding rate k_{off} , and binding occurring whenever two particles are within a distance $r_{\text{bind}} = 10^{-2}$, which corresponds to a physical distance of 1 nm. Binding distance roughly corresponds to a binding rate. See Supplemental Fig. 2.9 and Methods for details on simulation of reversible dimerization of surface molecules. We further assume that dimers diffuse more slowly by two-fold [30]. We explore the effect of varying k_{off} , i.e., of varying the equilibrium constant for the dimerization reaction.

Dimerization decreases the timescale of evacuation by orders of magnitude even for weak bi-molecular interaction strengths

At low k_{off} , evacuation times are decreased by over 5 orders of magnitude, as shown in Fig. 2.4A (towards left of horizontal axis). Order of magnitude changes in MFPT appear to track closely with corresponding steady-state monomer fraction, as a function of k_{off} (Fig. 2.4B). We can understand this heuristically as follows, making use of the asymptotic approximation (for monomer evacuation) in Eqn. 2.1 [74]. At low k_{off} , most molecules are in dimer form. Evacuation is nominally slowed by the reduction in diffusion coefficient, since the evacuation time scales as $T \propto D^{-1}$ in the asymptotic (infinite domain) limit. However, this effect is outweighed by the reduction in the number of independent particles, since $T \propto e^N/N^2$, which leads to an almost exponentially-lower MFPT [74]. Thus, the linear reduction in diffusion coefficient is dominated by the near-exponential dependence on the (linear) reduction in number of particles. Indeed when we compute MFPT as a function of the fraction of monomers in Fig. 2.4C, we observe an approximately exponential

relationship.

More surprisingly, this dramatic reduction in evacuation time occurs even for weak dimerization. When monomer fraction is as high as 80%, meaning only 20% of CD45 subunits are in dimers, the MFPT is already reduced by ten-fold.

These relatively high unbinding rates correspond to weak homodimerization affinities. We can compute an effective 2D dissociation constant, defined as the concentration of reactants at which half the reactants are in the product (dimer). We find that a reaction with $K_D^{2D} = 0.0058\text{nm}^{-2}$ would yield a ten-fold reduction in evacuation time. A reaction with $K_D^{2D} = 0.0009\text{nm}^{-2}$ would yield a 1000-fold reduction. Conversion of 2D chemical properties to the equivalent 3D properties is nontrivial [97, 56], but a lower-bound estimate can be obtained by dividing the 2D density by the confinement height of the reaction [97, 99, 11]. In this case, the upper bound for the confinement height is the height of CD45 [21]. Using this, we can compute a lower-bound affinity for 10-fold reduction of MFPT, $K_D^{3D} = 460\mu\text{M}$, which corresponds to a standard binding free energy of $\Delta G_{\text{bind}} = k_B T \ln K_D^{3D} = -19\text{kJ/mol}$. These binding strengths are an order of magnitude weaker than those measured for agonist TCR-peptide-MHC [53]. Note that these over-estimate the needed strength, since the confinement length we assumed to convert to a 3D affinity is an over-estimate.

Effects of higher-order oligomers on evacuation

This led us to wonder how evacuation times would be affected by the formation of higher-order complex molecular assemblies, for example as could be engineered using extracellular molecular linkers.

Full simulation of higher-order oligomerization was beyond our limits due to combinatorial complexity of number of molecular species and reactions between these. (This is algorithmically feasible, especially with rule-based modeling tools [5, 6] but would be difficult to parameterize, i.e., we would either need to estimate or explore a combinatorially large number of rate constants). How-

ever, we can use Weighted Ensemble to compute the evacuation time assuming k_{off} is sufficiently low, such that all molecules are in the highest-order oligomer, shown in Fig. 2.5. Here, again, we assume that diffusion coefficient is reduced proportional to the number of subunits in the complex. For oligomers larger than dimers, we assume the ensemble is homogenous and only made of the largest complex. Between oligomer size 1 and 2, we show heterogeneous mixtures using the same sweep of k_{off} from Fig. 2.4, but plotted as a function of the average number of subunits in each independently diffusing particle,

$$\text{Avg. \# subunits} = \frac{2[\# \text{ dimers}] + [\# \text{ monomers}]}{[\# \text{ dimers}] + [\# \text{ monomers}]} \quad (2.2)$$

Consistent with the result for dimers, these larger oligomers evacuate faster, despite diffusing more slowly. Indeed, hexamers with strong binding evacuate within 1 second.

We have assumed that the diffusion coefficient of oligomers is $D_n = D_1/n$, where n is the number of subunits. Note that for a single molecule with a transmembrane domain, or multiple transmembrane domains very close together, previous findings suggest a linear decrease in D with lateral radius of the diffusing membrane protein [43], while, other studies found that D depends more weakly on protein radius [78], and it has also been shown that the size-dependent decrease in D of membrane proteins depends on factors such as membrane crowding [55]. In our case, we assume the oligomers result from the attachment of multiple molecules each with a single transmembrane domain of unchanged size. So, the transmembrane domains are further apart than in [43, 78, 55], and therefore hydrodynamic interactions between the transmembrane domains is weaker. Previous studies suggest that, if separated by the size of a protein, the hydrodynamic interactions are weak and the scaling of drag coefficient with n becomes linear [30]. Nonetheless, our assumption of $D_n = D_1/n$ may overestimate the decrease in D upon oligomerization; in this case, the effect (decrease of MFPT upon dimerization/oligomerization) would be stronger than that shown in Figs. 2.4 and 2.5.

Close Contacts

Model for molecule behavior at a close contact

T cell receptor triggering can be induced by the formation of cellular protrusions called microvilli, which push against a surface, creating a region in close contact [40] between two surfaces, as shown schematically in Fig. 2.6A. If the close contact membrane-separation is smaller than the resting size of CD45, it has been hypothesized that this leads to dynamics in which CD45 can diffuse out but not back into the area of close contact [40]. Note that this is distinct from models in which the membrane deformation is tent-shaped or wedge-shaped, and therefore induces non-diffusive advection on compressed molecules [3]. We begin this section by exploring the model regime in which CD45 has an unspecified compressional resistance at a close contact size.

Although the compressional resistance of CD45 is a key property in kinetic segregation models, estimates are challenging. Efforts to measure a similar molecule [15] have yielded estimates around $k_{\text{spring}} = 0.1\text{pN/nm}$. If a close contact is held in place by T cell receptors bound to antigen, the height difference between the rest size of CD45 and the close contact size has been estimated to be $\Delta z = 6.6\text{nm}$ [21].

For large regions of close-contact, such as those induced by large microvilli, molecular size and compressibility imply an intermediate re-entry probability

Using a Boltzmann relationship between the compression energy and the re-entry probability (see Methods, Eq. 2.3), we can compute how these two molecular properties influence the ability of CD45 to enter the close contact, shown in Fig. 2.6B. Stiff or large molecules enter the ROI with near-zero probability, and soft or small molecules enter with high probability, but the physiologically estimated parameters lead to an entry probability of $p_{\text{entry}} \approx 0.6$. Our finding is in contrast to

previous models, e.g. [40], which assumed $p_{\text{entry}} = 0$. We perform simulations of the evacuation process with varying p_{entry} to investigate its effect on evacuation.

How well-approximated is the evacuation time by assuming zero re-entry, or by assuming free diffusion? It cannot be well-approximated by both, since the evacuation time we found above at $p_{\text{entry}} = 1$ is many orders of magnitude larger than the evacuation time in simulations from [40], who assumed $p_{\text{entry}} = 0$.

Physiological levels of molecular compressibility lead to significant reentry, leading in turn to significant delays in evacuation compared to purely one-way evacuation.

We use our WE-Smoldyn algorithm to simulate a density $\rho = 9$ of molecules undergoing diffusion, but with the assumption that a given molecule, after exiting the ROI, re-enters with probability p_{entry} . We find that the evacuation times, shown in Fig. 2.6C, indeed vary between the simple diffusion case $p_{\text{entry}} = 1$ and the no-entry case $p_{\text{entry}} = 0$. The evacuation time, given physiological estimates of k_{spring} and Δz , is around 10^5 seconds (pink vertical bar in Fig. 2.6C). Although this value is orders of magnitude faster than the evacuation time computed for the $p_{\text{entry}} = 1$ case (simple diffusion), it remains substantially longer than T Cell triggering times.

Note that estimates of k_{spring} have varied widely [16, 15]. As empirical uncertainty bounds are not presently available, we explored the effect of a potential two-fold increase or decrease of k_{spring} . Propagation of this uncertainty to p_{entry} implies a range from 0.36 to 0.77 (gray dashed lines in Fig. 2.6C), with a corresponding variation in the computed MFPT over approximately five orders of magnitude. These results demonstrate the high sensitivity of evacuation times on p_{entry} and, in turn, on the biophysical parameters that govern CD45 movement near the close contact.

We validate our results at $p_{\text{entry}} = 1$ by comparison with our simulations for simple, unhindered diffusion. We also solve for an analytic expression for the MFPT at $p_{\text{entry}} = 0$. This calculation is performed in the Supporting Text and shown in Supplemental Fig. 1.1. Agreement with

WE-Smoldyn is shown as the open green circle in Fig. 2.6C. We further performed brute force simulation for $0 < p_{\text{entry}} < 0.3$.

The dramatic effect of even small changes in p_{entry} led us to wonder about the relative importance of close contact size and gap size (which determined p_{entry}). Note that, so far, all figure panels have shown evacuation time and densities nondimensionalized by scaling with the radius of the ROI. Rescaling to physical units, at a fixed physical density $\tilde{\rho}$, would require simulations over a range of p_{entry} and ρ . Instead, we make use of our finding that evacuation time is an approximately exponential function of p_{entry} , and use this as an ansatz in Eq. 2.4 (Methods).

Evacuation times in physical units are shown in Fig. 2.6D, for a density $160 \times \mu\text{m}^{-2}$. Note this is lower than the physiological estimates used elsewhere by about 4-fold but at the limit of our current computational capability. At these parameters, a close contact region of radius less than $\sim 100\text{nm}$ evacuates spontaneously in sub-second time. Close contacts that are perfectly impenetrable also evacuate in sub-second time, up to at least 250nm, comparable to the size observed [17, 40]. However, close contacts larger than 100nm with 20% or more re-entry probability have significant slow-down in evacuation. Here, an observed change in the shape of the curves leads to an interesting prediction: At low $p_{\text{entry}} < 0.2$, relative changes in p_{entry} lead to more significant changes in evacuation time compared to the same relative change in close contact radius. At high $p_{\text{entry}} > 0.6$, relative changes in close contact radius R_{ROI} lead to more significant changes in evacuation time compared to the same relative change in entry probability. At the physiological estimate $p_{\text{entry}} = 0.6$, roughly, a 20% reduction in p_{entry} has the same effect in reducing MFPT as a roughly 20% reduction in size of the close contact.

Discussion

The paradigm of kinetic segregation — triggering a receptor by local depletion of its deactivating enzyme — has been proposed for a variety of surface receptors [34, 73]. The most developed example is T cell receptor triggering by CD45 depletion. In this work, we show that, first, simple diffusive motion of CD45 leads to spontaneous depletion extremely rarely, in agreement with previous results [74]. Spontaneous depletion is therefore not at risk of false positive receptor triggering in the absence of an external cue. Second, we show that oligomerization of CD45 dramatically increases the speed of depletion. And third, we show that a close contact may accelerate depletion, but depending on its gap size and the mechanical properties of CD45, depletion may nevertheless be extremely slow.

Our results on oligomerization make a prediction: that externally-induced oligomerization of CD45 into higher-order structures would lead to more rapid receptor triggering, and indeed sufficient oligomerization (e.g., dominant heptamers, Fig. 2.5) would lead to spontaneous receptor triggering. Such oligomerization could be performed on the extracellular regions of CD45. This could therefore provide a test of the prediction, particularly in a controlled system like a liposomal reconstitution [52]. It also predicts a mechanism through which CD45 modulation leads to orders-of-magnitude changes in a very proximal step in T cell function. It is intriguing to speculate the effect of oligomerization on overall T cell signaling, and on engineered T cell function, e.g., T cells used in therapeutics [47, 96]. However, the observation that CD45 has both positive and negative regulation of overall T cell signaling [91, 98?] suggests a highly nonlinear system, which demands more careful quantitation and quantitative modeling. Outside the context of CD45, galectins [39, 38] can cross-link surface proteins via their sugars, but their role is not clear. These galectins can generate higher-order oligomers. They also have weak specificity, and therefore could be a general mechanism to evacuate molecules with many sugars. It could also be that oligomerization is induced not by direct protein-protein interaction, but rather by a mobile raft of

lipid heterogeneity. Such oligomers would satisfy a similar mathematical model, perhaps with a different diffusion coefficient.

Calculation of the theoretical evacuation time at a close contact has implications for models of close-contact surface molecule dynamics. In particular, Fernandes et al. [40] made the assumption that once a CD45 leaves the close contact region, it cannot re-enter ($p_{\text{entry}} = 0$). We confirm here that this leads to depletion times on the order of seconds. However, using estimates of CD45 geometry and mechanics, we compute that if there is a 60% chance of re-entry, the depletion time increases to 10^5 seconds, much slower than observed timescales of receptor triggering [40, 54, 50]. Formally, there are several possible resolutions to this discrepancy: If the estimated geometry and mechanical properties of CD45 are accurate, there must be another phenomenon contributing to evacuation. Alternatively, the molecular spacing could be smaller than estimated in [21], or the molecules could be much stiffer. Indeed, according to our model, a factor of two change in estimate of the latter parameter (i.e., in the spring constant of CD45) induced a greater than hundred-fold change in the evacuation time. This sensitivity shows that close contacts under the current model could indeed result in short-enough (or nearly so) evacuation times, under different estimates of biophysical parameters. Thus, the question of what mechanisms drive evacuation, and particularly the estimation of p_{entry} , warrants further study.

The model we used is minimal in its assumptions and therefore subject to limitations. Our model focuses on the kinetic segregation mechanism, however, it has been proposed that kinetic segregation is just one of many mechanisms contributing to T cell receptor triggering [93]. Moreover, we focus on the inhibitory effect of CD45 on TCR signaling, whereas CD45 can both positively and negatively influence TCR signaling [29]. Within the kinetic segregation model, one limitation is our focus exclusively on total depletion, when the last molecule leaves the ROI. In reality, other steps in receptor triggering include ligand binding and receptor phosphorylation [22]. So, a more realistic model could be formulated in which the number of CD45 in the ROI determine a next-event rate. This rate would be high for total evacuation, and slower for partial evacuation.

The MFPT one would study would be the time until the next-event has occurred. Without further assumptions, it is possible this next-event could happen slower than the total evacuation time, since it adds a subsequent step, or faster, since it can be triggered when evacuation is not total. Another limitation is the assumption of a flat, two-dimensional membrane. In particular, our consideration of microvilli ignored the purely geometric effect of a microvillus, in which distances around the perimeter of the microvillus are smaller than distances around the ROI in our flat simulations. Simulating diffusion on such curved surfaces is computationally possible, but more expensive, in Smoldyn [5] and would require more characterization of the 3-dimensional shape of microvilli. Yet another limitation is our focus on motion of CD45, when in reality the receptor moves as well [22]. Further integrative models, at the cell scale, may also include multiple receptors, and therefore multiple opportunities for a T cell to activate. Future research may explore these directions.

Crowding is prevalent in biology [19, 37, 36, 28, 61, 89, 68]. For that reason, there are examples in which *un*-crowding may be important — that is, when molecules must evacuate from a region before a given process can occur, and so the problem of making space is of general interest. These include the many transient cell-cell contacts which occur during tissue development (e.g., the delta-notch system [70, 85]). There are also membrane-membrane contacts within cells, including between the endoplasmic reticulum and plasma membrane (where crowding could modulate interactions of molecules including Ora1 and Slim1 [62]). In 1D, an example is offered by transcriptional control in eukaryotes, which is achieved by the binding of many classes of proteins to DNA [88, 58]. Transcription factors (TFs) locate to binding sites within promoters and enhancers by 1D diffusion along the DNA and by attachment/detachment into the 3D cytoplasm [46, 69, 71]. The binding of larger structures, such as nucleosomes, which occupy ~ 150 base pairs (bp) of DNA, is inhibited by the presence of TFs, and therefore it is intriguing to wonder whether evacuation timescales are significant. Furthermore, enhancers, which are $\sim 200 - 1000$ bp stretches of DNA with 5-30 TF binding sites of various classes, may require evacuation of nucleosomes and transcriptional repressors to activate their target genes. Again in 1D, microtubules (inflexible polymers of the protein tubulin) are decorated by hundreds of microtubule-associated

proteins [2, 82]. These proteins exhibit significant crowding [28, 61] and lateral diffusion along the microtubule lattice [31, 48]. Large microtubule-binding molecules may therefore have to wait for a region to be clear before binding.

