Stochastic Spatiotemporal Simulation of a General Reaction System

Andrew J. Loza[a,*], Marc S. Sherman[b]

[a] *Yale University School of Medicine, New Haven, CT, United States of America*
[b] *Division of Gastroenterology, Massachusetts General Hospital, Boston, MA, United States of America*

Abstract

Biological systems frequently contain biochemical species present as small numbers of slowly diffusing molecules, leading to fluctuations that invalidate deterministic analyses of system dynamics. The development of mathematical tools that account for the spatial distribution and discrete number of reacting molecules is vital for understanding cellular behavior and engineering biological circuits. Here we present an algorithm for an event-driven stochastic spatiotemporal simulation of a general reaction process that bridges well-mixed and unmixed systems. The algorithm is based on time-varying particle probability density functions whose overlap in time and space is proportional to reactive propensity. We show this to be mathematically equivalent to the Gillespie algorithm in the specific case of fast diffusion. We develop a computational implementation of this algorithm and provide a Fourier transformation-based approach which allows for near constant computational complexity with respect to the number of individual particles of a given species. To test this simulation method, we examine reaction and diffusion limited regimes of a bimolecular association-dissociation reaction. In the reaction limited regime where mixing occurs between individual reactions, equilibrium numbers of components match the expected values from mean field methods. In the diffusion limited regime, however, spatial correlations between newly dissociated species persist, leading to rebinding events and a shift the in the observed molecular counts. In the final part of this work, we examine how changes in enzyme efficiency can emerge from changes in diffusive mobility alone, as may result from protein complex formation.

1. Introduction

Biological processes are often characterized by slowly diffusing components present in small numbers, as well as reactions that depend nonlinearly on substrate concentrations and may not be at equilibrium [1, 2, 3]. In this realm, fluctuations in the number and spatial distribution of the participating species become non-trivial and can lead to functional consequences [4]. This variation has been demonstrated to play a role in a wide variety of cellular processes, including fate in bacterial

*Corresponding author
Email address: andrew.loza@yale.edu (Andrew J. Loza)

populations following viral infection or nutrient restriction [5, 6], phenotypic variation in genetically identical populations [7, 8, 9], and cell differentiation in organism development [10].

Mathematical models are critical for understanding and predicting the behavior of biological systems, but deterministic models often break down in describing biological dynamics [11]. Molecular concentration becomes discontinuous in the spatial domain, due to slow diffusion, and in magnitude, due to countable fluctuations in the number of molecules. To understand the details of when and how fluctuations influence the behavior of biological systems, a variety of mathematical models have been developed to capture the discrete and stochastic nature of the systems involved [12, 13, 14, 15, 16, 17]. A computational approach was provided by Gillespie to simulate the chemical master equation [12]. In the original work, spatial distribution of molecules was assumed to be uniform, but an extension of the model using subvolumes allowed for simulation of spatial heterogeneity provided each subvolume fulfilled the well-mixed requirement [18, 14, 13]. Related work divided space into metastable compartments, simulating reactions within compartments and transitions between them [19]. A second class of simulation techniques relies on a molecular dynamics type approach, simulating the diffusive trajectories of species involved [15]. For simulations when the diffusive steps are small compared to inter-particle distance, this approach can become inefficient. To remedy this, methods referred to as Green's Function Reaction Dynamics (GFRD) and First-Passage Kinetic Monte Carlo (FPKMC) provided a strategy to jump forward in time to when interactions occur [16, 20, 17].

While these stochastic and discrete models extend deterministic and continuous approaches, they retain a few disadvantages. Subvolume methods may be influenced by the selection of subvolume size, with potential inaccurate results [21], and GFRD and FPKMC require careful treatment of certain spatial arrangements of molecules that, if not explicitly accounted for, can limit computational efficiency [17, 22, 23, 24]. Furthermore, processes that span timescales and length-scales may be coupled in biological systems, complicating the choices required to efficiently and accurately implement these algorithms. Therefore, further development of methods that address these shortcomings is warranted to advance the field.

Here we present the framework and implementation of an event driven stochastic spatiotemporal simulation of a generalized reaction process which bridges diffusion-limited stochastic kinetics with well-mixed reaction kinetics. The analytic framework of this algorithm is unique in its ability to span well-mixed and non-mixed systems, providing a seamless way to simulate complex biological processes. In the first part of this work, we develop the underlying reaction probability density function and the sampling procedure to determine reaction type, location, and participating molecules. We additionally demonstrate that this method is mathematically equivalent to the Gillespie algorithm in the specific limit of rapid diffusion. In the second part of this work, we provide an implementation for a bimolecular association-dissociation reaction and demonstrate the necessity

for reaction-diffusion modeling in predicting expected values when re-binding probabilities are high. We then provide strategies for increasing computational efficiency for systems with large numbers of particles. Lastly, we examine the influence of protein complex formation on enzyme efficiency solely through changes in diffusive mobility. This framework relies on general particle probability density functions without making assumptions about their forms. Here we focus on diffusive motion, however this framework could be used to simulate systems with more complex processes driving particle mobility.

2. Simulation Algorithm

2.1. System Definitions

Let the general system be defined as a set of M mobile species S_i that may interconvert through N reactions R_μ within a volume V. We are interested in systems with a discrete number of members of each species and therefore define X_i as the current number of molecules of S_i. Individual particle identity is accounted for through a second subscript j such that S_{ij} refers to the jth molecule of the ith species, where $1 \leq j \leq X_i$.

2.2. Derivation of the Reaction Probability Density Function

The Gillespie algorithm requires the computation of the reaction probability density function, $P(\tau, \mu)$, defined as the probability at time τ that the next reaction in V will occur in the differential time interval $\delta\tau$ and will be of type μ [12]. The goal of the present work is to extend this joint probability density function to include location and molecular identity by deriving $P(\tau, \mu, \vec{r}, S_{ij})$: The probability at time τ that the next reaction in V will occur in the differential interval $\delta\tau$; *and* will be of type μ; *and* will occur at position $\vec{r}$; *and* will involve the specific molecules S_{ij}.

