

Coupled mesoscopic and microscopic simulation of stochastic reaction–diffusion processes in mixed dimensions

Andreas Hellander¹, Stefan Hellander¹, Per Lötstedt^{1,*}

¹*Division of Scientific Computing, Department of Information Technology
Uppsala University, P. O. Box 337, SE-75105 Uppsala, Sweden
email: andreas.hellander@it.uu.se, stefan.hellander@it.uu.se, perl@it.uu.se*

Abstract

We present a new simulation algorithm that allows for dynamic switching between a mesoscopic and a microscopic modeling framework for stochastic reaction-diffusion kinetics. The more expensive and more accurate microscopic model is used only for those species and in those regions in space where there is reason to believe that a microscopic model is needed to capture the dynamics correctly. The microscopic algorithm is extended to simulation on curved surfaces in order to model reaction and diffusion on membranes. The accuracy of the method on and near a spherical membrane is analyzed and evaluated in a numerical experiment. Two biologically motivated examples are simulated in which the need for microscopic simulation of parts of the system arises for different reasons. First, we apply the method to a model of the phosphorylation reactions in a MAPK signaling cascade where microscale methods are necessary to resolve fast rebinding events. Then a model is considered for transport of a species over a membrane coupled to reactions in the bulk. The new algorithm attains an accuracy similar to a full microscopic simulation by handling critical interactions on the microscale, but at a significantly reduced cost by using the mesoscale framework for most parts of the biological model.

Keywords. multiscale model, stochastic simulation, chemical reaction, single molecule, diffusion, mesoscale, microscale, hybrid method, URDME

AMS subject classification. 65C05, 65C35, 92C05, 92C42

1 Introduction

A recurring topic in theoretical studies of the dynamic behavior of gene regulatory networks in biological cells is their robustness and response to stochastic fluctuations in the number of proteins or other macromolecules [21, 46, 47]. Small copy numbers of a majority of the key components in such networks render the classical deterministic, macroscopic description of the process unsuitable. Instead, intrinsic noise is incorporated in the models by the description of the reaction network as a stochastic process. Even though the majority of work in this area of molecular systems biology is staged in a coarse grained, well-stirred setting, there are recent examples where the interplay between reactions and molecular transport are important for the model’s dynamics [25, 48]. Advances in experimental techniques are making it possible to study not only the total copy number of proteins or mRNA in a cell, but also their position in the cell with single molecule resolution [20, 41]. As such techniques develop further, simulation methods capable of accurate and efficient stochastic simulation of spatially heterogeneous stochastic systems will continue to grow in importance.

*To whom correspondence should be addressed.

The most frequently used stochastic model in the context of modeling biochemical reaction networks in single cells is posed on the so called mesoscopic scale. It is based on a continuous time discrete space Markov process and a realization of the system can be obtained by the Stochastic Simulation Algorithm (SSA) [27]. In the generalization of the well-stirred mesoscopic model to the spatially dependent case, the cellular domain is subdivided into voxels or subvolumes. The state of the system is given by the number of molecules of each species in every voxel. The molecules in a voxel can react with each other and diffuse to adjacent voxels. Realizations of this process can be efficiently generated by adaptation of the SSA to space dependent problems as in [16, 22, 28].

In some cases the mesoscopic model is, being limited in spatial resolution to the mesh size, not capable of capturing critical features of the model [48] unless the mesh is extremely fine and special corrections and modifications of the model are made [24]. In such cases, a higher modeling accuracy is obtained by considering stochastic microscopic models where the positions of the particles are tracked continuously in space. In the microscopic model, individual molecules move by Brownian motion and when they are close they can react with each other. The molecules diffuse in the simulation either by taking small time steps in a solution of a Langevin equation [5, 32] or by sampling a probability distribution for the new position [29, 40, 52]. Such microscopic models are much more computationally demanding than a corresponding mesoscopic simulation, at least for reasonable mesh resolutions. On a macroscale, the model is deterministic and the evolution of the concentrations of the species is governed by a system of partial differential equations.

Many realistic models of biochemical reaction networks in a spatially resolved setting can be expected to display multiscale properties. Depending on the values of the reaction and diffusion constants and the copy number of partaking species, some reactions may require a microscopic simulation [48] while others can be modeled accurately on the mesoscale. The computational geometry may contain areas with fine scaled structures, e.g. near membranes, forcing a mesoscopic solver to use a very fine mesh. In both those cases, a purely microscopic simulation can be very expensive. Even if a purely mesoscopic model could be formulated despite the problems that arise due to small voxels [23, 24, 30], its simulation would require a very fine resolution dictated by the microscale features of the model, leaving degrees of freedom modeled accurately on the mesoscale grossly over-resolved, thereby causing unnecessarily costly simulations.

To address this problem, the mesoscopic and the microscopic levels of modeling are coupled in this paper by following the critical species as single molecules in critical regions of the domain, letting them react with the mesoscopic molecules when they are in the same voxel. In this way, species with high copy numbers can be treated mesoscopically. Another advantage with a micro-meso model is that for parts of the model, a mesoscopic description of the processes may not be known in detail and the solution is then to use a more basic microscopic model selectively. This is how binding, diffusion, unbinding, and reactions on a membrane are simulated in the computational results in Section 6.

The partitioning of the species is determined *a priori*. The microscopic description should be used with care only for those species and in those voxels were it is necessary. As it is now, the partitioning must be based on biological experience. More work is needed to derive computationally useful criteria for an automatic and dynamic partitioning.

There are a number of software packages for mesoscopic simulation e.g. [2, 22, 28, 50] and for microscopic simulation e.g. [4, 32, 48, 52] of cell biochemistry and diffusion. Some of them are compared in [14], but only a few integrate single molecule models with a meso or macro level model [6, 42]. Species are modeled microscopically or mesoscopically in Spatiocyte [6] in a Cartesian mesh or lattice. The molecules at the micro level diffuse in discrete jumps between the

voxels in the mesh. The molecules of some species are simulated with Brownian dynamics and the motion of the remaining species is modeled by a macroscopic diffusion equation in Cell++ [42]. Our contribution differs from these two methods in the following respects: Firstly, the mesoscopic mesh is unstructured as in [22] where a voxel shares faces with a variable number of other voxels. Secondly, assuming Brownian dynamics the microscopic part is based on Smoluchowski theory for diffusion and reactions [45] and is simulated by an efficient version from [29] of the GFRD algorithm [52] allowing longer microscopic time steps. Thirdly, the simulation in time is split in a time step such that the meso level is treated first with a frozen micro level and then the micro level is advanced with a frozen meso level.

The method we propose is made sufficiently general to simulate complex biochemical models in mixed dimensions by extending the method in [29] to handle an arbitrary 2D boundary. Our microscopic method simulates Brownian motion and reactions on any smooth 2D surface embedded in 3D by a locally planar approximation determined by the triangulation of the surface. A probability distribution is sampled for the next position of the molecule on the membrane. In this way, the randomness of the motion and reactions of single molecules bound to a membrane is modeled. Other software packages for microscopic simulations on membranes, such as Smoldyn [3–5] and MCell [32] are available. They are also based on the Smoluchowski model but differ from our approach. In both MCell and Smoldyn, a fixed discrete time step is used and the temporal accuracy in the time until reactions occur is limited by this time step. Consequently, the spatial resolution is not higher than the mean distance a molecule travels during a time step. Thus, to simulate a system as in [48] for which a high resolution is indeed needed, accuracy would require a very small time step making these methods inefficient. In GFRD the idea is to solve the Smoluchowski equation for pairs of molecules. Then the time for the next reaction in the system is sampled with arbitrary precision. We show that by using the techniques developed in [29] also in two dimensions we obtain an accurate and efficient algorithm for simulating diffusion and reactions on surfaces. The accuracy of the method is analyzed in the special case when a sphere is approximated by planar facets.

Algorithms coupling an atomistic description and a macroscopic continuum description discretized on a mesh are found in many other areas of physics and biology. For example, when the mesh size approaches the molecular scale in computational fluid dynamics, then the Navier-Stokes equations for the continuum break down and a particle model is necessary. An example is when the mean free path of gas molecules is large. Then the no-slip condition at solid surfaces is no longer valid. Another example is when the geometry in nanoscale and microscale devices has to be resolved at a molecular level [15, 26, 36]. Methods to couple a particle model and discretizations of a continuum model on a mesh are reviewed in [36]. The computational domain is split into subdomains with either a particle model or a continuum model [13, 15]. The difficulty is to handle the interaction at the interface between the subdomains. The particles and the continuous variables all live on the same mesh in the particle-in-cell method. Electrically charged particles interact with an electromagnetic field on a mesh in applications of the method in plasma physics [49]. In these methods, both the microscopic and macroscopic equations are integrated in the whole time interval of interest. This is in contrast to the idea in the heterogeneous multiscale method [17] and the “equation-free” method [33] where the microscopic level is visited only for short time intervals to feed the mesoscopic or macroscopic model with data.

The contents of the remaining sections of the paper are as follows. The model at the mesoscopic level is described in Section 2. In Section 3, we review the Smoluchowski model for diffusion and reactions of individual molecules and show how to adapt it to model the interaction between a particle and a curved membrane. The algorithm for coupling the mesoscale and the microscale is described in Section 4. To study the effect of our boundary representation, the surface of a sphere is represented locally by planes in Section 5 and the consequence of this

approximation on the accuracy of the microscale model is analyzed. The multiscale algorithm is applied to four different examples in Section 6. The diffusion of a molecule on a sphere and its annihilation by a trap is simulated in the first example and compared to an analytical result. In the second example, the simulation of a mitogen activated protein kinase (MAPK) cascade is compared to data from [48]. The algorithm is applied to the reversible adsorption of a molecule to a membrane in the third problem. Finally, results from a biologically motivated example are presented with species and reactions in the cytosol, adsorption of one species to the nuclear membrane, and its translocation into the nucleus.

Vectors and matrices are written in boldface. A vector \mathbf{u} has the components u_i and the elements of a matrix \mathbf{A} are A_{ij} or a_{ij} .

2 The mesoscopic scale

On the mesoscopic scale, the dynamics of the biochemical reaction network is modeled as a continuous time discrete space Markov process [31]. If the system is well-mixed such that transport of molecules can be ignored, the time dependent state will simply be the total copy number of each of the participating chemical species, and the stochastic process can be sampled exactly using the Stochastic Simulation Algorithm (SSA) due to Gillespie [27] or variants thereof [19].

In order to introduce spatial dependence and molecular transport, the domain Ω is partitioned into N non-overlapping voxels \mathcal{V}_i . For a system with M participating species, the state of the system can now be described by the $N \times M$ state matrix $\mathbf{X}(t)$ whose entry $x_{ij}(t)$ is the copy number of species j in voxel i at time t . The copy numbers of the species in voxel i are denoted by \mathbf{x}_i , and the components of \mathbf{x}_i are the number of molecules of species j in the voxels. Within each voxel \mathcal{V}_i we assume that the well-stirred assumption holds, and to each chemical reaction in the system we associate a stoichiometry vector $\boldsymbol{\nu}_{ir}, r = 1, \dots, R$. The reaction r in voxel i alters the state of the system according to $\mathbf{x}'_i = \mathbf{x}_i + \boldsymbol{\nu}_{ir}$. For example, a bimolecular reaction can be written



where the entries of $\boldsymbol{\nu}_{ir}$ are -1 for j, j' , 1 for j'' and zero otherwise. The propensity function $a_r(\mathbf{x}_i)$ is the probability per unit time for the reaction to fire in an infinitesimal time interval. In the Markov formalism, the time until the reaction occurs is an exponentially distributed random variable with mean $1/a_r(\mathbf{x}_i)$.