Simulations performed here would be unfeasible without recourse to an enhanced sampling algorithm. Weighted Ensemble has been applied to many different types of stochastic dynamics simulations, however ongoing challenges are present, e.g., in *a priori* selection of state-space binning strategies, hyperparameters, and weighting schemes to optimize convergence of desired observables [104]. Further systematic study of Weighted Ensemble hyperparameter selection and analysis methods should lead to further increases in efficiency and empower future rare event simulation studies. Our WE-Smoldyn code base was built on top of the Smoldyn dynamics engine, which is widely used, flexible and with a large user base. We anticipate the combination of Weighted Ensemble with spatial stochastic simulation, as highlighted by full-featured software like MCell-WESTPA [33], will open new avenues of research, including for the evacuation questions posed in the previous paragraph.

Methods

Model and Dynamics

Diffusion

Smoldyn is a time-stepping simulator with a continuous spatial domain (as opposed to a lattice-method). At each time step, molecule displacements are drawn from a Gaussian distribution whose width is determined by their diffusion coefficient D , with each chemical species having their own diffusion coefficient. At each Smoldyn timestep, Smoldyn tracks the location of each molecule, and stopping when it observes a complete evacuation of the ROI (Fig. 2.1A, right). Aside from

interactions with boundaries, barriers, and for molecular binding, all molecules diffuse independently and do not interact with each other. Boundary conditions were chosen to be reflective, though preliminary results did not show substantial differences between reflective and periodic boundary conditions.

As we are using a time-stepping based method, the determination of whether or not a molecule has evacuated the ROI within a timestep is based only on its starting and ending locations, and specific details of the trajectory between the time-steps are lost. This representation results in evacuation events that occur between two sequential time-steps that are not observed by the method. As evacuation events are lost but none are gained, this would result in estimates of the MFPT that are higher than the true value rather than an underestimate. To minimize the number of evacuation events lost from these missing trajectories, we chose to use a small Smoldyn timestep, $\Delta t = 10^{-6}$. To ensure this choice of Δt is small enough we confirmed that smaller timesteps give similar results, as shown in Fig. 2.8.

Intermolecular interactions and reversible dimerization

Smoldyn uses an algorithm that is qualitatively similar to the Collins-Kimball model of bi-molecular reactions and approaches the Smoluchowski model for short time steps [8]. The association reaction occurs when two monomers diffuse within a pre-defined distance of each other referred to as the binding radius, r_{bind} . Dissociation of dimers into two monomers is probabilistically determined at each time step, with probability determined by the detachment rate k_{off} [8].

Two-dimensional reactions are more complicated to analyze than their 3D counterparts [100?]. For example, there is no exact relationship between r_{bind} and a well-mixed k_{on} . We confirm that, for our choice of $\Delta t = 10^{-6}$ and ranges of r_{bind} and k_{off} , the steady-state unbound (monomeric) fraction is a smoothly increasing function of k_{off} and decreasing function of r_{bind} , as shown in Fig. 2.9A. The dissociation constant, meaning the value of k_{off} at which half of subunits are in

monomers, is a weakly increasing function of r_{bind} , as expected by previous theoretical treatments. [100?]. The unbinding radius, the distance between two monomers that dissociated in the previous step, was set such that the geminate recombination probability was 0.2, which gives r_{unbind} a nominal value of 0.0443134 (≈ 4 nm).

Close Contacts

By representing close contacts of the system as energy barriers caused from compression of molecules inside the ROI/close contact, we can then model these energy barriers by creating asymmetric behavior between molecules attempting to enter the ROI and those leaving it. Molecules attempting to leave the ROI are free to do so, while those attempting to enter are only allowed to do so probabilistically. The energy barrier between the ROI and rest of the domain is taken to be the energy required to compress a spring, $E_{\text{spring}} = 1/2k\Delta z^2$, where k is the physiological spring constant and Δz is the size of the compression. The thermodynamic relationship between this probability and the energy compression is given by

$$p_{\text{entry}} = e^{\frac{-E_{\text{spring}}}{k_b T}} \quad (2.3)$$

where k_b is Boltzmann's constant and T is the temperature. According to this definition, p_{entry} gives the probability that a molecule has energy greater than E_{spring} (and thus may gain entry to the ROI); it follows that p_{entry} also equals the ratio of concentrations of molecules inside versus outside the ROI at equilibrium. (Note, however, that p_{entry} is not the same as the instantaneous probability of entry of a single molecule during a collision event in the simulation. Particle-based simulation of this type of partial transmission was discussed previously by Andrews [4], and we made use of Smoldyn's built-in functionality for implementation.) In the absence of evidence otherwise, we assume the presence of the close contact does not influence diffusion coefficient D .

Parameter estimates and model nondimensional scaling

Estimates for R_{ROI} range from 100 – 220 nm [17, 40, 3], depending in part on the definition, e.g., whether it is the minimum region necessary for receptor triggering, or the observed depletion zone size. The diffusion coefficient D has been estimated to range from $0.01\mu\text{m}^2/\text{s}$ [18] to $0.3\mu\text{m}^2/\text{s}$ [40]. Roughly setting $R_{\text{ROI}} = 100\text{nm}$ and $D = 0.01\mu\text{m}^2/\text{s}$ conveniently sets the characteristic timescale $D/R_{\text{ROI}}^2 = 1\text{s}$.

Given the wide ranges of estimates, throughout this work we report times and distances in these scaled (nondimensional) units. Where appropriate, we report results in physical units, denoting these by explicitly including the unit (e.g. seconds or nanometers), and, if clarity necessitates, we use a tilde to denote the parameter with physical units. Domain size L and binding radius r_{bind} have units of R_{ROI} . The corresponding physical parameters $\tilde{L} = LR_{\text{ROI}}$ and $\tilde{r}_{\text{bind}} = r_{\text{bind}}R_{\text{ROI}}$ have units of nanometers or microns. The dynamics-engine timestep Δt has units of D/R_{ROI}^2 and the detachment rate k_{off} has units of R_{ROI}^2/D .

The scaled molecule density ρ has units of molecules per R_{ROI}^2 , and the physical molecular density $\tilde{\rho} = \rho/R_{\text{ROI}}^2$ has units of molecules per square nanometer. Another interchangeable quantity is the number of particles in the ROI in a uniform distribution, $\overline{n_{\text{in}}} = \pi\rho = \pi R_{\text{ROI}}^2\tilde{\rho}$. For CD45, estimates range from $482\mu\text{m}^{-2}$ [40] to $1000\mu\text{m}^{-2}$ [3]. For $R_{\text{ROI}} = 100\text{nm}$, this corresponds to a density ranging from $\rho = 4.82$ to $\rho = 10$. We use $\rho = 9$ as our focus [3] in all figures unless otherwise noted (e.g., in the ρ sweeps in Fig. 2.3 and the physical unit plot in Fig. 2.6D).

Ansatz for rescaling close contact evacuation time to physical units

To return to physical dimensions, a constant physical density $\tilde{\rho}$ requires varying scaled density ρ since $\tilde{\rho} = \rho/R_{\text{ROI}}^2$. In Fig. 2.6D, we plot the evacuation time in physical units over a range of R_{ROI} . Doing so would require a full exploration of both nondimensional ρ and p_{entry} , which is

outside of our computational capacity. So, as an approximation, we take the result in Fig. 2.6C, which suggests that evacuation time $T(p_{\text{entry}})$ is a simple exponential function of p_{entry} , that varies between the $p_{\text{entry}} = 0$ limit and $p_{\text{entry}} = 1$ limit. In other words,

$$T(p) \approx T(0) \left(\frac{T(1)}{T(0)} \right)^p \quad (2.4)$$

where $p = p_{\text{entry}}$. The MFPTs for $T(1)$ can be obtained from the simple diffusion simulations in Fig. 2.3A, while the MFPTs for $T(0)$ can be obtained analytically, see Fig. 1.1B.

Weighted Ensemble

Algorithm Overview

Weighted Ensemble obtains information of long timescale processes through multiple short timescale trajectories, hereafter denoted as replicas. A group of simulation replicas are initialized and attributed a probabilistic weight (see Fig. 2.2A). These replicas are allowed to evolve into a steady state distribution based on the ensemble space, which is then organized into bins based on location inside the state space. After the definition of these bins, the simulations are allowed to evolve for a fixed amount of time τ with periodic duplication and deletion of certain simulations (Fig. 2.2B). These duplications and deletions, dubbed splitting and merging, respectively, are done in such a way that preserves the total probabilistic weight of the system: splitting involves separating a simulation into two identical simulations each with half as much weight, and merging involves giving the weight from a deleted simulation to a simulation inside the same bin as the deleted simulation. While the total weight and its distribution between bins might change as the simulation evolves, the number of simulations inside each bin is manipulated so that computation power is evenly split between bins. Our binning order parameter is the number of subunits inside the ROI, n_{in} . So each monomer inside the ROI increased the order parameter by 1, while for simulations

with dimerization, each dimer inside the ROI increased the order parameter by 2.

If a replica reached the bin where $n_{in} = 0$, hereafter referred to as the flux bin, the replica is removed from memory, its probabilistic weight is recorded as outgoing flux, and then the weight is redistributed according to one of two different methods, see Fig. 2.7 and Reweighting Methods (below).

Model initialization

Each WE simulation is initialized with 1000 replicas of a Smoldyn simulation. In each replica, each of N molecules is randomly and uniformly placed throughout the entire domain. After this initialization, WE splitting and merging, and flux measurements are performed before each subsequent step of Smoldyn dynamics (Fig. 2.7C).

In simulations involving more than one molecular species, initialization is done with homogeneous molecular mixtures; either N monomers or $N/2$ dimers are uniformly distributed in the simulation, depending on which is closer to the steady state as found by brute force simulations in Fig. 2.9.

Reweighting Methods

There were two methods of redistributing weight removed from the system through flux into the flux bin ($n_{in} = 0$, see Fig 2.7). The first method, which was the method used unless otherwise noted, involves redistributing the weight by renormalization; the weight of all remaining replicas is scaled by the total weight remaining in the simulation. If replica i evacuates, it is removed and the weight for a replica j remaining in the simulation will scale according to

$$w_j \rightarrow \frac{w_j}{1 - w_i}. \tag{2.5}$$

This method works for calculating the mean transition time from the steady-state distribution (or small perturbation from steady-state) to a rare fluctuation. In other words, we are measuring the MFPT from $A \rightarrow B$, where B is defined as the bin $n_{in} = 0$ (which is rarely visited), and A is defined as encompassing all bins $n_{in} > 0$. Note that in many WE applications, significant time is required for the system to reach steady state, before which accurate MFPT estimates cannot be obtained from the averaged flux-to-target [12]. In our system, we know *a priori* the equilibrium distribution of particles undergoing simple diffusion. Our initialization of replicas according to the equilibrium distribution thus starts close to the non-equilibrium steady-state distribution (reached after some number of τ iterations), where the small weight entering the flux bin is continuously removed and returned to the remaining bins.

In some scenario, the evacuated state is not rare, and therefore the steady-state distribution is not well-approximated by the equilibrium distribution assuming no evacuation. Specifically, this occurs for simple diffusion when $\rho \leq 1$, at the left of Fig. 2.3A, and for close-contact simulations where $p_{\text{entry}} \approx 0$ (Fig. 2.6). In this latter scenario the steady state is the completely evacuated state, and we are seeking to compute a different transition time: from the uniform steady state (as if $p_{\text{entry}} = 1$) to the evacuated state. In these scenarios, a second reweighting method was used. This is described schematically in Fig 2.7B. In these method, weight from an evacuated replica is not redistributed to remaining replicas. Rather, each time a replica evacuates, a new replica is initialized as described above and given all of the weight from the evacuating replica. This ensures that the weight distribution throughout the state space remains statistically accurate, even when the flux of weight throughout the space is unidirectional.

Hyperparameters

The above-described Weighted Ensemble method requires the specification of hyperparameters τ , m_{targ} , and the max number of iterations. In principle, the selection of these hyperparameters should not impact the results of the WE simulation, but will impact the efficiency of convergence

to an accurate MFPT.

In an effort to maximize the observed number of flux events, m_{targ} was chosen to be high to maximize the number of replicas in bins nearby the flux bin (transient bins), but not higher than allowed by computer memory limitations. The values were $m_{\text{targ}} = 100$ for Fig. 2.3, Fig. 2.8, and 200 for Fig. 2.4, Fig. 2.5.

The WE step was chosen to be $\tau = 50\Delta t$, which we found to be large enough to give replicas time to change bins before the splitting and merging process began, while also avoiding being too long to ensure a high number of replicas inside the transient bins.

Implementation

Smoldyn simulations were executed through Smoldyn’s C library, libsmoldyndyn [5]. Combination of Smoldyn with Weighted Ensemble was written in C and is available to the public at

<https://github.com/dydtaylor/libsmoldynWE>

. Execution of libsmoldynWE simulations was done on UCI’s high performance cluster. Brute force Smoldyn simulations were executed in libsmoldyndyn, but outside of the libsmoldynWE weighted ensemble framework. All data was analyzed in MATLAB. Evaluation of analytical solution for $p_{\text{entry}} = 1$ was done in Wolfram Mathematica.

Analysis

To allot for burn-in time, i.e., an initial transient while the replicas approach a steady state, the first half of each run is discarded, and the mean flux measured in the second half of the run is taken to be a single measurement of the mean flux $\bar{\phi}$, averaged across WE steps. Each WE data point in this paper is calculated from 10 independent WE simulations; multiple independent runs are used

to minimize the effect of spurious correlations between iterations [104], which can cause inter-run variability in estimates (e.g., as seen in Fig. 2.2C). These 10 repeats are then arithmetically averaged to give an estimate of the mean flux, averaged across repeats, $\langle \bar{\phi} \rangle$, with error bars given by the standard error of the mean for the 10 WE simulations.

The average flux recorded from replicas reaching the flux bin, $\bar{\phi}$ is used to calculate the MFPT from the Hill relation [49, 12], $T = 1/\bar{\phi}$. The error bars of the flux $\delta\phi$ then propagate to MFPT error bars by $\delta T = |T|^2 \delta\phi$.

To calculate MFPT from brute force simulations, the arithmetic mean of 1000 (Fig. 2.3) or 500 (Fig. 1.1B) independent repeats was used, with error bars giving the standard error of the mean for those simulations. To calculate the single molecule first passage time distribution for Fig. 1.1A, an empirical CDF was created from the result of 20,000 brute force evacuations of a single molecule placed uniformly within the ROI.

To calculate monomer fractions from brute force simulations, purely monomeric Smoldyn simulations are initialized and run for 15 units of time. Afterwards, the monomer fraction we report is average measured over these last 5 time units.

Method validation

Several methods were used to validate the WE results. For MFPTs where achieving brute force results was computationally viable, brute force results were included along with WE results (see Figs. 2.3, 2.6C). Monomer and dimer fractions were verified with brute force Smoldyn simulations (Fig. 2.9). For each binding radius presented, the sigmoidal monomer fraction vs unbinding rate curve was executed for a range of timesteps to verify the timestep of $\Delta t = 10^{-6}$ was small enough.

Close-contact WE simulations were found to agree with brute force simulations for low p_{entry} (Fig. 2.6C). For $p_{\text{entry}} = 1$ we verified the endpoint with our simple diffusion WE simulations for $\rho = 9$.

An analytical solution for close-contact simulations when $p_{\text{entry}} = 0$ was obtained, see Supplemental. The analytic solution for a single molecule's first passage time distribution was compared with an empirical CDF obtained from 20000 brute force Smoldyn simulations (Fig. 1.1A) and the analytical form for the MFPT for a variety of evacuating molecules was compared with the results from brute force Smoldyn simulations, 500 for each data point (Fig. 1.1B). The analytic solution for the parameters used in the close-contact WE simulations is included in Fig. 2.6C.

An asymptotic solution for the MFPT for homogeneous monomeric solutions was obtained from previous work [74]. When applicable, these asymptotics were used to verify WE results (Figs. 2.3A, 2.8). However, as can be seen in figures 2.3, 2.8, agreement with the asymptotics at higher densities is dependent on the size of the domain used. For a range of densities, we did a sweep of domain sizes (Fig. 2.8). Weighted Ensemble estimates reach to within one order of magnitude of the asymptotics above $\approx L = 5$ at the highest simulated densities.

Convergence in time-step was done by comparing MFPT estimates for larger time-steps with the chosen time-step to ensure MFPT estimates did not undergo drastic differences in MFPT estimates. The time step of $\Delta t = 10^{-6}$, used for all simulations unless otherwise stated, is compared with with the timestep of $\Delta t = 5 \times 10^{-6}$ is shown in supplemental figure 2.8B-D.