We first derive the univariate distribution $P(\tau)$ and then decompose it into the multivariate distribution $P(\tau, \mu, \vec{r}, S_{ij})$, which can be sampled using Monte Carlo methods. Here we focus on a bimolecular reaction between two distinct species, but the derivation can be generalized for multiple reaction types. The reaction is:

$$S_1 + S_2 \longrightarrow S_3. \tag{1}$$

Two molecules have the potential to react if they collide with proper orientation and energy. A collision will take place whenever the centers of two spherical molecules come closer than:

$$d_{12} = (d_1 + d_2)/2, \tag{2}$$

Where d_{12} is the collision distance threshold, d_1 is the diameter of an S_1 molecule, and d_2 is the diameter of an S_2 molecule. If the average speed of an S_1 molecule relative to an S_2 molecule is $\langle v_{12} \rangle$, the S_1 molecule produces a potential collision volume over time that can be defined as:

$$\langle \delta V_{coll} \rangle = \pi {d_{12}}^2 \langle v_{12} \rangle \delta\tau. \tag{3}$$

By determining the probability that this collision volume is occupied by an S_2 molecule, the likelihood that the collision was successful, and the number of potential reactive collisions, Gillespie defines the total probability that a reaction R_μ will occur in V in the next time interval $\delta\tau$ as $h_\mu c_\mu \delta\tau$, where h_μ is the number of distinct combinations of reactant molecules and c_μ is the probability that a particular combination of reactant molecules will indeed react [12]. The present goal is to determine the value of $h_\mu c_\mu \delta\tau$ without making the assumption of spatial homogeneity. To begin, we define $z_\mu(\tau)\delta\tau$ as the probability that the next reaction R_μ will occur in V in the next time interval $\delta\tau$ beginning at time τ. Importantly, while $h_\mu c_\mu$ is constant with respect to time for any inter-reaction interval, z_μ may be a function of time.

A volume V can be broken into an infinite number of small sub-volumes in which the next reaction could occur. Because the probability of two reactions occurring in the time interval $\delta\tau$ vanishes as $\delta\tau$ goes to zero, the probabilities of this single next reaction being in each sub-volume are disjoint. This means, through the addition rule of disjoint events, that the sum of probabilities in all these sub-volumes is equal to $z_\mu \delta t$. In the limit of infinite sub-volumes, summation becomes integration, and we define a new variable, $\hat{z}_\mu$:

$$z_\mu(\tau) = \int_V \hat{z}_\mu(\tau, \vec{r}) d\vec{r}, \tag{4}$$

which is the probability density of the next reaction occurring at a specific position $\vec{r}$ in the differential interval beginning at time τ.

In a general sense, $\hat{z}_\mu$ describes a conditional probability, namely:

$$\hat{z}_\mu = P(A, B) = P(A|B) \cdot P(B), \tag{5}$$

Where event A is an R_μ reaction and event B is the presence of an S_1 molecule at a location. Computation of $P(B)$ will be addressed first, followed by $P(A|B)$.

The term $P(B)$, or the probability that an S_1 molecule is at a specific location, can be defined as follows. If the location of a molecule is known at a time $t = 0$, its location at $t \geq 0$ is unknown precisely, but it is describable in terms of probability if the driving forces behind motion are known. Here we assume that thermal fluctuations drive molecular motion, yielding a Brownian diffusion process, however this method can be generalized to any source of motion that can be described in terms of a probability density function for each member of each species: $P_{S_{ij}}(\vec{r}, t)$. When multiple S_i molecules are present, a pseudo-probability density function describing probabilities of finding S_i at a point can be written:

$$\mathbb{P}_{S_i}(\vec{r}, \tau) = \sum_{j=1}^{X_i} P_{S_{ij}}(\vec{r}, \tau). \tag{6}$$

This is not a true probability density function because while the different $P_{S_{ij}}$ are independent, they are not disjoint. Therefore, the union of the individual $P_{S_{ij}}$ probability distributions is not simply the sum of probabilities of individual events but rather must be adjusted by the intersection of these probabilities. However, the cost of neglecting the intersection terms depends on the absolute probabilities of events, and as the absolute probabilities of the distinct events approach zero, the fractional error from neglecting their intersection becomes negligible (the probability of individual events is $\mathrm{O}(\delta)$, the probability of intersection events is $\mathrm{O}(\delta^2)$ and higher). Therefore, if we consider a region where all the individual $P_{S_{ij}}$ are high and query the probability that any S_i molecules are in that region, the integration of $\mathbb{P}_{S_i}$ over the spatial limits of that region will be a poor approximation of the true probability (integrating $\mathbb{P}_{S_i}$ over V makes this readily apparent, as a value greater than 1 is seen). If, however, the region of interest is very small, such as δV_{coll}, then the probability of individual molecules being within the region is small (computed by integrating $P_{S_{ij}}$ over the limits of the region δV_{coll}), and as $\delta\tau$ drops to zero, the probability of two particles being within the collision volume is negligible. Thus, $P(B)$ in (5), or the probability that an S_1 molecule is within an infinitesimal volume in space, can be well approximated by $\mathbb{P}_{S_1}$:

$$P(B) = \mathbb{P}_{S_1}. \tag{7}$$

The first term on the right hand side of (5), $P(A|B)$, is the probability that a reaction will occur given there is an S_1 molecule at a specific location. Following the logic used by Gillespie, we break this into the probability that an S_2 molecule is within the collision volume of a given S_1 molecule and the probability that such a collision is productive. Using θ_R as the probability a collision is productive, $P(A|B)$ becomes:

$$P_{A|B}(\vec{r}) = \int_{\delta V_{coll}} \mathbb{P}_{S_2}(\vec{r}) d\vec{r} \cdot \theta_R, \tag{8}$$

Where $\vec{r}$ is the location of a given S_1 molecule and δV_{coll} is the collision volume as defined in (3). The sum rule approximation employed here by the use of $\mathbb{P}_{S_2}$ is valid because the collision volume is vanishingly small. The probability of any single molecule being within it is small, so any higher order terms that result from the product of individual probabilities will be negligible. To simplify this expression, we recognize that δV_{coll} is small such that $\mathbb{P}_{S_2}$ is essentially constant throughout this domain. Thus (8) becomes:

$$P_{A|B} = \mathbb{P}_{S_2} \delta V_{coll} \cdot \theta_R. \tag{9}$$

We now can re-write (5) substituting in (3), (7), (9):

$$\hat{z}_\mu \delta t = \mathbb{P}_{S_1} \mathbb{P}_{S_2} \left(\pi {d_{12}}^2 \langle v_{12} \rangle \right) \theta_r \cdot \delta t. \tag{10}$$

Integrating $\hat{z}_\mu \delta t$ with respect to its spatial domain over a volume V gives the probability that a reaction occurs in the volume, $z_\mu \delta t$, as defined in (4). The only terms that depend on space are the pseudo-probability distributions, $\mathbb{P}_{S_1}$ and $\mathbb{P}_{S_2}$, allowing for simplification:

$$z_\mu(\tau)\delta t = \int_V \mathbb{P}_{S_1}\mathbb{P}_{S_2}\left(\pi {d_{12}}^2 \langle v_{12}\rangle\right)\theta_R d\vec{r}\cdot\delta t, \tag{11}$$

$$z_\mu(\tau)\delta t = \left(\pi {d_{12}}^2 \langle v_{12}\rangle\right)\theta_R \int_V \mathbb{P}_{S_1}\mathbb{P}_{S_2} d\vec{r}\cdot\delta t. \tag{12}$$