The domain Ω is covered by an unstructured primal mesh consisting of tetrahedra with a triangulated boundary $\partial\Omega$. A dual mesh is constructed with the voxels \mathcal{V}_j centered around the nodes or vertices of the mesh.

A primal mesh and a dual mesh are depicted in 2D to the left in Figure 1. The edges of the voxel in the center are the lines from the centroid of a triangle to the midpoint of an edge. A dual mesh in 3D is obtained by generalizing the construction in 2D. Quadrilateral parts of three triangular planes form the faces of a voxel intersecting a tetrahedron. A triangle has its base on an edge opposing the vertex of the voxel and a corner at the midpoint of the opposite edge. For every tetrahedron, the intersection of these triangles will define the quadrilaterals that separates the vertices, see Figure 1b. For example, the face that separates the leftmost and top vertices in the tetrahedron in the figure is defined by the centroid of the tetrahedron, the centroid of the front left triangular face, the midpoint of the edge between the vertices and the centroid of the rear face.




Fig. 1: A primal mesh (solid) and the dual mesh (dashed) in 2D (a). A tetrahedron in 3D with a triangular plane, the upper left part of which is a face of the voxel with the top vertex in the center and the voxel with lower left vertex in the center (b).

The molecules move between the voxels by diffusion. The diffusive motion of molecules is modeled as a first order event

$$X_{ij} \xrightarrow{q_{ik}x_{ij}} X_{kj}, \quad (2)$$

where one molecule of species j diffuses from voxel i to the adjacent voxel k with propensity $q_{ik}x_{ij}$. The coefficients q_{ik} are determined by the diffusion coefficient D and a finite element discretization of the Laplacian on the primal mesh in [22], where the properties of this approximation are discussed in detail. Complicated geometries have a simpler representation in an unstructured mesh composed of tetrahedra compared to a Cartesian grid with cubic voxels. In particular, the representation of curved membranes is simplified considerably. The membrane is triangulated by the faces of the tetrahedra on the membrane. A discretization of a membrane of such quality is difficult to achieve in a Cartesian grid where the membrane surface cuts the cubic voxels irregularly. Realizations of the system on unstructured meshes can be obtained in an efficient manner by the Next Subvolume Method (NSM) [19] as implemented in the URDME software package [12].

The advantage with a mesoscopic model is that only the copy number of the species in each voxel is updated and diffusion is modeled as jumps between discrete voxels. In a microscopic model, the paths and reactions of individual molecules are tracked requiring much more computing time for the same number of molecules. If the mesh is kept fixed, the cost of simulating a trajectory on the mesoscale scales linearly with the number of molecules in a fixed domain while it scales quadratically for the microscopic method [29].

3 The microscopic scale

The microscopic model for the motion of the molecules is here Brownian dynamics with the Smoluchowski model [45] for reacting molecules. The probability density function (PDF) for the position of a molecule satisfies a parabolic partial differential equation. For N molecules, we have a N -body problem which cannot be solved analytically. One approach to simulate such a system is to choose a small time step and then update the position of each molecule by sampling from a normal distribution. After each time step one determines whether a reaction has occurred or not. This approach is taken in Smoldyn [5] and MCell [32]. The problem with this approach is that for an accurate result the time step must be chosen very small and as a consequence the method may become inefficient. If the system does not have properties where the finest scales have to be resolved these methods are a good choice.

An efficient algorithm for simulating systems that require high accuracy allowing longer time steps and well defined reaction dynamics is the Green’s Function Reaction Dynamics (GFRD) proposed in [52, 53] for 3D problems. Improvements have been developed in [48] and [29]. The GFRD idea is to decompose the problem into one-body and two-body problems by choosing a time step Δt such that molecules are unlikely to react with more than one other molecule or the boundary of the domain $\partial\Omega$. Details how to choose Δt can be found in [29, 40, 48, 52]. If the molecule is located at \mathbf{z}^n at time $t = t^n$, the new position \mathbf{z}^{n+1} at $t = t^{n+1} = t^n + \Delta t$ is sampled from the cumulative distribution function (CDF) derived from the corresponding PDF.

The GFRD algorithm is here extended to reactions and diffusion on and in the vicinity of arbitrary curved surfaces. A surface is approximated locally by the tangential plane reconstructed from the triangulation of the surface. The new position of a molecule on or near the surface is sampled from the distribution with the PDF satisfying the Smoluchowski equation in a Cartesian coordinate system. Other methods for simulation of molecular diffusion and reactions on biological membranes have been developed and applied to simulations on e.g. the cell membrane [38, 39] and the endoplasmic reticulum [43]. An overview of modeling difficulties and simulation techniques for membranes is found in [9]. The Laplace-Beltrami equation for macroscopic diffusion of molecules on curved surfaces is solved in [39, 43, 44]. The surface is reconstructed from a Cartesian grid in [39, 44] and the difference approximation on the surface is embedded in 3D in [43, 44]. A microscopic model for realization of the reaction-diffusion process on plane membranes is proposed in [38] and in [11] a general method for diffusion on curved surfaces including anisotropy is described. In [3] molecule–surface interaction probabilities for different scenarios are proposed for use with particle-based simulators that use a fixed time step such as Smoldyn [4].

3.1 Molecules in three dimensions

A single molecule at the position \mathbf{z}^n at t^n moves according to Brownian motion and the new location at t^{n+1} is sampled from a normal distribution in 3D with mean \mathbf{z}^n and variance $2D\Delta t$. The PDF for the particle’s position is

$$p_{\mathbf{z}}(\mathbf{z}, t^{n+1} | \mathbf{z}^n, t^n) = \frac{1}{(4\pi D\Delta t)^{3/2}} \exp\left(-\frac{\|\mathbf{z} - \mathbf{z}^n\|^2}{4D\Delta t}\right), \quad (3)$$

where $\|\cdot\|$ is the Euclidean norm.

Consider a pair of molecules with two different diffusion coefficients D_1 and D_2 and positions \mathbf{z}_1^n and \mathbf{z}_2^n at t^n . The new positions \mathbf{z}_1 and \mathbf{z}_2 are determined by introducing two new coordinates

$$\mathbf{z}_{12} = \sqrt{D_2/D_1}\mathbf{z}_1 + \sqrt{D_1/D_2}\mathbf{z}_2, \quad \mathbf{y} = \mathbf{z}_2 - \mathbf{z}_1. \quad (4)$$

The new \mathbf{z}_{12}^{n+1} is sampled as in (3) from a 3D normal distribution with mean \mathbf{z}_{12}^n and variance $2D_{12}\Delta t$ where D_{12} is the combined diffusion $D_{12} = D_1 + D_2$. The relative distance \mathbf{y}^{n+1} between the molecules is sampled from a distribution with a PDF satisfying the Smoluchowski equation. In a spherical coordinate system with $\mathbf{r} = (r, \theta, \phi)$, the PDF $p_{\mathbf{r}}(\mathbf{r}, t | \mathbf{r}^n, t^n)$ solves

$$\partial_t p_{\mathbf{r}} = D_{12} \left(\frac{\partial^2 p_{\mathbf{r}}}{\partial r^2} + \frac{2}{r} \frac{\partial p_{\mathbf{r}}}{\partial r} + \frac{1}{r^2 \sin \theta} \frac{\partial}{\partial \theta} \left(\sin \theta \frac{\partial p_{\mathbf{r}}}{\partial \theta} \right) + \frac{1}{r^2 \sin^2 \theta} \frac{\partial^2 p_{\mathbf{r}}}{\partial \phi^2} \right), \quad (5)$$

with the boundary condition at the reaction radius $r = \sigma$ between the two molecules

$$4\pi\sigma^2 D_{12} \frac{\partial p_{\mathbf{r}}}{\partial r} \Big|_{r=\sigma} = k_r p_{\mathbf{r}}(r = \sigma, t | \mathbf{r}^n, t^n) \quad (6)$$

and initial condition $p_{\mathbf{r}}(\mathbf{r}, t|\mathbf{r}^n, t^n) = \delta(\mathbf{r} - \mathbf{r}^n)$. The probability of the molecules to react with each other is controlled by the association coefficient k_r . When $k_r = 0$ then they never react and when $k_r \rightarrow \infty$ the boundary is absorbing and an association always occurs at $r = \sigma$. The same equation is also valid for the binding of a molecule to a spherical membrane but then σ is the radius of the sphere and very much larger than the reaction radius.

There is an analytical solution to (5) with boundary condition (6) in [10, p. 382] but the solution is complicated and expensive to evaluate. Instead a scheme based on operator splitting is proposed in [29] simplifying the computations of the required PDFs. Approximate samples of the distribution with the PDF in (5) can be computed in the following way:

1. Solve in the radial direction for $p_r(r, t|r^n, t^n), r \geq \sigma, t \geq t^n$, in

$$\partial_t p_r = D_{12} \left(\frac{\partial^2 p_r}{\partial r^2} + \frac{2}{r} \frac{\partial p_r}{\partial r} \right), \quad (7)$$

with initial condition $p_r(r, t^n|r^n, t^n) = \delta(r - r^n)$ and boundary condition

$$4\pi\sigma^2 D_{12} \frac{\partial p_r}{\partial r} \Big|_{r=\sigma} = k_r p_r(\sigma, t|r^n, t^n). \quad (8)$$

Sample the distribution for r^{n+1} . There is an analytical solution to (7) and (8) in [35] and also the CDF is integrated analytically there.

2. Solve in the polar and azimuthal directions for $p_\theta(\theta, t|r^{n+1}, 0, \phi^n, t^n), \theta \in [0, \pi], \phi \in [0, 2\pi]$, in

$$\partial_t p_\theta = \frac{D_{12}}{(r^{n+1})^2} \left(\frac{1}{\sin \theta} \frac{\partial}{\partial \theta} \left(\sin \theta \frac{\partial p_\theta}{\partial \theta} \right) + \frac{1}{\sin^2 \theta} \frac{\partial^2 p_\theta}{\partial \phi^2} \right), \quad (9)$$

with initial condition $p_\theta(\theta, t^n|r^{n+1}, 0, \phi^n, t^n) = \delta(\theta)/((r^{n+1})^2 \sin \theta)$. There is an analytical formula for the PDF, see e.g. [29], and the CDF can be integrated numerically and tabulated from that formula. The CDF is sampled for θ^{n+1} . The solution in (9) is independent of ϕ and ϕ^{n+1} is sampled from the uniform distribution in $[0, 2\pi)$.

3.2 Molecules near and on surfaces

Here, a curved membrane embedded in 3D space is represented by planar facets. A molecule in the cytosol can attach to a planar facet and a molecule on a facet is allowed to diffuse on it and react with other molecules on the the same facet or other adjacent facets. More details of this geometrical approximation can be found in Section 4.2.

