

MULTISCALE MODELING OF DIFFUSION PROCESSES IN DENDRITES AND DEDRITIC SPINES

by

Fredrik Eksaa Pettersen

THESIS

for the degree of

MASTER OF SCIENCE



Faculty of Mathematics and Natural Sciences
University of Oslo

June 2014

Acknowledgements

This masters thesis marks the end of a five year education, which I could not have completed without the support of my parents.

I would like to thank my counselor, Hans Petter Langtangen, for convincing me to rewrite the entire thesis. It was hard, but worth it. Thanks to Geir Halnes, and Gaute Einevoll for help with the applications part of the thesis, and interesting discussions.

Thanks to Torbjørn Seland, for giving me valuable input on formulations.

And finally, a special thanks to my fantastic girlfriend, Nina, who has put up with my complaining throughout this project!

Contents

1	Introduction	5
1.1	Outline	7
2	Theory	9
2.1	Random Walks	10
2.1.1	An error estimate for a diffusion RW solver	12
2.1.2	Random number generator	12
2.2	Backward Euler schemes in two or more spatial dimensions . .	12
2.3	Combining micro and macro scale models	17
2.3.1	The algorithm	17
2.3.2	Conversion between length scales	18
2.3.3	Coupling the models through step length	19
2.3.4	Boundary conditions for the random walk	20
2.4	Potential problems	20
2.5	Other possible coupling methods	22
3	Analysis	25
3.1	Computation of the error	26
3.2	Verification techniques	26
3.3	Testing the PDE solver	28
3.3.1	Verification by manufactured solutions	28
3.3.2	Verification by convergence tests	29
3.3.3	Verification of FE scheme by exact numerical solution .	32
3.3.4	Verification of BE scheme by exact numerical solution .	34
3.4	Verification of the RW solver	36
3.5	Verification of the hybrid solver	38
4	Physical application	41
4.1	Physical scope	42
4.2	Implementation	43
4.3	Parameters and details	44

5	Results	47
5.1	Results of verification	48
5.1.1	The FE scheme	48
5.1.2	The BE scheme and the block tridiagonal solver	48
5.1.3	The RW solver	49
5.1.4	The hybrid diffusion solver	49
5.2	Results for physical application	49
5.3	Using other particle dynamics	50
6	Discussion and conclusion	53
6.1	Discussion	54
6.1.1	Application	54
6.1.2	Possible extensions	55
6.1.3	Other work on the topic	56
6.2	Concluding remarks	56
A	Appendix	59
A.1	Debugging	60
A.1.1	Compiler/syntax errors	60
A.1.2	Segmentation faults	60
A.1.3	Finite difference methods	61
A.1.4	Random Walk and Monte Carlo methods	63
A.1.5	The developed software	64
A.1.6	Last resorts	64

List of Figures

2.1	A general volume contained by a surface.	10
2.2	Illustration of $C(x)$	12
2.3	Reason for RW update each time step	19
2.4	Workaround for negative concentrations, illustration	21
2.5	Idea behind averaging solutions	23
2.6	Combination by polynomial regression	24
3.1	Error plots FE	29
3.2	Error plots BE	29
3.3	Convergence tests in time FE scheme	30
3.4	Convergence tests for BE scheme	31
3.5	Spatial error tests	31
3.6	Numerical exact error plots FE	34
3.7	Numerical exact error plots for BE scheme	35
3.9	Test RW	37
3.10	Error test for BE combined with RW in 1D	39
3.11	Effects of increasing relative size of walk area	39
4.1	Pyramidal neuron	42
4.2	Difference between hybrid diffusion solver and dendrite - spine diffusion model	44
4.3	Initial condition used in the application	45
5.1	Diffusion time for long necked spines	50
5.2	Diffusion times with least squares fit	51
5.3	Hybrid solver with DSMC particle model	52
A.1	Visual debugging	62

List of Abbreviations

BE Backward Euler

DSMC Direct Simulation Monte Carlo

FE Forward Euler

FLOP Floating point Operation

MC Monte Carlo

PDE Partial Differential Equation

RW Random Walk

Chapter 1

Introduction

Diffusion processes are extremely important in modern science, and have so many applications that it is challenging to list them all. Brownian motion of particles, momentum in liquids, and atomic diffusion are just some examples of processes described by the diffusion equation.

Both in general transport processes, and particularly in diffusion processes, there are cases in which parts of the process cannot be described by a continuum model, but the rest of the process can. Examples of such processes are fluid flow in nanoporous media and diffusion in the extracellular space of the brain. There are also many examples in materials design.

These types of processes can be called multi scale processes in the sense that more than one model is required to describe the entire system. Usually these are on different length scales. There are three obvious ways to handle multi scale problems. The first alternative is to ignore the problem, and assume that the chosen model can accurately solve the problem. For diffusion processes, this approach often works a lot better than it should. A more realistic approach is to create an intermediate scale and model. In the limit between continuum and particle dynamics this is often called a *meso scale* model. An example of a meso scale model is dissipative particle dynamics [14], where clusters of particles are modeled as individual particles. The particle clusters have different properties than the individual atoms or molecules which make up the substance. A final alternative is a *hybrid* model that uses more than one model of the problem in order to describe it. Some hybrid models exist today, but these are mostly aimed at specific problems, like dendrite solidification [11], or hybrid fluid flow models which combine molecular dynamics and Navier-Stokes solvers [10]. Other hybrid solvers for diffusion processes have also been developed [6], but they are closed in the sense that the computer code is not commonly available.

The aim of this masters thesis is to develop and implement a simple, yet flexible hybrid diffusion solver from the ground and up where all parts of the theory and implementation are fully understood and transparent. In principle, any particle dynamics model could be used, but for the sake of verification a stochastic model has been implemented. This is no limitation, as the interface to the lower scale model is simple, and the lower scale model works as a standalone unit.

A large emphasis has been put on verification of all parts of the hybrid model. Each individual solver has tested both as a standalone unit, and as a part of the combined, hybrid solver in order to verify that the code functions properly.

1.1 Outline

The thesis in your hands has the following structure:

Chapter 2 contains a detailed review of the coupling between a diffusion equation and a lower scale Random Walk model, as well as a look at band diagonal linear systems.

Chapter 3 defines the error estimate and shows a thorough analysis of all parts of the developed software in order to verify that it functions properly.

In *chapter 4* the developed software is modified to describe a physical application in which a hybrid diffusion solver is demanded.

Chapter 5 presents and discusses the results from the verification of the computer code and the physical application.

Finally, *chapter 6* discusses the thesis as a whole and looks at possible improvements and extensions that can be done.

The *appendix* is a general guide to debugging the methods that have been implemented during this project.

Chapter 2

Theory

This chapter will concern the theory needed to implement a hybrid diffusion solver that combines a diffusion equation and Random Walk solver as a lower scale model. A Random Walk is a Monte Carlo (MC) method, which is a large group of algorithms that use random sampling to produce numerical results.

2.1 Random Walks

A random walker is a point object which after a certain amount of time jumps a fixed step length either right or left with equal probability for jumping either way. In d spatial dimensions the axis on which the jump happens is chosen at random before the direction of the jump is decided.

Say N random walkers are placed at the same position within a volume, V of arbitrary shape, to make a spatial distribution of walkers, $C = N\delta(x)$. Here, $\delta(x)$ is the Dirac Delta function [3], defined by its properties

$$\delta(x - a) = \begin{cases} 0, & x \neq a \\ \infty, & x = a \end{cases} \quad (2.1)$$

V is contained by a surface S and the relation is shown in Figure 2.1.




Figure 2.1: A general volume contained by a surface.

A flux of walkers is a vector size describing the number of walkers passing a surface element dS per unit time, per unit area.

As the random walkers start to move, some of them will leave V , which results in a change in C . The number of walkers leaving V is described by

the change in C integrated over the entire volume,

$$\int_V \frac{\partial C}{\partial t} dV. \quad (2.2)$$

The number of walkers leaving V must be equal to the number of walkers crossing S per unit time and area. Integrating the outward normal component of the flux of walkers, \mathbf{J} over the entire surface results in the number of walkers leaving V ,

$$\int_S \mathbf{J} \cdot \mathbf{n} dS. \quad (2.3)$$

Since the number of walkers is conserved, the change in C per time must be equal to the number of walkers leaving V plus any walkers added to V :

$$\int_V \frac{\partial C}{\partial t} dV = \int_S \mathbf{J} \cdot \mathbf{n} dS + s(x, t), \quad (2.4)$$

where $s(x, t)$ denotes a source term. Through Greens theorem, (2.3) can be reformulated

$$\int_V \frac{\partial C}{\partial t} dV = \int_V \nabla \cdot \mathbf{J} dV + s(x, t). \quad (2.5)$$

The flux \mathbf{J} can be expressed by Fick's first law as the diffusive flux

$$\mathbf{J} = D\nabla C, \quad (2.6)$$

where D is the diffusion coefficient. Since V was chosen as an arbitrary volume (2.5) is independent of the integration and the integrals can be dropped. The diffusive flux \mathbf{J} (2.6) is also inserted to yield the diffusion equation

$$\frac{\partial C}{\partial t} = D\nabla^2 C + s(x, t). \quad (2.7)$$

Throughout this thesis C will denote a spatial distribution of random walkers while u will denote the solution to the diffusion equation, which is a concentration distribution. The difference between u and C is best illustrated in Figure 2.2.




Figure 2.2: Illustration of the difference between $u(x, t)$ (in (a)) which is the solution to the diffusion equation, and $C(x, t)$ (in (b)) which is the number of random walkers within an area of $\pm \frac{\Delta x}{2}$ around each mesh point.

2.1.1 An error estimate for a diffusion RW solver

In general a Monte Carlo method will fluctuate around the correct (steady state) solution. The fluctuations have an amplitude related to the number of samples, N used by the relation

$$\epsilon \propto \frac{1}{\sqrt{N}}. \quad (2.8)$$

Equation (2.8) will be used as an error estimate for the solution of diffusion processes by a random walk (RW) model and is assumed to hold for all parts of the simulation, not only the steady state.

2.1.2 Random number generator

To produce a random walk one must have random numbers. The random numbers in this thesis are produced by a five seeded xor-shift algorithm copied from George Marsaglia [9]. This generator is chosen because of its very large period ($\sim 10^{48}$) compared to the computational cost.