Defining the quantity of $z_\mu \delta t$ allows us to determine the reaction probability distribution function $P(\tau,\mu)$, the probability that the next reaction occurs at time τ and is a μ reaction. This probability can be decomposed into two components. The first component is $P_0(\tau)$, the probability that no reaction occurs over time τ. The second component is $z_\mu(\tau)d\tau$, the probability that a reaction occurs in an infinitesimal window right after τ and has already been defined in (12). Because reactions are independent, random events that occur with an underlying rate, the Poisson distribution describes their occurrences. $P_0(\tau)$ is the zero-event probability density function resulting from the inhomogeneous Poisson process governed by the time-dependent event rate $z_\mu(\tau)$. Therefore the probability that no reactions (of any type, not just μ) occur in a time interval τ is given by the following equations:

$$P_0(\tau) = P(n_R(\tau)=0) = \frac{\Lambda(\tau)^0}{0!}e^{-\Lambda(\tau)}, \tag{13}$$

$$\Lambda(\tau) = \int_0^\tau \left(\sum_{\nu=1}^{N} z_\nu(\xi)\right) d\xi, \tag{14}$$

Where n_R is the number of reactions and ξ is a dummy time-variable. Combining (12), (13), and (14), we can now express $P(\tau,\mu)$ in terms of the two parts described above:

$$P(\tau,\mu) = z_\mu(\tau)\exp\left(-\int_0^\tau \left[\sum_{\nu=1}^{N} z_\nu(\xi)\right] d\xi\right). \tag{15}$$

The probability that a specific R_μ reaction happens is a function of all other possible reactions due to the summation term in the exponential. For the next reaction to be of type R_μ, the reactions $R_{\nu\neq\mu}$ and the specific reaction R_μ must not yet have occurred. The final step in computing $P(\tau)$ is to sum over all possible reaction types:

$$\begin{aligned} P(\tau) &= \sum_{\mu=1}^{N} P(\tau,\mu) \\ &= \left(\sum_{\mu=1}^{N} z_\mu(\tau)\right)\exp\left(-\int_0^\tau \left[\sum_{\nu=1}^{N} z_\nu(\xi)\right] d\xi\right). \end{aligned} \tag{16}$$

This is the probability density function describing the likelihood of the wait-time until the next reaction. To verify that this probability density function is properly normalized over its defined domain, we must show that $\int_0^\infty P(\tau) = 1$. We recognize that (16) takes the following form:

$$P(\tau) = f'(\tau)e^{-\int_0^\tau f'(\xi)d\xi}, \tag{17}$$

Where $f'(\tau) = \left(\sum_{\mu=1}^{N} z_\mu(\tau)\right)$. For the set of functions f where $f(0) = 0$ and $f(\tau > 0)$ integrating yields:

$$\int_0^\infty P(\tau) = \int_0^\infty f'(\tau)e^{-(f(\tau)-f(0))} = 1 - e^{-(f(\infty)-f(0))}. \tag{18}$$

For the set of functions where $f'(\tau) \geq 0$ and therefore $f(\tau) \geq f(0)$, (18) is bounded with the following inequality:

$$0 \leq 1 - e^{-(f(\infty)-f(0))} \leq 1. \tag{19}$$

This is true here because for all R_μ reactions, $z_\mu(\tau) \geq 0$. The circumstances where $\int_0^\infty P(\tau) \neq 1$ describe situations where there is a finite probability that no further reactions occur and provide insight on the reactive process: for a reaction to occur, $f(\infty)$ must not remain bounded and therefore f' must not decay to 0.

2.3. Gillespie Algorithm Equivalence

We are now at a point to ask if this reaction probability density function is equivalent to the Gillespie derivation in the limit of a fast diffusion rates relative to reaction rates. We begin by demonstrating the following equality:

$$z_\mu \delta t = h_\mu c_\mu \delta t. \tag{20}$$

In a bounded system governed by the limit of fast diffusion, each molecule has equal probability of being anywhere within V, so the pseudo-probability distributions defined in (6) are sums of the uniform distribution over V:

$$\mathbb{P}_{S_1} = \frac{X_1}{V}, \mathbb{P}_{S_2} = \frac{X_2}{V}. \tag{21}$$

From reference [12], the right hand side of (20) is equal to:

$$h_\mu c_\mu \delta t = X_1 X_2 \left(V^{-1}\pi {d_{12}}^2 \langle v_{12}\rangle \delta t\right) \theta_R. \tag{22}$$

Substituting the values from (21) into (12) gives:

$$z_\mu(\tau)\delta t = \left(\pi {d_{12}}^2 \langle v_{12}\rangle)\right) \theta_R \int_V \left[\frac{X_1}{V} \cdot \frac{X_2}{V}\right] d\vec{r} \cdot \delta t. \tag{23}$$

All the terms within the integral are constants with respect to position, therefore (23) reduces to:

$$z_\mu(\tau)\delta t = \left(\pi {d_{12}}^2 \langle v_{12}\rangle)\right) \theta_R \left(X_1 X_2 V^{-1}\right) \cdot \delta t. \tag{24}$$

Which is exactly equal to $h_\mu c_\mu \delta t$ in (22). To establish equivalence of $P(\tau)$ in this limit, we recognize that z_μ defined in (24) is time independent and the inhomogenous Poisson process used to derive (16) becomes a homogenous one:

$$P(\tau) = \sum_{\mu=1}^{N} z_\mu \exp\left(\sum_{\nu=1}^{N} z_\nu \tau\right), \tag{25}$$

Which, using the demonstrated relationship in (20), is equivalent to $P(\tau)$ of the Gillespie algorithm [12].

2.4. Sampling $P(\tau, \mu, \vec{r}, S_{ij})$ by decomposition of $P(\tau)$

The joint probability distribution $P(\tau, \mu, \vec{r}, S_{ij})$ may be sampled through the use of conditional probabilities in a series of steps:

1. What is the wait time τ?
2. Given a wait time τ, what reaction R_μ took place?
3. Given a wait time τ and a specific reaction R_μ, where did it take place?
4. Given a wait time τ, specific reaction R_μ, and location $\vec{r}$, which molecules S_{ij} were involved?

These questions are the four terms on the right-hand side of the following equation and may be answered by sampling each distribution using Monte Carlo methods:

$$P(\tau, \mu, \vec{r}, S_{ij}) = P(\tau)P(\mu|\tau)P(\vec{r}|\tau, \mu)P(S_{ij}|\tau, \mu, \vec{r}). \tag{26}$$

In this section, we proceed through the sampling of each distribution using a computational implementation for a simple bimolecular reaction as a guide (Fig 1A).

What is the wait time τ? The reaction wait-time τ can be directly obtained by inversion sampling of the probability density function $P(\tau)$ shown in (16), providing the time at which the next reaction occurs (Fig 1B). Inversion sampling is a general method for random sampling of a probability distribution given its cumulative distribution function. Cumulative distribution functions are monotonic and map the domain of the probability distribution to a range between 0 to 1. Inversion sampling selects a number on the interval 0 to 1 from a uniform random distribution. The cumulative distribution function maps this value from the range back to a unique value of the domain.