As is shown in [29], the splitting technique in Section 3.1 can be used for pairs of molecules on a plane in 2D. After a change of coordinates as in (4), the PDF of the relative position between two molecules in cylindrical coordinates $\mathbf{r} = (r, \phi)$ satisfies on a membrane facet

$$\partial_t p_{\mathbf{r}} = D_{12} \left(\frac{\partial^2 p_{\mathbf{r}}}{\partial r^2} + \frac{1}{r} \frac{\partial p_{\mathbf{r}}}{\partial r} + \frac{1}{r^2} \frac{\partial^2 p_{\mathbf{r}}}{\partial \phi^2} \right), \quad (10)$$

with the boundary condition at the reaction radius $r = \sigma$

$$2\pi\sigma D_{12} \frac{\partial p_{\mathbf{r}}}{\partial r} \Big|_{r=\sigma} = k_r p_{\mathbf{r}}(r = \sigma, t|\mathbf{r}_n, t_n). \quad (11)$$

A new position (r, ϕ) is sampled in two steps using an approximation of the PDF in (10):

1. Solve in the radial direction for $p_r(r, t|r^n, t^n)$, $r \geq \sigma$, $t \geq t^n$, in

$$\partial_t p_r = D_{12} \left(\frac{\partial^2 p_r}{\partial r^2} + \frac{1}{r} \frac{\partial p_r}{\partial r} \right), \quad (12)$$

with initial condition $p_r(r, t^n|r^n, t^n) = \delta(r - r^n)$ and boundary condition (11). The analytical solution to (12) is relatively complicated involving an integral of a rational expression in Bessel functions [10, p. 370]. We solve (12) numerically in this paper and sample the corresponding numerically integrated CDF for r^{n+1} .

2. Sample a normal distribution for ϕ^{n+1} in the azimuthal direction with the PDF

$$p_\phi(\phi, t^{n+1}|0, t^n) = \frac{1}{(4\pi D\Delta t)^{1/2}} \exp\left(-\frac{\phi^2}{4D\Delta t}\right). \quad (13)$$

The diffusion equation in a Cartesian coordinate system for a molecule above a facet with the normal in the z -direction is needed for the binding of a molecule to a membrane. Assume that there is no other molecule in the neighborhood. With $\mathbf{r} = (x, y, z)$, the PDF $p_{\mathbf{r}}(\mathbf{r}, t|r^n, t^n)$ satisfies

$$\partial_t p_{\mathbf{r}} = D \left(\frac{\partial^2 p_{\mathbf{r}}}{\partial x^2} + \frac{\partial^2 p_{\mathbf{r}}}{\partial y^2} + \frac{\partial^2 p_{\mathbf{r}}}{\partial z^2} \right), \quad (14)$$

with the association coefficient k_r and the boundary condition at $z = 0$

$$D \frac{\partial p_{\mathbf{r}}}{\partial z} \Big|_{z=0} = k_r p_{\mathbf{r}}(z = 0, t|r^n, t^n). \quad (15)$$

Using operator splitting, the sampling of the distribution defined by (14) is performed as follows:

1. Solve in the z direction for $p_z(z, t|z^n, t^n)$, $z \geq 0$, $t \geq t^n$, in

$$\partial_t p_z = D \frac{\partial^2 p_z}{\partial z^2}, \quad (16)$$

with initial condition $p_z(z, t^n|z^n, t^n) = \delta(z - z^n)$ and boundary condition (15). The analytical solution is derived in [10, p. 358]. A new z^{n+1} is obtained by sampling the numerically integrated CDF.

2. Sample a normal distribution in 2D for the position parallel to the plane if $z^{n+1} > 0$ from the previous splitting step or on the approximating plane if $z^{n+1} = 0$. The PDF for x and y is with $x^n = y^n = 0$

$$p_{xy}(x, y, t^{n+1}|0, 0, t^n) = \frac{1}{4\pi D\Delta t} \exp\left(-\frac{x^2 + y^2}{4D\Delta t}\right). \quad (17)$$

In this section we have described how a micro particle can bind to and diffuse on a planar surface in 2D. In the next section, we will describe a hybrid simulation strategy that combines the mesoscale and microscale. The mesoscale model requires a spatial discretization, and each surface embedded in the 3D domain is approximated by a surface mesh composed of triangles. This boundary mesh representation is used to approximate the surface locally with facets providing a natural surface representation for the microscopic part of the solver. This strategy is described in detail in the next section, and the approximations made in the special case of a spherical surface \mathcal{S}^2 are analyzed in detail in Section 5.

4 Coupling the mesoscale and the microscale

In this section we describe the algorithm used to simulate a hybrid system. The coupling of the meso- and microscales is based on an operator splitting scheme over species and space. Species requiring a more detailed description are simulated at the microscale level while the remaining species are simulated on the mesoscale. The algorithm is flexible: species can be simulated on the mesoscale in some parts of the domain and on the microscale in other parts.

4.1 Splitting the system into two subsystems

In the beginning of every time step, the species in the voxels are divided into one mesoscopic part and one microscopic part. A species can be simulated on the mesoscale in some voxels and on the microscale in other voxels. As illustrated in Section 6, this flexible structure is important to facilitate partitioning of systems where the need to resolve parts of the model on the microscale arises from different properties of the model.

Given a partitioning in species and space, the chemical reactions of the original system have to be divided into three parts: one purely mesoscopic part involving only mesoscopic variables, one purely microscopic part and one part that involves both mesoscopic and microscopic degrees of freedom. The last set of reactions are bimolecular reactions where one of the reactants is a mesoscopic variable and the other a microscopic variable in the same voxel. Such mixed reactions are executed by the more detailed microscopic solver, and the mesoscopic variable is simply assumed to be uniformly distributed in the voxel according to the local well-mixed assumption on the mesoscale.

Suppose that species A is simulated at the microscopic level and that B is a mesoscopic species and let them react with each other and form the product C . If the microscopic rate constant for the reaction is k_{micro} in



then the mesoscopic rate constant is $k_{\text{meso}} = 4\pi\sigma D_{12}k_{\text{micro}}/(4\pi\sigma D_{12} + k_{\text{micro}})$ where D_{12} is the sum of the diffusion constants for A and B and σ is the reaction radius. The mesoscopic propensity a_{AB} for the reaction is $k_{\text{meso}}[B]$ where $[B]$ is the number B molecules in the voxel. The time to the next reaction is sampled from an exponential distribution with mean equal to $1/a_{AB}$. When a reaction has occurred then the microscopic A disappears, the number of mesoscopic B molecules is lowered by one and a new microscopic C appears at the old position of A .

The algorithm advances the system a time step $\Delta\tau$ forward in time from τ^n to $\tau^{n+1} = \tau^n + \Delta\tau$ using a splitting scheme of first order accuracy in time. The steps in the algorithm are:

1. Freeze all microscopic state variables and update the mesoscopic part of the system as described in Section 2 during the time step $\Delta\tau$. The chemical reactions that are executed by the mesoscale solver only involve mesoscale reactants.
2. Freeze the mesoscopic state variables and update the microscopic part of the system $\Delta\tau$ in time using the method described in Section 3. Mixed reactions involving both mesoscopic and microscopic variables are executed by the microsolver as in (18). For every local time step $\Delta t < \Delta\tau$ taken by the microscopic solver, we track the voxel in the mesh containing the molecule by the procedure in Section 4.3.
3. Update the state at τ^{n+1} by determining if any meso- or microvariables have been created or destroyed in the time step and repeat from Step 1 until the final simulation time has been reached.

While simple in structure, in practice the splitting algorithm requires that information about the geometry and boundaries as defined by the unstructured mesh is communicated to the microsolver. The position of a microscopic molecule is easy to determine in a Cartesian mesh in a cubic Ω but such a mesh does not possess the geometric flexibility of an unstructured mesh in a general domain.

4.2 Boundary approximation for the microscale solver

The boundary of the domain $\partial\Omega$ is covered by non-overlapping triangles. Each vertex of the triangle is also a vertex in a tetrahedron, and is thus a center vertex in a voxel in the dual mesh used by the mesoscopic solver and defined in Section 2. For a vertex with coordinate ζ_i on the boundary, we associate the volume element \mathcal{V}_i with *one* plane or facet \mathcal{P}_i approximating the boundary. This vertex is a vertex of many triangles \mathcal{T}_{ij} , $j = 1, \dots, N_i$, in the surface mesh. Denote the outward normal to each of these triangles by \mathbf{n}_{ij} . The plane \mathcal{P}_i that a molecule in voxel \mathcal{V}_i can bind to or diffuse on is defined by the point ζ_i and the averaged normal $\mathbf{n}_i = N_i^{-1} \sum_{j=1}^{N_i} \mathbf{n}_{ij}$ scaled to unit length. In principle, this approximation is not necessary, but it simplifies the implementation considerably compared to letting the molecules move directly on the triangles in surface mesh. The bookkeeping of reactions between molecules in space and on the intersection of $\partial\Omega$ and a voxel would be more complicated, since we would have to keep track of which triangle the molecule is closest to and choose a time step such that we only have to consider reactions between a molecule and one triangle at a time.

In the hybrid algorithm, all these facets are precomputed in an initialization step and stored in an extended mesh data structure. A data structure corresponding to \mathcal{T}_{ij} is also generated in a preprocessing step as this information will be needed frequently in the parts of the algorithm described in the next subsection. An analogous data structure for the tetrahedra in the mesh is also formed in this way.

4.3 Mapping between microscopic and mesoscopic degrees of freedom

For the microscopic simulation the underlying mesoscopic mesh is in principle not needed, apart from representing the boundaries. However, in order to allow mixed reactions across the two scales it is necessary to keep track of which voxel in the mesh each micro-molecule is located in. Also, to avoid having to compare the molecule’s position to each planar surface approximating the boundary in every step of the microscopic solver, we need to know if the molecule is close to one of the approximating boundary planes, i.e. if a molecule is close to a vertex on an exterior or interior boundary of the domain. For consistency between the two models, a molecule on the microscale cannot simply be assigned to the voxel corresponding to its closest vertex in the primal mesh. There is no guarantee for a general unstructured mesh that this will be the correct voxel in the dual mesh. In general, it would be computationally expensive for the hybrid algorithm to find and assign the correct voxel to a microscopic molecule after each local time step Δt . However, as explained in the next section, we will restrict Δt in the microscopic solver in such a way that the probability of a molecule to diffuse across several voxels in the mesh is small. This means that the procedure outlined below will often exit early already in the first step. Knowing that the voxel containing a molecule at time t^{n+1} is expected to be close to its voxel at t^n we proceed as follows:

1. Starting at the last voxel the molecule was known to belong to, loop over all tetrahedra in the primal mesh having the center vertex in the voxel as one of its corners and determine if the molecule is inside one of those tetrahedra.

2. If we do not find a tetrahedron containing the molecule in Step 1, we continue by finding the closest vertex to the current position of the molecule.
3. Once we have found the nearest vertex in the mesh, the search procedure in Step 1 above is repeated starting from that vertex.
4. When the tetrahedron containing the particle has been found in the primal mesh, we determine which dual element the molecule belongs to by checking its position with respect to the planes defined by the quadrilateral surfaces shown in the right panel of Figure 1.

The decision in Step 1 is based on the barycentric coordinates of the molecule with respect to the tetrahedron. If all coordinates are positive, the molecule is found inside the tetrahedron. The matrices that map the Cartesian coordinates of the molecule to the barycentric coordinates are precomputed and stored before starting the hybrid algorithm. The cost of determining if a molecule is inside a tetrahedron is then that of performing matrix vector multiplication with a 3×3 matrix.

The vertex in Step 2 is found by first comparing the distance from the vertex of its previous voxel to the neighboring vertices. If the molecule is found to be closer to one of the neighbors, then focus is shifted to that vertex and the procedure is repeated. In this way we will "slide" in the mesh towards the nearest vertex. A limited number of distances between vertices have to be computed and compared in this step.