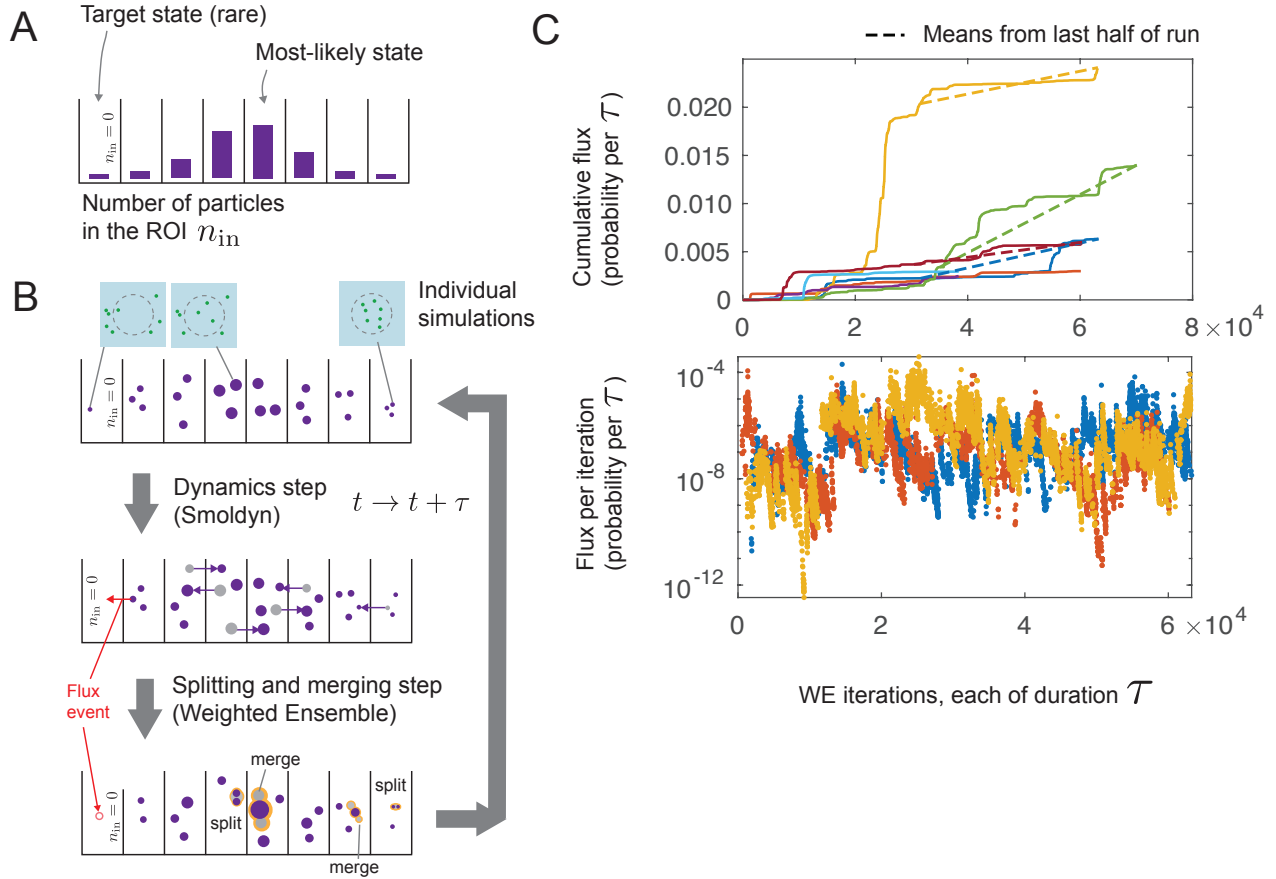


Figure 2.2: Weighted Ensemble algorithm with Smoldyn spatial dynamics simulations. (A) Partitioning of simulation state-space into bins based on the number of molecules in the ROI, $0 < n_{\text{in}} < N$, where N is the total number of molecules. Bar heights represent the relative probability of the system to be found in the corresponding bin at equilibrium. In the simplest case in which molecules experience diffusion-only dynamics, this is a binomial distribution. (B) Description of the algorithm. Individual simulations (replicas) are shown as purple circles, and have statistical weights that can vary from replica to replica, represented here through circle size. Simulations are allowed to propagate according to dynamics simulated by the dynamics engine Smoldyn [7] for a period of time τ . After the dynamics step has completed, the replicas are reexamined to see their new bin locations. The number of replicas in each bin is then compared to m_{targ} , the desired number of replicas in each bin. Bins with more than m_{targ} replicas have a “merging” event, where the replicas with the smallest individual weights are removed and their weight is redistributed to another replica within the same bin. Bins with fewer than m_{targ} (but still > 0) replicas have “splitting” events where the replicas with the most weight are duplicated into 2 daughter replicas, with the weight from the parent being redistributed equally to the daughters. Flux events, representing complete evacuations of the ROI, (red) have their replicas deleted and weight redistributed to replicas outside of the flux bin. For details on this redistribution of weights from flux events, see 2.7. (C) Flux measurements from multiple, independent runs displayed in two different ways: total flux accumulated (top, y-axis scaled linearly) and flux accumulated per WE iteration (bottom, y-axis scaled logarithmically; fewer runs shown for clarity). For each WE run, the measurements in the first half of the run are discarded to exclude the fluxes that might be measured during weight redistribution between the initial simulation state and the simulation state after many WE steps have passed. The mean fluxes measured during this period from multiple independent runs (slopes of dashed lines) are then used to estimate the mean first-passage time to the evacuated state. Cumulative fluxes for 7 runs are shown in the top figure, with the time series for 3 of those runs being shown in the bottom plot. Since the bottom plot is on a logarithmic scale, WE iterations where no flux was measured (in this case, $\sim 90\%$ of iterations) do not appear.




Figure 2.3: Mean evacuation time from the ROI for molecules only experiencing diffusion shows evacuation times well outside the experimentally observed timescales at physiological densities. Despite this drastic increase in evacuation times, use of Weighted Ensemble (WE) allows calculation on computationally-feasible timescales. (A) Evacuation times for densities ranging up to $\rho = 10$ molecules per R_{ROI}^2 . Estimate of physiological density of CD45 assuming an ROI of $\sim 100\text{nm}$ in radius and $\sim 9 \times 10^{-4}$ molecules per nm^2 (see main text) marked with a vertical line. Evacuation times from WE for domain size $L = 5$ (blue) agree with brute force simulations (red) at low ρ and with an asymptotic (infinite domain) approximation at low to intermediate ρ [74]. (The asymptotics overestimate evacuation times for both methods at very low ρ). Each WE data point uses 10 independent runs. Error bars for WE runs are computed by taking the average flux measured from the second half of each WE run as a single measurement, computing the 95% confidence interval of these flux measurements, and propagating these errors. Brute force error bars come from the standard error of the mean for 1000 runs. (B) Total computational costs for the data points in (A). WE data points come from summing the computation time of 10 independent WE simulations, with $m_{\text{targ}} = 100$ replicas per bin. WE runs were programmed to run until flux measurements from the final 1/3 of WE iterations and the middle 1/3 of WE met the following criteria: (i) Non-zero flux measurements in each third numbered at least 500 (ii) The KS-statistic between the two thirds reached a value of 0.02 or lower (iii) The KS-statistic between the two thirds with zeros removed reached a value of 0.3 or lower. Oftentimes, one or more of these requirements would not be met prior to reaching computation limits of an individual run, and in that case WE simulation would automatically stop after 2.5 days of computational time spent on WE and Smoldyn dynamics. Brute force data points come from summing the computation time of 1000 independent brute force evacuation events. WE simulations used reweighting method 2 for $\rho \leq 1$ and reweighting method 1 for $\rho > 1$ (see supplemental Fig. 2.7). Parameters were $\tau = 50\Delta t$, $m_{\text{targ}} = 100$, $\Delta t = 10^{-6}$.




Figure 2.4: Dimerization of evacuating molecules causes several-order reduction in evacuation time (MFPT). (A) Right: Schematic of dimerization model used in Smoldyn. Molecules, diffusing at a rate D_1 , that get within a certain distance, $r_{\text{bind}} = 0.001$, of each other combine into a dimer that diffuses at a rate of $D_2 = D_1/2$. These dimers have an unbinding rate, k_{off} , which varies across simulations. Left: Reduction in unbinding rate k_{off} (moving right-to-left on axes) leads to several-order reduction in evacuation time. At low dimerization fraction (high k_{off}), evacuation time is the same as for non-interacting diffusing molecules with $\rho = 9$, $D = D_1$, shown by the upper dashed line. At high dimerization fraction (low k_{off}), evacuation time is the same as for non-interacting particles but with half the density and half the diffusion coefficient, i.e., $\rho = 4.5$, $D = D_2$. (B) Average monomer fraction for the simulations given in (A) vs k_{off} . (C) Same data as (A) and (B) showing the relationship between monomer fraction and evacuation time. A 10-fold reduction in evacuation time (vertical line in (A) and (C)) occurs when only 20% of the individual molecules are in a dimer, suggesting even weak binding can have drastic impacts on evacuation events.




Figure 2.5: Formation of higher-order oligomers further reduces evacuation time, including times below ~ 1 second for hexamers. Assuming high-affinity (e.g., low k_{off}) oligomerization would lead to homogenous ensemble of oligomers of a particular size (horizontal axis). Oligomer diffusion coefficients were taken to scale reciprocally with the number of subunits $D_n = D_1/n$ [30]. Intermediate affinity oligomerization (red) would lead to heterogeneous monomeric and dimeric mixtures. We compute these intermediate evacuation times from dimerization data in Fig. 2.4 and Eq. 2.2.

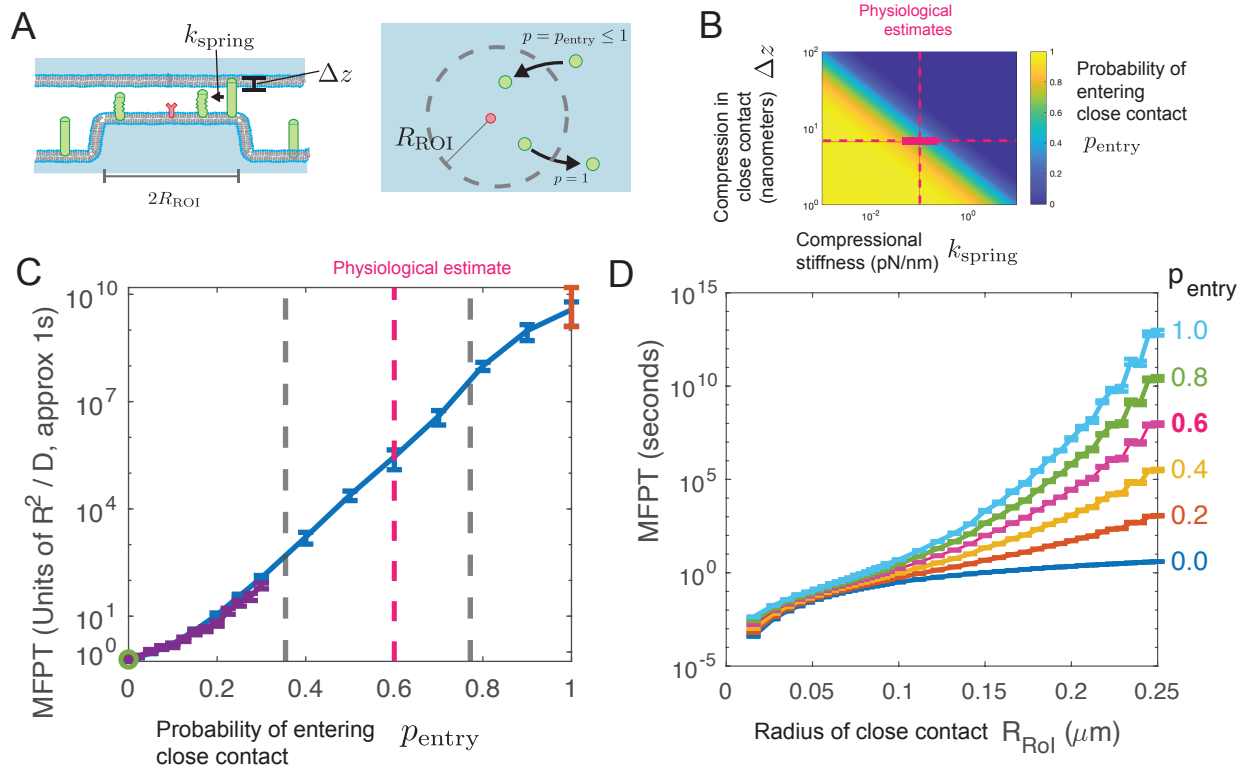


Figure 2.6: Formation of close contacts creates an energetic barrier for entry into the ROI. Even moderate energy barriers can create order of magnitude reduction in evacuation time at physiological densities. (A) Schematic of close contacts and how we chose to model them in Smoldyn. Close contacts, such as those from microvilli, can cause compression of large transmembrane particles such as CD45. The energy associated with this compression is modeled as a compressed spring $E_{\text{spring}} = 1/2k\Delta z^2$, where k is the compressional spring constant, and Δz is the distance of compression. This leads to dynamics in which movement into the ROI is reduced. We model the reduction as a probability p_{entry} from the thermodynamic relation Eq. 2.3. (B) Heatmap of p_{entry} for a variety of compressional stiffnesses k_{spring} and the compression size Δz . Pink dashed lines represent physiological estimates of both. Solid pink rectangle represents two-fold increase and decrease of k_{spring} . (C) Evacuation time (MFPT) versus p_{entry} , with the physiological estimate for p_{entry} given by the vertical dashed, pink line. Included are results from WE simulations (blue) brute force simulations (purple). At $p_{\text{entry}} = 1$, we include the diffusion-only simulation result from Fig. 2.3 (red). At $p_{\text{entry}} = 0$, we compute an analytic expression for the MFPT in Eq. A.29 and Fig. 1.1, shown here in green. The inset shows, with a linear y-axis, close agreement between brute force and WE for $p_{\text{entry}} < 0.1$. For WE runs, an alternative method of redistributing weight in the flux bin is used, see 2.7. Gray dashed lines show p_{entry} and MFPT if the spring constant were increased or decreased by two-fold. (D) Evacuation time versus radius of close contact, R_{ROI} , re-scaled to physical units (seconds, microns) to show impact of physical parameters, assuming the ansatz Eq. 2.4. Here we use lower density of CD45 of $160\mu\text{m}^{-2}$ (lower than central estimate used elsewhere) to achieve simulation at larger ROIs.

Supplemental material

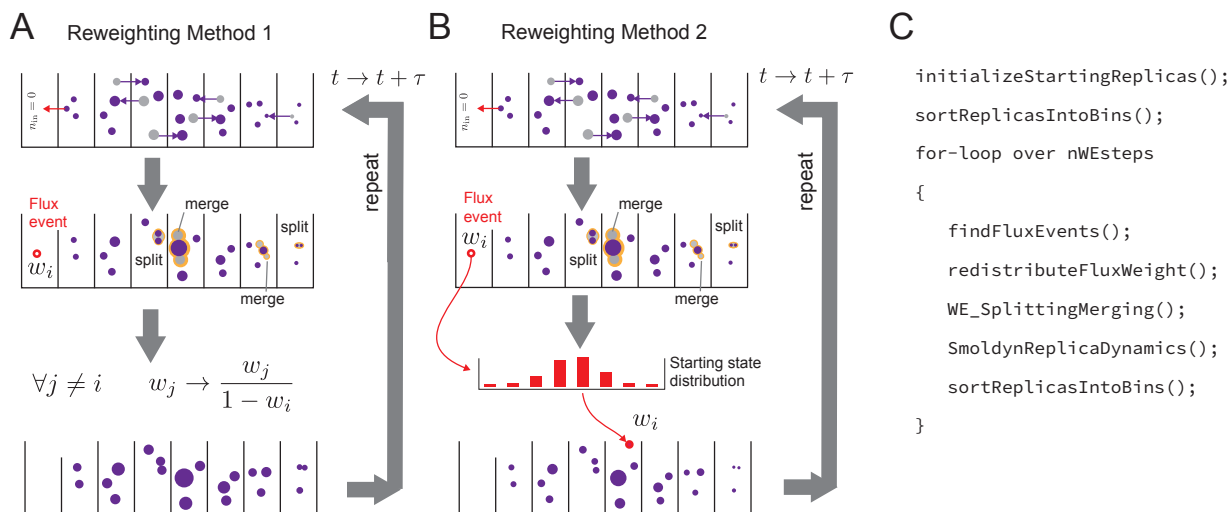


Figure 2.7: Different reweighting methods used for WE simulations. Method A is used for figures 2.2C, 2.3, 2.4, and 2.5. Method B is used for the WE runs in 2.6C and D. (A) Upon occurrence of a flux event, the weight from the flux event, w_i is returned to the simulation by scaling the weight of the remaining simulations by their remaining weight, $w_j \rightarrow w_j/1 - w_i$. Note how weight distribution changes if replica movement between the bins is limited, e.g. in figure 2.6 with $p_{\text{entry}} = 0$. As bins are based on the number of molecules inside the ROI, no weight is allowed to re-enter a bin after it leaves, which has a significant impact on steady-state weight distribution. (B) Upon occurrence of a flux event, the weight is given to a newly initialized replica drawn from the starting state distribution (all molecules are placed randomly and uniformly throughout the domain). (C) Pseudo-code of WE algorithm described in 2.2, including reweighting method “redistributeFluxWeight()”.