Given a wait time τ, what reaction R_μ took place? The conditional probability $P(\mu|\tau)$ can be expressed using (15) and (16) as:

$$P(\mu|\tau) = \frac{P(\tau,\mu)}{P(\tau)} = \frac{z_\mu(\tau)}{\sum_{\mu=1}^{N} z_\mu(\tau)}. \tag{27}$$

This is the the proportion of the total event probability that each reaction contributes. This is a discrete distribution with respect to the reaction variable μ and can be sampled to obtain the type of reaction that occurs (Fig 1B inset).

Given a wait time τ and a specific reaction R_μ, where did it take place? The third question is answered by sampling the probability distribution $P(\vec{r}|\tau,\mu)$, which represents the relative contribution of each point in space towards the occurrence of the reaction R_μ and is given by the following equation:

$$P(\vec{r}|\tau,\mu) = \frac{P(\tau,\mu,\vec{r})}{P(\tau,\mu)} = \frac{\hat{z}_\mu(\vec{r},\tau)}{z_\mu(\tau)}. \tag{28}$$

Additionally, because $\hat{z}_\mu$ itself is a joint probability distribution over the number of spatial dimensions in the system, further decomposition needs to be done to sample each dimension (Fig 1C).

Given wait time τ, specific reaction R_μ, and location $\vec{r}$, which molecules S_{ij} were involved? To answer the final question, we need to sample the probability distribution $P(S_{ij}|\tau,\mu,\vec{r})$. That is we need to answer the question: given a reaction, time and position, which of the involved molecules participated? This quantity can be expressed by the following:

$$P(S_{ij}|\tau,\mu,\vec{r}) = \frac{P_{S_{ij}}(\vec{r})}{\mathbb{P}_{S_i}(\vec{r})}. \tag{29}$$

This is a discrete probability distribution where the probability of the jth molecule being picked is the proportion of $\mathbb{P}$ it contributes at point $\vec{r}$. Sampling of this probability distribution is again performed using the inversion method as described in Question 2. This computation is performed for each species participating in the reaction (Fig 1D and E).

The reaction step has now been fully characterized. The reacting particles can be removed from the list of S_{ij} and the created particle can be added to S_{ij} at the location of the reaction, updating the quantities X_i. The final step is to update the positions of the molecules that did not react this step but have moved due to diffusion, by sampling the individual probability density functions $P_{S_{ij}}(\vec{r})$. This completes one round of the simulation and leaves the system ready for the next (Fig 1F).

3. Results

3.1. Diffusion-Influenced Bimolecular Reactions

To verify that this algorithm produces expected behavior from reaction-limited to diffusion-limited regimes, we simulated a simple bimolecular association-dissociation reaction on a two dimensional domain. In two dimensions, effects due to local concentration differences as might arise in a reaction-diffusion system are exaggerated because the probability of two given molecules reacting is much more strongly dependent on separation distance than in three dimensions [11]. The bimolecular reaction is defined as $S_1 + S_2 \rightleftarrows^{k_f}_{k_r} S_3$, where k_f and k_r are the microscopic rate constants for association and dissociation, respectively. The algorithm was implemented using the MatLab programming language (Natick, MA).

This reaction was simulated across a range of values for k_r/k_f and diffusion constants for species S_1 and S_2. Under conditions of fast diffusion where mixing occurred between reactions, the mean count of S_1 closely matched the mean-field expected value Fig 2 (blue points). In the setting of slow diffusion, the effective association rate is expected to depart from the microscopic value due to the persistence of spatial correlations in molecular position created by the dissociation reaction. The dissociation of an S_3 molecule produces an S_1 molecule and an S_2 molecule separated by only a short distance. These two molecules can either diffuse away from each other or re-associate. Under conditions of fast diffusion, re-association is less likely and the temporary spatial correlation between the two reactants is more quickly lost. Under conditions of slow diffusion, re-association frequently occurs, shifting the expected values for species counts. Indeed as the diffusion constants for S_1 and S_2 molecules were lowered, the observed S_1 count shifted to the right, indicating that fewer S_1 molecules were present than expected given the rate constants alone Fig 2 (red and yellow points).

3.2. Efficient Computational Strategies

A fundamental question for all reaction-diffusion simulations is how computational time scales as the number of molecules in the system increase. An ideal algorithm would have low start-up computational costs and minimal increases in computational time as the number of molecules increases. In situations where the number of molecules is large, this scaling property may be significantly more important than any start-up costs. The proposed algorithm has the potential to be computationally intensive, especially in regards to the generation of the pseudo-probability distributions $\mathbb{P}_{S_i}(\vec{r}, \tau)$, which requires the summation of individual probability distributions for each member of a species. However, though the use of specific computational strategies, the algorithm can be implemented in an efficient manner with highly desirable scaling in computational time as molecular number increases. This relies on the fact that $P_{S_{ij}}$ evolves identically for a single species except for a shift in position, providing an alternative strategy for computation employing convolution. For particles

whose positions evolve according to a general probability distribution function $f(\vec{r}, t)$, the function $\mathbb{P}_{S_i}(\vec{r}, \tau)$ can be re-defined as:

$$\mathbb{P}_{S_i}(\vec{r}, t) = \sum_j f(\vec{r}, t) * \delta(\vec{r} - \vec{r_j}). \tag{30}$$

Where δ is the Dirac delta function. This convolution may be implemented directly or through use of Fourier transforms:

$$\mathbb{P}_{S_i}(\vec{r}, t) = \mathcal{F}^{-1}\left[\mathcal{F}(f(\vec{r}, t)) \cdot \mathcal{F}\left[\sum_j \delta(\vec{r} - \vec{r}_j)\right]\right]. \tag{31}$$

Where $\mathcal{F}$ is the Fourier transform operator. In this approach, the summation of complex functions is avoided by converting calculations to a single multiplication step in frequency-space.

To demonstrate the computational efficiency of the convolution strategy, the time required to compute $\mathbb{P}_{S_i}(\vec{r}, \tau)$ for an increasing number of particles was performed. The direct summation method was compared to direct convolution and Fourier transform based convolution where the total number of particles, N, ranged from 10 to 10^4. While computation time scaled linearly with N for the direct sum method, computation time for both convolution based approaches was independent of particle number (Fig 3). This reflects optimum scaling that may or may not be realizable for all computational implementations of the simulation. It may be desirable to scale the numerical resolution with the particle number to maintain the same strength in the assumption underlying (6). In this case, the number of elements in the simulated domain scales with the number of particles. Therefore the efficiency becomes bounded by that of the Fast Fourier Transform algorithm at $N \log(N)$. Regardless of which approach is taken when increasing the number of particles within a simulation, a computational efficiency that scales for N particles between a lower bound of zero-order and upper bound of $N \log(N)$ is a highly advantageous feature of this algorithm.