The molecule can always be found in Step 3 within two layers of tetrahedra around the vertex with minimal distance to the molecule, unless it is outside the domain. This can happen if the molecule is close to a boundary since it will be bounded by the facets approximating the surface described in the previous section and not by the triangulated surface mesh. In that case we assign the molecule to its closest tetrahedron measured by the absolute value of the sum of negative components in the molecule's barycentric coordinates. The computational work in this step as the same as in Step 1.

The normals of the quadrilateral faces of the voxels are precomputed and stored initially for use in Step 4. Note that this is the same necessary preprocessing step as in the assembly of a standard finite volume discretization starting from the primal mesh. At most three such comparisons are needed to assign the molecule to its correct dual, i.e. at most three scalar products of vectors with three elements need to be computed.

These four algorithmic steps are simplified considerably in a Cartesian mesh. On a 2D surface, the algorithm is analogous but since molecules will move on the averaged planes and not on the boundary triangles directly, a position on the actual surface mesh is first determined in Steps 1 and 3 above by projecting orthogonally onto the triangles \mathcal{T}_{ij} in the boundary mesh. When the correct voxel \mathcal{V}_i has been found, the molecule is projected to the facet \mathcal{P}_i .

After each splitting time step the system has to be updated according to the partitioning of the meso- and microscales. A mesoscopic molecule may have entered a voxel in which it should be simulated on the microscale. It is then assigned a position in space, and since we assume that molecules on the mesoscale are well-mixed, the position is given by sampling from a uniform distribution over the voxel. A microscopic molecule may have entered a voxel in which it should be simulated on the mesoscale, and in this case we simply increment the mesoscopic state vector accordingly. If there are no chemical reactions, then the total number of molecules of the species is preserved but the identity of a microscopic molecule is lost when it is converted to the mesoscale.

4.4 Time step selection

The splitting time step $\Delta\tau$ should be appropriately chosen, neither too large nor too small. If $\Delta\tau$ is too small when a microscopic molecule has entered a new voxel it is likely that it has not had time to become well-mixed in the new voxel and may diffuse back to its previous voxel with high probability. On the other hand, if $\Delta\tau$ is too large the error introduced by the operator splitting will be large since the error is proportional to $\Delta\tau$. The splitting time step should be chosen such that the distance a microscopic molecule can move during $\Delta\tau$ is about the length of a voxel, i.e. $\sqrt{D\Delta\tau} \sim h$ or $\Delta\tau \sim h^2/D$, where h is a measure of the mesh size. The macroscopic partial differential equation modeling diffusion is the diffusion equation. Discretize the time and space derivatives. The time step in an explicit time marching method to solve that equation will depend in the same manner on the spatial step for a stable solution.

The time step Δt in the microscopic simulations satisfies $\Delta t < \Delta\tau$ and should be so small that a collision between freely moving single molecules or pairs of molecules or the boundary is improbable. These restrictions on Δt are discussed in detail in [29]. A curved surface is approximated by planar facets. Let h be a measure of the length of a facet. The distance a diffusing molecule on a surface moves in a time step is of the order of $\sqrt{D\Delta t}$ and it is shown in the next section that also here $\sqrt{D\Delta t} \sim h$ or $\Delta t \sim h^2/D$ for good accuracy.

5 Approximations near and on the membrane

In this section the analytical solutions of the densities in (7), (16), (9) and (17) in the microscopic model will be compared near and on a spherical membrane \mathcal{S}^2 with radius σ . The sphere is approximated by the planes or facets \mathcal{P}_i generated in Section 4.2. We will show that the PDFs are close if the model constants fulfill certain conditions and the distance a molecule moves in a microscopic time step Δt is small compared to σ .

Assume that the diffusion coefficient is D in space and D_s on the surface and that the initial position of the molecule is above the north pole at radius r' with the polar angle $\theta = 0$, see Figure 2a, or on the surface of one of the facets \mathcal{P}_1 approximating the sphere, see Figure 2b. For the same new radial position r and the same new polar angle θ , the exact PDF for the sphere is compared to the approximate PDF from the triangulated surface. The algorithm for association and diffusion on a membrane can handle any smooth surface by a local approximation with the set of \mathcal{P}_i but an analysis is not possible for such surfaces in the same detail as it is for \mathcal{S}^2 .

5.1 Density outside the sphere

The PDFs of the new position of a molecule above a sphere are determined in a spherical coordinate system in (7) and in the Cartesian coordinate system of the tangent plane in (16). The angle between the normal of the tangent plane \mathcal{P}_1 and the direction of the polar axis is ϕ in Figure 2a.

Given that the location of the molecule is at r' at $t = t^n$, the PDF solving (7) and (8) for the position in the radial direction r outside a sphere at $t^{n+1} = t^n + \Delta t$ is (see [10, p. 368] or [34])

$$\begin{aligned}
 p_r(r, \Delta t | r') &= \frac{1}{8\pi r r' \sqrt{D\pi\Delta t}} (\exp(-a_r^2) + \exp(-b_r^2)) \\
 &\quad - \frac{1}{4\pi r r'} \cdot \frac{k_r + 4\pi D\sigma}{4\pi D\sigma^2} \exp(-b_r^2) \exp((b_r + c_r)^2) \operatorname{erfc}(b_r + c_r), \\
 a_r &= \frac{r - r'}{\sqrt{4D\Delta t}}, \quad b_r = \frac{r + r' - 2\sigma}{\sqrt{4D\Delta t}}, \quad c_r = \frac{k_r + 4\pi D\sigma}{4\pi D\sigma^2} \sqrt{D\Delta t},
 \end{aligned} \tag{19}$$




Fig. 2: Local coordinate systems with one tangential plane \mathcal{P}_1 for approximation in the radial direction (a). Local coordinate systems with two tangential planes \mathcal{P}_1 and \mathcal{P}_2 for approximation in the polar direction (b).

where erfc is the complementary error function [1, Ch. 7]. The boundary condition at $r = \sigma$ is given by (8)

$$4\pi D\sigma^2 \frac{\partial p_r}{\partial r} = k_r p_r. \quad (20)$$

The coordinates of r' in the Cartesian system are $\rho' = r' \sin \phi$ and $z' = r' \cos \phi - \sigma$. A new position z_r is sampled in the z -direction perpendicular to the tangent plane. The PDF for z_r to be at the same radius r as in (19) is compared to p_r . The PDF for the position z is

$$p_z(z, \Delta t | z') = \frac{1}{2\sqrt{D\pi\Delta t}} (\exp(-a_z^2) + \exp(-b_z^2)) - \frac{k_z}{D} \exp(-b_z^2) \exp((b_z + c_z)^2) \text{erfc}(b_z + c_z), \quad (21)$$

$$a_z = \frac{z_r - z'}{\sqrt{4D\Delta t}}, \quad b_z = \frac{z_r + z'}{\sqrt{4D\Delta t}}, \quad c_z = k_z \sqrt{\frac{\Delta t}{D}},$$

see [10, p. 358]. The boundary condition at $z = 0$ is (cf. (15))

$$D \frac{\partial p_z}{\partial z} = k_z p_z. \quad (22)$$

Let the relation between k_r and k_z be

$$k'_r = \frac{k_r}{4\pi D\sigma^2}, \quad \frac{k_z}{D} = \frac{k_r + 4\pi D\sigma}{4\pi D\sigma^2} = \frac{k_r}{4\pi D\sigma^2} + \frac{1}{\sigma} = k'_r \left(1 + \frac{1}{k'_r \sigma}\right). \quad (23)$$

If $\sigma \gg 1/k'_r$ then $k_z/D \approx k'_r$ and $c_r \approx c_z$ and the boundary condition is approximately the same for both p_r and p_z . The distance between r' at t^n and r at t^{n+1} is at most a small factor times $\sqrt{D\Delta t}$. Then the quotient r/r' satisfies

$$\left|1 - \frac{r'}{r}\right| = \left|\frac{r - r'}{r}\right| \leq \left|\frac{r - r'}{\sigma}\right| \sim \frac{\sqrt{D\Delta t}}{\sigma}. \quad (24)$$

In Figure 2a, we find that $(\rho')^2 + (z_r + \sigma)^2 = r^2$. Consequently, for small ϕ

$$\begin{aligned} z_r - z' &= r\sqrt{1 - \left(\frac{r'}{r}\right)^2 \sin^2 \phi} - r' \cos \phi = r - r' + \frac{r'}{2}\phi^2 \left(1 - \frac{r'}{r}\right) + \mathcal{O}(\phi^4), \\ z_r + z' &= r\sqrt{1 - \left(\frac{r'}{r}\right)^2 \sin^2 \phi} + r' \cos \phi - 2\sigma = r + r' - 2\sigma - \frac{r'}{2}\phi^2 \left(1 + \frac{r'}{r}\right) + \mathcal{O}(\phi^4). \end{aligned} \quad (25)$$

To be able to resolve the geometry of the sphere, the size of the facet $h \approx \sigma\phi_{\max}$ should be a small factor times $\sqrt{D\Delta t}$. Thus, $\phi \leq \phi_{\max} \sim \sqrt{D\Delta t}/\sigma$. A comparison between a_r and b_r in (19) and a_z and b_z in (21) using the expressions in (25) and (24) shows that

$$\begin{aligned} a_z &= a_r + a_\phi\phi^2 + \mathcal{O}(\phi^4), \quad a_\phi = \frac{r'}{2\sqrt{4D\Delta t}} \left(1 - \frac{r'}{r}\right) \sim \frac{r'}{\sigma}, \\ b_z &= b_r + b_\phi\phi + \mathcal{O}(\phi^3), \quad b_\phi \approx -\frac{r'\phi}{\sqrt{4D\Delta t}} \sim -\frac{r'}{\sigma}. \end{aligned} \quad (26)$$

By Taylor expansion one can show that a perturbation ε in the arguments a_z, b_z , and c_z will result in a perturbation of $\mathcal{O}(\varepsilon)$ in p_z .

Multiplying p_r in (19) by the surface element $4\pi r^2$ and comparing with p_z in (21), we conclude by (23), (24), and (26) that the relation between them is

$$4\pi r^2 p_r \approx p_z \quad (27)$$

if $\sigma \gg \sqrt{D\Delta t}$, $\phi \leq \phi_{\max} \sim \sqrt{D\Delta t}/\sigma$, and $k_z = k'_r$ and $\sigma \gg 1/k'_r$.

The probability for the molecule to bind to the sphere at t^{n+1} with the initial position $r' > \sigma$ is derived in [34]

$$\begin{aligned} p_r(*, \Delta t | r') &= \frac{\sigma}{r'} \frac{k_r}{k_r + 4\pi D\sigma} \left(\operatorname{erfc}(d) - \exp(-d^2) \exp((d+f)^2) \operatorname{erfc}(d+f) \right), \\ d &= \frac{r' - \sigma}{\sqrt{4D\Delta t}}, \quad e = k'_r \sqrt{D\Delta t}, \quad f = e + \frac{\sqrt{D\Delta t}}{\sigma} = k_z \sqrt{\frac{\Delta t}{D}}. \end{aligned} \quad (28)$$

Since

$$\frac{\sigma}{r'} \approx 1 - \frac{r' - \sigma}{r'}, \quad \frac{k_r}{k_r + 4\pi D\sigma} = \frac{k'_r}{k'_r + 1/\sigma} \approx 1 - \frac{1}{\sigma k'_r}$$

and $r' - \sigma$ is a small factor times $\sqrt{D\Delta t}$, the first two factors in (28) are close to one when $\sigma \gg \sqrt{D\Delta t}$ and $\sigma \gg 1/k'_r$. The probability to react with the tangent plane in the Cartesian system is found in [18]

$$p_z(*, \Delta t | z') = \operatorname{erfc}(d) - \exp(-d^2) \exp((d+e)^2) \operatorname{erfc}(d+e). \quad (29)$$

By comparing (28) and (29), we see that $p_r(*, \Delta t | r')$ is well approximated by $p_z(*, \Delta t | z')$ when σ is large compared to $1/k'_r$, and $\sqrt{D\Delta t}$.