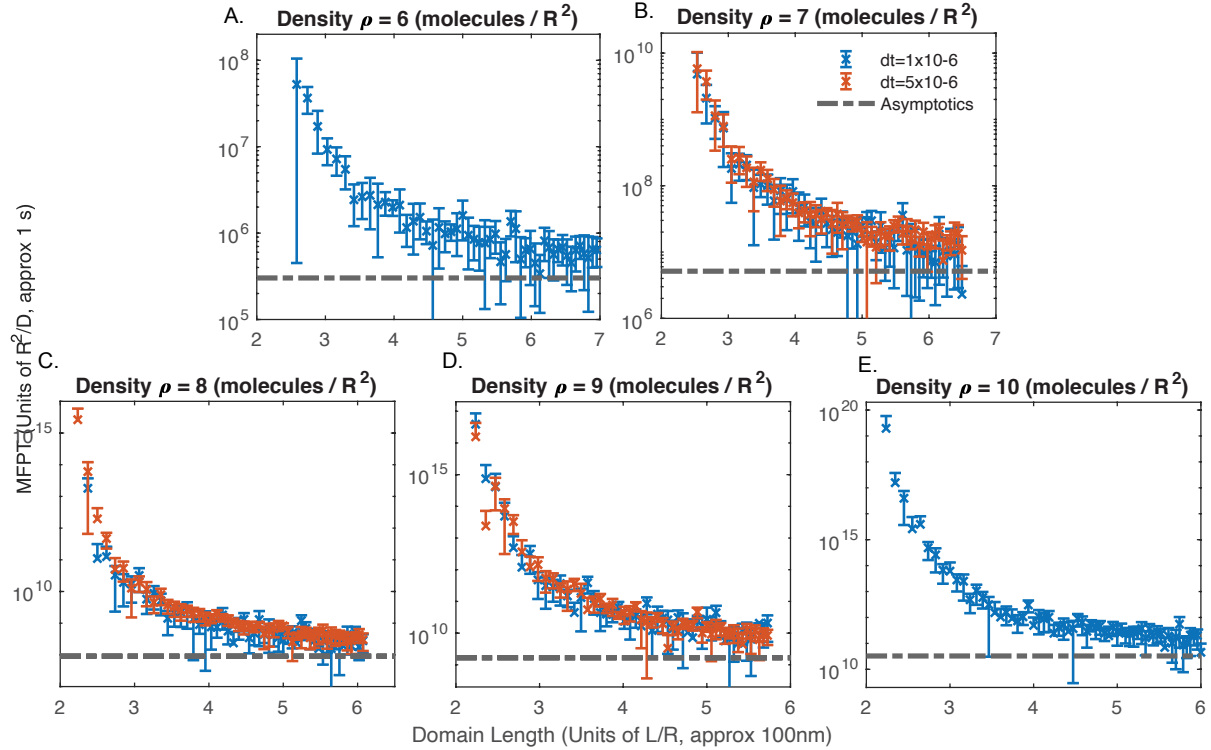


Figure 2.8: Evacuation time (MFPT) dependence on domain size L for given densities. (A-E) MFPT versus L for increasing values of density ρ . Horizontal dashed lines give the value of the infinite-domain limit from [74]. Weighted Ensemble estimates of evacuation time approach the asymptotic estimate as L increases. At density $\rho = 9$ (D), the Weighted Ensemble estimate reaches to within about one order of magnitude of the asymptotic near domain size $L = 5$. In light of this, dimerization simulations in the main text were chosen to have a domain size of $L = 5.333R_{\text{ROI}}$ (leading to total number of monomeric subunits $N = 256$ at $\rho = 9$). To investigate impact of time step on convergence, time steps of $\Delta t = 10^{-6}$ and $\Delta t = 5 \times 10^{-6}$ were used for panels B,C and D. For all runs included, WE parameters were $\tau = 50\Delta t$, $m_{\text{targ}} = 100$. Each data point had 10 WE runs, with error bars representing standard error of the mean as calculated in Methods. Error bars were calculated from the 95% confidence interval of the mean flux, and propagating this error to the MFPT.

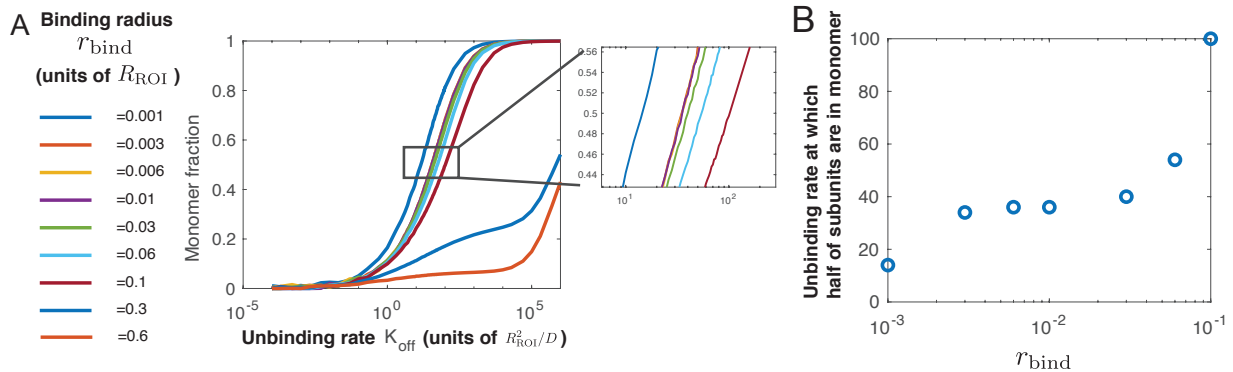


Figure 2.9: Particle-based reaction diffusion simulations of reversible binding kinetics on a 2D domain. Simulations were done in Smoldyn with 256 monomeric subunits in a square domain with side length $L = 5.333$ and run until reaching a steady state fraction of monomers. Each curve was independently found to be insensitive to decreasing the simulation time-step Δt . (A) Monomer fraction at steady state for a range of k_{off} and r_{bind} . (B) The effective dissociation constant, i.e., the unbinding rate k_{off} for which 50% of the monomeric subunits are in the monomeric form, as a function of binding radius. For binding radii close to the estimated physiological binding radius of $r_{\text{bind}} = 10^{-2}$ (corresponding to roughly 1nm), the steady-state monomer we see weak insensitivity between the monomer fraction and the binding radius, as previously reported, due to the complicated logarithmic nature of two-dimensional diffusion [101, 60]. At binding radii larger than 0.3, the system approaches the crowding regime, resulting in large impact to monomer fraction.

Table 2.1: Table of model variables, model parameters, Weighted Ensemble and simulation hyperparameters. For discussion of values of model parameters, see Results and Methods.

Symbol	Description	Value unless specified
Model parameters		
R_{ROI}	Radius of region of interest	100 nm
D	Diffusion coefficient of single-subunit molecule	$0.01 \mu\text{m}^2/\text{s}$
ρ	Surface density of molecule (units of molecules per R_{ROI}^2)	$9R_{\text{ROI}}^{-2}$
$\tilde{\rho}$	Surface density of molecule (units of molecules per μm^2)	$900\mu\text{m}^{-2}$ (note exception in Fig. 2.6D)
\bar{n}_{in}	Mean number of molecules in the ROI assuming uniform	$\bar{n}_{\text{in}} = \pi\rho \approx 30$
Model variables		
n_{in}	Number of particles in ROI	
ϕ	Rate of flux into the evacuated state	
T	Mean time to the evacuated state (units of R_{ROI}^2/D)	
\tilde{T}	Mean time to the evacuated state (units of seconds)	
Simulation hyperparameters		
Δt	Dynamics timestep	$10^{-6}R_{\text{ROI}}^2/D \approx 10^{-6}\text{s}$
N	Total number of particles	$N = L^2\rho = 256$
L	Domain size	$5.333R_{\text{ROI}} \approx 533 \text{ nm}$
Weighted Ensemble hyperparameters		
τ	Time interval between WE merging/splitting	$50\Delta t \approx 5 \times 10^{-5}\text{s}$
m_{targ}	Number of replicas per bin	50-200
(No symbol)	Number of repeated WE runs	10

Chapter 3

Quantifying Uncertainty in Weighted Ensemble Simulations and Optimizing Weighted-Ensemble Specific Parameters

Introduction

Error Quantification

While much work has been done making use of Weighted Ensemble since its first introduction by Huber and Kim in 1995 [51, 84, 81, 102, 103, 105, 32, 104], including work on binning methods [1], there is an ambiguity with error quantification. WE simulations are useful for estimating the order of magnitude of quantities [79, 13]. Many WE simulations produce high log-variance data sets [13], and error bars are typically large relative to the quantity measured. A consequence of this is the production of unphysical confidence intervals using standard uncertainty methods (i.e. standard error of the mean) [72]. In the context of the prior chapter on cell surface MFPTs, these

can produce negative lower-bounds on the confidence intervals.

Many of these properties of high log-variance data sets have been studied before, see, e.g. [13, 45, 72, 1]. We approach this uncertainty from two different directions: within-run variation and between-run variation. The former comes from looking at the flux trace, weight distribution, and other quantities from a single WE simulation, with the hope being that properties from these can be used to make determinations of uncertainty. The latter involves using multiple independent and identically distributed (iid) WE simulations to estimate mean and standard error of the mean; it is a somewhat crude approach, as improving uncertainty estimated from multiple iid WE simulations is best done by simply throwing more computational power at the problem, and requires little insight to create uncertainty estimates regardless of the specifics of the WE simulations.

While each replica inside a WE simulation has internal dynamics that function independently from each other, the creation of daughter replicas through the splitting process necessitates that there are correlations between successive measurements. As such, each successive flux measurements from a single WE simulation are not statistically independent [104, 13, 45]. While treating each successive measurement as independent is inappropriate, if one can determine the correlation time of the flux measurements then one could split an individual WE simulation into multiple independent measurements, helping to improve the precision of the simulations. However, this correlation time cannot always be easily determined.

In a WE simulation, if one is interested in measuring a MFPT for a given process, one needs to define a target state, as well as the bin weight distribution associated with starting state. While the initial configuration of the starting state may be arbitrary, the weight distribution must be allowed to converge to a non-equilibrium steady state in order to obtain accurate flux measurements. As such, it makes sense to discard portions of the WE dataset that are made prior to reaching this non-equilibrium steady state. This is consistent with practices used for similar molecular simulations, even those not implementing WE, see [45]. Determining how much of the data should be discarded depends on being able to distinguish the steady-state weight distribution from the initial weight

distribution. It is worth noting the general difficulty with determining convergence for a stochastic system when the true value of the measurements are not known a priori [24]. Additionally, several of the flux traces we observed in the previous chapter had distinguishable trends and clear visible correlations, see Fig. 3.4. As a consequence, we began investigating potential metrics that can be used to evaluate how well a WE simulation has converged with itself as well as accuracy.

Investigating quantitative differences between Weighted Ensemble metaparameters.

In addition to simulation parameters associated with the physical dynamics of a system, e.g. diffusion coefficient, chemical binding rates, etc., WE simulations have specific metaparameters associated with the execution of the WE process. The metaparameters of primary focus in this paper are τ , the period of WE splitting / merging, and m_{targ} , the number of replicas maintained in each WE bin in the splitting and merging process, the number of WE iterations to perform, and lastly the bin definitions for a WE simulation. For the system we studied in chapter 2, our ability to create more bins is limited by the discrete nature of our order parameter; we will focus primarily on τ and m_{targ} , while the number of WE iterations is dependent on computational resources.

Choosing an appropriate τ and m_{targ} is “somewhat of an artform” [13]. The field would benefit from a method of determining an ideal set of metaparameters, and some recent progress has been made in this area [9]. Our goal is to see if there were clear indications of certain metaparameter sets performing measurably better than others and to discover properties that can predict simulation accuracy.

Measuring and Understanding Uncertainty in WE

Before discussing uncertainty quantification in WE, it is worth first noting two different ways to use sets of WE simulations to calculate a first passage time. One approach involves treating each WE simulation as a measurement of the average flux, using the full set of WE simulations to get a combined measurement of the average flux, then using that flux estimate to calculate the first passage time. Another approach involves treating each WE simulation as a measurement of the MFPT, then averaging the individual MFPT estimates to calculate the first passage time. For a set of n WE simulations, each WE simulation can be used to get an estimate of the flux ϕ_i or of the MFPT $T_i = \frac{1}{\phi_i}$. The averaging methods therefore amount to:

$$T_{\text{arithmetic}} = \frac{1}{\bar{\phi}} = \left(\frac{1}{n} \sum_i (\phi_i) \right)^{-1} \quad (3.1)$$

$$T_{\text{harmonic}} = \bar{T} = \frac{1}{n} \sum_i T_i = \frac{1}{n} \sum_i \left(\frac{1}{\phi_i} \right). \quad (3.2)$$

Where the subscripts indicate the method of averaging the fluxes. Note, the harmonic mean of $1/x$ is equal to the reciprocal of the arithmetic mean of x , thus $T_{\text{arithmetic}}$ corresponds to using the arithmetic mean of the fluxes and the harmonic mean of the T_i , while T_{harmonic} , which involves taking the arithmetic mean of the T_i , corresponds to using the harmonic mean of the fluxes. It is worth noting that the harmonic mean is always less than or equal to the arithmetic mean, which implies that $T_{\text{arithmetic}} \leq T_{\text{harmonic}}$ for our reciprocals of these means. Indeed, this is confirmed in Fig. 3.1 A, B.

Moving on to error quantification, we note multiple ways of calculating errors from sets of WE simulations. Firstly, for a set of n iid WE simulations, one can use the n WE simulations to calculate a standard error of the mean, which was done to create the error bars in Fig. 3.1 A, B. Additionally, one can consider variance about a trend that is known a priori, or the “wobble”

about the known behavior. For instance, in Fig. 3.1B, the large-domain behavior of the simulations implies an eventual horizontal trend. If we make the generous assumption that we have reached convergence at $L > 6$, we can calculate the SEM amongst the T_L in that region. In the figure, this corresponds to a SEM of 16 data points, each one corresponding to a different L . This wobble-implied error is included in Fig. 3.1C as horizontal grey lines. Furthermore, one can instead consider that each of the 10 WE simulations that goes into the 16 data points as a measure of this horizontal asymptote, and instead calculate the SEM of the conglomerate 160 WE simulations. As can be seen in Fig. 3.1 C, these values are similar to the wobble-implied error but clearly distinct from it. Also included are the error bars from each of the 16 data points. It is comforting to note that these error bars are all within an order of magnitude from the other two errors, suggesting that despite the method chosen to calculate the error bars, the broader impact on conclusions made from WE simulations is minor.

The Search for Key Metrics in Weighted Ensemble

In all real uses the true values and trends of quantities estimated from WE are not known, and the only way to come up with more accurate estimates is to throw more computation time at the problem. There are many different ways of doing so; one can make individual WE simulations more taxing or run them for longer, or one can simply run more independent WE simulations. The ambiguity in the nature of how many data points can be collected from an individual WE run has inspired our search for specific metrics that might be used to understand when a WE simulation has converged to its steady-state. In our search for these metrics, we investigated statistical properties of flux traces to see if there were any trends between these metrics and accuracies of the measurements. In the following tests, we assumed that we knew the true value of the answer from an asymptotic value given by [74], and then attempted to find trends between our metrics of interest and the accuracies of the simulations. In particular, we investigated Kolmogorov-Smirnov (KS) test statistics determining convergence of WE simulations within-run, our own modified KS




Figure 3.1: A. WE Evacuations for a particle density of 6 molecules per unit area vs total domain size. Each data point has 10 identical WE simulations averaged together, with the two different curves being associated with how the average was taken. Error bars are obtained from the SEM of the 10 WE simulations. The expected behavior of this curve is a smooth decay towards the asymptotic value for this MFPT [74]. However, sometimes nearby points will display behavior contrary to the general trend. For each data point, note that the MFPT estimate from the arithmetic mean of the fluxes lies below the equivalent estimate from the harmonic mean of the fluxes. B. The latter 16 data points from A. At larger L , the curve takes a generally flat shape, suggesting the measurements should be for the same MFPT. In spite of that, we still see a “wobble” of the curve, and in this regime you can see pairs of points that do not have overlapping error bars. This wobble suggests an error distinct from the SEM measured for each data point. C. For the final 16 data points, we compare the magnitudes of errors calculated from 3 different methods: the error suggested from the wobble (SEM of the 16 data points plotted in B), the SEM from the 10 runs of an individual data point (error bars from A, B), and the SEM from combining all 160 runs together, with error defined as in 3.3. We note that the wobble error and the 160 run total error are similar in magnitudes in both cases.

statistic (see below and Fig. 3.2), the fraction of non-zero measurements in our flux trace, and within run variance.