3.3. Diffusive Mobility and Enzyme Efficiency

The formation of protein complexes is a ubiquitous part of biological systems [25, 26]. Much work has focused on understanding how conformational changes induced by protein-protein interactions can modulate intrinsic reaction rates through allostery [27]. These interactions also lead to the formation of macromolecular complexes or subcellular compartments where mobility of proteins differs from free diffusion. These structures, like lipid rafts and scaffolds, are known to impact cell signalling[28], but the direct effect of changes in mobility on reaction kinetics are less well understood. Using a molecular dynamics approach, Soula et. al show that changes in mobility can impact apparent binding affinities, but recognize computational challenges in using a molecular dynamics framework for modeling these processes [29]. The algorithm proposed here provides an efficient avenue to study these systems. In this section we examine a model reaction-diffusion system

motivated by quantifying the effect that protein complex formation may have on enzymatic activity simply through the change in diffusivity that results from forming a macromolecular assembly.

Consider a system where a substrate is produced in a spatially restricted manner and acted on by an enzyme that is capable of forming complexes with other proteins. The reaction scheme used in the simulation is shown in Fig 4A. A molecule S is produced by an essentially fixed point source and diffuses away subject to degradation, creating a spatial gradient in the probability of encountering S. An enzyme A binds S and converts it to P. Alternatively, while bound to S, A can be altered in a manner that leads to a change in mobility but does not affect intrinsic enzymatic efficiency. This scenario represents a generalization of a number of specific mechanisms that could lead to spatial localization of an enzyme near a source of substrate. This may reflect a change in phosphorylation state or conformation that leads to protein complex formation, increased lipid-bilayer binding, or any alteration that reduces the diffusion rate of A. As a result, the mobility of A is reduced in a region where S is likely to be found.

This system provides a generalized scenario for examining the effect that protein complex formation has on enzymatic efficiency even if intrinsic catalytic rates remain constant. Figure 4B and C show a snapshot of the enzyme positions at the conclusion of 8000 reaction time-steps for the case of no mobility reduction and a 1000-fold reduction. As expected, enzymes where mobility reduction is high are concentrated in the vicinity of the source of S, indicating successful recruitment near the source of substrate due to the reduction in mobility. The production of P depends on the relative mobility of A^* and is higher in cases where reduction is the greatest but changes as the system evolves in time (Fig 4D-F). For large reductions in mobility, the production rate of P is higher than if no reduction occurs and remains so for thousands of reaction steps. In this case, although the enzymes search space less quickly, they remain in a region with an increased probability of encountering S, leading to higher overall efficiency. For intermediate drops in mobility, however, the production rate of P initially exceeds but eventually falls below the control level. Here, diffusion is not reduced enough to maintain the high degree of initial spatial correlation with the source of S to balance the effect on search-times.

4. Discussion

In this work, we use probabilistic techniques to simulate a general reaction system in both space and time. The approach is driven by reaction events, is capable of spanning diffusion limited to well-mixed regimes, and is shown to be mathematically identical to the Gillespie algorithm in the limit of rapid diffusion. Using a computational implementation of this algorithm, we demonstrate that it captures equilibrium values in well-mixed systems and re-binding events when diffusion is slow. Lastly, in a model of protein complex formation, we find tradeoffs may exist between spatial localization and search-time due to macromolecular assembly. For simplicity, this algorithm will be

referred to as the Spatiotemporal Gillespie algorithm (STG).

The STG method relies on probabilistically jumping between events in order to skip details that do not change the progression of the system in a manner similar to GFRD and FPKMC methods [16, 20, 17]. The primary difference between these methods and STG is the definition of an event and the construction of the underlying wait time distribution. GFRD and FPKMC are event driven, but the event definition is related to probability of multi-way-interactions rather than successful reactions. In the regime of highly diffusion-limited reactions, this is equivalent, but when interactions do not guarantee reactions, these definitions differ. In this case, both GFRD and FPKMC increment through time by simulating many unsuccessful interaction events before a reaction occurs. By focusing on reactive events, the STG algorithm avoids this potential inefficiency that may arise when a near-diffusion-limited process or even a reaction-limited process is part of the system. As a result, the STG algorithm can simultaneously simulate reaction-limited (highlighted by its Gillespie algorithm equivalence) and diffusion-limited processes that share components. This hybrid functionality is a fundamental strength of the STG algorithm.

To achieve this, the STG algorithm relies on fundamental assumptions in both the analytical definition and computational implementation. The primary assumption of the analytical definition is that non-reactive collisions with labeled particles (those that have been given an identity in the system) do not significantly alter the random-walk statistics that arise from collisions with all other unlabeled particles in the environment. This assumption may not hold if electrostatic interactions are involved that disrupt diffusive motion [30]. This is a challenge not unique to this approach, but may be handled here by defining the diffusion-altering event as a separate reaction in the system and creating an intermediate species for more precise treatment.

The primary assumption in the numerical implementation of the STG algorithm is in the numerical resolution used for the pseudo-probability distribution $\mathbb{P}_S$, which approximates the probability that a single molecule resides at a position. In theory, collision volumes are vanishingly small and the probability that two molecules simultaneously fall within it approaches zero. In this limit, the union probability is well approximated by summing individual probabilities and neglecting the intersection probabilities. In practice, the minimum value that individual probabilities may take depends on the resolution chosen for numerical approximation of the individual probability density functions. Therefore, to maintain the accuracy of this assumption, the numerical resolution must be kept high enough that the probability of a single particle of a given type residing in a numerical discretization of space is small such that the probability of two remains negligible. This approach, however, makes the STG approach highly flexible. Any motion that can be probabilistically described can be incorporated into the framework, allowing for the study of systems where body forces play a role and processes in 1, 2, and 3 dimensions. Furthermore, systems with mixed dimensionalities, for example a 1D DNA strand within a 3D environment, could be modeled with the STG framework.

The analysis on computational efficiency demonstrates that the algorithm exhibits efficient scaling as the number of particles increases when convolution-based methods are employed. In the best-case scenario, when the resolution of numerical implementation does not need to be changed, scaling is zero-order: additional particles do not effect the computational time for a given reaction. However, even if the numerical resolution is scaled with the number of particles, the computational complexity follows that of the method chosen for convolution. While the computations required to implement the STG algorithm are more involved than simulating diffusive trajectories or subvolume transitions, its reaction-event driven nature and scaling properties are highly advantageous features as parallelized CPU and GPU computing power continues to grow.

The simple model of protein complex formation presented here demonstrates two fundamental competing factors in reaction diffusion systems: control of protein localization and effects on search times. In all cases, the enzymes A are likely to undergo a reduction in mobility where S molecules are more prevalent and do so with the same kinetics. Such a change may occur when substrate binding influences covalent modification of the enzyme itself, leading to changes in binding partner affinity as is seen in Rho GTPase enzymes [31]. This localization in a zone of higher levels of S leads to increased production rates. As time progresses, enzymes that only underwent a moderate reduction in mobility begin to spread out and production drops. These enzymes are slower to search space and are no longer restricted to a region with high local concentrations of S and as a result, production falls below that of the case where no reduction in mobility occurs. In summary, protein complex formation increases effective enzyme efficiency through localization near a source of substrate. However, maintenance of spatial correlation with the source becomes of increased importance to avoid the disadvantageous lengthening of search-times that results from mobility reduction.