The conclusion from the analysis is that *the error in the PDF due to the approximation of the sphere by planar facets is small when the size of the facet h is of the same order as $\sqrt{D\Delta t}$ and the radius of the sphere σ fulfills $\sigma \gg \sqrt{D\Delta t}$ and $\sigma \gg 1/k'_r$.*

5.2 Density on the sphere

The PDFs of a new polar position at time t of a molecule diffusing on the sphere or the approximation of the sphere are determined and compared.

Consider a molecule initially at the north pole $\theta = 0$ of the sphere. It diffuses on the sphere for $t > t^n$. The PDF for the polar angle, the solution of (9), is

$$p_\theta = \frac{1}{4\pi\sigma^2} \sum_{\ell=0}^{\infty} (2\ell+1) \exp(-\ell(\ell+1)\gamma) P_\ell(\cos(\theta)), \quad (30)$$

where P_ℓ is a Legendre polynomial and $\gamma = D_s \Delta t / \sigma^2$. Since $\cos(\theta) \approx 1 - \theta^2/2$ for small θ and $P_\ell(1) = 1$, there exist coefficients c_ℓ such that $P_\ell \approx 1 - c_\ell \theta^2$. By the definition of the Legendre polynomials and their recursion formula [1, Ch. 8], $c_0 = 0, c_1 = 1/2$, and for $\ell \geq 1$ the coefficients satisfy

$$c_{\ell+1} = \frac{2\ell+1}{\ell+1} \left(c_\ell + \frac{1}{2} \right) - \frac{\ell}{\ell+1} c_{\ell-1}. \quad (31)$$

The solution to (31) is $c_\ell = \ell(\ell+1)/4$. Hence, the first two terms in the Taylor expansion of p_θ about $\theta = 0$ are

$$\begin{aligned} p_\theta &= \frac{1}{4\pi\sigma^2} \left(S_0 - \frac{\theta^2}{4} S_2 + \mathcal{O}(\theta^4) \right), \\ S_0 &= \sum_{\ell=0}^{\infty} (2\ell+1) \exp(-\ell(\ell+1)\gamma), \quad S_2 = \sum_{\ell=0}^{\infty} \ell(\ell+1)(2\ell+1) \exp(-\ell(\ell+1)\gamma). \end{aligned} \quad (32)$$

The sums in (32) can be approximated by integrals I_0 and I_2 defined by

$$\begin{aligned} I_0 &= \int_0^{\infty} (2x+1) \exp(-x(x+1)\gamma) dx = \frac{1}{\gamma}, \\ I_2 &= \int_0^{\infty} x(x+1)(2x+1) \exp(-x(x+1)\gamma) dx = \frac{1}{4\gamma^2}. \end{aligned} \quad (33)$$

It is shown in the Appendix that S_0 is bounded from below and above by

$$\begin{aligned} &\gamma^{-1} (1 - \exp(-(2 - 4\sqrt{2\gamma} + 3\gamma))(1 - \exp(-\sqrt{\gamma}(\sqrt{2} + \sqrt{\gamma})))) \\ &< S_0 < \gamma^{-1} (1 + \sqrt{2\gamma} \exp(-(1 - \gamma/2)/2)). \end{aligned} \quad (34)$$

Thus, for small γ , the upper bound is $\gamma^{-1}(1 + \mathcal{O}(\sqrt{\gamma}))$, the lower bound is $\gamma^{-1}(1 - \mathcal{O}(\sqrt{\gamma}))$, and $S_0 \approx 1/\gamma$. An approximation of p_θ in (32) is then for small θ

$$p_\theta = \frac{1}{4\pi\sigma^2\gamma} \left(1 - \frac{\theta^2}{16\gamma} + \mathcal{O}(\theta^4) \right). \quad (35)$$

Assume that the normals of two adjacent tangent planes and the polar axis lie in the same plane. The molecule is initially at a point on the left plane \mathcal{P}_1 in Figure 2b with a polar angle $\theta = 0$. The angle between the normal of \mathcal{P}_1 and the polar axis is ϕ_1 . The intersection of the polar axis and $z = 0$ in the (ρ, z) coordinate system is at $\rho_1 = \sigma \tan(\phi_1)$. At $\theta = \theta_c$ there is an edge in the approximating surface and to the right we have the second plane \mathcal{P}_2 . The distance between the origin and the edge is $h = \sigma \tan(\phi_1 + \theta_c)$. This is a measure of the size of the facets on the sphere. The angle between the planes \mathcal{P}_1 and \mathcal{P}_2 is ϕ_2 . To determine a polar angle on

the approximation of the sphere, ρ is first sampled from a Gaussian distribution centered at $(\sigma \tan(\phi_1), 0)$ on \mathcal{P}_1 with the PDF

$$p_\rho(\rho, \Delta t) = \frac{1}{4\pi D_s \Delta t} \exp\left(\frac{-\rho^2}{4D\Delta t}\right) = \frac{1}{4\pi\gamma\sigma^2} \exp\left(-\frac{1}{4\gamma} \left(\frac{\rho}{\sigma}\right)^2\right). \quad (36)$$

Then the position θ on the sphere is obtained by projecting the sampled ρ in the radial direction on the sphere. Let the first sampled value be $\rho_1 \leq h - \sigma \tan(\phi_1)$. Then the new position of the molecule is on \mathcal{P}_1 and the corresponding sampled angle θ_1 satisfies

$$\tan(\theta_1 + \phi_1) = (\rho_1 + \sigma \tan(\phi_1))/\sigma. \quad (37)$$

The ρ_1 corresponding to a particular θ_1 is

$$\rho_1 = \sigma(\tan(\theta_1 + \phi_1) - \tan(\phi_1)). \quad (38)$$

If the sampled value is ρ_2 with $\rho_2 > h - \sigma \tan(\phi_1)$, then the position is projected on \mathcal{P}_2 along its normal. Then the new coordinates will be (see Figure 2b)

$$(\rho_p, z_p) = (\rho_* \cos(\phi_2)^2 + h \sin(\phi_2)^2, -(\rho_* - h) \sin(\phi_2) \cos(\phi_2)), \quad \rho_* = \rho_2 + \sigma \tan(\phi_1). \quad (39)$$

In this case, the new polar angle θ_2 after projection satisfies

$$\tan(\theta_2 + \phi_1) = \rho_p/(\sigma + z_p), \quad \theta_1 \leq \theta_c < \theta_2. \quad (40)$$

The sampled point $(\rho_2 + \sigma \tan(\phi_1), 0)$ corresponding to a given θ_2 is by (39) and (40)

$$\rho_p = \sigma \left(\frac{\tan(\theta_2 + \phi_1) + \sigma^{-1} h \sin(\phi_2)(\cos(\phi_2) \tan(\theta_2 + \phi_1) - \sin(\phi_2))}{\cos(\phi_2)(\cos(\phi_2) + \tan(\theta_2 + \phi_1) \sin(\phi_2))} - \tan(\phi_1) \right). \quad (41)$$

If the geometric resolution of the sphere is fine, i.e. $\sigma \gg h$, then ϕ_1 and ϕ_2 are small and after MacLaurin expansion in ϕ_2 we have

$$\rho_2 = \sigma \tan(\theta_2 + \phi_1)(1 + \phi_2(h/\sigma - \tan(\theta_2 + \phi_1)) + \mathcal{O}(\phi_2^2)) - \sigma \tan(\phi_1). \quad (42)$$

If the time step is small such that $\sigma \gg \sqrt{D_s \Delta t}$ then $\theta_2 = \mathcal{O}(\phi_2)$ and we always have $\phi_1 < \phi_2$. Hence, $h/\sigma - \tan(\theta_2 + \phi_1) = \sin(\phi_1) - \tan(\theta_2 + \phi_1) = \mathcal{O}(\phi_2)$ and ρ_p can be written

$$\rho_2 = \sigma(\tan(\theta_2 + \phi_1)(1 + \mathcal{O}(\phi_2^2)) - \tan(\phi_1)). \quad (43)$$

The expansion of the leading terms of the expression inside the parenthesis in (38) and (43) is now for $j = 1, 2$

$$\tan(\theta_j + \phi_1) - \tan(\phi_1) = \tan(\theta_j)(1 + \phi_1 \tan(\theta_j) + \mathcal{O}(\phi_1^2)). \quad (44)$$

The approximate PDF for the polar angle is then by (36), (38), (43), and (44)

$$p_\rho(\theta, \Delta t) = \frac{1}{4\pi\gamma\sigma^2} \exp\left(\frac{-\tan(\theta_j)^2(1 + \mathcal{O}(\phi_j^2))}{4\gamma}\right), \quad j = 1, 2. \quad (45)$$

For small angles, the Taylor expansion is

$$p_\rho = \frac{1}{4\pi D_s \Delta t} \left(1 - \frac{\theta_j^2}{4\gamma} + \mathcal{O}(\theta_j^2 \phi_j^2) \right), \quad j = 1, 2, \quad (46)$$

when the new location is on \mathcal{P}_j , $j = 1, 2$. Since $\theta_j < h/\sigma$, the quotient $\theta_j^2/\gamma < (h/\sqrt{D_s\Delta t})^2 \sim 1$. The quotient p_θ/p_ρ of p_θ in (30) and p_ρ in (36) for $\phi_1 = 0$ follows from (35) and (46) when γ is sufficiently small such that $S_0 \approx 1/\gamma$

$$\frac{p_\theta}{p_\rho} \approx 1 + \frac{3\theta^2}{16\gamma} + \dots \quad (47)$$

The conclusion from this analysis is as in Section 5.1. *The error in the PDF due to the approximation of the sphere by planar facets is small when h is of the same order as $\sqrt{D_s\Delta t}$ and the radius of the sphere σ fulfills $\sigma \gg \sqrt{D_s\Delta t}$.*

6 Results

In this section we consider three examples to illustrate the accuracy and the flexibility of the proposed method. First, we show that diffusion and reactions on a curved surface are computed accurately by comparing an analytical result and a simulation of two molecules diffusing and reacting on a sphere. In the two remaining examples, the need for a microscale solver for parts of the model arises for two conceptually different reasons.

In the second example, taken from [51], one reaction is highly diffusion limited and needs to be resolved in a detailed manner. If association is fast compared to diffusion, fast re-association reactions are a determining factor for the behavior of the model. To capture that phenomenon by a purely mesoscopic model, an extremely fine space discretization approaching the interaction radius of the molecules is required. For very small voxels, special corrections and modifications of the model are necessary [23, 24]. As the molecules are free to move in the entire domain, the mesh has to be uniformly fine. If the reaction pair to be resolved on the microscale is only one of many possible events in the system, most parts of the model are going to be vastly over-resolved, causing unnecessarily long simulation times. As an alternative to this approach, we show that a moderately fine mesh is sufficient if a small subsystem is simulated with the microscopic model.