2.2 Backward Euler schemes in two or more spatial dimensions

The reader is assumed to know the basics of solving partial differential equations (PDEs) by finite difference methods. For the sake of clarity the two

schemes used to solve the diffusion equation (2.7) in this thesis are written out using the theta rule for finite differences:

$$\frac{u^{k+1} - u^k}{\Delta t} = \theta D \frac{\partial^2 u^{k+1}}{\partial x^2} + (1 - \theta) D \frac{\partial^2 u^k}{\partial x^2}. \quad (2.9)$$

Setting $\theta = 0$ gives the Forward Euler (FE) scheme which is explicit and conditionally stable with stability criterion

$$\Delta t \leq \frac{\Delta x^2}{2dD_{max}}, \quad (2.10)$$

where d is the spatial dimension and D_{max} denotes the maximum value of the diffusion coefficient. Setting $\theta = 1$ gives the Backward Euler (BE) scheme which is unconditionally stable.

This section will present a simple, but elegant way of efficiently solving the diffusion equation by the implicit BE scheme in more than one spatial dimension. The reason for using the BE scheme is that we want to use a large time step and still have a large spatial resolution. First, let us look at the BE discretization of the diffusion equation in 1D:

$$u_i^{k+1} = \frac{\Delta t}{2\Delta x^2} ((D_{i+1} + D_i)(u_{i+1}^{k+1} - u_i^{k+1}) - (D_i + D_{i-1})(u_i^{k+1} - u_{i-1}^{k+1})) + u_i^k. \quad (2.11)$$

Here u is the unknown function which solves the diffusion equation. Note that

$$u_i^k = u(t_0 + k \cdot \Delta t, x_0 + i \cdot \Delta x),$$

and not u to the k 'th power.

Neumann boundary conditions will be used, these are described in (2.12).

$$\frac{\partial u}{\partial n} = 0|_{\text{on boundary}} \quad (2.12)$$

Solving the BE scheme by hand for a very small mesh consisting of 4 mesh points will reveal the structure of the BE scheme.

$$\begin{aligned} u_0^{k+1} &= \frac{\Delta t}{2\Delta x^2} (2(D_0 + D_1)(u_1^{k+1} - u_0^{k+1})) + u_0^k \\ u_1^{k+1} &= \frac{\Delta t}{2\Delta x^2} ((D_2 + D_1)(u_2^{k+1} - u_1^{k+1}) - (D_1 + D_0)(u_1^{k+1} - u_0^{k+1})) + u_1^k \\ u_2^{k+1} &= \frac{\Delta t}{2\Delta x^2} ((D_3 + D_2)(u_3^{k+1} - u_2^{k+1}) - (D_2 + D_1)(u_2^{k+1} - u_1^{k+1})) + u_2^k \\ u_3^{k+1} &= \frac{\Delta t}{2\Delta x^2} (2(D_2 + D_3)(u_3^{k+1} - u_2^{k+1})) + u_3^k \end{aligned}$$

Rearranging this and setting $\alpha = \frac{\Delta t}{2\Delta x^2}$ results in a normal system of linear equations:

$$\begin{aligned} u_0^{k+1} (1 + 2\alpha(D_0 + D_1)) - 2\alpha u_1^{k+1}(D_1 + D_0) &= u_0^k \\ u_1^{k+1} (1 + \alpha(D_2 + 2D_1 + D_0)) - \alpha u_2^{k+1}(D_2 + D_1) - \alpha u_0^{k+1}(D_1 + D_0) &= u_1^k \\ u_2^{k+1} (1 + \alpha(D_3 + 2D_2 + D_1)) - \alpha u_3^{k+1}(D_3 + D_2) - \alpha u_1^{k+1}(D_2 + D_1) &= u_2^k \\ u_3^{k+1} (1 + 2\alpha(D_3 + D_2)) - 2\alpha u_2^{k+1}(D_3 + D_2) &= u_3^k \end{aligned}$$

which is arranged as a coefficient matrix multiplied by the unknown next time step

$$\begin{pmatrix} 1 + 2\alpha(D_0 + D_1) & -2\alpha(D_1 + D_0) & 0 & 0 \\ -\alpha(D_1 + D_0) & 1 + \alpha(D_2 + 2D_1 + D_0) & -\alpha(D_2 + D_1) & 0 \\ 0 & -\alpha(D_2 + D_1) & 1 + \alpha(D_3 + 2D_2 + D_1) & -\alpha(D_3 + D_2) \\ 0 & 0 & -2\alpha(D_3 + D_2) & 1 + 2\alpha(D_3 + D_2) \end{pmatrix} \mathbf{u}^{k+1} = \mathbf{u}^k \quad (2.13)$$

Denote the coefficient matrix in (2.13) by \mathbf{M}

$$\mathbf{M}\mathbf{u}^{k+1} = \mathbf{u}^k. \quad (2.14)$$

If a linear system like the one in (2.14) has a solution it can always be found by Gaussian elimination. However, for a sparse linear system a lot of the calculations in a Gaussian elimination will be done with zeros. A dramatically more efficient way to solve the linear system in (2.14) is by exploiting the fact that it is tridiagonal. In order to be more consistent with the algorithm to be used, (2.14) will be rewritten as

$$\mathbf{M}\mathbf{u} = \mathbf{u}_p, \quad (2.15)$$

where $\mathbf{u}_p = \mathbf{u}^k$ denotes the solution at the previous time step, and $\mathbf{u} = \mathbf{u}^{k+1}$ denotes the solution at the current time step. Tridiagonal systems can be solved extremely efficiently by the “tridiag” algorithm listed below.

```
void tridiag(double *u, double *up, int n, double *a,
double *b, double *c){
double *H = new double[n];
double *g = new double[n];
for(int i=0;i<n;i++){
H[i] = 0;
g[i] = 0;
}
g[0] = up[0]/b[0];
H[0] = c[0]/b[0];
for(int i=1; i<n; i++){
//forward substitution
```

```

    H[i] = -c[i]/(b[i] + a[i]*H[i-1]);
    g[i] = (up[i] - a[i]*g[i-1])/(b[i] + a[i]*H[i-1]);
  }
  u[n-1] = g[n-1];
  for(int i=(n-2); i>=0; i--){
    //Backward substitution
    u[i] = g[i] - H[i]*u[i+1];
  }
}

```

Algorithm 2.1: The tridiag algorithm

Algorithm 2.1 uses three vectors a , b and c to describe the coefficient matrix, \mathbf{M} . Vector b contains the diagonal elements, a contains the sub-diagonal elements, and c contains the elements above the diagonal.

In two spatial dimensions the BE discretization gives the scheme shown below,

$$u_{i,j}^k = - \underbrace{\frac{D\Delta t}{\Delta x^2}}_{\alpha} (u_{i+1,j}^{k+1} + u_{i-1,j}^{k+1}) + \underbrace{\left(1 + \frac{2D\Delta t}{\Delta x^2} + \frac{2D\Delta t}{\Delta y^2}\right)}_{\gamma} u_{i,j}^{k+1} - \underbrace{\frac{2D\Delta t}{\Delta y^2}}_{\beta} (u_{i,j+1}^{k+1} + u_{i,j-1}^{k+1}). \quad (2.16)$$

If (2.16) is solved by hand on a 3×3 mesh by the same steps as in 1D it gives the following linear system

$$\left(\begin{array}{ccc|ccc|ccc} \gamma & -2\beta & 0 & -2\alpha & 0 & 0 & 0 & 0 & 0 \\ -\beta & \gamma & -\beta & 0 & -2\alpha & 0 & 0 & 0 & 0 \\ 0 & -2\beta & \gamma & 0 & 0 & -2\alpha & 0 & 0 & 0 \\ \hline -\alpha & 0 & 0 & \gamma & -2\beta & 0 & -\alpha & 0 & 0 \\ 0 & -\alpha & 0 & -\beta & \gamma & -\beta & 0 & -\alpha & 0 \\ 0 & 0 & -\alpha & 0 & -2\beta & \gamma & 0 & 0 & -\alpha \\ \hline 0 & 0 & 0 & -2\alpha & 0 & 0 & \gamma & -2\beta & 0 \\ 0 & 0 & 0 & 0 & -2\alpha & 0 & -\beta & \gamma & -\beta \\ 0 & 0 & 0 & 0 & 0 & -2\alpha & 0 & -2\beta & \gamma \end{array} \right) \mathbf{u} = \mathbf{u}_p \quad (2.17)$$

The dashed lines confine a total of nine 3×3 matrices which can be used to reformulate the $2n + 1 = 7$ band diagonal system as a block tridiagonal system as shown below.

$$\left(\begin{array}{c|c|c} B_0 & C_0 & 0 \\ \hline A_1 & B_1 & C_1 \\ \hline 0 & A_2 & B_2 \end{array} \right) \mathbf{u} = \mathbf{u}_p \quad (2.18)$$

All entries in eq. (2.18) are 3×3 matrices. In the general case where the diffusion equation is discretized and solved on a mesh of size $n \times n$ the resulting block tridiagonal system will consist of n^2 matrix entries which all are $n \times n$ matrices.

One should note that the $n \times n$ matrix \mathbf{u} which solves the 2D diffusion equation has be rewritten as a vector \mathbf{u} of size n^2 .

In order to solve this block tridiagonal linear system Algorithm 2.1 must be generalized to support matrices as entries in \mathbf{M} . Luckily this generalization is trivial and all that is needed is to replace divisions by matrix inverses [13]. An updated *pseudo coded* scheme is listed in (2.19)

There is a forward substitution

$$\begin{aligned} H_0 &= -B_0^{-1}C_0 \\ \mathbf{g}_0 &= B_0^{-1}\mathbf{u}_{p0} \\ H_i &= -(B_i + A_iH_{i-1})^{-1}C_i \\ \mathbf{g}_i &= (B_i + A_iH_{i-1})^{-1}(\mathbf{u}_{pi} - A_i\mathbf{g}_{i-1}) \end{aligned} \quad (2.19)$$

Followed by a backward substitution

$$\begin{aligned} \mathbf{u}_{n-1} &= \mathbf{g}_{n-1} \\ \mathbf{u}_i &= \mathbf{g}_i + H_i\mathbf{u}_{i+1} \end{aligned}$$

It is possible to solve the BE scheme in 3D by the block tridiagonal solver as well. The linear system will be $2n^2 + 1$ band diagonal and must first be reduced to a block matrix which is $2n + 1$ band diagonal like the one in eq. (2.17) before it is further reduced to a block tridiagonal form. Matrix entries will be $n^2 \times n^2$ matrices.