The methods proposed here provide a rigorous framework for assessing the effects of slow diffusion and small numbers of molecular species as well as the conditions where these considerations become important. Furthermore, this unique approach unifies reaction-limited and diffusion-limited regimes within a single framework, allowing of straightforward simulation of biological systems where reaction of both regimes are coupled. We anticipate our framework will help decipher the complex interplay between reaction and diffusion in biological signaling.

5. Acknowledgments

We would like to greatly thank Elliot Elson for his input and many useful discussions. We would also like to acknowledge Mauricio Del Razo Sarmina and Hong Qian for their valuable and constructive feedback. This work was supported by the National Institutes of Health Grants T32 EB018266 and T32 GM007200 to A.J.L.

References

References

[1] B. Huang, H. Wu, D. Bhaya, A. Grossman, S. Granier, B. K. Kobilka, R. N. Zare, Counting low-copy number proteins in a single cell, Science 315 (5808) (2007) 81–84.

[2] E. M. Ozbudak, M. Thattai, H. N. Lim, B. I. Shraiman, A. Van Oudenaarden, Multistability in the lactose utilization network of escherichia coli, Nature 427 (6976) (2004) 737–740. URL `http://www.nature.com/nature/journal/v427/n6976/pdf/nature02298.pdf`

[3] J. Elf, J. Paulsson, O. G. Berg, M. Ehrenberg, Near-critical phenomena in intracellular metabolite pools, Biophysical journal 84 (1) (2003) 154–170. URL `http://ac.els-cdn.com/S0006349503748395/1-s2.0-S0006349503748395-main.pdf?_tid=7aaa82be-60f6-11e6-b072-00000aacb35e&acdnat=1471052421_b46ab528a21c1ab24961bc0ac3d4e8e8`

[4] O. G. Berg, A model for the statistical fluctuations of protein numbers in a microbial population, Journal of theoretical biology 71 (4) (1978) 587–603.

[5] A. Arkin, J. Ross, H. H. McAdams, Stochastic kinetic analysis of developmental pathway bifurcation in phage λ-infected escherichia coli cells, Genetics 149 (4) (1998) 1633–1648.

[6] H. Maamar, A. Raj, D. Dubnau, Noise in gene expression determines cell fate in bacillus subtilis, Science 317 (5837) (2007) 526–529.

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

[8] W. J. Blake, M. Kærn, C. R. Cantor, J. J. Collins, Noise in eukaryotic gene expression, Nature 422 (6932) (2003) 633–637.

[9] W. J. Blake, G. Balázsi, M. A. Kohanski, F. J. Isaacs, K. F. Murphy, Y. Kuang, C. R. Cantor, D. R. Walt, J. J. Collins, Phenotypic consequences of promoter-mediated transcriptional noise, Molecular cell 24 (6) (2006) 853–865.

[10] M. F. Wernet, E. O. Mazzoni, A. Çelik, D. M. Duncan, I. Duncan, C. Desplan, Stochastic spineless expression creates the retinal mosaic for colour vision, Nature 440 (7081) (2006) 174–180.

[11] A. Mahmutovic, D. Fange, O. G. Berg, J. Elf, Lost in presumption: stochastic reactions in spatial models, Nature methods 9 (12) (2012) 1163–1166. URL `http://www.nature.com/nmeth/journal/v9/n12/pdf/nmeth.2253.pdf`

[12] D. T. Gillespie, A general method for numerically simulating the stochastic time evolution of coupled chemical reactions, Journal of computational physics 22 (4) (1976) 403–434.

[13] J. Hattne, D. Fange, J. Elf, Stochastic reaction-diffusion simulation with mesord, Bioinformatics 21 (12) (2005) 2923–2924.

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

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

[16] J. S. van Zon, P. R. Ten Wolde, Simulating biochemical networks at the particle level and in time and space: Green's function reaction dynamics, Physical review letters 94 (12) (2005) 128103.

[17] T. Oppelstrup, V. V. Bulatov, A. Donev, M. H. Kalos, G. H. Gilmer, B. Sadigh, First-passage kinetic monte carlo method, Physical Review E 80 (6) (2009) 066701.

[18] M. Malek-Mansour, J. Houard, A new approximation scheme for the study of fluctuations in nonuniform nonequilibrium systems, Physics Letters A 70 (5) (1979) 366–368.

[19] S. Winkelmann, C. Schütte, The spatiotemporal master equation: Approximation of reaction-diffusion dynamics via markov state modeling, The Journal of Chemical Physics 145 (21) (2016) 214107.

[20] J. S. van Zon, P. R. Ten Wolde, Green's-function reaction dynamics: a particle-based approach for simulating biochemical networks in time and space, The Journal of chemical physics 123 (23) (2005) 234910.

[21] M. Chen, F. Li, S. Wang, Y. Cao, Stochastic modeling and simulation of reaction-diffusion system with hill function dynamics, BMC systems biology 11 (3) (2017) 1–11.

[22] K. Takahashi, S. Tănase-Nicola, P. R. Ten Wolde, Spatio-temporal correlations can drastically change the response of a mapk pathway, Proceedings of the National Academy of Sciences 107 (6) (2010) 2473–2478.

[23] D. Fange, O. G. Berg, P. Sjöberg, J. Elf, Stochastic reaction-diffusion kinetics in the microscopic limit, Proceedings of the National Academy of Sciences 107 (46) (2010) 19820–19825.

[24] I. Hepburn, W. Chen, S. Wils, E. De Schutter, Steps: efficient simulation of stochastic reaction–diffusion models in realistic morphologies, BMC systems biology 6 (1) (2012) 1.

[25] N. J. Krogan, G. Cagney, H. Yu, G. Zhong, X. Guo, A. Ignatchenko, J. Li, S. Pu, N. Datta, A. P. Tikuisis, Global landscape of protein complexes in the yeast saccharomyces cerevisiae, Nature 440 (7084) (2006) 637–643.

[26] E. D. Levy, E. B. Erba, C. V. Robinson, S. A. Teichmann, Assembly reflects evolution of protein complexes, Nature 453 (7199) (2008) 1262–1265.

[27] J. Kuriyan, D. Eisenberg, The origin of protein interactions and allostery in colocalization, Nature 450 (7172) (2007) 983–990.

[28] H. E. Grecco, M. Schmick, P. I. Bastiaens, Signaling from the living plasma membrane, Cell 144 (6) (2011) 897–909.

[29] H. Soula, B. Caré, G. Beslon, H. Berry, Anomalous versus slowed-down brownian diffusion in the ligand-binding equilibrium, Biophysical journal 105 (9) (2013) 2064–2073.