In the final example, the geometry and molecular interaction of a certain species with the boundary motivate a more detailed microscopic model to be used locally in space. Instead of dividing only the species into a mesoscopic and a microscopic part as in the second example, we here partition both species and space and show that high accuracy is retained while significantly reducing the number of molecules that are simulated at the microscopic level.

6.1 Diffusion and reactions on a membrane

To show that our treatment of surfaces is accurate we consider a simple example with two molecules A and B diffusing on a sphere and reacting instantly upon collision. For a single molecule diffusing on a sphere, one can compute the average time for the molecule to be captured by a trap defined by a cone with half-angle θ and a spherical cap [8, 37]. Since we are considering molecules with small radii compared to the radius of the sphere, the time for the two molecules to react can be well approximated by the case of a single molecule and a trap. Thus, for two molecules with diffusion constants D_A and D_B and reaction radii σ_A and σ_B the average time until they react is close to

$$\tau_r = \frac{\sigma^2}{D_A + D_B} \left(\frac{2 \ln \left(\frac{2}{1 - \cos(\sigma_A + \sigma_B)} \right)}{1 + \cos(\sigma_A + \sigma_B)} - 1 \right), \quad (48)$$

where σ is the radius of the sphere. In Figure 6.1, we compare the result of microscopic simulations for three different mesh resolutions to τ_r in (48). The error bars correspond to 95% confidence intervals computed by sampling 6500 realizations of the system. The parameters are $D_A = D_B = 10^{-14}\text{m}^2$, $\sigma_A = \sigma_B = 10^{-9}\text{m}$ and the radius of the sphere is $9.23 \cdot 10^{-7}\text{m}$. The time step is chosen so small that the error from the discretization of the sphere is larger than that due to the discretization of time.






Fig. 3: The average time until a reaction between two molecules diffusing on a sphere with different triangulations compared to τ_r in (48). Even for the coarsest mesh resolution, the analytical solution falls in the 95% confidence interval.

As we see in the figure, the approximation improves as we refine the discretization of the sphere as expected. A finer mesh requires a smaller time step, and since the molecules on average diffuse a distance proportional to \sqrt{t} , the time step has to decrease as h^2 where h is the length parameter of the voxels, cf. Sections 4.4 and 5. It would be desirable to be able to estimate the error more precisely for general surfaces *a priori* or *a posteriori*.

6.2 Simulating highly diffusion limited bimolecular reactions

In [48] a model of one layer of a mitogen-activated protein kinase (MAPK) pathway is studied.



In the first two reactions in (49) K is phosphorylated in two steps by the kinase KK to become K_p and K_{pp} . The dephosphorylation of K by the phosphatase P occurs in two steps in (50) from K_{pp} to K_p and finally to K . The enzymes are inactive, KK^* and P^* , after the reactions in (49) and (50) and are activated in (51) to become KK and P . They are now prepared

to act as enzymes in new phosphorylation and dephosphorylation reactions. The parameters k_i , $i = 1, \dots, 6$, have typical values and $k_7 = \ln(2)/\tau_{\text{rel}}$, see [48].

The response time τ_{res} is defined as the average time to reach 50% of the steady state level of the doubly phosphorylated substrate K_{pp} . It depends on spatio-temporal correlations which cannot be captured on the macroscale or the mesoscale when k_7 is large and the diffusion is slow. The difference between the microscopic and mesoscopic simulations is due to the possibility of fast rebindings. In particular, if k_7 is large, then KK^* and P^* will be reactivated quickly. Thus, on the microscale, the dissociation of $KK - K$ and $P - K_{pp}$ could be followed by a fast reactivation and then a rebinding to K or K_p or to K_{pp} or K_p . If this process is sufficiently fast, it cannot be captured on the macro- or mesoscale unless the spatial resolution is very high. In [24] it is shown how to compute correct mesoscopic rate constants and how to account for reactions over neighboring voxels in the RDME to be able to simulate the MAPK model in (49), (50), and (51) accurately. This, however, requires a uniformly fine mesh down to scales of just a few molecular radii.

Since most of the system can be simulated accurately on the mesoscale, we propose a splitting of the system where the complexes $KK - K$ and $P - K_{pp}$ are simulated on the microscale and the rest of the system is simulated on the mesoscale. When $KK - K$ or $P - K_{pp}$ dissociates the products are simulated on the microscale in the remaining part of the splitting time step $\Delta\tau$. Thus for a reasonable $\Delta\tau$ that time will be sufficient for the molecules to become well-mixed in a voxel, see Section 4.4. After a splitting time step, the products are again considered as mesoscopic.

In Figure 4 we compare the results of our hybrid method to those obtained with GFRD in [48] and with the macroscale by solving the reaction rate equations. The parameter in the definition of k_7 is $\tau_{\text{rel}} = 1\mu\text{s}$. The other parameters in (49), (50), and (51) are as in [48].

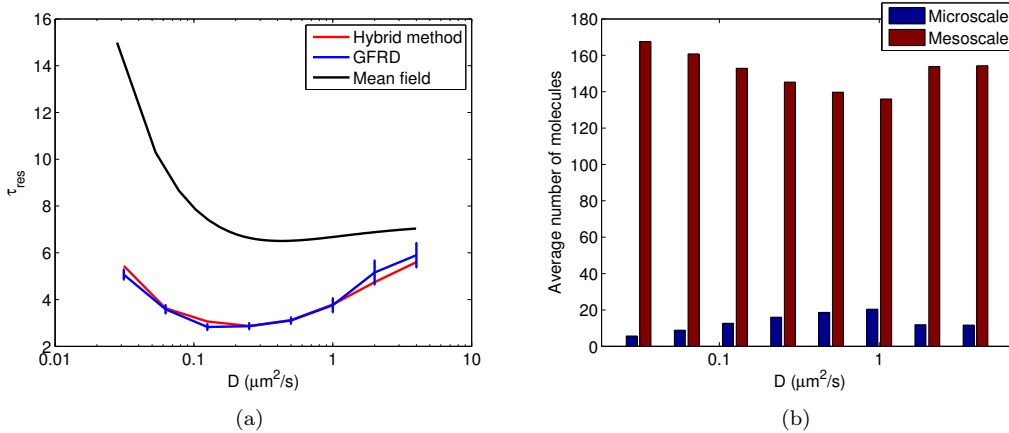


Fig. 4: (a) The average response time τ_{res} of the doubly phosphorylated substrate K_{pp} as a function of the diffusion constant D . (b) The average number of molecules simulated on the microscale in the beginning of a time step compared to the average number of molecules simulated on the mesoscale.

The response is significantly faster on the microscale than on the macroscale. The hybrid method agrees very well with the results from [48] as shown in Figure 4a, but the fraction of the total number of molecules simulated on the microscale is significantly reduced, see Figure 4b. The cost of simulating molecules with the microscopic method grows almost quadratically with the number of molecules [29]. Although some overhead is introduced in the hybrid scheme, we

still gain orders of magnitude in computing time with the mixed method without compromising the accuracy.

6.3 Molecular transport from the cytosol to the nucleus

The flexibility and full capability of our algorithm is demonstrated in an example with more than one subdomain, reactions both in the bulk and on a membrane, as well as reversible adsorption of a molecule to a membrane. The model we consider will be a prototypical model for transport from the cytosol to the nucleus via nuclear pore complexes (NPCs). It will not contain all known components in that process; we will rather focus on the key interactions between a species in the cytoplasm and the nuclear membrane and how accurately that interaction can be simulated by the mesoscopic, microscopic and hybrid methods respectively.

The model is defined in Table 1. It consists of three logical modules: a set of reactions in the cytosol (1–4) that result in the creation of a species C , C 's reversible binding to the nuclear membrane (5–6), and its subsequent reaction with the pore P on the membrane (7–9) leading to its translocation into the nucleus. The proteins A and B are here abstract variables, but as an example we can envision B being the end product of the signal cascade modeled in the previous section while A could be modeling an importin, a helper protein that needs to bind to a protein targeted for nuclear transport in order to link it to the pore. Reaction (2) is then the coarse-grained creation of B , and (1) would implicitly model the recycling of the importin A from the nucleus to the cytoplasm.

Reactions in the bulk	
(1)	$\emptyset \xrightarrow{\mu_1} A$
(2)	$\emptyset \xrightarrow{\mu_2} B$
(3, 4)	$A + B \xrightleftharpoons[k_2]{k_1} C$
Reversible adsorption of C to the membrane	
(5, 6)	$C \xrightleftharpoons[k_d]{k_a} C_m$
Translocation of C into nucleus	
(7, 8)	$C_m + P \xrightleftharpoons[k_p^-]{k_p} C_m-P$
(9)	$C_m-P \xrightarrow{k_r} P + C_n$

Table 1: The reactions of the full translocation model. Reactions (1–4) take place in the bulk, (5) near the reactive boundary, the membrane, and (6–9) on the membrane. After B is released in the nucleus, A can be recycled to the cytoplasm to complete another transport cycle.

6.3.1 Adsorption to a membrane

Before simulating the full system we illustrate the accuracy of the hybrid approach for reactions between molecules and a surface in Figure 5. The reversible adsorption of C to the membrane (5,6) is simulated with the mesoscopic, microscopic and hybrid methods in Figure 5 using a space discretization with 6706 voxels. The domain is taken to be a box with sides of length $4.6 \cdot 10^{-6} m$ and $6 \cdot 10^{-6} m$ at the bottom and the height $4.8 \cdot 10^{-6} m$. The boundary is reactive

at the bottom, $z = 0$. The total number of molecules in the simulation is 10. The reaction rates are $k_a = 10^{-7} \text{ ms}^{-1}$ and $k_d = 1.0 \text{ s}^{-1}$ and the diffusion constant is $D = 10^{-12} \text{ m}^2\text{s}^{-1}$.

The rate constant for mesoscopic binding of a molecule to the surface has unit ms^{-1} . It is thus necessary to divide with a local length scale in the mesoscopic propensity. How to choose that length parameter in the case of reactions with surfaces or between molecules on surfaces is unknown for an unstructured, tetrahedral mesh. In the case of a structured Cartesian mesh, correct mesoscopic rate constants for reactions occurring on a plane surface are derived in [24].

For the unstructured mesh, using two different naïve choices for the length scale parameter l in the mesoscopic propensities for membrane binding, the mesoscopic and microscopic models differ for the steady-state distribution of the total number C on the membrane. Use as l the local mesh size parameter h provided by the mesh generator and the radius of a sphere, r_{sph} with the same volume as the voxel sharing triangular faces with the membrane. The mesoscopic results differ from each other, from the microscopic simulations and the mean field prediction. Using the radius of a sphere improves the result, but there is still an error. However, simulating all C molecules in voxels with centers ζ within a slice of the domain defined by $\mathcal{Z}_{\text{micro}} = \{\zeta = (x, y, z) | z < z_{\text{micro}}\}$ yields good results also for relatively small values of z_{micro} (a). The number of molecules at the micro level increases slowly when z_{micro} increases as seen in Figure 5b. When $z_{\text{micro}} = 1.5 \mu\text{m}$ the hybrid simulations show excellent agreement with pure microscale simulations in the whole domain, but only $\sim 40\%$ of the molecules are updated with the microscopic solver.