Efficiency of the block tridiagonal algorithm

In general the FE scheme will solve a discrete PDE in d spatial dimensions with n mesh points in each dimension using

$$\mathcal{O}(n^d)$$

floating point operations (FLOP). This is the maximum efficiency possible for a PDE in d dimensions (without using some form of symmetry argument) and will be the benchmark for the BE scheme as well.

Implicit schemes like BE result in linear systems which, as mentioned, can always be solved by Gaussian elimination by some $\mathcal{O}(n^3)$ FLOPs for a

$n \times n$ matrix. In d spatial dimensions the matrices are $n^d \times n^d$, meaning that Gaussian elimination will require $\mathcal{O}(n^{3d})$ FLOPs which is extremely inefficient.

The block Tridiagonal solver requires inverting n matrices of size $n^{d-1} \times n^{d-1}$ as well as some matrix-matrix multiplications. All of these steps require $\mathcal{O}(n^{3(d-1)})$ FLOPs, but as long as the coefficient matrix is constant they need only be done once. In other words, the H_i terms from (2.19) can be stored. The remaining calculations include three matrix vector multiplications, one of which demands n FLOPs because the matrix is diagonal. The other two matrices are dense, so (for now) the remaining two matrix - vector multiplications require n^2 FLOPs. In summary, the block tridiagonal solver requires $\mathcal{O}(n^{2n-1})$ FLOPs, which is an order more efficient than a general LU decomposition, and in fact any other direct solver the author knows of. In 2D the block tridiagonal solver is one order slower than the FE scheme, which means it is still very much usable for small meshes.

Another point which should be mentioned is that the block tridiagonal solver is extremely memory efficient compared to other linear system solvers. Since the algorithm only loops through the non zero entries, there is no need to assemble the entire matrix. For a 2D simulation on a 100×100 mesh, the full assembled matrix will have $100^2 \times 100^2 = 10^8$ matrix entries, all of which are double precision floats. The total size of the matrix in RAM will be $8 \cdot 10^8$ bytes, which is getting close to the available RAM of a normal computer at $\sim 8 \cdot 10^9$ bytes. In comparison, storing only the non zero entries requires three block vectors of size 100 with entries that are 100×100 matrices. A total of $24 \cdot 10^6$ bytes, or a factor 100 less. Effectively, the memory impact is also reduced from $8 \cdot n^{2d}$ bytes to $8 \cdot n^{2d-1}$ bytes by switching to the block tridiagonal solver.

2.3 Combining micro and macro scale models

2.3.1 The algorithm

After setting an initial condition and diffusion constant the diffusion problem is solved on both the microscopic and macroscopic meshes and combined into a common solution by the following steps.

- The result from previous PDE time step, \mathbf{u}_p , is converted to a distribution of random walkers and sent to the RW solver.
- The RW solver does a predefined number of micro scale time steps

which correspond to one PDE time step.

- The result from the RW solver is converted back to a concentration and this replaces the PDE solution, \mathbf{u}_p .
- \mathbf{u}_p is then used as input to calculate the next time step.

2.3.2 Conversion between length scales

As the previous section states, the result from the last PDE time step is converted to a distribution of random walkers. This is done by specifying a conversion rate denoted Hc (as was done by Plapp and Karma [11]) which is a single, real integer defined by

$$C_{ij} = Hc \cdot u_{ij} \tag{2.20}$$

As before u_{ij} is the solution to the 2D diffusion equation evaluated at $x_0 + i\Delta x$ and $y_0 + j\Delta y$, and C_{ij} is the number of random walkers located in the rectangle defined by $x_i \pm \frac{\Delta x}{2}$ and $y_j \pm \frac{\Delta y}{2}$. Walkers are given a random position within this rectangle at the beginning of each PDE time step.

The reason to not simply use the positions from last time step is that the concentration in the RW area could change from one time step to the next. This will result in a different distribution of walkers which must be calculated at each time step. This is illustrated in Figure 2.3.




Figure 2.3: This figure shows an exaggerated case of how the concentration in an area changes over the course of a simulation. The distribution of walkers is directly proportional to the integral of $u(x, t)$, illustrated by a colored area. There is clearly a larger concentration in the area after 1000 time steps than at the beginning of the simulation, meaning that the number of walkers should be recalculated at each time step. This figure was made using the FE scheme with $\Delta x = \frac{1}{300}$, $D = 1$ and $\Delta t \approx 5.56 \cdot 10^{-6}$.

Equation (2.20) effectively states that one "unit" of the solution u will directly correspond to Hc random walkers. Depending on the application and the units of the PDE solution, $Hc \approx 20 - 50$.

2.3.3 Coupling the models through step length

To ensure that the τ time steps performed by the RW solver sums up to one time step for the PDE solver a limitation must be imposed on the RW solver. This limitation can only be imposed on the step length of the random walkers, so the task at hand is to relate the step length to the PDE time step, Δt , and τ .

Through an Einstein relation the variance in position is coupled to the diffusion constant.

$$\langle \tilde{\Delta x}^2 \rangle = 2Dd\tilde{\Delta t} \quad (2.21)$$

Where $\tilde{\Delta}t$ is the time step for the random walkers which is exactly $\tilde{\Delta}t = \frac{\Delta t}{\tau}$. Similarly $\tilde{\Delta}x$ is the spatial resolution on the micro scale which is also the step length, l , for the random walkers. Rewriting equation (2.21) gives the final expression for the step length

$$l = \sqrt{2dD\frac{\Delta t}{\tau}} \quad (2.22)$$

2.3.4 Boundary conditions for the random walk

Because the RW area must correspond to an area on the PDE mesh which has a finite, constant size the RW model must have boundary conditions imposed on it. In order to make absolutely sure that no walkers disappear and by extension that the amount of concentration is conserved, perfectly reflecting boundaries are chosen for the RW model. Reflecting boundaries correspond to zero flux Neumann boundaries

$$\frac{\partial C}{\partial n} = 0 \quad (2.23)$$

Neumann boundaries might seem unphysical, but the redistribution of walkers at each PDE time step must also to some extent be considered as a boundary condition in which fluxes are exchanged between the two models. In hindsight, a possibly better choice of boundary conditions for the RW model would have been perfect exchange of flux between the two solvers.

$$\frac{\partial u}{\partial n} = \frac{\partial C}{\partial n} \quad (2.24)$$

Since time does not allow for this to be implemented and tested it will be left as a reference to possible future work.

2.4 Potential problems

This section will list and discuss some of the problems which were encountered while working on the thesis and present the solutions or workarounds.

Negative concentration of walkers

Physically a negative concentration does not make sense, but if an initial condition with negative values is imposed on the system, the software will try to allocate a negative number of walkers. Trying to handle this is more of an oddity than anything else, but a workaround has been found. If the absolute

value of the negative concentration is taken as input to the RW solver and the sign at each mesh point is stored while the RW solver advances the system, the resulting solution can be multiplied by the stored sign to give back a negative concentration. The workaround will at least keep the simulation from crashing, but has a fundamental problem.

Figure 2.4a shows how a negative concentration is converted to a positive concentration by taking the absolute value. This results in an abrupt change where the original function, $u(x, t)$, crossed the x-axis. Seeing as a diffusion process very efficiently evens out abrupt changes, the point where $u(x, t)$ crossed the x-axis will get an increased value. As Figure 2.4b shows, this point should remain equal to zero, but does not.

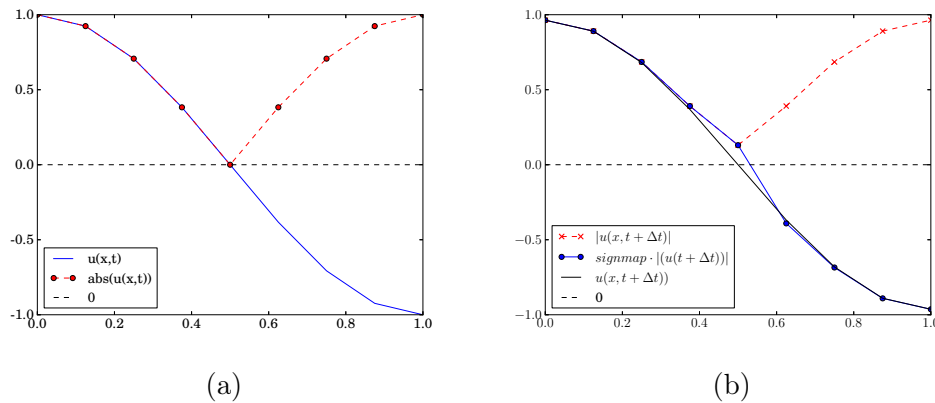


Figure 2.4: An illustration of the proposed workaround for negative concentrations and an illustration of how it performs (b). The fundamental problem is that a diffusion process will even out discontinuities such as the one in (a). The result of this evening out is shown in (b) as an increase in concentration in the middle mesh point which should remain equal to zero.

Smooth solutions

A diffusion process is very effective when it comes to dampening fast fluctuations, and so any solution of the diffusion equation will be smooth. Random walks are stochastic and though a RW will describe a diffusion process it will fluctuate around the solution. In this case a dilemma arises; on the one hand there is the smoothness of the solution to consider, on the other hand the stochastic term was introduced believing that it adds detail to the model. This is partly the reason for solving the RW model before the PDE model. The fluctuations will be reduced by the PDE solver and the solution, though still fluctuating, will be smoother.

Section 2.5 describes some other methods which were attempted in order to reduce the fluctuations from the RW model.

Number of time steps on the random walk level

The reason for introducing a random walk model in the diffusion solver was to include an area with more advanced dynamics than regular diffusion on a smaller time-, and length scale. This means that the RW solver must take τ time steps for each PDE time step. τ could in principle be any integer, but should be chosen large enough so the central limit theorem is satisfied. This usually means $\tau \geq 50$.