[30] Y. Liu, E. Fratini, P. Baglioni, W.-R. Chen, S.-H. Chen, Effective long-range attraction between protein molecules in solutions studied by small angle neutron scattering, Physical review letters 95 (11) (2005) 118102.

[31] M. Liu, F. Bi, X. Zhou, Y. Zheng, Rho gtpase regulation by mirnas and covalent modifications, Trends in cell biology 22 (7) (2012) 365–373.

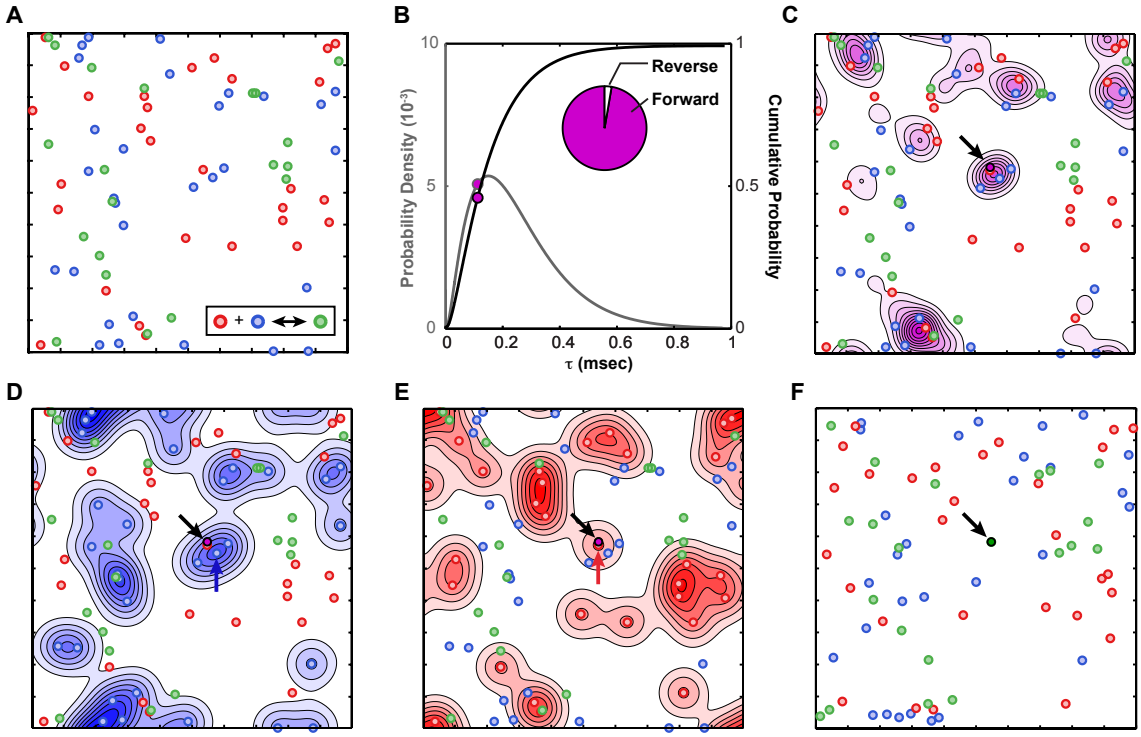

Figure 1: **Simulation of a single reaction step.**

A. Initial positions of molecules. Molecules are randomly distributed over the environment. B. Reaction probability density function, $P(\tau)$ shown in gray and the cumulative probability is shown in black. The cumulative probability distribution is sampled via the inversion method, and the resulting value of τ is the wait time (magenta dots). The relative probabilities of the forward and reverse reactions occurring at this wait time is shown in the inset; random sampling selects the forward reaction. C. The spatial distribution of the reaction likelihood, $\hat{z}_\mu$, and the sampled reaction location (magenta dot). D, E. The location density functions $\mathbb{P}_{S_{ij}}$ for S_1 and S_2 respectively. The S_1 and S_2 molecules likely to have participated are selected from these distributions. F. The updated system with removal of reacting S_1 and S_2, new S_3, and displacement by diffusion of non-reacting molecules.

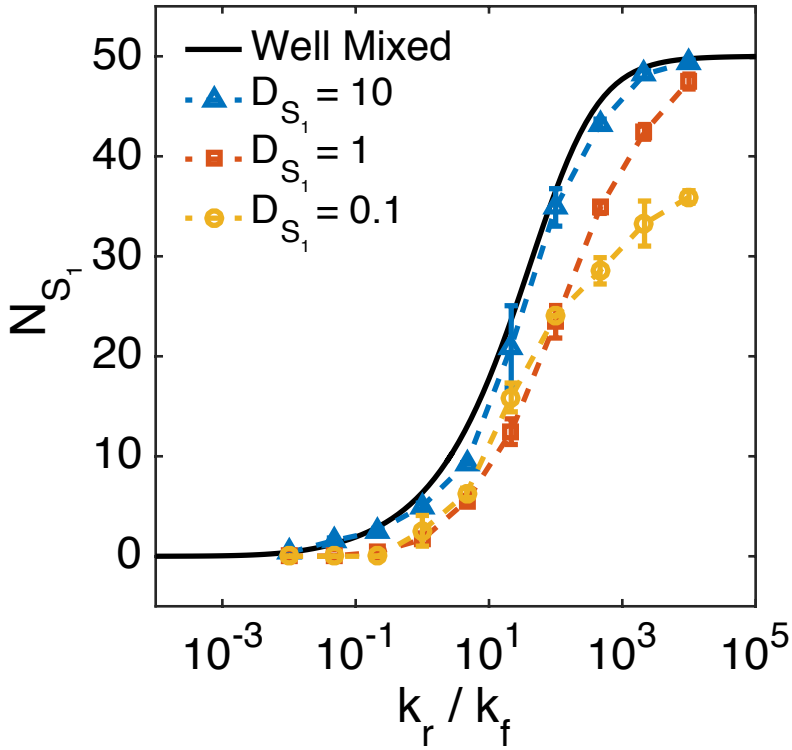

Figure 2: **Association-dissociation reaction with re-binding effects.**

Expected number of S_1 molecules computed under a mean-field approximation or simulated as a function of reactant diffusion constant. The results of the mean-field approach are shown in the black line, and simulation results are shown by the dashed lines and symbols. The diffusion constant for S_1 ranged from $10\mu\text{m}^2\text{s}^{-1}$ to $0.1\mu\text{m}^2\text{s}^{-1}$. Additional parameter values used in the simulation are as follows: initial $S_1 = 30$, initial $S_2 = 30$, initial $S_3 = 20$, $D_{S_1} = D_{S_2}$, $D_{S_3} = D_{S_1}/\sqrt{2}$, $k_f = 10\mu\text{m}^2\text{s}^{-1}$ and k_r is set by k_f and the ratio given by the x-axis. The simulation size was 1μm by 1μm, solved on a 200 by 200 point numerical domain. D_{S_1}, D_{S_2}, and D_{S_3} are diffusion coefficients of the species given by the subscripts.