Because WE is best used for order of magnitude estimates, we choose to measure accuracy in

terms of logarithmic differences.

$$\log a = |\log \bar{\phi} - \log G| \tag{3.3}$$

Where $\log a$ is our logarithmic accuracy estimator, $\bar{\phi}$ is the average measured flux, and G is the average flux predicted from the asymptotics given by [74]. Using this quantity, a value of x indicates agreement to within $x - 1$ orders of magnitude between simulation and asymptotics, e.g. a value of 1 indicates complete agreement.

Our first investigation was of the KS statistic on accuracy measurements, with the logic being that once steady-state has been achieved, flux values will be sampled from the same probability distribution. We therefore calculated a KS statistic for each WE simulation by comparing the distribution of flux values in the third quarter of the run with the flux values in the fourth quarter of the run. However, as is shown in Fig. 3.3 A, oftentimes there are a large number of null flux measurements, resulting in KS values that can become very small very quickly regardless of behavior of the non-null values. Our solution was to create a dual-distribution KS statistic (DKS),

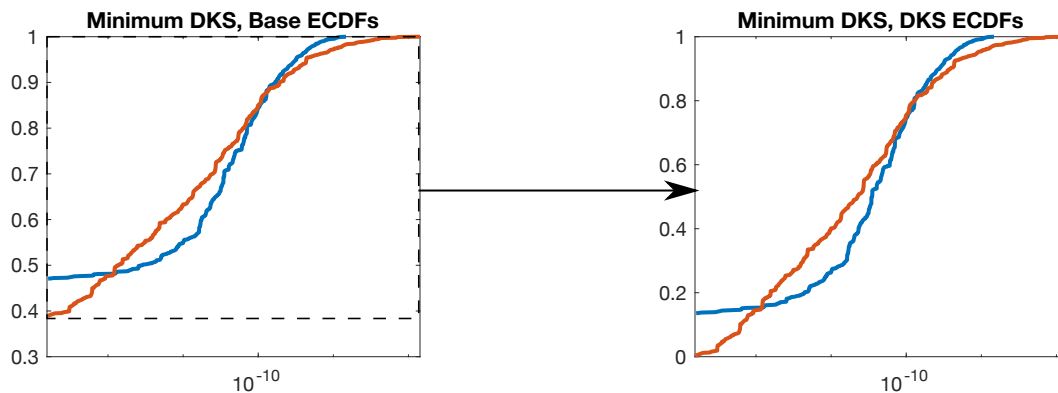


Figure 3.2: Visual representation of differences between the ECDFs used for the KS statistic and the DKS statistic. For the DKS statistic, zeros in both of the KS distributions are removed, and a new KS statistic is calculated from the new DKS distributions.

which we visually represent in Fig. 3.2, with the intention being to compare the non-null values of the fluxes. To do so, we counted the number of null measurements in both the third quarter and fourth quarter of data. Null measurements that were duplicated in both quarters were removed




Figure 3.3: Candidate metrics of flux measurements do not reveal a correlation with estimated accuracy. A. Comparing 3rd quarter and 4th quarter ECDFs of the flux measurements (left, showing the minimum recorded KS statistic) reveals that even fairly large KS statistics can have lower errors than even some of the smallest KS statistics, suggesting the statistic is inappropriate for an accuracy predictor. In general, a high frequency of zeros in the flux measurements can cause this statistic to get very small. B. The Dual distribution KS statistic, one comparing 3rd and 4th quarter ECDFs with zeroes contained in both distributions removed. Similar to the KS statistic, the scatter plot shows that this DKS statistic is not indicative of achieving a certain error magnitude. C. A flux time series with the largest fraction of non-zero measurements. A high fraction of non-zero measurements does seem to be indicative of lower errors, but in general a high fraction of zero measurements doesn't preclude a high accuracy. D. In run variation of flux from average flux. Blue is the cumulative flux measured from a WE simulation, while red is the average cumulative flux measurement from the run. A high value indicates large deviations between the in-run behavior, something associated with large weight distribution deviations, which might be associated with the weight distribution significantly varying from the non-equilibrium steady-state distribution. The scatter plot, however, reveals no trend between this statistic and overall accuracy.

completely, while a quarter with more null measurements would retain the extras. These two dual-distributions, with the nulls removed, then had a KS test performed on them, resulting in the DKS

statistic. However, similar to the KS statistic, investigating trends between this DKS statistic and the error (Fig. 3.3B) again showed no notable trends; in fact, some of the lowest error values were obtained even for DKS values of nearly 1.

Our next investigation was in the fractional number of non-zero flux measurements (fractional NNZ). Prior to performing a detailed investigation, we understood that a high fractional NNZ was not exclusive to accurate simulations; we had observed very low NNZ fractions giving some of the more accurate results. Nevertheless, the scatter plot in Fig. 3.3 C shows some of the strongest trends of the four scatter plots included. While a low NNZ fraction was still capable of giving accurate measurements, the largest errors were eliminated as NNZ fraction increased, suggesting that NNZ fraction is indicative of a good WE setup where weight is allowed to travel more freely.

Lastly, we created a metric for within-run variation from the run's average values. The total flux measured at any one WE step is compared to the expected flux measured at that time step (given that the total amount of flux measured is already known), with the logic being that large variations between the cumulative flux measured and the expected flux measured is indicative of large chunks of weight traveling throughout the system rather than an idealized steady, continuous weight at each time step. However, the scatter plot of this metric shown in Fig. 3.3D again shows a lack of any significant trends.

The Sawtooth Effect

One key feature observed in WE simulations, albeit hard to quantify, was sawtooth-shaped flux traces, see Figs. 3.4, 3.5. This appears to be a general feature of the WE algorithm, e.g., see Bogetti et al. [13]. Though never directly observed, our hypothesis is that these sawtooth traces are the result of large amounts of weight entering into a bin in one step followed by slow decay as the weight in the bin propagates outside of the bin, with each spike being associated with weight from a new replica first entering the bin. To investigate this, flux time series were compared




Figure 3.4: Weight time-series for a variety of bins. A. The bin weight vs WE step for the same WE simulation. Fine-grained bins of the WE simulation are bins 0 through 8, with bin 9 being the first of 4 coarse grained bins. While the coarse grained bin doesn't show significant correlations between it and the weight of the fine-grained bins, the fine-grained bins show distinctive saw-tooth behavior that propagates through all of them, though some are lost as the distance from the coarse-grained bins increases. Note, weight is scaled by T^*/τ , hence values > 1 . B. Bin weight ratios between several fine-grained bins and the coarse-grained bin with its fine-grained neighbor. Note in the fine-grained bins these ratios are mostly flat, though there are spikes that appear to correlate with the beginning of a new saw-tooth.

with bin weight time series of neighboring bins, see Fig 3.4 A. Our weighted ensemble bins were divided such that bins 0-8 were much more finely-grained than the remaining bins of the system. A side-by-side comparison revealed that these sawtooth patterns existed in each of the fine-grained bins, though detail is lost in bins furthest away from the coarse-grained bins. This suggests that the weight in these bins is much less dependent on how the weight propagates between the fine-grained bins and more dependent on how weights get introduced into the bins from the coarse-grained bins.

To investigate this further, time series of the ratio of neighboring bin weights were created. Interestingly, the fine-grained bin ratios appeared to somewhat frequently get restored to a constant value, with the break from these constant values being spikes that temporally align to the start of a new sawtooth. This suggests that amongst the fine-grained bins, weight propagates very quickly throughout them, though not instantaneously. Furthermore, the ratio between a fine-grained bin and its coarse grained neighbor is continuously increasing aside from spikes downwards associated with the start of a sawtooth. This suggests that in fact, the coarse grained bins are ultimately the main drivers of flux measurements. A future direction that has recently been investigated by [9] would be in creating a set of different m_{targ} values for each bin, with few replicas existing in the coarsest grained bins and a very large m_{targ} in the coarse grained bins neighboring the fine-grained

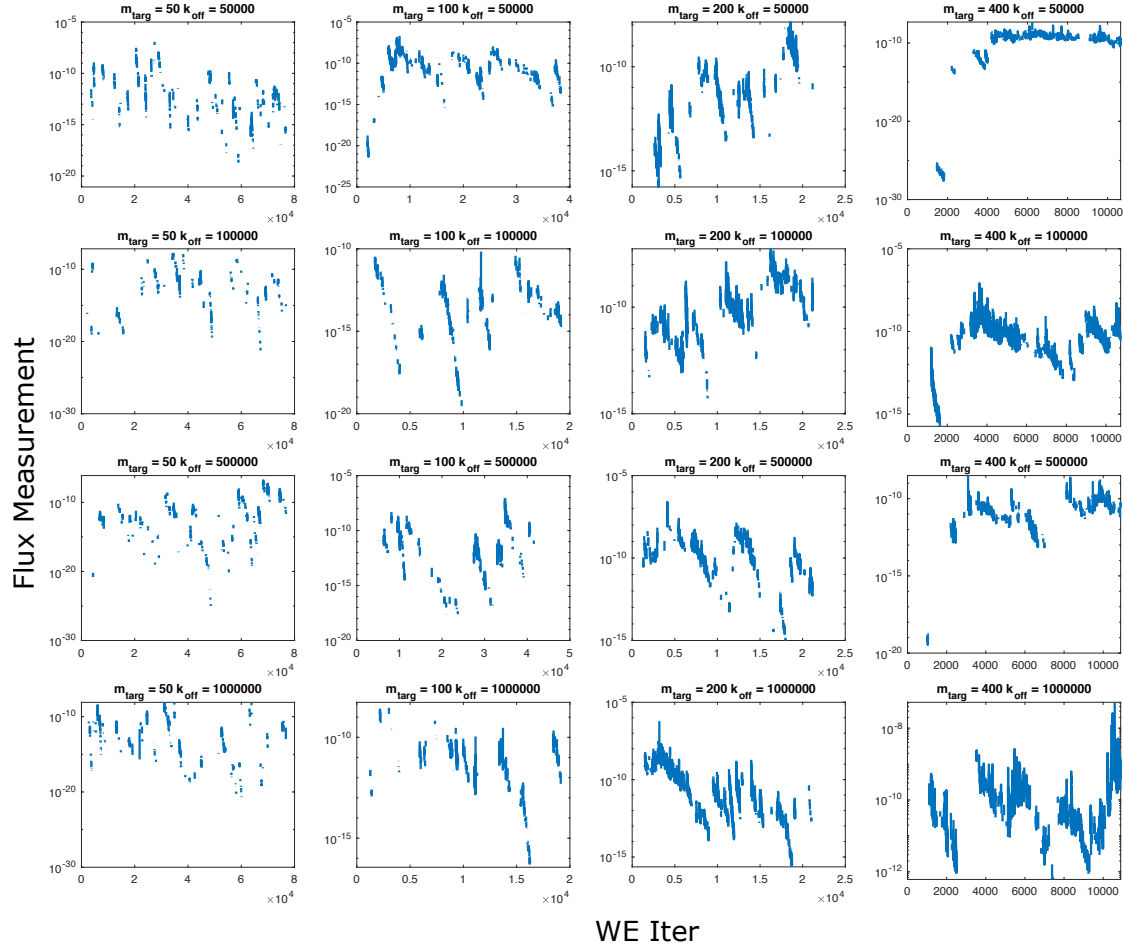


Figure 3.5: Flux time series for a variety of dynamics parameters k_{off} and metaparameters m_{targ} . Higher m_{targ} is correlated with fewer zeros and less dramatic sawtooth behavior.

bins.

Metaparameter Dependence

After investigating the properties of many flux traces, we moved on to running identical dynamics simulations with different WE metaparameters and comparing the results. In Fig. 3.5 we see that for several different simulations, there can still be sawtooth like behavior even at relatively large m_{targ} values such as for $k_{\text{off}} = 50000$ and $m_{\text{targ}} = 200$. However, it is clear that these increases in m_{targ} do visibly have an impact on the one metric above associated with higher accuracies, the NNZ fraction. Unfortunately, this higher m_{targ} limits the number of WE iterations that can

finish. Each WE simulation was set up to use 2.5 days of computation time, and we can see that the number of iterations completed seems to scale inversely with m_{targ} , thus counter balancing the increased NNZ fraction with lower flux measurements total. Rather than comparing NNZ fractions across WE metaparameter sets, instead we chose to do a large sweep showcasing the impact of metaparameters on accuracy measurements for a dynamically simpler system, allowing us to use brute force simulation results as a test for accuracy. The results are given in Fig. 3.7. Each square of the heatmap is associated errors given from 3 iid WE simulations, while simultaneously implementing a computation time limit of total number of dynamics steps that can be executed, $dt_{\text{max}} = \tau * n_{\text{bins}} * m_{\text{targ}}$. Unfortunately, rather than achieving a desired outcome of clear regions of τ and m_{targ} , we see that most metaparameter sets resulted in good accuracies within an order of magnitude (heatmap values of 2 or less).

Discussion

Our goal of this investigation was to help clear up ambiguities of WE simulations involving simulation accuracy and optimization of metaparameters with regards to simulation accuracy. In the process, we desired to create a method of selecting metaparameters that would minimize error for a fixed amount of computational resources, with an ideal method being one that is easily adaptable to a wide range of WE simulations rather than the 2d surface dynamics evacuations that were the primary focus of WE simulations in this dissertation. While we were unsuccessful in our goal of creating a method for metaparameter selection, our investigation revealed several illuminating properties of WE simulations which we will summarize here.

Firstly, there is ambiguity in how best to calculate the MFPT, whether done from taking an arithmetic mean of first passage times or arithmetic mean of fluxes. There is a relationship between the two, as the arithmetic mean of the MFPT / fluxes is the harmonic mean of the fluxes / MFPT. In our cases, the arithmetic mean of the fluxes makes the most intuitive sense, while also agreeing




Figure 3.6: Influence of m_{targ} on MFPT estimates, at constant number of WE iterations, shows that higher m_{targ} is associated with both better estimates for the MFPT as well as lower spread between MFPT estimates. For each of the 4 panels, 2 WE runs were done for a variety of unbinding rates k_{off} and replicas per bin m_{targ} . WE iterations per panel are approx. 22000 for $\tau = 20$, 10000 for $\tau = 50$, 5000 for $\tau = 100$, and 2800 for $\tau = 200$. Unbinding rates were chosen such that monomerization fraction was approximately 1 for each of the runs, so that MFPTs for each run should be identical (black horizontal line).

more closely with the predictions from asymptotics 2.1.

Secondly, while quantification of error can be done by comparing to a priori trends or through conglomeration of many runs, we see that in general these different methods of quantifying error give similar results on an order of magnitude scale.

Thirdly, identifying a metric that consistently predicts accuracy of WE simulations is difficult, especially if one desires that those that do not pass a certain threshold of the metric will fail. Of the metrics investigated, only NNZ fraction seemed to have any sort of correlation between accuracy and the metric, and even then it only implied that you would be consistently within 2




Figure 3.7: Heatmap of logarithmic accuracy defined in Eq. 3.3 across a wide range of τ and m_{targ} values for given computation times. The overall trend is relatively minor and hard to distinguish, as most of the simulations pass a basic accuracy threshold of order-of-magnitude accuracy (values 2 and below).

orders of magnitude, which might not be satisfactory. We furthermore see that metaparameter selection based around this NNZ metric might result in other limitations, such as the number of WE iterations that actually get completed (Fig. 3.5).

Lastly, we identified how weights in fine-grained bins can be dependent on incoming weights from relatively few replicas in the coarse grained bins. While this is somewhat obvious, the fast propagation of these weights throughout the fine-grained bins suggests that even more computation time might be ideally placed on the transitional states. As a consequence, we hypothesize that careful choice of binning methods will optimize computational resources. While the purpose of Weighted Ensemble is to better allocate computational resources to rare regions, ensuring a smooth entry into those rare regions is perhaps equally important, if not more so.