Alternatively τ can be calculated from the assumption that the error term with respect to time arising from the RW model is proportional to the square root of the time step,

$$\epsilon \propto \sqrt{\tilde{\Delta}t}. \quad (2.25)$$

This assumption comes from the fact that the step length is proportional to the square root of the time step. In order to make the error of the same size as the error from the PDE model we use

$$\begin{aligned} \tilde{\Delta}t &= \frac{\Delta t}{\tau} \\ \implies \epsilon &\propto \sqrt{\frac{\Delta t}{\tau}} \\ \mathcal{O}(\Delta t) &\geq \sqrt{\frac{\Delta t}{\tau}} \\ \implies \tau &\geq \frac{1}{\Delta t} \end{aligned}$$

Note that the error associated with the number of walkers introduced is of much higher significance than the error from the time step.

2.5 Other possible coupling methods

Running the solution from the PDE model through a RW model gives two solutions for a part of the mesh. As it turns out it is sufficient to simply replace the PDE solution with the RW solution in this area, which is nice seeing as it is the simplest possible method. There are, however, some other ways to do this which have been tested, and found inadequate or simply been disregarded. Three of these are mentioned below.

Averaging the two solutions

Taking the arithmetic mean of the two solutions will reduce the magnitude of the fluctuations, as shown in Figure 2.5. This method is simple and reduces fluctuations by approximately one half (approximately because fluctuations will be present after the first time step).




Figure 2.5: This plot shows the PDE solution before and after the stochastic solver has been called, and the arithmetic mean of the two solutions which has fluctuations of smaller amplitude.

The reason an average was scrapped is simply that equally satisfactory results were achieved by only replacing the solution.

Polynomial regression

In order to smoothen out the solution from the RW model a polynomial regression function was tested. The principle was to look at the solution from the RW model as a set of data points which have some dependency on their position on the PDE mesh. Figure 2.6 shows how this method worked for a slightly exaggerated case. This approach turns out to have two problems; first and foremost the amount of concentration is no longer conserved, secondly the endpoints are usually way off.

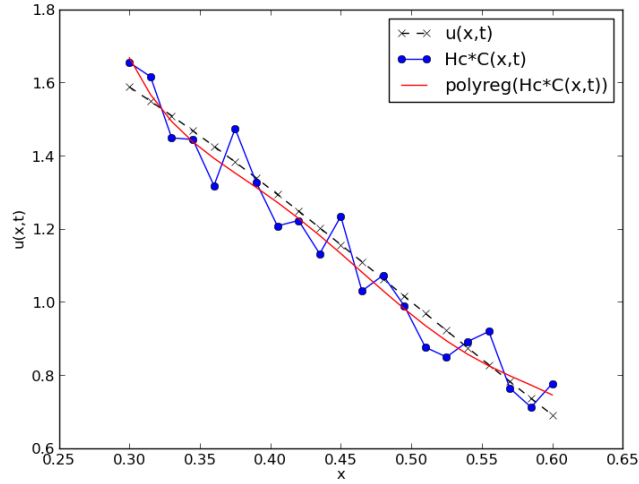


Figure 2.6: This plot shows the PDE solution before and after the stochastic solver has been called, and a polynomial regression of degree 6 using $C(x, t)$ as data points.

Cubic spline interpolation

Doing an interpolation will force the combined solution to be equal to the RW solution at some mesh points. In practice this means choosing which points are correct, and which are not. The author has not found any reasonable way to do this.

In addition, cubic spline interpolation has the same problem as polynomial regression when it comes to conserving energy.

Chapter 3

Analysis

3.1 Computation of the error

Solving PDEs numerically will result in errors because derivatives are approximated by finite differences and the PDE is required to be fulfilled only on the mesh points.

The error, ϵ , is measured by comparing the result from a numerical simulation to an exact solution, u_e and taking the norm of the difference. Specifically ϵ is measured by the L2 norm which is defined in (3.1):

$$\begin{aligned}\epsilon(t^k) &= \|u(t^k) - u_e(t^k)\|_2 \\ &= \iint \sqrt{(u(t^k, x, y) - u_e(t^k, x, y))^2} dx dy \\ &\approx \sqrt{\Delta x \Delta y \sum_{i=0}^k \sum_{j=0}^k (u(t^k, x_i, y_j) - u_e(t^k, x_i, y_j))^2}.\end{aligned}\tag{3.1}$$

A time dependent error estimate has been chosen because it allows for investigation of the evolution of the error over the course of a simulation. Some of the error tests will require a single number as an error measure. In these cases the norm of $\epsilon(t)$, shown below, is used:

$$\epsilon = \sqrt{\Delta t \sum_{k=0}^T \epsilon(t^k)^2}.\tag{3.2}$$

3.2 Verification techniques

All parts of the hybrid diffusion solver are subject to testing. First, the PDE and RW solver will be tested individually, then the hybrid solver is tested. This thesis will focus on three verification techniques which are described below. The aim for all of these techniques is to make sure that the implementation is correct by comparing simulations to an exact solution.

For the PDE solver, there are errors arising from both the spatial derivative and from the time derivative. Since an incorrect implementation of the spatial derivative will cause the solution to be unstable, which is easily noticed through visual inspection, the tests will focus on verifying the time derivative. In some cases, isolating the contribution to $\epsilon(t)$ from the time derivative will be necessary, and is ensured by setting $\Delta t \gg \Delta x^2$. A time step of this size violates the stability criterion for the FE scheme, and so some of the tests are slightly inaccurate for this scheme.

The verification techniques that will be carried out are:

- Manufactured solutions

By adjusting the source term, $s(x, t)$, one can, in principle, choose any solution to the diffusion equation. Should we, for example, want the solution to be $u(x, t) = \frac{1}{x+1}$, all that is needed is to set the source term equal to $-D \frac{\partial^2 u}{\partial x^2}$;

$$s(x, t) = \frac{2D}{(x+1)^3},$$

and the problem is solved (as long as the boundary conditions are correct). For the verification that follows, the chosen solution is

$$u(x, y, t) = e^{-\pi^2 t} \cos(\pi x) \cos(\pi y) + 1, \quad (3.3)$$

which requires no source term and fulfills the zero flux boundary conditions naturally. The point of these tests is to verify that $\epsilon \sim \Delta t$ and the tests can be done for a time step which fulfills the stability criterion.

- Convergence tests

For the schemes used in this thesis the error term is described by

$$\epsilon = C_x \Delta x^2 + C_t \Delta t^1, \quad (3.4)$$

where the coefficients C_x and C_t are unknown. Notice that there is one term arising from the time derivative and one from the spatial derivative and that they are of different order.

Where possible, convergence will be tested by isolating one of the error terms. For example, setting $\Delta t \gg \Delta x^2$, implies $C_x \Delta x^2 \approx 0$ and

$$\epsilon \propto \Delta t^r,$$

where $r = 1$. As Δt tends to zero, r is expected to converge to one, and that is the purpose of these tests. By comparing the error from two simulations with different time steps, the convergence rate, r , can be measured from

$$r \simeq \frac{\log(\epsilon_1/\epsilon_2)}{\log(\Delta t_1/\Delta t_2)}.$$

Notice that a single number must be used to describe the error, as shown in (3.2).

- Exact numerical solutions

The numerical schemes are actually reformulations of the PDE as difference equations, which have their own exact solutions. These will be called the numerical exact solutions, and they are slightly different from the exact solutions to the PDE. The reason for finding the numerical exact solutions is that the scheme theoretically will represent this solution with no error. In practice there will always be round off errors and other factors, but an error term close to machine precision is expected.

3.3 Testing the PDE solver

3.3.1 Verification by manufactured solutions

Equation (3.3) solves the diffusion equation, so using $u(x, y, t = 0)$ as the initial condition for a simulation will give us both the numerical solution and the exact solution. An important property of both the FE and BE discretization schemes is that the difference between these solutions is of the same order as the time step.

$$\epsilon(t) \sim \mathcal{O}(\Delta t)$$

Figures 3.1 and 3.2 show that for both the FE and BE scheme in both 1D and 2D the error is of the expected magnitude, which is roughly Δt . Another interesting property of the error plots is that the error tends to zero after a large number of time steps. By inserting the limit $t \rightarrow \infty$ in (3.3) we observe that the error is expected to tend to zero because the limit value of one can be exactly recreated by both schemes.

$$\lim_{t \rightarrow \infty} e^{-\pi^2 t} \cos(\pi x) \cos(\pi y) + 1 = 1$$




Figure 3.1: Error plot for the FE scheme in 1D (a) and 2D (b). All simulations have been run using $\Delta t = \frac{\Delta x^2}{5}$ which fulfills the stability criterion in 1D and 2D.




Figure 3.2: Error plot for the BE scheme in 1D (a) and 2D (b). Notice that the error is roughly of the same order as the time step. Figures 3.1b and 3.2b show the same simulation using FE and BE schemes in 2D. The error is slightly larger in the BE discretization by a factor of 4.

3.3.2 Verification by convergence tests

The convergence tests are done by isolating the error term from either the time derivative or the spatial derivative, and refining the relevant discretization parameter over several simulations. However, for the FE scheme it is impossible to isolate the error from the time derivative without violating the

stability criterion. To cope with this problem a single discretization parameter, h , is introduced by ensuring

$$\frac{\Delta t}{\Delta x^2} = k,$$

where k is a constant. The error term from (3.4) can now be rewritten by setting $\Delta t = h$;

$$\begin{aligned}\epsilon(t) &= C_t h + C_x \frac{h}{k} \\ &= (C_t + C_x/k)h \\ &= C_k h^r,\end{aligned}$$

where r is still expected to converge to one.

Figures 3.3 and 3.4 verify that the error associated with the time derivative is of the expected order.

In Figure 3.5 the error from the spatial derivative has been isolated, and tested for both schemes. The error term associated with the spatial derivative is of second order, as seen in (3.4). A second order error term is expected to have second order convergence.




Figure 3.3: Convergence tests with respect to the single discretization parameter, h , for the FE scheme in 1D (a) and 2D (b).




Figure 3.4: Convergence tests with respect to the time derivative for the BE scheme in 1D (a) and 2D (b).