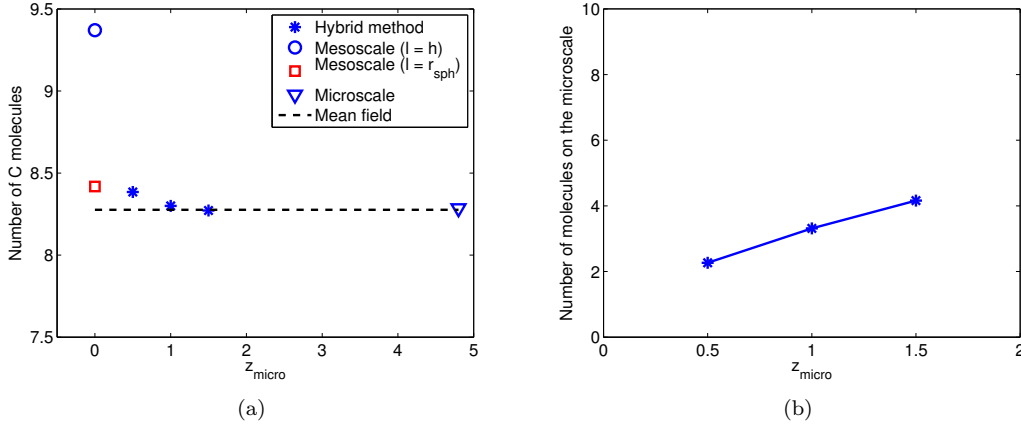


Fig. 5: (a) The result of the hybrid method with z_{micro} in μm is compared with the mesoscopic model with the length parameter $l = h$ and $l = r_{sph}$ and with the pure microscopic model. The dashed line corresponds to the mean field prediction. (b) The time averaged number of microscopic molecules in a simulation with the hybrid method for different z_{micro} .

6.3.2 The full translocation model

Figure 6a displays a snapshot of the full system in Table 1 in a purely microscopic simulation. The radius of the cell (the outer sphere) is taken to be $2.58\mu\text{m}$ and that of the nucleus (the inner sphere) $0.98\mu\text{m}$. These are realistic sizes for a yeast cell in an early stage of the cell cycle.

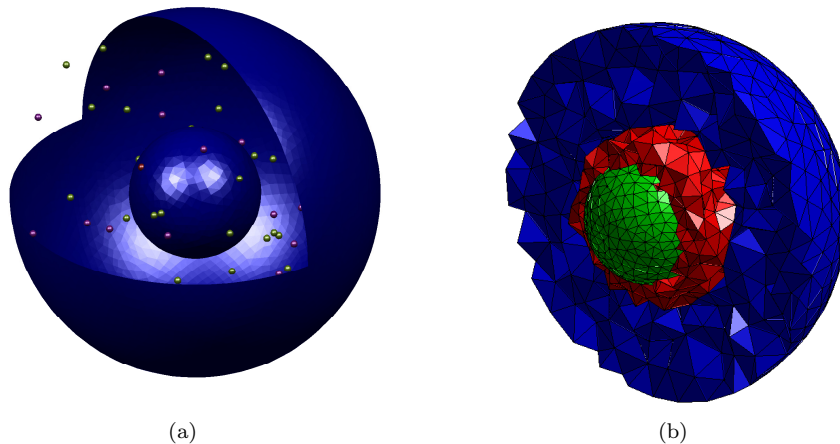


Fig. 6: (a) A snapshot of a purely microscopic simulation. The size of molecules is larger in the figure than in the actual simulation for illustrative purposes. (b) The partial adsorption of C molecules to the nuclear membrane (green) at the microscopic level is resolved by letting C be a microscopic degree of freedom in a layer of tetrahedral elements around the nucleus (red). In the remaining part of the domain (blue), the system is updated in time by the mesoscopic model.

If species A corresponds to an importin protein, it will be present in higher copy numbers than the other species in the model. All A molecules are therefore simulated mesoscopically as well as all other molecules outside a sphere with radius $1.3 \cdot 10^{-6}m$, the blue domain in Figure 6b. All molecules (except for A molecules) are simulated on the microscale close to and on the membrane, the red and green domains in Figure 6b. We are using the more expensive microscopic simulation only near and on the nuclear membrane, where it is most needed for the accuracy.

The reaction rates in Table 1 are $\mu = 0.03 s^{-1}$, $k_1 = 2.5336 \cdot 10^{-20} m^3s^{-1}$, $k_2 = 0.02 s^{-1}$, $k_a = 10^{-6} ms^{-1}$, $k_d = 0.02 s^{-1}$, $k_p = 10^{-11} m^2s^{-1}$, $k_p^- = 0.04 s^{-1}$ and $k_r = 0.1 s^{-1}$. The diffusion constant D for all molecules in space is $10^{-12} m^2s^{-1}$ and on the membrane D_s is $10^{-14} m^2s^{-1}$. The reaction radius σ for all molecules is $10^{-9} m$ and the splitting time step $\Delta\tau$ is 0.5. In Figure 7, the average number of molecules during a simulation of the full system is plotted. Initially, there are 1000 A molecules in the cytoplasm, while B and C are present in low copy numbers.

In this example the mesoscopic diffusion of the A molecules is dominating the execution time, and the time spent on the microscale simulations is negligible. Hence, the cost of the hybrid simulation is close to that of a purely mesoscopic simulation, but with higher accuracy for reactions near and on the membrane. The cost of the microscale simulation and the overhead introduced by the coupling routines in Section 4 depends on the problem and the parameters. In the above examples, the overhead ranges between 0 and 30 per cent of the total execution time. The microscale simulation could be dominant if a major part of the molecules are bound to a membrane as the interaction between molecules on a surface is expensive to simulate.




Fig. 7: The time evolution of the system in Table 1 is simulated.

7 Discussion

For certain systems it has been shown that the classical, purely mesoscopic description for reaction-diffusion systems fails to capture important dynamics predicted by a corresponding microscopic description [23, 24, 48, 52]. This is due to the fact that the assumption of molecules being well stirred in a voxel is not satisfied at fine scales. Thus, in order to resolve the model a high mesh resolution is needed and this leads to a breakdown of the mesoscopic model in its classical formulation [30]. In [24] this issue is addressed by computing corrected mesoscopic rate constants and considering reactions involving neighboring voxels.

In this work, we have taken another approach. Instead of using a very fine mesh in a mesoscopic simulation to resolve fast association reactions, sharp gradients, or geometric details, we combine the mesoscopic model with a microscopic model in a hybrid method. The main advantages are our ability to resolve fine scaled structures in the model without increasing the stiffness in the parts of the model that can be treated accurately at the mesoscopic level as well as the possibility to switch to a microscopic model when the corresponding mesoscopic model cannot easily be defined. Numerical examples illustrate both these capabilities of the method. Ultimately, how the approach taken in this paper will compare to that in [24] in terms of computational efficiency will depend on the nature of the particular model under study. The extension of [24] to unstructured meshes is non-trivial and further work is needed in order to be able to make a fair efficiency comparison, but in general, the more scale separation the model displays the more likely it is that the hybrid method will be faster.

An alternative is to use adaptive mesh refinement [7] to increase the accuracy locally in parts of the domain. This is possible provided that all parts of the model can be formulated on the mesoscale and that the RDME can be corrected to account for errors due to small voxels. However, adaptive mesh refinement is complicated to implement in three dimensions with preserved mesh quality for unstructured meshes, and we believe that our approach is simpler, more accurate and more flexible.

In Sections 4.4 and 5 we show that it is necessary to restrict the time step Δt in the microscopic solver such that the probability is small that the molecules traverse more than one voxel

in Δt . That time step scales as the mesh size squared. In a purely microscopic simulation with GFRD style algorithms, it is not necessary to restrict the time step in such a way. For a fine mesh this could of course be adding a significant overhead to the microscale part of the hybrid simulation. However, the basic idea underlying our algorithm is that instead of using a very fine mesh where high accuracy is needed we use a microscopic simulation. Hence, for an appropriate mesh and partitioning, the potential overhead from this time step restriction should be small compared to the gain due to the reduction of the size of the microscale part of the system in relation to a full microscopic simulation.

In our method, a curved surface embedded in 3D is approximated by planar facets derived from the unstructured surface mesh. For problems where membranes are important, an unstructured mesh can represent the membrane geometry accurately. The precision of this approximation is analyzed for diffusion on a sphere. In numerical experiments with reactions and diffusion the agreement is very good between simulations and an analytical result.

In the examples, we made use of a static partitioning determined before the simulation was started. The implementation of the algorithm itself is very flexible, and would allow for an extension to dynamic partitioning of the system. The development of good criteria to base such adaptive splittings on is a natural next step to make the algorithm fully automatic.

Acknowledgement

This work has been supported by the Swedish Research Council and the Graduate School in Mathematics and Computing at Uppsala University. Computing time has been provided by Uppsala Multidisciplinary Center for Advanced Computational Science (UPPMAX). Comments by Brian Drawert and Emil Kieri helped us improve the paper.

Appendix




Fig. 8: The upper and lower bounds on the terms of the sum S_0 with $\gamma = 0.8 \cdot 10^{-2}$, $x_- = 7$ and $x_+ = 8$.

The upper and lower bounds on S_0 in (32) are derived by integral estimates as illustrated in Figure 8.

Let $f(x)$ be the continuous function such that the terms in S_0 are $f(\ell)$, $\ell = 0, 1, \dots$. The maximum of

$$f(x) = (2x + 1) \exp(-x(x + 1)\gamma)$$

is $\sqrt{2} \exp(-(1 - \gamma/2)/2) / \sqrt{\gamma}$ at $x_{\max} = (\sqrt{2/\gamma} - 1)/2$. Let $x_- = \lfloor x_{\max} \rfloor$ and $x_+ = \lceil x_{\max} \rceil$

denote the largest and smallest integers such that $x_- \leq x_{\max} \leq x_+$. Then by Figure 8

$$\int_{-1}^{x_- - 1} f(x) dx = \int_0^{x_-} f(x-1) dx < \sum_{\ell=0}^{x_- - 1} f(\ell) < \int_0^{x_-} f(x) dx,$$

$$\int_{x_+}^{\infty} f(x) dx < \sum_{\ell=x_+}^{\infty} f(\ell) < \int_{x_+}^{\infty} f(x-1) dx = \int_{x_+ - 1}^{\infty} f(x) dx.$$

By combining the above inequalities and observing that $\int_{x_- - 1}^{x_-} f(x) dx < f(x_-)$, the bounds on the sum are

$$\int_{-1}^{x_-} f(x) dx + \int_{x_+}^{\infty} f(x) dx < \sum_{\ell=0}^{\infty} f(\ell) < \int_0^{\infty} f(x) dx + f(x_-).$$

A lower bound on the integrals in the left hand side is

$$\begin{aligned} & \int_{-1}^{x_-} f(x) dx + \int_{x_+}^{\infty} f(x) dx \\ &= \gamma^{-1}(1 - \exp(-x_-(x_- + 1)\gamma) + \exp(-x_+(x_+ + 1)\gamma)) \\ &\geq \gamma^{-1}(1 - \exp(-x_{\max}(x_{\max} - 1)\gamma)(1 - \exp(-2(x_- + 1)\gamma))) \\ &\geq \gamma^{-1}(1 - \exp(-x_{\max}(x_{\max} - 1)\gamma)(1 - \exp(-2(x_{\max} + 1)\gamma))) \\ &= \gamma^{-1}(1 - \exp(-(2 - 4\sqrt{2}\gamma + 3\gamma))(1 - \exp(-\sqrt{\gamma}(\sqrt{2} + \sqrt{\gamma}))))), \end{aligned}$$

and an upper bound on the sum is $\gamma^{-1} + f(x_{\max})$. These are the bounds in (34).