Moving forwards, we wish to see if there is a consistent way to find the “worst” metaparameters, or specifically metaparameter sets that are clearly and consistently worse than other metaparameter sets. In general, accuracy was relatively static across different sets of metaparameters, even in those that impacted the NNZ metric. As the landscape of the heatmap in Fig. 3.7 is relatively flat and remained relatively flat across a variety of computation times, it is possible that the specific metaparameters chosen aren’t important so long as there is enough opportunity for replicas to make key transitions between bins.

Chapter 4

Discussion

In this dissertation, we investigated rare events in the context of cell surface dynamics, specifically the evacuation of CD45 from T-cell receptors during antigen discrimination. While there are many potential causes of directed motion on the cell surface, a minimalist approach of starting with simply thermal diffusion reveals a clear fundamental difference in timescale separating evacuation from a purely thermal evacuation and the evacuations observed physiologically. To attempt to bridge this gap in timescales, we investigated the impact of several common features of cell-surface dynamics to see their impact on evacuation timescales. We found that dimerization / oligomerization and energy barriers, which correspond to the necessary compression of CD45 molecules in order for a close contact to form, both can help to bring the timescale closer to the physiological limit.

In the process of studying these surface dynamics, we created software that combines the biochemical simulator Smoldyn with the rare event sampling algorithm Weighted Ensemble. Our foray into Weighted Ensemble caused us to look in a greater detail at properties of the simulation traces output by the method. This investigation into the broader properties of Weighted Ensemble was inspired by the desire to help create a standard for both understanding uncertainty estimates

from a single Weighted Ensemble trace as well as the impact of Weighted Ensemble-specific simulation metaparameters on these uncertainty estimates as well as the overall accuracy of Weighted Ensemble simulations. We will discuss the impact of each of these projects below.

Impact of oligomerization and close contacts on receptor triggering

The activation of TCRs and evacuation of CD45 surrounding them during antigen discrimination is a widely observed but not fully explained phenomenon. Of the many models of T-cell activation, such as kinetic proofreading, receptor scanning, and serial triggering, only the kinetic segregation model attempts to explain the observed evacuations of CD45 from the regions surrounding the TCR. In this work, we created a simple model for kinetic segregation, starting from a thermal-only regime. In this sense, while the triggering of TCRs is clearly the motivation for the work, the main focus of this work is the study of how these additional features impact rare evacuations in diffusive processes rather than in creating a model as close to the real-world scenario as possible. In that regard, our work on dimerization and close contacts reveal that the relatively simple effect of partial limitations on entry into the evacuation region and oligomerization of the evacuators can have drastic impacts on the speed of the evacuation events which in turn can influence the feasibility of the evacuation being observed on a timescale useful to an individual cell.

The impact of oligomerization on these evacuation events might seem obvious upon first observation of the asymptotic estimates for diffusion-only evacuations and their approximately exponential dependence on the density of the evacuators. However, the ability of this self-binding to make an impact on the feasibility of an event cannot be understated. For CD45, for instance, evacuation times of a purely monomeric solution can decrease from hundreds of years to 10s of hours while still remaining explicitly non-homogeneous. For instance, evacuation times of around 10^5 seconds, or ~ 1 day were observed for a solution with a monomerization fraction of 0.2. While this sort of timescale is not applicable to the phenomenon of TCR activation, this brings the timescale of

the event from one that is virtually never observable to one that can be observed within the lifetime of the cell. Furthermore, the extension of a dimerization process to one of oligomerization can further reduce these timescales to sub-second processes in extreme cases. In fact, the near-exponential dependence of these evacuation times on the density of the evacuating particles has shown how even small changes in the binding affinity can have order of magnitude impacts on these evacuation times, as can be seen from our CD45 estimate of a 10-fold reduction in MFPT at a monomerization fraction of 0.8. Even at the lowest densities, the differences between monomeric and dimeric solutions can mean reduction from timescales of minutes to seconds, a difference that cannot be understated when considering binary outcomes such as the triggering of a receptor. We would very much like to see future experimental research that attempts to control the triggering of a process by changing the binding affinity of molecules that evacuate the nearby region.

Similar to oligomerization, the statement that an energy barrier can cause drastic decreases in these evacuation times might be seen as a prediction that is meant to be self-fulfilling. Indeed, our work on evacuation times for diffusing particles in the one-way case reveals that when these energy barriers are in place, evacuation times grow incredibly slowly with density. However, the connection with the physiological estimates for the spring constant of CD45 suggest that this energy difference alone is not enough to explain the evacuation of the region on the timescales observed physiologically. However, this does potentially provide a means of testing the causal relationship of kinetic segregation on TCR activation; whether or not the evacuation triggers the activation or the activation (or rather, the initialization of the close-contact by the T-cell and antigen providing cell). If kinetic segregation is necessary simply because of the physical barrier that CD45 creates between the two cells, rather than some biochemical property of CD45 such as binding affinity, then replacing CD45 with proteins of varying length and stiffness might give insight into how these mechanical properties influence CD45 evacuation and TCR triggering.

It is worth mentioning that while segregation between activated TCRs and CD45 has been observed experimentally, the resolution of these observations has not determined to what extent these

evacuations actually need to be achieved on a per-molecule basis. Complete evacuation of the region is a very large obstacle to overcome for diffusive processes alone. However, the possibility that only partial evacuation is necessary might have significant impacts on the overall evacuation time, especially in cases where dimerization and close contacts are present. If partial evacuation is all that is required for a triggering event, then the MFPTs we found in this work would be overestimates of the time to TCR triggering. As our current software is well situated to investigate the evacuation times for partial evacuations using the Weighted Ensemble method, investigations into the timescale of partial evacuations for diffusion-only as well as the other simple models in this document is certainly within our technological, though not temporal, reaches. Furthermore, while the work in this document suggests that the processes investigated alone are not sufficient to explain CD45 activation in the context of TCR activation, perhaps the introduction of partial evacuation alleviates the need for active processes to achieve these evacuations.

Conversely, an evacuation event occurring does not ensure that an antigen presenting cell is able to utilize the evacuation. With the high surface density of CD45, there exists the possibility that the evacuation events to be short-lived relative to the time it takes for an antigen presenting cell to take advantage of the evacuation. In that case, the time for the first evacuation would be an underestimate of the true triggering time. While this paper does not provide resolution between the simultaneous over/underestimations mentioned here, we reiterate that the time to achieve partial evacuations can be measured from changing the Smoldyn runtime commands programmed in LibsmolWE to stop simulations upon partial evacuations instead of full evacuations in combination with changing how the flux bin is defined. Similarly, the triggering event could be modeled as a Poisson rate that occurs when the evacuation has occurred, potentially even adjusting the rate to distinguish between partial and full evacuations. Nonetheless, the conclusions from this work on the impact of dimerization and close contacts remains valid for complete evacuation events, though we encourage awareness of the above mentioned effects when considering evacuation events as a condition for receptor triggering.

Uncertainty quantification and metaparameter selection in Weighted Ensemble rare event simulation

Weighted ensemble simulations can often result in data sets that have a high log-variance and correlated data, and complications that result from these favor approaching conclusions made from these data sets in “atypical” ways compared to the more “typical” methods such as standard error of the mean. Firstly, when combining data from iid WE runs one could compute the mean in two different ways, choosing to either take an arithmetic mean for individual MFPT estimates or an arithmetic mean for individual flux estimates. Furthermore, it is known that the result calculated from the arithmetic mean of flux estimates will always be less than the arithmetic mean of MFPT estimates. The work of chapter 2 resulted in frequent over estimations of an asymptotic solution, and as a result lent itself towards using arithmetic means of individual flux estimates, though this trend need not always be true.

There are multiple potential ways we identified to determine uncertainty from combinations of WE runs: using the SEM from iid WE runs, or comparison of sets of similar runs and using regression analysis. Regardless of which of these methods is used to calculate error, resulting error bars remain mostly the same to within order of magnitude. Metrics created to serve as accuracy predictors in general had little correlation between the metric and known accuracy to ground truth. The one metric investigated that did show some sort of correlation was the fraction of non-zero flux measurements for a WE simulation. In spite of this correlation, we found that low NNZ fractions did not preclude low errors. The ability of low NNZ fractions to have very high ground truth accuracies was somewhat surprising, and we believe that this information is useful in situations where the user’s ability to create more bins is limited, (e.g. discrete order parameters). When fixing the total computation time allowed for WE simulations, we did not see clear trends between m_{targ} , τ , and overall accuracy relating to ground truth.

WE traces often show “sawtooth” behavior in flux bins as well as neighboring fine-grained bins.

Furthermore, the appearance of these sawtooths in one bin appeared to be correlated with similar behavior in its neighbors. We saw indications that higher m_{targ} values were associated with higher non-zero fractions as well as less dramatic sawtooth behavior, while higher m_{targ} values also showed better agreement with ground truth and less spread between estimates. We hypothesize that the sawtooth behavior remains a possible candidate as an accuracy predictor, though we did not succeed at finding a way to quantify this behavior that agreed with ad hoc visual inspection. We believe that future work investigating properties of WE simulations should investigate these sawtooths, specifically their influence on accuracy and self-convergence. We hypothesize that these sawtooths are indicative of larger overall trends in weight distribution spawning from relatively few parent replicas, resulting in highly correlated data. As such, a desire for the weight distribution to come from many different parent replicas and minimize correlations suggest that minimizing the severity of these sawtooths can be useful for WE simulation optimization should these hypotheses be true.

Bibliography

- [1] J. L. Adelman and M. Grabe. Simulating rare events using a weighted ensemble-based string method. *J. Chem. Phys.*, 138(4):44105, jan 2013.
- [2] B. Alberts, A. Johnson, J. Lewis, M. Raff, and K. Roberts. *Molecular Biology of the Cell*. Garland, 2014.
- [3] J. F. Allard, O. Dushek, D. Coombs, and P. A. van der Merwe. Mechanical modulation of receptor-ligand interactions at cell-cell interfaces. *Biophys. J.*, 102(6):1265–1273, mar 2012.
- [4] S. S. Andrews. Accurate particle-based simulation of adsorption, desorption and partial transmission. *Phys. Biol.*, 6(4):046015, nov 2009.
- [5] S. S. Andrews. Smoldyn: particle-based simulation with rule-based modeling, improved molecular interaction and a library interface. *Bioinformatics*, 33(5):710–717, 2017.
- [6] S. S. Andrews. Modeling Biomolecular Site Dynamics, Methods and Protocols. *Methods Mol. Biol. (N. Y.)*, 1945:179–202, 2019.
- [7] S. S. Andrews, N. J. Addy, R. Brent, and A. P. Arkin. Detailed Simulations of Cell Biology with Smoldyn 2.1. *PLoS Comput. Biol.*, 6(3):e1000705, mar 2010.
- [8] S. S. Andrews and D. Bray. Stochastic simulation of chemical reactions with spatial resolution and single molecule detail. *Phys. Biol.*, 2004.
- [9] D. Aristoff and D. M. Zuckerman. Optimizing weighted ensemble sampling of steady states. *Multiscale Model. Simul.*, 18(2):646–673, 2020.
- [10] N. Bag, D. A. Holowka, and B. A. Baird. Imaging FCS Delineates Subtle Heterogeneity in Plasma Membranes of Resting Mast Cells. *Mol. Biol. Cell*, pages mbc.E19–10–0559, jan 2020.
- [11] G. Bell, M. Dembo, and P. Bongrand. Cell adhesion. Competition between nonspecific repulsion and specific bonding. *Biophys. J.*, (6):1051–1064, jun.
- [12] D. Bhatt, B. W. Zhang, and D. M. Zuckerman. Steady-state simulations using weighted ensemble path sampling. *J. Chem. Phys.*, (1):014110, jul.

- [13] A. T. Bogetti, B. Mostofian, A. Dickson, A. Pratt, A. S. Saglam, P. O. Harrison, J. L. Adelman, M. Dudek, P. A. Torrillo, A. J. DeGrave, et al. A suite of tutorials for the westpa rare-events sampling software [article v1. 0]. *LiveCoMS*, 1(2), 2019.
- [14] S. S. A. Bray and Dennis. Stochastic simulation of chemical reactions with spatial resolution and single molecule detail. *Phys. Biol.*, 1(3):137, 2004.
- [15] N. J. Burroughs, K. Köhler, V. Miloserdov, M. L. Dustin, P. A. van der Merwe, and D. M. Davis. Boltzmann energy-based image analysis demonstrates that extracellular domain size differences explain protein segregation at immune synapses. *PLoS Comput. Biol.*, 7(8), 2011.
- [16] N. J. Burroughs and C. Wülfing. Differential Segregation in a Cell-Cell Contact Interface: The Dynamics of the Immunological Synapse. *Biophys. J.*, (4):1784–1796, oct.
- [17] E. Cai, K. Marchuk, P. Beemiller, C. Beppler, M. G. Rubashkin, V. M. Weaver, A. Gérard, T.-L. Liu, B.-C. Chen, E. Betzig, F. Bartumeus, and M. F. Krummel. Visualizing dynamic microvillar search and stabilization during ligand detection by T cells. *Science*, (6338):eaal3118, may.
- [18] C. W. Cairo, R. Das, A. Albohy, Q. J. Baca, D. Pradhan, J. S. Morrow, D. Coombs, and D. E. Golan. Dynamic Regulation of CD45 Lateral Mobility by the Spectrin-Ankyrin Cytoskeleton of T Cells. *J. Biol. Chem.*, 285(15):11392–11401, apr 2010.
- [19] M. Candotti and M. Orozco. The Differential Response of Proteins to Macromolecular Crowding. *PLoS Comput. Biol.*, 12(7):e1005040—18, jul 2016.
- [20] C. B. Carbone, N. Kern, R. A. Fernandes, E. Hui, X. Su, K. C. Garcia, and R. D. Vale. In vitro reconstitution of T cell receptor-mediated segregation of the CD45 phosphatase. *Proc. Natl. Acad. Sci. U. S. A.*, 114(44):E9338–E9345, oct 2017.
- [21] V. T. Chang, R. A. Fernandes, K. A. Ganzinger, S. F. Lee, C. Siebold, J. McColl, P. Jönsson, M. Palayret, K. Harlos, C. H. Coles, E. Y. Jones, Y. Lui, E. Huang, R. J. C. Gilbert, D. Klenerman, A. R. Aricescu, and S. J. Davis. Initiation of T cell signaling by CD45 segregation at close contacts. *Nat. Immunol.*, 17:574, mar 2016.
- [22] K. Y. Chen, E. Jenkins, M. Körbel, A. Ponjavic, A. H. Lippert, A. M. Santos, N. Ashman, C. O’Brien-Ball, J. McBride, D. Klenerman, and S. J. Davis. Trapping or slowing the diffusion of T cell receptors at close contacts initiates T cell signaling. *Proc. Natl. Acad. Sci. U. S. A.*, 118(39):e2024250118, sep 2021.
- [23] W.-X. Chew, K. Kaizu, M. Watabe, S. V. Muniandy, K. Takahashi, and S. N. V. Arjunan. Reaction-diffusion kinetics on lattice at the microscopic scale. *Phys. Rev. E*, 98(3):32418, sep 2018.
- [24] J. D. Chodera. A Simple Method for Automated Equilibration Detection in Molecular Simulations. *J. Chem. Theory Comput.*, 12(4):1799–1805, apr 2016.