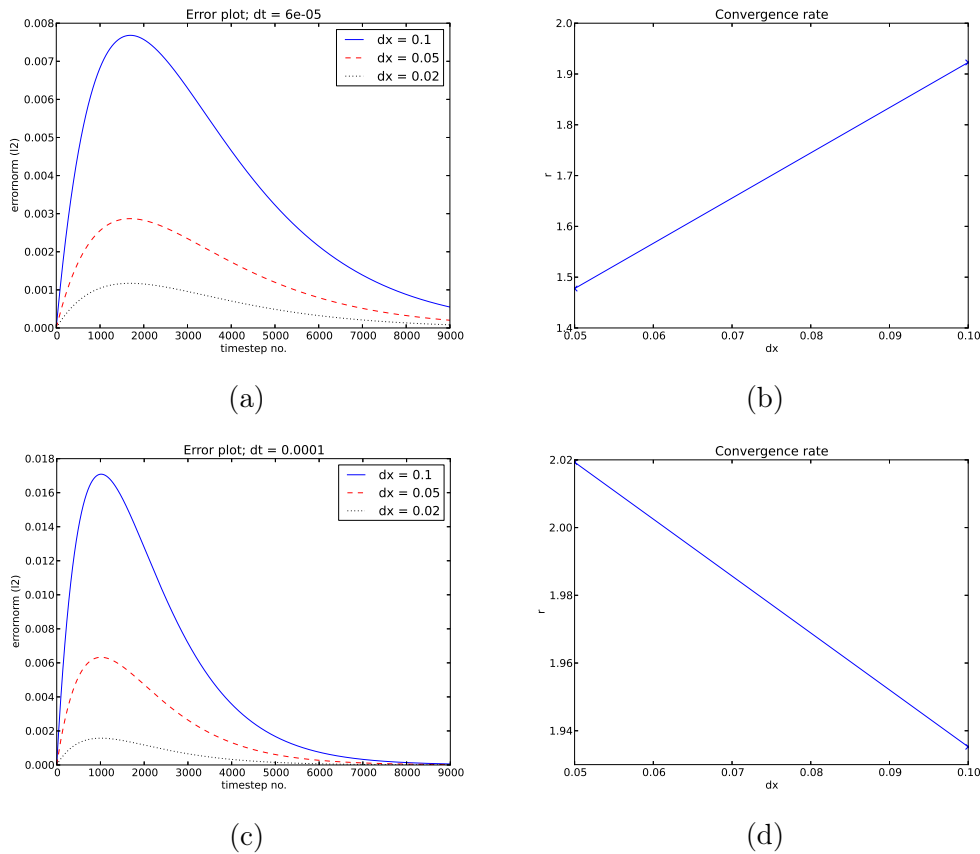


Figure 3.5: These plots show the error and convergence from the spatial derivative using BE (a & b) and FE (c & d) discretization schemes. The expected convergence rate for both discretizations is $r = 2$, which is recreated to a decent accuracy. Testing has been done using 1D simulations for the BE scheme, and 2D simulations for the FE scheme. This should not make any difference for the results.

3.3.3 Verification of FE scheme by exact numerical solution

Discretization of the diffusion equation by the FE scheme yields the following numerical scheme in 1D:

$$u^{k+1} = D\Delta t u_{2x}^k + u^k, \quad (3.5)$$

where u_{2x} denotes the double derivative of u with respect to x . Equation (3.5) is a difference equation, which is most easily solved by writing out the first four iterations and finding a pattern

$$\begin{aligned} u^1 &= D\Delta t u_{2x}^0 + u^0 \\ u^2 &= D\Delta t u_{2x}^1 + u^1 \\ &= D\Delta t [D\Delta t u_{4x}^0 + u_{2x}^0] + D\Delta t u_{2x}^0 + u^0 \\ &= (D\Delta t)^2 u_{4x}^0 + 2D\Delta t u_{2x}^0 + u^0 \\ u^3 &= D\Delta t u_{2x}^2 + u^2 \\ &= D\Delta t [(D\Delta t)^2 u_{6x}^0 + 2D\Delta t u_{4x}^0 + u_{2x}^0] + (D\Delta t)^2 u_{4x}^0 + 2D\Delta t u_{2x}^0 + u^0 \\ &= (D\Delta t)^3 u_{6x}^0 + 3(D\Delta t)^2 u_{4x}^0 + 3D\Delta t u_{2x}^0 + u^0 \\ u^4 &= D\Delta t u_{2x}^3 + u^3 = \dots \\ &= (D\Delta t)^4 u_{8x}^0 + 4(D\Delta t)^3 u_{6x}^0 + 6(D\Delta t)^2 u_{4x}^0 + 4D\Delta t u_{2x}^0 + u^0 \end{aligned}$$

From the iterations above a pattern can be recognized for iteration $k + 1$:

$$u^{k+1} = \sum_{i=0}^k \binom{k}{i} (D\Delta t)^i u_{2ix}^0, \quad (3.6)$$

where u^0 is the initial condition;

$$u^0 = \cos(\pi x). \quad (3.7)$$

The spatial derivatives are found by inserting the initial condition into the centered difference approximation to the second derivative :

$$\begin{aligned} u_{2x}^0 &= \frac{1}{\Delta x^2} (\cos(\pi(x + \Delta x)) - 2\cos(\pi x) + \cos(\pi(x - \Delta x))) \\ &= \frac{2}{\Delta x^2} (\cos(\pi\Delta x) - 1) \cos(\pi x) \\ u_{4x}^0 &= [u_{2x}^0]_{2x} \frac{1}{\Delta x^2} \left[\frac{u_{2x}^0}{\cos(\pi x)} (\cos(\pi(x + \Delta x)) - 2\cos(\pi x) + \cos(\pi(x - \Delta x))) \right] \\ &= \frac{4}{\Delta x^4} (\cos(\pi\Delta x) - 1)^2 \cos(\pi x) \\ &\dots \end{aligned}$$

The pattern continues allowing the final numerical exact solution to be expressed below:

$$u^{k+1} = \sum_{i=0}^k \binom{k}{i} (D\Delta t)^i \frac{2^i}{\Delta x^{2i}} (\cos(\pi\Delta x) - 1)^i \cos(\pi x). \quad (3.8)$$

By construction the Neumann boundary conditions are fulfilled since $\frac{\partial \cos(\pi x)}{\partial x} = -\pi \sin(\pi x)$ and $\sin(0) = \sin(\pi) = 0$.

Although the FE scheme is expected to reproduce (3.8) to machine precision ($\epsilon \approx 10^{-16}$) there are two problems with the solution which will have an effect on the error:

- Δx^{2i} will quickly tend to zero, and the computer will interpret it as zero. This will cause division by zero, which again results in “Not a number” (nan) and ruins the simulation. Testing if $\Delta x^{2i} > 0$ and returning zero if the test fails will fix the problem. The argumentation for ignoring the troublesome terms is given below.
- $\binom{k}{i}$ goes towards infinity for large k and i . The computer can only represent numbers up to $\sim 10^{308}$, which limits the number of steps to 170 since $k! > 10^{308}$ for $k > 170$.

As a side note, (3.8) illustrates how the stability criterion for the FE scheme comes into place. In the numerical exact solution, the exponential which is found in the exact solution to the PDE (eq. 3.3) is replaced by an amplification factor A^k . This amplification factor can be found in equation (3.8) as

$$A^k = \left(\frac{2D\Delta t}{\Delta x^2} \right)^i. \quad (3.9)$$

Inserting a time step larger than $\frac{\Delta x^2}{2D}$ will make the amplification factor, A , larger than one, which in turn makes the solution blow up.

The stability criterion also illustrates why the terms where

$$\frac{1}{\Delta x^{2i}} \rightarrow \infty$$

can be dropped. By the stability criterion, the time step will cancel out Δx^2 , and the result will be a number smaller than 1 raised to a rather large power, resulting in a number comparable to zero.

The results from comparing a 1D simulation to the numerical exact is shown in Figure 3.6a. As expected the error is larger than machine precision by at most two orders of magnitude because of accumulating error terms from the dropped terms in (3.8).

Using the same method as in the 1D case, a numerical exact solution can be found to the 2D FE scheme;

$$u^{k+1} = \sum_{i=0}^k \binom{k}{i} (D\Delta t)^i \left[2^{i-1} \cos(\pi x) \cos(\pi y) \left(\frac{(\cos(\pi \Delta x))^i}{\Delta x^{2i}} + \frac{(\cos(\pi \Delta y))^i}{\Delta y^{2i}} \right) \right]. \quad (3.10)$$

The same problems as in the 1D case will apply to (3.10) with the same solutions. Figure 3.6b shows how the 2D simulation compares to the numerical exact solution. As was the case in 1D, the error is larger than machine precision, but much smaller than Δt .




Figure 3.6: Error plots for the FE scheme in 1D (a) and 2D (b) when comparing with the numerical exact solution. Some of the terms in the numerical exact solutions are ignored to prevent overflow and these are responsible for the increasing error which is slightly larger than machine precision.

3.3.4 Verification of BE scheme by exact numerical solution

The exact numerical solution to the BE scheme is found by solving the linear system which arises from the discretization (see Section 2.2). Given the k 'th time step, the next time step is found by

$$\begin{aligned}
 \mathbf{M}\mathbf{u}^{k+1} &= \mathbf{u}^n \\
 \mathbf{u}^{k+1} &= \mathbf{M}^{-1}\mathbf{u}^k \\
 &= \mathbf{M}^{-1}(\mathbf{M}^{-1}\mathbf{u}^{k-1})
 \end{aligned}$$

Doing the separation to the end relates the n 'th time step to the initial condition

$$\mathbf{u}^k = (\mathbf{M}^{-1})^k \mathbf{u}^0, \quad (3.11)$$

where $(\mathbf{M}^{-1})^k$ is the inverse of \mathbf{M} to the k 'th power.

Taking the inverse of \mathbf{M} will result in a dense matrix where a lot of the entries are close to zero (e.g. 10^{-20}). Doing calculations with such a matrix results in a lot of round-off errors which will reduce the accuracy of the numerical exact solution. The error should theoretically be machine precision, but is expected to at least, be much smaller than Δt . Errors from both 1D and 2D simulations are shown in Figure 3.7. Though computationally intensive, the numerical exact solution for the BE scheme has none of the limitations that were found for the FE scheme.