Fig. 9: The quotient (solid) between the sum S_0 and $1/\gamma$ and the upper and lower bounds on γS_0 (dashed) for different values of γ .

The bounds and S_0 divided by $1/\gamma$ are compared for varying γ in Figure 9. The gap between the upper and lower estimates closes rapidly for decreasing γ .

References

- [1] M. Abramowitz and I. A. Segun. *Handbook of Mathematical Functions*. Dover, New York, 1965.
- [2] M. Ander, P. Beltrao, B. D. Ventura, J. Ferkinghoff-Borg, M. Foglierini, A. Kaplan, C. Lemerle, I. Tomas-Oliveira, and L. Serrano. SmartCell, a framework to simulate cellular processes that combines stochastic approximation with diffusion and localisation: analysis of simple networks. *Syst. Biol.*, 1:129–138, 2004.

- [3] S. S. Andrews. Accurate particle-based simulation of adsorption, desorption and partial transmission. *Phys. Biol.*, 6:046015, 2009.
- [4] S. S. Andrews, N. J. Addy, R. Brent, and A. P. Arkin. Detailed simulations of cell biology with Smoldyn 2.1. *PLoS Comput. Biol.*, 6(3):e1000705, 2010.
- [5] S. S. Andrews and D. Bray. Stochastic simulation of chemical reactions with spatial resolution and single molecule detail. *Phys. Biol.*, 1:137–151, 2004.
- [6] S. N. V. Arjunan and M. Tomita. A new multicompartmental reaction-diffusion modeling method links transient membrane attachment of E. coli MinE to E-ring formation. *Syst. Synth. Biol.*, 4:35–53, 2010.
- [7] B. Bayati, P. Chatelin, and P. Koumoutsakos. Adaptive mesh refinement for stochastic reaction-diffusion processes. *J. Comput. Phys.*, 230:13–26, 2011.
- [8] V. A. Bloomfield and S. Prager. Diffusion-controlled reactions on spherical surfaces. *Biophys. J.*, 27:447–453, 1979.
- [9] K. Burrage, J. Hancock, A. Leier, and D. V. Nicolau Jr. Modelling and simulation techniques for membrane biology. *Brief. Bioinf.*, 8(4):234–244, 2007.
- [10] H. S. Carslaw and J. C. Jaeger. *Conduction of Heat in Solids*. Oxford University Press, London, 2nd edition, 1959.
- [11] M. Christensen. How to simulate anisotropic diffusion processes on curved surfaces. *J. Comput. Phys.*, 201:421–438, 2004.
- [12] J. Cullhed, B. Drawert, S. Engblom, and A. Hellander. Urdme 1.1: User’s manual. Technical Report 003, Department of Information Technology, Division of Scientific Computing, Uppsala University, 2010.
- [13] R. Delgado-Buscalioni, K. Kremer, and M. Praprotnik. Coupling atomistic and continuum hydrodynamics through a mesoscopic model: Application to liquid water. *J. Chem. Phys.*, 131:244107, 2009.
- [14] M. Dobrzyński, J. V. Rodríguez, J. A. Kaandorp, and J. G. Blom. Computational methods for diffusion-influenced biochemical reactions. *Bioinformatics*, 23(15):134–155, May 2007.
- [15] A. Donev, J. B. Bell, A. L. Garcia, and B. J. Alder. A hybrid particle-continuum method for hydrodynamics of complex fluids. *Multiscale Model. Simul.*, 8:871–911, 2010.
- [16] B. Drawert, M. J. Lawson, L. Petzold, and M. Khammash. The diffusive finite state projection algorithm for efficient simulation of the stochastic reaction-diffusion master equation. *J. Chem. Phys.*, 132(7):074101, 2010.
- [17] W. E, B. Engquist, X. Li, W. Ren, and E. Vanden-Eijnden. Heterogeneous multiscale methods: A review. *Comm. Comput. Phys.*, 2(3):367–450, 2007.
- [18] A. Edelman and N. Agmon. Brownian dynamics simulations of reversible reactions in one dimension. *J. Chem. Phys.*, 99:5396–5404, 1993.
- [19] J. Elf and M. Ehrenberg. Spontaneous separation of bi-stable biochemical systems into spatial domains of opposite phases. *Syst. Biol.*, 1:230–236, 2004.

- [20] J. Elf, G.-W. Li, and X. S. Xie. Probing transcription factor dynamics at the single-molecule level in a living cell. *Science*, 316(5828):1191–1194, 2007.
- [21] M. B. Elowitz, A. J. Levine, E. D. Siggia, and P. S. Swain. Stochastic gene expression in a single cell. *Science*, 297(5584):1183–1186, 2002.
- [22] S. Engblom, L. Ferm, A. Hellander, and P. Lötstedt. Simulation of stochastic reaction-diffusion processes on unstructured meshes. *SIAM J. Sci. Comput.*, 31:1774–1797, 2009.
- [23] R. Erban and J. Chapman. Stochastic modelling of reaction-diffusion processes: algorithms for bimolecular reactions. *Phys. Biol.*, 6:046001, 2009.
- [24] D. Fange, O. G. Berg, P. Sjöberg, and J. Elf. Stochastic reaction-diffusion kinetics in the microscopic limit. *Proc. Natl. Acad. Sci. USA*, 107(46):19820–19825, 2010.
- [25] D. Fange and J. Elf. Noise induced Min phenotypes in *E. coli*. *PLoS Comput. Biol.*, 2(6):e80, 2006.
- [26] A. L. Garcia, J. B. Bell, W. Y. Crutchfield, and B. J. Alder. Adaptive mesh and algorithm refinement using direct simulation Monte Carlo. *J. Comput. Phys.*, 154:134–155, 1999.
- [27] D. T. Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J. Comput. Phys.*, 22(4):403–434, 1976.
- [28] J. Hattne, D. Fange, and J. Elf. Stochastic reaction-diffusion simulation with MesoRD. *Bioinformatics*, 21:2923–2924, 2005.
- [29] S. Hellander and P. Lötstedt. Flexible single molecule simulation of reaction-diffusion processes. *J. Comput. Phys.*, 230:3948–3965, 2011.
- [30] S. A. Isaacson. The reaction-diffusion master equation as an asymptotic approximation of diffusion to a small target. *SIAM J. Appl. Math.*, 70:77–111, 2009.
- [31] N. G. van Kampen. *Stochastic Processes in Physics and Chemistry*. Elsevier, Amsterdam, 5th edition, 2004.
- [32] R. A. Kerr, T. M. Bartol, B. Kaminsky, M. Dittrich, J.-C. J. Chang, S. B. Baden, T. J. Sejnowski, and J. R. Stiles. Fast Monte Carlo simulation methods for biological reaction-diffusion systems in solution and on surfaces. *SIAM J. Sci. Comput.*, 30(6):3126–3149, 2008.
- [33] I. G. Kevrekidis, C. W. Gear, J. M. Hyman, P. G. Kevrekidis, O. Runborg, and C. Theodoropoulos. Equation-free, coarse-grained multiscale computation: Enabling microscopic simulators to perform system-level analysis. *Comm. Math. Sci.*, 1(4):715–762, 2003.
- [34] H. Kim and K. J. Shin. Exact solution of the reversible diffusion-influenced reaction for an isolated pair in three dimensions. *Phys. Rev. Lett.*, 82(7):1578–1581, 1999.
- [35] H. Kim, M. Yang, and K. J. Shin. Dynamic correlation effect in reversible diffusion-influenced reactions: Brownian dynamics simulation in three dimensions. *J. Chem. Phys.*, 111(3):1068–1075, 1999.
- [36] P. Koumoutsakos. Multiscale flow simulations using particles. *Ann. Rev. Fluid Mech.*, 37:457–487, 2005.

- [37] J. Lindermann and D. Lauffenburger. Analysis of intracellular receptor/ligand sorting - calculation of mean surface and bulk diffusion times within a sphere. *Biophys. J.*, 50:295–305, 1986.
- [38] M. I. Monine and J. M. Haugh. Reactions on cell membranes: Comparison of continuum theory and Brownian dynamics simulations. *J. Chem. Phys.*, 123:074908, 2005.
- [39] I. L. Novak, F. Gao, Y.-S. Choi, D. Resasco, J. C. Schaff, and B. M. Slepchenko. Diffusion on a curved surface coupled to diffusion in the volume: Application to cell biology. *J. Comput. Phys.*, 226:1271–1290, 2007.
- [40] T. Oettel, V. V. Bulatov, A. Donev, M. H. Kalos, G. H. Gilmer, and B. Sadigh. First-passage kinetic Monte Carlo method. *Phys. Rev. E*, 80:066701, 2009.
- [41] A. Raj, P. van den Bogaard, S. A. Rifkin, and A. van Oudenaarden. Imaging individual mRNA molecules using multiple singly labeled probes. *Nat. Meth.*, 5(10):877–879, 2008.
- [42] C. Sanford, M. L. K. Yip, C. White, and J. Parkinson. Cell++ – simulating biochemical pathways. *Bioinformatics*, 22:2918–2925, 2006.
- [43] I. F. Sbalzarini, A. Hayer, A. Helenius, and P. Koumoutsakos. Simulations of (an)isotropic diffusion on curved biological surfaces. *Biophys. J.*, 90:878–885, 2006.
- [44] P. Schwarz, D. Adalsteinsson, P. Colella, A. P. Arkin, and M. Onsum. Numerical computation of diffusion on a surface. *Proc. Natl. Acad. Sci. USA*, 102(32):11151–11156, 2005.
- [45] M. v. Smoluchowski. Versuch einer mathematischen Theorie der Koagulationskinetik kolloider Lösungen. *Z. phys. Chemie*, 92:129–168, 1917.
- [46] P. S. Swain. Efficient attenuation of stochasticity in gene expression through post-transcriptional control. *J. Mol. Biol.*, 344(4):965 – 976, 2004.
- [47] P. S. Swain, M. B. Elowitz, and E. D. Siggia. Intrinsic and extrinsic contributions to stochasticity in gene expression. *Proc. Natl. Acad. Sci. USA*, 99(20):12795–12800, 2002.
- [48] 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. USA.*, 107(6):2473–2478, 2010.
- [49] D. Tskhakaya, K. Matyash, R. Schneider, and F. Taccaogna. The Particle-In-Cell method. *Contrib. Plasma Phys.*, 47:563–594, 2007.
- [50] S. Wils and E. D. Schutter. STEPS: modeling and simulating complex reaction-diffusion systems with Python. *Front. Neuroinf.*, 3:15, 2009.
- [51] J. S. van Zon and M. J. Morelli and S. Tănase-Nicola and P. R. ten Wolde. Diffusion of transcription factors can drastically enhance the noise in gene expression. *Biophys. J.*, 91:4350–4367, 2006.
- [52] J. S. van Zon and P. R. ten Wolde. Green’s-function reaction dynamics: A particle-based approach for simulating biochemical networks in time and space. *J. Chem. Phys.*, 123:234910, 2005.
- [53] J. S. van Zon and P. R. ten Wolde. Simulating biochemical networks at the particle level and in time and space: Green’s-function reaction dynamics. *Phys. Rev. Lett.*, 94:128103, 2005.