- [25] K. Choudhuri, D. Wiseman, M. H. Brown, K. Gould, and P. A. van der Merwe. T-cell receptor triggering is critically dependent on the dimensions of its peptide-MHC ligand. *Nature*, 436(7050):578–582, 2005.
- [26] L. Clemens, O. Dushek, and J. Allard. Intrinsic Disorder in the T Cell Receptor Creates Cooperativity and Controls ZAP70 Binding. *Biophys. J.*, (2):379–392, jan.
- [27] A. Connolly, R. Panes, M. Tual, R. Lafortune, A. Bellemare-Pelletier, and E. Gagnon. TMEM16F mediates bystander TCR-CD3 membrane dissociation at the immunological synapse and potentiates T cell activation. *Sci. Signaling*, (675), mar.
- [28] L. Conway, D. Wood, E. Tuzel, and J. L. Ross. Motor transport of self-assembled cargos in crowded environments. *Proc. Natl. Acad. Sci. U. S. A.*, 109(51):20814–20819, dec 2012.
- [29] A. H. Courtney, A. A. Shvets, W. Lu, G. Griffante, M. Mollenauer, V. Horkova, W.-L. Lo, S. Yu, O. Stepanek, A. K. Chakraborty, and A. Weiss. CD45 functions as a signaling gatekeeper in T cells. *Sci. Signaling*, 12(604):eaaw8151, oct 2019.
- [30] E. N. Cytrynbaum, Y. D. Li, J. F. Allard, and H. Mehrabian. Estimating the bending modulus of a FtsZ bacterial-division protein filament. *Phys. Rev. E*, (1):011902, jan.
- [31] R. Dixit, J. L. Ross, Y. E. Goldman, and E. L. F. Holzbaur. Differential Regulation of Dynein and Kinesin Motor Proteins by Tau. *Science (New York, NY)*, 319(5866):1086–1089, feb 2008.
- [32] R. M. Donovan, A. J. Sedgewick, J. R. Faeder, and D. M. Zuckerman. Efficient stochastic simulation of chemical kinetics networks using a weighted ensemble of trajectories. *J. Chem. Phys.*, 139(11):115105, sep 2013.
- [33] R. M. Donovan, J.-J. Tapia, D. P. Sullivan, J. R. Faeder, R. F. Murphy, M. Dittrich, and D. M. Zuckerman. Unbiased Rare Event Sampling in Spatial Stochastic Systems Biology Models Using a Weighted Ensemble of Trajectories. *PLoS Comput. Biol.*, 12(2):e1004611, feb 2016.
- [34] O. Dushek, J. Goyette, and P. A. Van Der Merwe. Non-catalytic tyrosine-phosphorylated receptors. *Immunol. Rev.*, 250(1):258–276, nov 2012.
- [35] O. Dushek, P. van der Merwe, and V. Shahrezaei. Ultrasensitivity in Multisite Phosphorylation of Membrane-Anchored Proteins. *Biophys. J.*, 100(5):1189–1197, 2011.
- [36] L. Dyson and R. E. Baker. The importance of volume exclusion in modelling cellular migration. *J. Math. Biol.*, 71(3):1–21, apr 2015.
- [37] A. H. Elcock. Models of macromolecular crowding effects and the need for quantitative comparisons with experiment. *Curr. Opin. Struct. Biol.*, 20(2):196–206, apr 2010.
- [38] S. Farhadi, R. Liu, and G. Hudalla. Tuning Multivalent Signaling of Extracellular Galectin-3. *FASEB J.*, (S1):fasebj.2021.35.S1.03143, may.

- [39] S. A. Farhadi, R. Liu, M. W. Becker, E. A. Phelps, and G. A. Hudalla. Physical tuning of galectin-3 signaling. *Proc. Natl. Acad. Sci. U. S. A.*, (19):e2024117118, may.
- [40] R. A. Fernandes, K. A. Ganzinger, J. C. Tzou, P. Jönsson, S. F. Lee, M. Palayret, A. M. Santos, A. R. Carr, A. Ponjavic, V. T. Chang, C. Macleod, B. C. Lagerholm, A. E. Lindsay, O. Dushek, A. Tilevik, S. J. Davis, and D. Klenerman. A cell topography-based mechanism for ligand discrimination by the T cell receptor. *Proc. Natl. Acad. Sci. U. S. A.*, 116(28):14002 LP – 14010, jul 2019.
- [41] T. Fujiwara, K. Ritchie, H. Murakoshi, K. Jacobson, and A. Kusumi. Phospholipids undergo hop diffusion in compartmentalized cell membrane. *J. Cell Biol.*, 157(6):1071–1082, jun 2002.
- [42] T. K. Fujiwara, K. Iwasawa, Z. Kalay, T. A. Tsunoyama, Y. Watanabe, Y. M. Umemura, H. Murakoshi, K. G. N. Suzuki, Y. L. Nemoto, N. Morone, and A. Kusumi. Confined diffusion of transmembrane proteins and lipids induced by the same actin meshwork lining the plasma membrane. *Mol. Biol. Cell*, 27(7):1101–1119, apr 2016.
- [43] Y. Gambin, R. Lopez-Esparza, M. Reffay, E. Sierrecki, N. S. Gov, M. Genest, R. S. Hodges, and W. Urbach. Lateral mobility of proteins in liquid membranes revisited. *Proc. Natl. Acad. Sci. U. S. A.*, 103(7):2098–2102, 2006.
- [44] D. Goswami, K. Gowrishankar, S. Bilgrami, S. Ghosh, R. Raghupathy, R. Chadda, R. Vishwakarma, M. Rao, and S. Mayor. Nanoclusters of GPI-Anchored Proteins Are Formed by Cortical Actin-Driven Activity. *Cell*, (6):1085–1097, dec.
- [45] A. Grossfield, P. N. Patrone, D. R. Roe, A. J. Schultz, D. W. Siderius, and D. M. Zuckerman. Best Practices for Quantification of Uncertainty and Sampling Quality in Molecular Simulations [Article v1.0]. *LiveCoMS*, 1(1):5067, 2018.
- [46] P. Hammar, P. Leroy, A. Mahmutovic, E. G. Marklund, O. G. Berg, and J. Elf. The lac Repressor Displays Facilitated Diffusion in Living Cells. *Science (New York, NY)*, 336(6088):1595–1598, jun 2012.
- [47] M. J. Harris, M. Fuyal, and J. R. James. Quantifying persistence in the T-cell signaling network using an optically controllable antigen receptor. *Mol. Syst. Biol.*, (5):e10091, may.
- [48] J. Helenius, G. Brouhard, Y. Kalaidzidis, S. Diez, and J. Howard. The depolymerizing kinesin MCAK uses lattice diffusion to rapidly target microtubule ends. *Nature*, 441(7089):115–119, may 2006.
- [49] T. L. Hill. *Free Energy Transduction and Biochemical Cycle Kinetics*. Dover Books on Chemistry. Dover Publications, 2005.
- [50] J. Huang, M. Brameshuber, X. Zeng, J. Xie, Q.-j. Li, Y.-h. Chien, S. Valitutti, and M. M. Davis. A single peptide-major histocompatibility complex ligand triggers digital cytokine secretion in CD4+ T cells. *Immunity*, 39(5):846–857, 2013.

- [51] G. A. Huber and S. Kim. Weighted-ensemble Brownian dynamics simulations for protein association reactions. *Biophys. J.*, 70(1):97–110, jan 1996.
- [52] E. Hui. Understanding T cell signaling using membrane reconstitution. *Immunol. Rev.*, 291(1):44–56, 2019.
- [53] J. B. Huppa, M. Axmann, M. A. Mörtelmaier, B. F. Lillemeier, E. W. Newell, M. Brameshuber, L. O. Klein, G. J. Schütz, and M. M. Davis. TCR–peptide–MHC interactions in situ show accelerated kinetics and increased affinity. *Nature*, (7283):963–967, feb.
- [54] D. J. Irvine, M. A. Purbhoo, M. Krogsgaard, and M. M. Davis. Direct observation of ligand recognition by T cells. *Nature*, 419(6909):845–849, 2002.
- [55] M. Javanainen, H. Martinez-Seara, R. Metzler, and I. Vattulainen. Diffusion of integral membrane proteins in protein-rich membranes. *J. Phys. Chem. Lett.*, 8(17):4308–4313, 2017. PMID: 28823153.
- [56] M. E. Johnson, A. Chen, J. R. Faeder, P. Henning, I. I. Moraru, M. Meier-Schellersheim, R. F. Murphy, T. Prüstel, J. A. Theriot, and A. M. Uhrmacher. Quantifying the roles of space and stochasticity in computer simulations for cell biology and cellular biochemistry. *Mol. Biol. Cell*, 32(2):186–210, 2021.
- [57] Y. Jung, I. Riven, S. W. Feigelson, E. Kartvelishvily, K. Tohya, M. Miyasaka, R. Alon, and G. Haran. Three-dimensional localization of T-cell receptors in relation to microvilli using a combination of superresolution microscopies. *Proc. Natl. Acad. Sci. U. S. A.*, 113(40):E5916 LP – E5924, oct 2016.
- [58] T. B. Kepler and T. C. Elston. Stochasticity in transcriptional regulation: Origins, consequences, and mathematical representations. *Biophys. J.*, 81(6):3116–3136, dec 2001.
- [59] M. Klann and H. Koepl. Spatial simulations in systems biology: from molecules to cells. *Int. J. Mol. Sci.*, 13(6):7798–7827, 2012.
- [60] S. D. Lawley and J. P. Keener. Rebinding in biochemical reactions on membranes. *Phys. Biol.*, 14(5):56002, 2017.
- [61] C. Leduc, K. Padberg-Gehle, V. Varga, D. Helbing, S. Diez, and J. Howard. Molecular crowding creates traffic jams of kinesin motors on microtubules. *Proc. Natl. Acad. Sci. U. S. A.*, 109(16):6100–6105, apr 2012.
- [62] L. Li, X. Shi, X. Guo, H. Li, and C. Xu. Ionic protein–lipid interaction at the plasma membrane: what can the charge do? *Trends Biochem. Sci.*, 39(3):130–140, mar 2014.
- [63] Q.-J. Li, A. R. Dinner, S. Qi, D. J. Irvine, J. B. Huppa, M. M. Davis, and A. K. Chakraborty. CD4 enhances T cell sensitivity to antigen by coordinating Lck accumulation at the immunological synapse. *Nat. Immunol.*, 5(8):791–799, 2004.
- [64] J. Lin and A. Weiss. The tyrosine phosphatase CD148 is excluded from the immunologic synapse and down-regulates prolonged T cell signaling. *J. Cell Biol.*, 162(4):673–682, 2003.

- [65] T. Lukeš, D. Glatzová, Z. Kvíčalová, F. Levet, A. Benda, S. Letschert, M. Sauer, T. Brdička, T. Lasser, and M. Cebeauer. Quantifying protein densities on cell membranes using super-resolution optical fluctuation imaging. *Nat. Commun.*, 8(1):1731, 2017.
- [66] Y. Ma, Y. J. Lim, A. Benda, J. Lou, J. Goyette, and K. Gaus. Clustering of the ζ -Chain Can Initiate T Cell Receptor Signaling. *Int. J. Mol. Sci.*, (10):3498, may.
- [67] T. T. Marquez-Lago, A. Leier, and K. Burrage. Anomalous diffusion and multifractional Brownian motion: simulating molecular crowding and physical obstacles in systems biology. *IET Syst. Biol.*, 6(4):134–142, 2012.
- [68] S. R. McGuffee and A. H. Elcock. Diffusion, Crowding & Protein Stability in a Dynamic Molecular Model of the Bacterial Cytoplasm. *PLoS Comput. Biol.*, 6(3):e1000694, mar 2010.
- [69] K. McKinney, M. Mattia, V. Gottifredi, and C. Prives. p53 Linear Diffusion along DNA Requires Its C Terminus. *Mol. Cell*, 16(3):413–424, nov 2004.
- [70] L. Meloty-Kapella, B. Shergill, J. Kuon, E. Botvinick, and G. Weinmaster. Notch Ligand Endocytosis Generates Mechanical Pulling Force Dependent on Dynamin, Epsins, and Actin. *Dev. Cell*, 22(6):1299–1312, jun 2012.
- [71] 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. *J. Phys. A: Math. Theor.*, 42(43):434013, 2009.
- [72] B. Mostofian and D. M. Zuckerman. Statistical uncertainty analysis for small-sample, high log-variance data: Cautions for bootstrapping and bayesian bootstrapping. *J. Chem. Theory Comput.*, 15(6):3499–3509, 2019.
- [73] E. Natkanski, W. Y. Lee, B. Mistry, A. Casal, J. E. Molloy, and P. Tolar. B Cells Use Mechanical Energy to Discriminate Antigen Affinities. *Science (New York, NY)*, 340(6140):1587–1590, jun 2013.
- [74] J. Newby and J. Allard. First-Passage Time to Clear the Way for Receptor-Ligand Binding in a Crowded Environment. *Phys. Rev. Lett.*, 116(12):128101, mar 2016.
- [75] F. W. J. Olver, N. I. of Standards, T. (U.S.), D. W. Lozier, R. F. Boisvert, and C. W. Clark. NIST Handbook of Mathematical Functions Hardback and CD-ROM. 2011.
- [76] R. H. Pullen and S. M. Abel. Catch Bonds at T Cell Interfaces: Impact of Surface Reorganization and Membrane Fluctuations. *Biophys. J.*, (1):120–131, jul.
- [77] V. Rajani, G. Carrero, D. Golan, G. de Vries, and C. Cairo. Analysis of Molecular Diffusion by First-Passage Time Variance Identifies the Size of Confinement Zones. *Biophys. J.*, 100(6):1463–1472, mar 2011.
- [78] S. Ramadurai, A. Holt, V. Krasnikov, G. van den Bogaart, J. A. Killian, and B. Poolman. Lateral diffusion of membrane proteins. *J. Am. Chem. Soc.*, 131(35):12650–12656, 2009. PMID: 19673517.

- [79] D. Ray, T. Gokey, D. L. Mobley, and I. Andricioaei. Kinetics and free energy of ligand dissociation using weighted ensemble milestoning. *J. Chem. Phys.*, 153(15):154117, oct 2020.
- [80] Y. Razvag, Y. Neve-Oz, J. Sajman, O. Yakovian, M. Reches, and E. Sherman. T Cell Activation through Isolated Tight Contacts. *Cell Rep.*, 29(11):3506–3521.e6, 2019.
- [81] A. Rojnuckarin, D. R. Livesay, and S. Subramaniam. Bimolecular reaction simulation using Weighted Ensemble Brownian dynamics and the University of Houston Brownian Dynamics program. *Biophys. J.*, 79(2):686–693, aug 2000.
- [82] R. Rouzier, R. Rajan, P. Wagner, K. R. Hess, D. L. Gold, J. Stec, M. Ayers, J. S. Ross, P. Zhang, T. A. Buchholz, H. Kuerer, M. Green, B. Arun, G. N. Hortobagyi, W. F. Symmans, and L. Pusztai. Microtubule-associated protein tau: a marker of paclitaxel sensitivity in breast cancer. *Proc. Natl. Acad. Sci. U. S. A.*, 102(23):8315–8320, jun 2005.
- [83] J.-F. Rupprecht, O. Bénichou, D. Grebenkov, and R. Voituriez. Exit time distribution in spherically symmetric two-dimensional domains. *Journal of Statistical Physics*, 158(1):192–230, 2015.
- [84] A. S. Saglam and L. T. Chong. Highly Efficient Computation of the Basal kon using Direct Simulation of Protein-Protein Association with Flexible Molecular Models. *J. Phys. Chem. B*, 120(1):117–122, jan 2016.
- [85] B. Shergill, L. Meloty-Kapella, A. A. Musse, G. Weinmaster, and E. Botvinick. Optical Tweezers Studies on Notch: Single-Molecule Interaction Strength Is Independent of Ligand Endocytosis. *Dev. Cell*, 22(6):1313–1320, jun 2012.
- [86] X. Shi, Y. Bi, W. Yang, X. Guo, Y. Jiang, C. Wan, L. Li, Y. Bai, J. Guo, Y. Wang, X. Chen, B. Wu, H. Sun, W. Liu, J. Wang, and C. Xu. Ca²⁺ regulates T-cell receptor activation by modulating the charge property of lipids. *Nature*, 493(7430):111, apr 2014.
- [87] J. A. Siller-Farfán and O. Dushek. Molecular mechanisms of T cell sensitivity to antigen. *Immunol. Rev.*, (1):194–205, sep.
- [88] K. Sneppen, M. A. Micheelsen, and I. B. Dodd. Ultrasensitive gene regulation by positive feedback loops in nucleosome modification. *Mol. Syst. Biol.*, 4:1–5, apr 2008.
- [89] J. C. Stachowiak, E. M. Schmid, C. J. Ryan, H. S. Ann, D. Y. Sasaki, M. B. Sherman, P. L. Geissler, D. A. Fletcher, and C. C. Hayden. Membrane bending by protein–protein crowding. *Nat. Cell Biol.*, 14(9):944–949, aug 2012.
- [90] K. Takahashi, S. Tănase-Nicola, and P. R. ten Wolde. Spatio-temporal correlations can drastically change the response of a MAPK pathway. *Proc. Natl. Acad. Sci. U. S. A.*, 107(6):2473 LP – 2478, feb 2010.
- [91] A. Takeda, A. Matsuda, R. M. J. Paul, and N. R. Yaseen. CD45-associated protein inhibits CD45 dimerization and up-regulates its protein tyrosine phosphatase activity. *Blood*, (9):3440–3447, may.