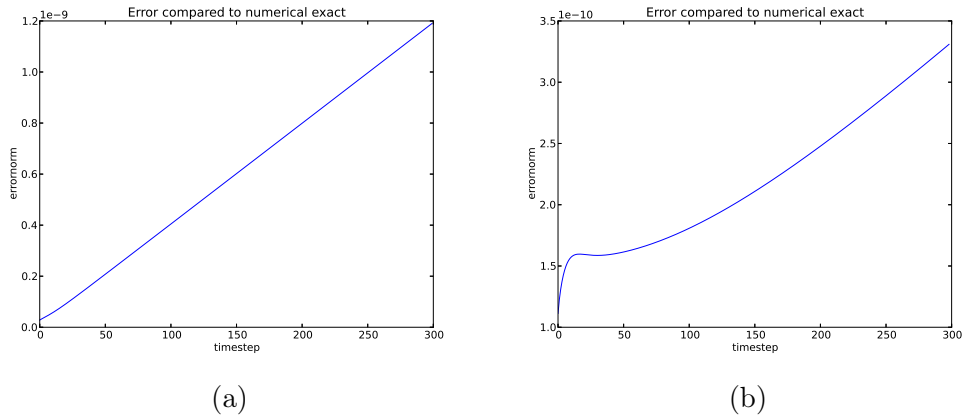


Figure 3.7: Plots showing the error for the BE scheme in 1D (a) and 2D (b) compared to the numerical exact solution. The error is not machine precision, but significantly smaller than Δt which for these simulations is $\Delta t = 0.01$. This increased error originates in the many round-off errors in the inverted matrix where a lot of terms are of the order 10^{-16} and smaller.

3.4 Verification of the RW solver

The aim of this section is to verify that the implemented RW model will solve the diffusion equation on the correct time scale, and to verify that the error term from the RW solver is dependent on the conversion factor, Hc . In order to carry out the tests, a new initial condition that can easily be recreated by random walkers will be introduced.

$$u(t = 0, x) = H\left(x - \frac{1}{2}\right) \quad (3.12)$$

Where $H(x)$ denotes the Heaviside step function [1], which is defined by its properties

$$H(x - a) = \begin{cases} 1 & x > a \\ \frac{1}{2} & x = a \\ 0 & x < a \end{cases} \quad (3.13)$$

The new initial condition gives a new exact solution which is found by separation of variables and Fourier series as

$$u(x, t) = \frac{1}{2} + 2 \sum_{n=\text{odd}}^{\infty} \frac{(-1)^{\frac{n-1}{2}}}{n\pi} e^{-(n\pi)^2 t} \cos(n\pi x) \quad (3.14)$$

Figure 3.9 verifies that the error term improves by adding more walkers. And Figure 3.8 shows that the RW model will fulfill the chosen time step on the PDE level. The convergence rate is not very good, but this is expected. Stochastic effects causes the error to fluctuate around a certain value rather than converge to zero, making it very difficult to get a good error measure.




Figure 3.8: Error plot (a) and convergence test (b) for 1D RW solver using a Heaviside step function as initial condition. In these tests both the time step and the conversion factor are changed for each simulation, and the conversion factor follows the previously proposed limit $Hc \geq \frac{1}{\Delta t^2}$. For each Δt the RW simulation does 250 steps with a step length calculated from (2.22). The expected convergence rate is 0.5, and to some extent this is achieved here. Note, however, that due to fluctuations in the solution, getting a good error measure is difficult and beyond our control.

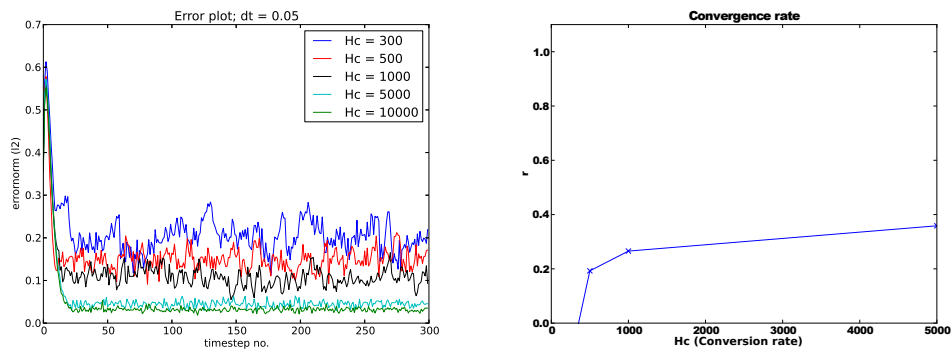


Figure 3.9: As a comparison to Figure 3.8, this test has been done for the same 1D Heaviside step function as initial condition, but keeps the time step fixed at $\Delta t = 0.05$ and increases the conversion factor. The convergence rate (fig. b) is worse than for the time convergence test, and seems to reach a limit where increasing the number of walkers has little effect on the error.

3.5 Verification of the hybrid solver

This section aims to achieve first order convergence in time for the hybrid solver by introducing a sufficient number of walkers. Effects of varying the number of mesh points affected by the RW solver will also be illustrated.

The hybrid diffusion solver is a combination of the PDE and RW solvers, which means that the error term for the hybrid model is described by

$$\epsilon(t) = C_t \Delta t + C_x \Delta x^2 + \frac{C_{RW}}{\sqrt{Hc}}, \quad (3.15)$$

where C_{RW} is an unknown stochastic coefficient.

From (3.15) it is clear that the error term from the RW solver is dominant, as it is of order $r = 0.5$. In order to achieve first order convergence, a limitation must be placed on the conversion rate. Assuming that the spatial error term is negligible:

$$\epsilon(t) = C_t \Delta t + \underbrace{C_x \Delta x^2}_{\approx 0} + \frac{C_{RW}}{\sqrt{Hc}} \quad (3.16)$$

$$\mathcal{O}(\Delta t) \approx C_t \Delta t + \frac{C_{RW}}{\sqrt{Hc}} \quad (3.17)$$

$$\implies \frac{1}{\sqrt{Hc}} \leq \Delta t. \quad (3.18)$$

Which finally gives the restriction on the conversion rate

$$Hc \geq \frac{1}{\Delta t^2}. \quad (3.19)$$

By (3.19) the time step should be as large as possible in order to limit the number of walkers and achieve a reasonable computational performance. The FE scheme has a strict limitation on the time step from the stability criterion, making the BE scheme desirable. For this reason all tests will be done with the BE scheme.

Figure 3.10 shows first order convergence for the hybrid model.


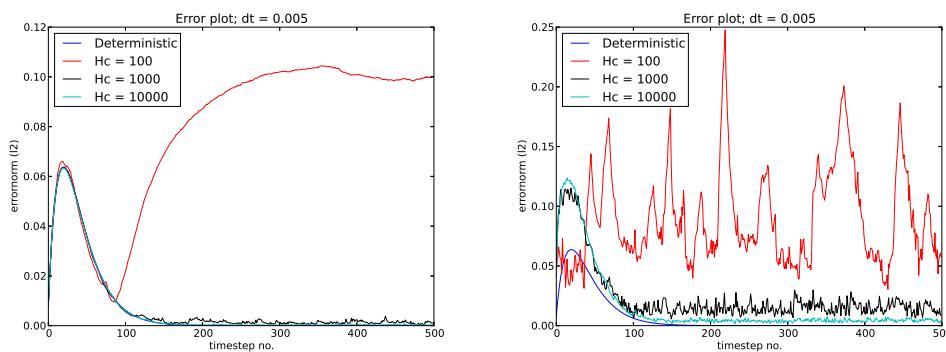


Figure 3.10: 3.10a shows the error plot for a test where Δx was fixed at $\Delta x = \frac{1}{100}$ and Δt was reduced from 0.05 to 0.01 and finally to 0.005. The conversion rate, Hc was updated for each simulation to have the value $Hc = \frac{1}{\Delta t^2}$, meaning the error from the walkers should be smaller than the error from the time derivative. Walkers are placed on 10% of the mesh from $x = 0.4$ to $x = 0.5$. 3.10b shows the convergence rate in time for the same test.

Decreasing or increasing the relative size of the microscopic model has serious effects on the error. Figure 3.11a shows that the error term is comparable to the error from a simulation without walkers when only 5% of the mesh points are affected by walkers. As a comparison, Figure 3.11b shows that the error will not be comparable to a simulation without walkers, even for a large conversion rate, when as much as 35% of the mesh points are affected by walkers.



(a) Walkers on 5% of the mesh points. (b) Walkers on 35% of the mesh points.

Figure 3.11: The effect of increasing the size of the walk area for a fixed $\Delta t = 0.05$ and $\Delta x = 0.01$ using the BE discretization.

Chapter 4

Physical application

The following chapter is included to show that the developed hybrid diffusion solver can be applied to real life problems with only minor modifications.

4.1 Physical scope

“To understand the central nervous system of even the simplest animals requires an understanding of the nervous system at many different levels, over a wide range of both spatial and temporal scales” - David Sterratt [7].

There are potentially many applications of a hybrid diffusion solver within the field of neuroscience, which is the study of the central nervous system. This chapter will look at a small part of a neuron located in the brain of a rat.




Figure 4.1: A pyramidal neuron with an apical dendrite (starting by the arrow). The inset shows an amplification of the apical dendrite and illustrates the size of dendritic spines (small outgrowths). The arrow points to the cell body, called the Soma. Note that this image does not show how neurons are located within the brain. Image from www.neurolex.org.

Figure 4.1 shows a pyramidal neuron with an apical dendrite, and an inset with dendritic spines. Pyramidal neurons take their name from the triangular shape of the Soma, and an apical dendrite is a dendrite springing from the apex of the Soma. Spines are (often) the receiving end of a chemical synapse, which is a junction between two neurons that allows for communication between the two cells.

This application will look at diffusion of $\text{PKC}\gamma$ in a segment of an apical dendrite and into dendritic spines. $\text{PKC}\gamma$ is a protein found in neurons which is associated with memory storage and associative learning [12]. Upon activation it will diffuse out of the Soma, through a dendrite, and into a dendritic spine to reinforce or reduce the absorption of neurotransmitters.

Effectively we will investigate the diffusion time for random walkers through spine necks which are very narrow ($\leq 0.5\mu\text{m}$). Spine necks are thought to act as diffusion barriers which hinder, but do not stop diffusion [4]. The results will be compared to results by Craske et.al [4], who in 2005 studied a similar problem in rodent pyramidal neurons harvested from the hippocampus.