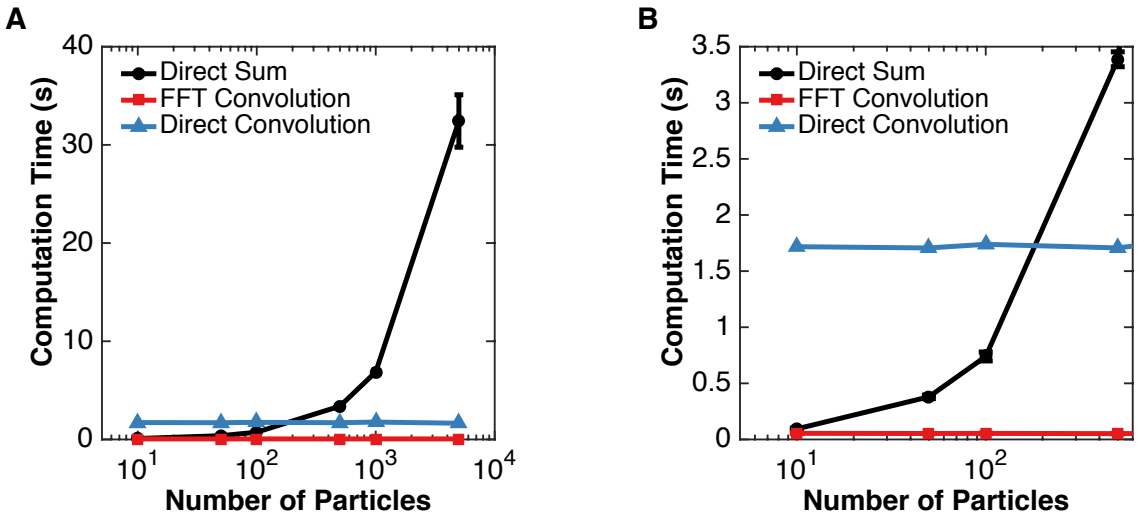

Figure 3: **Comparison of computational efficiency for summation and convolution based approaches.** A. Time to compute $\mathbb{P}_S$ as a function of number of S molecules for summation, direct convolution, and Fourier transform convolution. B. Magnification of cases between 10 and 500 molecules to highlight relative efficiencies for small numbers of molecules. Parameter values used for testing computational efficiency were: $D = 1\mu\text{m}^2\text{s}^{-1}$, $\tau = 10^{-3}$s, where D is the diffusion coefficient and τ is as defined in the main text. The domain size was 1μm by 1μm discretized on a 1024 by 1024 grid.

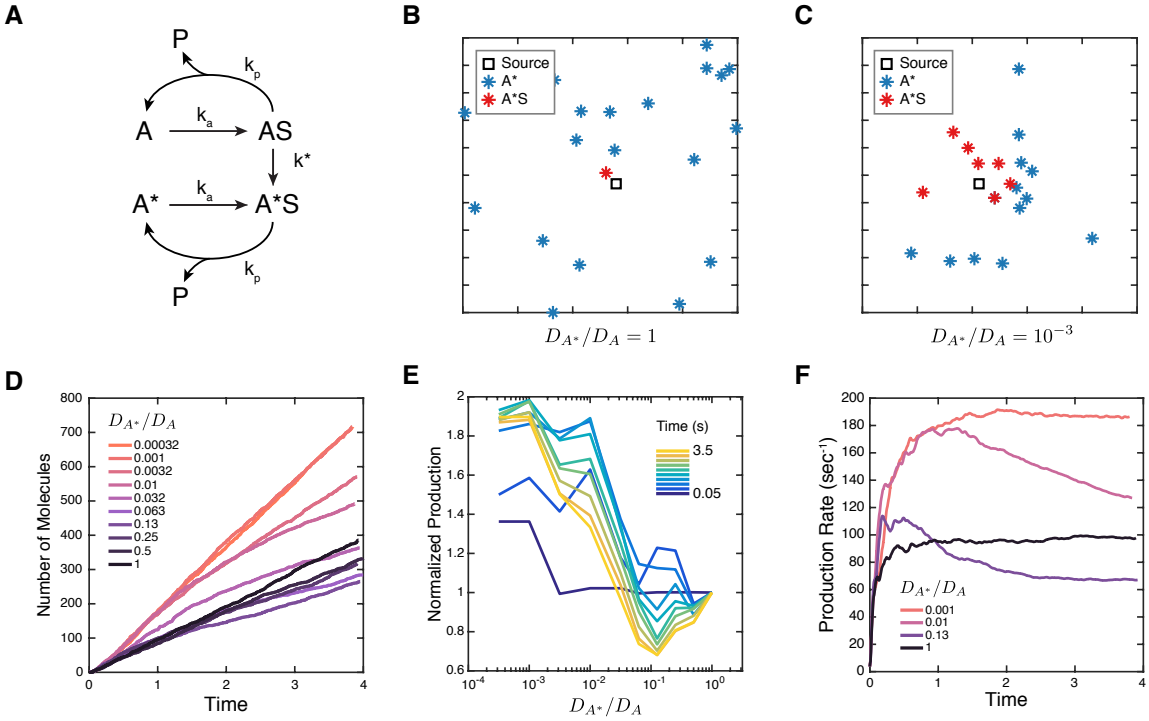

Figure 4: **Changes to enzyme mobility alter effective production rates.**

A. Diagram of reaction. Substrate S is produced by a relatively stationary point source. Enzyme A binds this substrate and converts it to a product P. Alternatively, when bound to the substrate, A can be converted to a form with lower mobility, A* with identical intrinsic catalytic ability to turn S to P. B. Snapshot of enzyme positions and occupancy after 8000 reaction time-steps with no reduction in mobility. The source of S is denoted by the black square. C. Equivalent to B but with a thousand-fold reduction in mobility. Enzymes are clustered near the substrate source and are more likely to be bound to S. D. Total number of P molecules produced as a function of time for a range of reductions in diffusivity. E. Equivalent data to D showing total production as a function of diffusion constant reduction at different times. For intermediate reductions in diffusion, production initially exceeded but then fell below the control case. F. Production rate for selected reductions in mobility spanning three orders of magnitude, computed by a sliding linear-regression to find the slope of curves in D. For high reductions in mobility, production rate doubles the control case. For lesser reductions, production rate initially exceeds the control case but diminishes with time. Parameter values used in the simulation were: $D_S = D_A = 1\mu\text{m}^2\text{sec}^{-1}$, substrate production rate $= 1000$ sec^{-1}, substrate decay rate $= 300$ s^{-1}, $k_a = 10\mu\text{m}^2\text{s}^{-1}$, $k^* = 100\mu\text{m}^2\text{s}^{-1}$, $k_p = 20\mu\text{m}^2\text{s}^{-1}$. The domain size was 1μm by 1μm discretized on a 200 by 200 grid. Here D variables denote diffusion coefficients and k variables denote reaction rates.