- [92] A. TAKEDA, J. J. WU, and A. L. MAIZEL. Evidence for monomeric and dimeric forms of CD45 associated with a 30-kDa phosphorylated protein. *J. Biol. Chem. (Print)*, 267(23):16651–16659, 1992.
- [93] P. A. van der Merwe and O. Dushek. Mechanisms for T cell receptor triggering. *Nat. Rev. Immunol.*, 11(1):47–55, 2011.
- [94] R. Varma, G. Campi, T. Yokosuka, T. Saito, and M. L. Dustin. T cell receptor-proximal signals are sustained in peripheral microclusters and terminated in the central supramolecular activation cluster. *Immunity*, 25(1):117–127, 2006.
- [95] M. von Smoluchowski. Versuch einer mathematischen Theorie der Koagulationskinetik kolloider Lösungen. *Z. Phys. Chem. (Leipzig, Ger.)*, 92:129–168, 2010.
- [96] A. D. Waldman, J. M. Fritz, and M. J. Lenardo. A guide to cancer immunotherapy: from T cell basic science to clinical practice. *Nat. Rev. Immunol.*, (11):651–668, nov.
- [97] Y. Wu, J. Vendome, L. Shapiro, A. Ben-Shaul, and B. Honig. Transforming binding affinities from three dimensions to two with application to cadherin clustering. *Nature*, (7357):510–513, jul.
- [98] Z. Xu and A. Weiss. Negative regulation of CD45 by differential homodimerization of the alternatively spliced isoforms. *Nat. Immunol.*, (8):764–771, aug.
- [99] O. N. Yogurtcu and M. E. Johnson. Cytosolic proteins can exploit membrane localization to trigger functional assembly. *PLoS Comput. Biol.*, (3):e1006031, mar.
- [100] O. N. Yogurtcu and M. E. Johnson. Theory of bi-molecular association dynamics in 2d for accurate model and experimental parameterization of binding rates. *J. Chem. Phys.*, 143(8):084117, 2015.
- [101] O. N. Yogurtcu and M. E. Johnson. Theory of bi-molecular association dynamics in 2D for accurate model and experimental parameterization of binding rates. *J. Chem. Phys.*, 143(8):08B614_1, 2015.
- [102] B. W. Zhang, D. Jasnow, and D. M. Zuckerman. Efficient and verified simulation of a path ensemble for conformational change in a united-residue model of calmodulin. *Proc. Natl. Acad. Sci. U. S. A.*, 104(46):18043–18048, nov 2007.
- [103] B. W. Zhang, D. Jasnow, and D. M. Zuckerman. The ”weighted ensemble” path sampling method is statistically exact for a broad class of stochastic processes and binning procedures. *J. Chem. Phys.*, 132(5):54107, feb 2010.
- [104] D. M. Zuckerman and L. T. Chong. Weighted Ensemble Simulation: Review of Methodology, Applications, and Software. *Annu. Rev. Biophys.*, 46(1):43–57, may 2017.
- [105] M. C. Zwier, A. J. Pratt, J. L. Adelman, J. W. Kaus, D. M. Zuckerman, and L. T. Chong. Efficient Atomistic Simulation of Pathways and Calculation of Rate Constants for a Protein-Peptide Binding Process: Application to the MDM2 Protein and an Intrinsically Disordered p53 Peptide. *The J. Phys. Chem. Lett.*, 7(17):3440–3445, sep 2016.

Appendix A

Analytic derivation of $p_{\text{entry}} = 0$ FPT distribution

Here we compute the mean first evacuation time from a disc of a collection of molecules undergoing diffusion, when reentry is prohibited ($p_{\text{entry}} = 0$). To solve this problem, we first find expressions involving evacuation time for a single molecule placed at a specific location inside the region to be evacuated. We then use that to find the mean first evacuation time for a molecule placed arbitrarily in the domain by averaging its expression across the domain. Lastly, we adapt that expression to systems of N non-interacting molecules placed randomly across the domain as an expression of the mean final evacuation time. Note that the single-particle exit problem has been solved previously, see [83].

Single Molecule Evacuation

Let $R = R_{\text{ROI}}$ and $p = p(\mathbf{r}, t | \mathbf{r}_0)$ be the probability density of a single molecule given the molecule's initial position $\mathbf{r}_0 = (r_0, \theta_0)$. This probability obeys the two-dimensional diffusion

equations

$$\frac{1}{D}\partial_t p = \partial_r^2 p + \frac{1}{r}\partial_r p + \frac{1}{r^2}\partial_\theta^2 p \quad (\text{A.1})$$

with boundary conditions

$$p(\mathbf{r}, 0) = \delta(\mathbf{r} - \mathbf{r}_0) \quad (\text{A.2})$$

$$p(R, \theta, t) = 0. \quad (\text{A.3})$$

This equation is separable, $p(r, \theta, t) = T(t)\Theta(\theta)R(r)$. The separated equations are

$$\partial_t T + D\lambda^2 T = 0 \quad (\text{A.4})$$

$$\partial_\theta^2 \Theta + m^2 \Theta = 0. \quad (\text{A.5})$$

$$r^2 \partial_r^2 R + r \partial_r R + (r^2 \lambda^2 - m^2) R = 0 \quad (\text{A.6})$$

The temporal and angular solutions are

$$T(t) = C e^{-\lambda^2 D t} \quad (\text{A.7})$$

$$\Theta(\theta) = A \sin(m\theta) + B \cos(m\theta) \quad (\text{A.8})$$

for unknown constants A, B, C . The periodicity requires that $\Theta(\theta) = \Theta(\theta + 2\pi)$, which demands m be an integer. The radial equation is an example of Bessel's equation, and given the boundedness of the solution at the origin it is appropriate to limit the possible solutions to Bessel functions of the first kind. This also defines the constant λ to be one of the zeroes of the Bessel function for some corresponding m , with $J_m(\lambda_{m,n} R) = 0$, where the m subscript specifies the order of Bessel function and n subscript refers to one of the (infinitely many) zeroes of the function. Combining

arbitrary constants our solution may be written generally as

$$p(\mathbf{r}, t|\mathbf{r}_0) = \sum_m \sum_n \left(e^{-(\lambda_{m,n}/R)^2 Dt} J_m(\lambda_{m,n} r/R) [A_{m,n} \cos(m\theta) + B_{m,n} \sin(m\theta)] \right). \quad (\text{A.9})$$

Survival Probability and Passage Time Distributions

We introduce the survival probability $G(t)$, the probability that the molecule remains in the evacuation region at time t . This can be expressed in terms of an integral over the position

$$G(t|\mathbf{r}(0) = \mathbf{r}_0) = \int_{\Omega} d\mathbf{r} p(\mathbf{r}, t|\mathbf{r}_0), \quad (\text{A.10})$$

where Ω is the spatial domain of the evacuation of radius R . This leads to

$$G(t|\mathbf{r}(0) = \mathbf{r}_0) = 2\pi \int_0^R r dr \sum_n A_n e^{-(\lambda_n/R)^2 Dt} J_0(\lambda_n r/R) \quad (\text{A.11})$$

$$= 2\pi R^2 \sum_n \frac{A_n}{\lambda_n} J_1(\lambda_n) e^{-(\lambda_n/R)^2 Dt}. \quad (\text{A.12})$$

Eq. A.11 by recognizing that the integral vanishes when $m > 0$ and relabeling the indices to only count the zeros of J_0 . Eq. A.12 the property of Bessel functions $\int r^{\nu+1} J_{\nu}(r) dr = r^{\nu+1} J_{\nu+1}(r)$ [75].

Using the orthogonality relation

$$\int_0^R J_0(\lambda_m r/R) J_0(\lambda_n r/R) r dr = \frac{1}{2} R^2 [J_1(\lambda_n)]^2 \delta_{mn}, \quad (\text{A.13})$$

allows us to find the coefficients

$$A_n = \frac{1}{\pi R^2 J_1(\lambda_n)} \int_{\Omega} J_0(\lambda_n r/R) \delta(\mathbf{r} - \mathbf{r}_0) d\mathbf{r} \quad (\text{A.14})$$

$$= \frac{J_0(\lambda_n r_0/R)}{\pi R^2 [J_1(\lambda_n)]^2}, \quad (\text{A.15})$$

and our final expression for the survival probability G is

$$G(t|\mathbf{r}_0) = 2 \sum_n e^{-(\lambda_n/R)^2 Dt} \frac{J_0(\lambda_n r_0/R)}{J_1(\lambda_n) \lambda_n}. \quad (\text{A.16})$$

The survival probability can also be written in terms of the distribution of first passage times, $\varrho(\tau|\mathbf{r}_0)$, as

$$G(t|\mathbf{r}_0) = \int_t^\infty \varrho(\tau|\mathbf{r}_0) d\tau, \quad (\text{A.17})$$

which implies

$$\varrho(\tau|\mathbf{r}_0) = -\partial_\tau G(\tau|\mathbf{r}_0). \quad (\text{A.18})$$

Integration by parts then gives the mean first passage time of a single diffusing molecule in terms of its positional probability distribution,

$$\bar{\tau} = \int_0^\infty \tau \varrho d\tau = \int_0^\infty d\tau G = \int_0^\infty d\tau \int_{\Omega} d\mathbf{r} p(\mathbf{r}, \tau|\mathbf{r}_0). \quad (\text{A.19})$$

To find the mean first passage time for a molecule randomly placed throughout the domain, rather than at a specific initial condition \mathbf{r}_0 , we average across all possible initial conditions

$$\langle \bar{\tau} \rangle = \frac{1}{\pi R^2} \int_0^\infty \int_{\Omega} \tau \varrho d\tau d\mathbf{r}_0 = \frac{1}{\pi R^2} \int_0^\infty d\tau \int_{\Omega} d\mathbf{r} \int_{\Omega} d\mathbf{r}_0 p(\mathbf{r}, \tau|\mathbf{r}_0) = \int_0^\infty d\tau \langle G(\tau) \rangle. \quad (\text{A.20})$$

where

$$\langle G(t) \rangle = \frac{2}{\pi R^2} \int_{\Omega} d\mathbf{r}_0 \sum_n e^{-(\lambda_n/R)^2 Dt} \frac{J_0(\lambda_n r_0/R)}{J_1(\lambda_n) \lambda_n} = 4 \sum_n \left(\frac{e^{-(\lambda_n/R)^2 Dt}}{\lambda_n^2} \right). \quad (\text{A.21})$$

We confirm this result by comparison with brute-force simulation in Fig. 1.1A.

As an aside, noting that the molecule should always remain in the region at $t = 0$, $\langle G(t) \rangle$ should evaluate to 1 at $t = 0$, implying the following property of the zeroes of the order 0 Bessel function of the first kind:

$$\sum_{n=1}^{\infty} \frac{1}{\lambda_n^2} = \frac{1}{4}, \quad (\text{A.22})$$

which is an interesting relation that this work confirms numerically.

Generalization to N Molecules

To generalize this solution for N distinct, non-interacting molecules consider the joint probability that all molecules evacuate prior to some time t' :

$$P(\tau_{N,\text{All}} \leq t') = P(\tau_1 \leq t', \tau_2 \leq t', \dots, \tau_N \leq t') \quad (\text{A.23})$$

$$= \prod_i P_i(\tau_i \leq t'), \quad (\text{A.24})$$

where $P(\tau_{N,\text{All}} \leq t')$ is the probability that the evacuation time for every molecule is less than t' . As the molecules are independent and non-interacting, we express it as a product of their individual probabilities of evacuating prior to t' , P_i . These P_i are the complement to each molecule's associated $G(t|\mathbf{r}_i)$ given in the previous section,

$$P_i(\tau_i \leq t') = 1 - G(t'|\mathbf{r}_i). \quad (\text{A.25})$$

We are interested in the survival probability of the entire system of N particles, which we can substitute into Eq. A.19 and Eq. A.20 to get the mean first passage time of the entire system. We designate the total survival probability $G_{N,\text{Any}}(t)$, i.e. the probability that at least one molecule remains inside the evacuation region at some time t , as the complement to the probability of every molecule evacuating.

$$G_{N,\text{Any}}(t) = 1 - P(\tau_{N,\text{All}} \leq t) . \quad (\text{A.26})$$

Combining Eq. A.26 into Eq. A.23 yields

$$G_{N,\text{Any}}(t|\mathbf{r}_1, \dots, \mathbf{r}_N) = 1 - \prod_i (1 - G(t|\mathbf{r}_i)) . \quad (\text{A.27})$$

As the molecules are non-interacting, we can use $\text{Cov}(G(t|\mathbf{r}_1), G(t|\mathbf{r}_2)) = 0$ to obtain $\langle G(t|\mathbf{r}_1)G(t|\mathbf{r}_2) \rangle = \langle G(t|\mathbf{r}_1) \rangle \langle G(t|\mathbf{r}_2) \rangle = \langle G(t) \rangle^2$ and find an expression for the spatial average of the probability at least one molecule remains in the region of interest. The spatial average of $G_{N,\text{Any}}$ can then be written as

$$\langle G_{N,\text{Any}}(t) \rangle = 1 - (1 - \langle G(t) \rangle)^N = - \left(\sum_{k=1}^N \binom{N}{k} (-\langle G(t) \rangle)^k \right) \quad (\text{A.28})$$

The mean first passage time for the entire system is then given by

$$\langle \tau_{N,\text{All}} \rangle = \int_0^\infty d\tau \langle G_{N,\text{Any}}(\tau) \rangle \quad (\text{A.29})$$

We confirm this result by comparison with brute-force simulation in Fig. 1.1B.

Placement Outside of Evacuation Region To account for placement outside of the evacuation region, if we have n molecules and $\langle \tau_{m,\text{All}} \rangle$ is the evacuation time for exactly m molecules inside the ROI, with $m \leq N$, we can use the fact that for each molecule, the probability of being placed inside the evacuation region is $p(\mathbf{r}_j \in \Omega_{\text{ROI}}) = A_{\text{ROI}}/A_{\text{Domain}}$, the ratio of areas of the ROI and the

whole domain. The overall evacuation then can be measured as a mean of the possible evacuation times across the Bernoulli trials of successful molecule placement inside the evacuation region,

$$\langle \tau_{\text{All}} \rangle = \sum_i \left(\binom{N}{i} \left(\frac{A_{\text{ROI}}}{A_{\text{Domain}}} \right)^i \left(\frac{A_{\text{Domain}} - A_{\text{ROI}}}{A_{\text{Domain}}} \right)^{n-i} \langle \tau_{i,\text{All}} \rangle \right). \quad (\text{A.30})$$

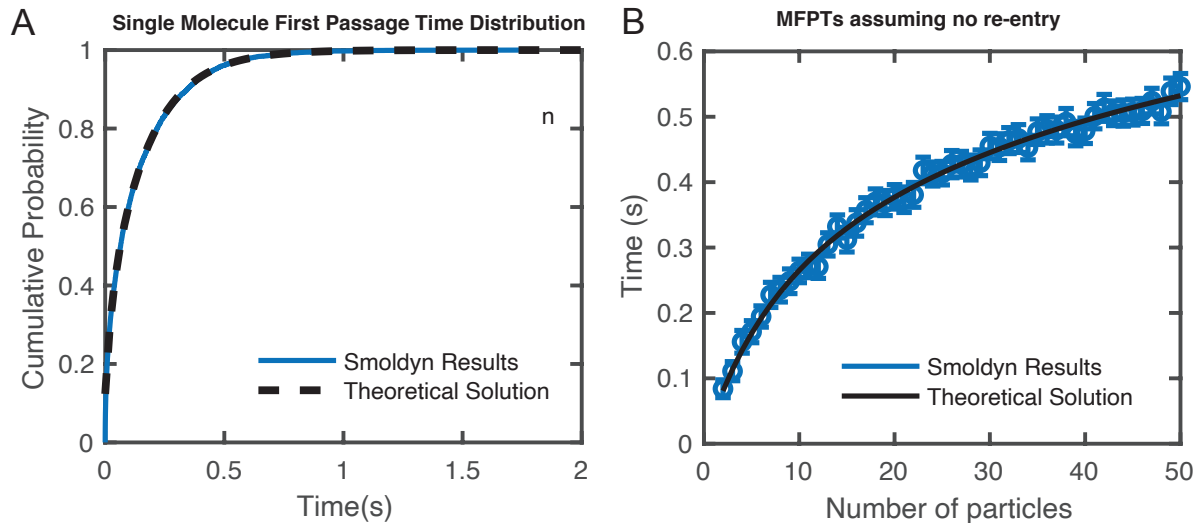


Figure 1.1: Evaluation of analytical solution for MFPT with $p_{\text{entry}} = 0$ shows strong agreement with brute force simulation. (A) Cumulative distribution for the first passage time of a single molecule to leave the ROI from analytical expression in Eq. A.21 (black) and Smoldyn results come as the result of 20000 individual simulations (blue). (B) MFPT for multiple molecules inside of the domain, analytically from Eq. A.29 (black) and computationally from Smoldyn (blue). In addition to the N values shown, an analytic value for $N=256$ is given in Fig. 2.6C. Error bars for Smoldyn results are given as the standard error of the mean for 500 Smoldyn simulations per data point.