4.2 Implementation

There is a difference between the approach of the developed hybrid solver and the dendrite - spine system with respect to geometry which is best summarized in Figure 4.2.

The dendrite segment is approximated as a cylinder with diameter $\sim 10\mu\text{m}$ and length $\sim 50\mu\text{m}$. Furthermore, little of interest is assumed to happen in the radial direction. All things considered, diffusion in the dendrite will be modeled as a 1D process.

Spines have a wide variety of geometries, with various properties [2]. For this application, however, only the neck length and width of a spine is of interest. The spines will therefore be modeled as two dimensional objects with rectangular necks and trapezoidal heads.

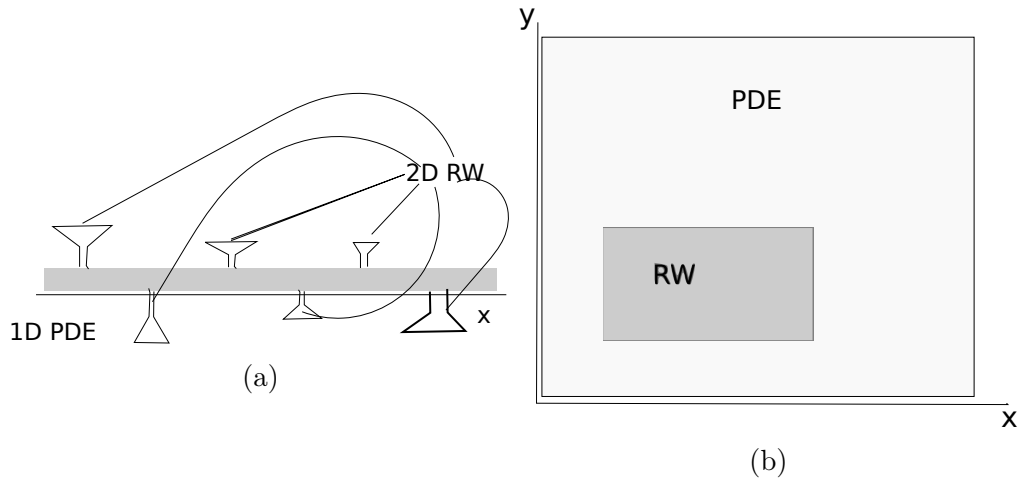


Figure 4.2: The geometric difference between the original hybrid diffusion solver (b) and the numerical setup to model $\text{PKC}\gamma$ diffusing into dendritic spines (a). Spines are attached by a few mesh points on the PDE mesh, which also determines the neck width of the spine.

The hybrid diffusion solver has been slightly modified in order to recreate the new geometry. Mathematically the largest difference is that the random walkers are left alone, meaning that the diffusion equation is not solved in the spines. There are also some differences in the coupling; at each PDE time step there is a probability for a protein to diffuse into a spine if the concentration at the beginning of the spine neck is large enough. This probability decreases with the number of proteins which have already diffused into the spine. The width of a spine neck is determined by the number of PDE mesh points it connects to, which, for this purpose, is chosen at random. All other length parameters are determined randomly, but required to fit with measurements from Arellano et.al [2].

4.3 Parameters and details

Within the first three seconds after $\text{PKC}\gamma$ is released from Soma, a portion of the released amount will bind to the dendrite wall. This process is not governed by the diffusion equation, and seems to increase the speed with which $\text{PKC}\gamma$ diffuses into the spine [4]. For this application, these dynamics are not of interest. In order to still have a realistic model, the initial condition has been altered:

Assuming that the actual release of $\text{PKC}\gamma$ from Soma can be modeled as a delta pulse, the solution will be a Gaussian function. The modified initial

condition places a small amount of PKC γ over the entire dendrite, and adds some random fluctuations. This is illustrated in Figure 4.3.




Figure 4.3: Initial condition used in the application

The result of using this new initial condition is simply that approximately 3 seconds can be added to all the measured times.

As Figure 4.3 shows, the concentration of PKC γ has units $\frac{\text{nMol}}{\text{L}}$. Craske et.al. [4] measured increases in concentration of around $5\frac{\text{nMol}}{\text{L}}$ in spine heads, which corresponds to 1 – 2 proteins once the volume of a spine is taken into account. In fact, by calculating the volume of the dendrite segment we are simulating, the conversion factor has been estimated at $Hc = 5 - 25$.

The model should have a few properties in order to be a realistic representation. First of all, if a large number of proteins are located at the junction to a spine, it is more likely that a protein diffuses into the spine. This is ensured by making the probability for a protein to diffuse into a spine to be dependent on the amount of PKC γ present at the beginning of the spine neck.

In order for a spine to be two dimensional, it must cover at least two points of the PDE mesh. This limits the minimal spatial resolution to the minimal neck diameter of a spine, which is $0.22\mu\text{m}$ [2].

Real particles will “feel” a concentration gradient, but random walkers do not. As spine heads are filled with proteins, new proteins must therefore be hindered from entering a spine if it is full. This is achieved by reducing the probability for a particle to diffuse into a specific spine each time a particle is absorbed in the head of that spine.

Chapter 5

Results

This chapter will summarize and discuss the results from chapters 3 and 4.

5.1 Results of verification

The plots that are discussed in this section can be found in chapter 3.

5.1.1 The FE scheme

In general, the results from testing the FE implementation are very good. Especially the results from testing the exact numerical solution are informative. The numerical exact solution test is a verification of the implementation of the FE scheme. A successful test indicates that the scheme is correct and will solve the diffusion equation to the expected accuracy. Both the error plots from section 3.3.1 and the spatial convergence test are very close to perfect, further suggesting a correct implementation.

Considering all of the three tests are close to perfect, the FE scheme is deemed correctly implemented within the limits of the applied tests.

5.1.2 The BE scheme and the block tridiagonal solver

Two of the tests done on the BE scheme are slightly off target. These are the spatial convergence test and the numerical exact solution.

Though the numerical exact solution test is expected to give an error close to machine precision, the results are five orders of magnitude larger than machine precision. On the other hand, the error term is also eight orders of magnitude smaller than the time step and there are a lot of possible round-off errors in the inverted coefficient matrix which can explain the lack of accuracy.

The spatial convergence test demands that the error associated with the spatial derivative is much larger than the error from the time derivative. As the spatial resolution is increased, the spatial error becomes less dominant because the time step is held constant. Because the spatial convergence is still clearly larger than one, and starts out at two (which is the value it should converge to), this is considered to be adequate.

Effectively, the manufactured solution test and convergence tests verify the assembly of the coefficient matrix, since an incorrect assembly would produce a larger error term. The numerical exact solution test verifies the implementation and accuracy of the block tridiagonal solver, since an incorrectly implemented solver would give results comparable to the numerical exact

solution.

In summary, the tests done on the BE scheme verifies that it, and the block tridiagonal solver, are implemented correctly.

5.1.3 The RW solver

Due to fluctuations in the solution it is hard to get a good convergence rate for the RW solver. However, it is measured to be close to 0.5 both with respect to the time step and with respect to the number of walkers introduced. The result of verifying the RW solver is that it is deemed correctly implemented within the limits of the applied tests.

5.1.4 The hybrid diffusion solver

For this thesis it is considered very important that the hybrid diffusion solver can have first order convergence in time given a high enough conversion rate. As Figure 3.10b shows this was achieved, and the hybrid solver is therefore considered adequately implemented for the time being.

Effects of varying the relative size of the RW area were also illustrated. The result of increasing the portion of the PDE mesh affected by an RW model is clearly a larger error term overall, while decreasing the size of the RW area significantly improves the error term. All of this is expected since more fluctuations are introduced by the RW solver.

5.2 Results for physical application

Craske et.al. [4] suggest that the neck of spines act as diffusion barriers which slow down, but don't completely stop the diffusion of PKC γ into spines. The function of this barrier is a bit unclear, but the presence of it is undisputed. In their measurements they found a delay of around 5 – 10 seconds from elevated concentration levels in the dendrite until a similarly elevated concentration level occurred in spines with necks longer than $0.5\mu\text{m}$. Using parameter values which resemble the values found in actual (rodent) neurons and neurites in the developed software, the observed delay-times have been recreated. Figure 5.2 shows plots of the observed diffusion times into spines. This figure shows a clear trend for longer diffusion times as the neck length of the spine increases. Figure 5.1 further supports this claim and implies the average diffusion time for PKC γ into long necked spines to be in accordance

with the results from Craske et.al. Seeing as there are no additional complexities added to the RW model we can assume that the spine neck does in fact function as a diffusion barrier.

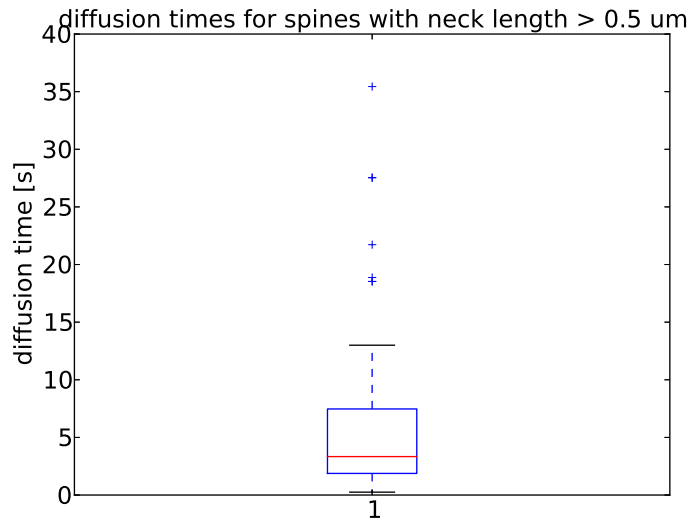


Figure 5.1: Boxplot of the relative diffusion times (time between elevated concentration in dendrite and elevated concentration in spine head) into spines with necks longer than $0.5\mu\text{m}$. Similar studies were done by Craske et.al. [4] and found diffusion time (unclear whether relative or not) to be somewhere around 5-10 seconds.

Through the simulations it became apparent that there must be some sort of limiting factor which limits the number of $\text{PKC}\gamma$ particles that are let into the spine. In real life this is achieved by a concentration gradient which tends to zero (or negative values), meaning that no particles will diffuse into the spine after it is “filled” up. A random walker will not feel this concentration gradient unless it is explicitly told so. The alternative solution then, is to reduce the probability for particles to diffuse into a spine for each particle that gets caught in the spine head by a factor $\frac{1}{10}$.

5.3 Using other particle dynamics

The aim of this project has always been to create a flexible framework where especially the particle dynamics can be substituted. A test was therefore executed where Direct Simulation Monte Carlo (DSMC) code developed by Anders Hafreager for his masters thesis [8], was used in stead of the RW

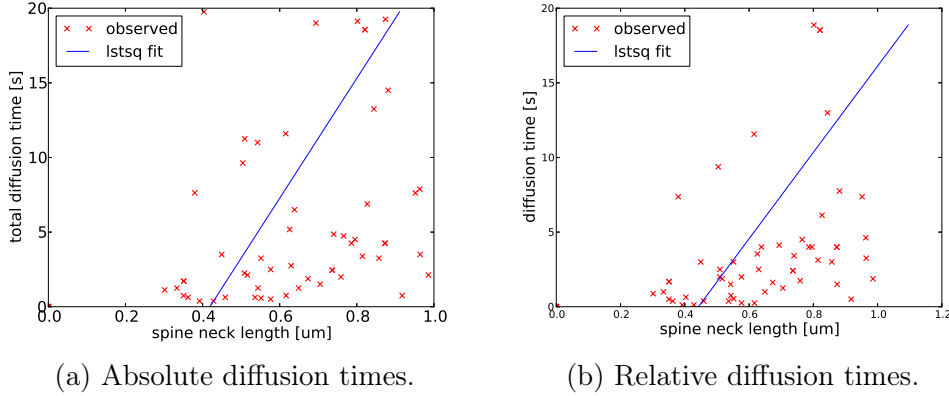


Figure 5.2: Absolute (a) and relative (b) diffusion times into spines with outliers removed from the data. Outliers are the observations illustrated by crosses in Figure 5.1. The lines represent a least squares fit of the results. For (a) the least squares fit is $y \approx 0.0346x + 0.4424$, where y is the diffusion time and x is the neck length of the spine. For (b) the least squares fit is $y \approx 0.0249x + 0.4189$.

model. DSMC is a particle dynamics approach where the velocities of individual particles are modeled, but some particle interactions are modeled by a stochastic process. Typically, particle collisions are modeled by a Monte Carlo method, but boundary conditions are not.

A problem arises with respect to the coupling described in this thesis: In the DSMC model, there is no fixed step length for each particle. It is possible to measure a diffusion coefficient which, to some extent, can be controlled by adjusting the temperature of the system, but this has not been done.

As a result, this must be viewed as an experiment in which the aim is for the two softwares to communicate as intended. In this sense, the experiment was successful. Figure 5.3 shows one time step of a simulation where the DSMC model was used as particle dynamics model. As is apparent from the figure, the coupling is not very good.




Figure 5.3: This figure shows the hybrid diffusion solver with a DSMC particle model. The coupling between the two modes is not very good, but the communication between them works.

Chapter 6

Discussion and conclusion

6.1 Discussion

A large portion of the work that has been put into this thesis has been on code implementation. Because it is very demanding to read computer code on paper, and it does not necessarily provide extra clarity, the code is not included in this text. The complete computer code is published at *github.com* under the following address

<https://github.com/fepettersen/thesis>

In summary, this project contains

- Implementation of both explicit and implicit solvers for the (anisotropic) diffusion equation using finite difference methods in one and two spatial dimensions. The implicit solver can also handle three dimensions, but this requires writing code that assembles the coefficient matrix.
- Implementation of a direct solver for banded linear systems which rewrites the system as block tridiagonal. To my knowledge, this block tridiagonal solver is one order more efficient than any other direct solver.
- Implementation of a RW model for (anisotropic) diffusion in one and two spatial dimensions, with the grounds for three dimensions laid.
- Combination of the PDE and RW models into a hybrid diffusion model.
- Thorough testing and verification of all the implemented solvers.
- An application of a slightly modified version of the hybrid diffusion model on diffusion of PKC γ from thick apical dendrites into very narrow dendritic spines.

6.1.1 Application

The model gives fairly good results for the application on PKC γ into spines which largely are in agreement with the results by Craske et.al [4]. However, the mean diffusion times found from simulations seem to consistently lie towards the lower limits of the experimental results. One possible extension to the model which might fix this is to introduce an absorption probability in the spine head which is fairly large (say 80% per second), or even increases with the amount of time spent in the spine head. A setup like this should increase the average diffusion time by a few seconds. It does not, however,

reflect the physical process to a better accuracy seeing as a concentration increase in a spine head will be measured quickly.

6.1.2 Possible extensions

As the project stands now it is mostly a proof of concept for hybrid diffusion solvers as well as a first attempt at a flexible framework for similar problems. This means that there are quite a few possible extensions that can be done to the project, perhaps in conjunction with another masters thesis.

A simple RW was chosen for the lower scale model because it fulfills the diffusion equation and is therefore easy to work with. The idea, however, was always to create a software in which the lower scale model can easily be substituted for a better one. By letting the lower scale model work as a standalone unit which communicates with the rest of the software through a file containing the positions of all the walkers (or particles), this is ensured. All that is needed to switch lower scale model is to put the new solver in the correct place with respect to the rest of the software, and make sure it can communicate in the described manner. As a test, the DSMC code developed by Anders Hafreager [8] was used as a lower scale model for one simulation. Naturally, a few problems arose, but from a strictly programming point of view it worked.

Another, rather simple extension would be to find an application which requires modeling anisotropic diffusion. This has already been implemented in the PDE and RW solvers. The anisotropic RW implementation is based on the article by Farnell and Gibson [5].

As mentioned in section 2.3.4, perfect flux exchange between the higher and lower scale model might have been a better boundary condition for the RW model than zero flux boundary conditions. In principle, changing boundary conditions on the RW solver is simple, seeing as it is completely separate from the rest of the software. The reason the boundary conditions have not been changed is that it requires a complete workover of the coupling between the two models as well, which is a too large job to complete at this point.

Perhaps the biggest weakness of the software, as it stands, is the limitation in mesh geometry. Implementing a mesh geometry which is not square in a finite difference method turns out to be very complicated, and requires transforming the PDE to a new set of coordinates. Ultimately, one must solve an entirely different equation. Alternatively, a finite element method

can be used. Finite element software will already have support for new mesh geometries built in, making the suggested transformations unnecessary.

There are, of course, a lot of physical problems which the developed software can be applied to. One possibility is modeling sodium ions that diffuse through ion channels in the cell membrane of neurons into the extracellular space of the brain.

Ion channels are very narrow, often only allowing one ion through at a time, and they are (often) specific to one ion. Two 1D PDEs can describe the intra-, and extra-cellular concentrations of ions, and a lower scale model can describe the ion channel. This problem would also need to consider drift terms arising from Coulomb forces and using a modified particle model which can describe the dynamics inside an ion channel.

In principle, any diffusion process where a small portion of a large number of particles diffuse into narrow passages could be modeled by the developed hybrid diffusion solver.

6.1.3 Other work on the topic

As the project was being finished, I came across an article by Flekkøy et.al.[6] describing the same problems that are addressed in this thesis. In this article, the authors try to combine a simple diffusion solver with a simple random walk solver and end up concluding that this is possible. This thesis has been done completely independently of said article, and takes a slightly different approach to the problem as well.

6.2 Concluding remarks

In this thesis a hybrid diffusion solver in which parts of the process can be modeled by a particle dynamics description has been developed. All parts of the solver have been verified to work properly, including the hybrid model.

The developed software mainly relies on the implicit BE scheme to solve the diffusion equation, both in 1D and in 2D. Like any implicit scheme, the BE scheme results in a system of linear equations of size $n^d \times n^d$ which must be solved at each time step. In order to do so, a block tridiagonal solver has been implemented, with an efficiency of $\mathcal{O}(n^{2d-1})$. To my knowledge, this is the most effective direct solver available, with alternatives like LU decomposition and Levinson recursion using $\mathcal{O}(n^{2d})$ FLOPs. The limiting factors of

the block tridiagonal solver are two matrix-vector multiplications which will use $\mathcal{O}(n^{2(d-1)})$ FLOPs. If a faster matrix-vector multiplication scheme exists it will reduce the computational work to $\mathcal{O}(n^d)$.

Appendix A

Appendix

A.1 Debugging

In any project which involves programming one is bound to do some debugging. This project is no exception. Debugging can be extremely frustrating because no one sees all the hours that go into finding the bugs. Should a bug not be found, however, it will be noticed immediately. This appendix is an introductory guide to debugging finite difference (FD) solvers and RW implementations. It starts out with some general tips on how to handle error messages, and moves on to deal with FD and RW implementations. Finally, some words on how to debug the developed software are included.

A.1.1 Compiler/syntax errors

If you are programming in a compiled language like Fortran or C/C++ it is very easy to forget or misspell some syntax. Usually, the compiler will tell you what is wrong. Should it not do so, there are compiler options which provide extra information about syntactically correct but questionable looking code (-Wall). Some times one error results in several error messages. Start with the messages that are understandable, and recompile before checking out the other messages.

If you are building a larger project which requires linking, remember that packages must be linked in the correct order. For example; the Armadillo linear algebra library is backened by LAPACK and BLAS. In addition to linking the Armadillo library, both LAPACK and BLAS must be linked, and they must be linked in the correct order:

```
g++ myprog.cpp -o myprog -O2 -larmadillo -llapack -lblas
```

Anything else will give very cryptic compiler errors.

Interpreted languages, like Python or MatLab, mostly include debuggers which provide extensive information about any errors in the code, read them thoroughly.

A.1.2 Segmentation faults

Segmentation fault is a very common abort message to get from C++. It means that the program has tried to access a part of the computer memory (RAM) which it is not allowed to access by the operating system. Interpreted languages will abort with a more informative error message, saying

something about where the error was encountered. In a compiled language, “Segmentation fault” is all you get, which basically means that there is something wrong somewhere in your code. As an illustration, the terminal output below shows the output from a python- and C++- program (both named “segfault”) for the same error:

```
fredrik@Workstation:~/uio$ python segfault.py
Traceback (most recent call last):
  File "segfault.py", line 3, in <module>
    b = sys.argv[1]
IndexError: list index out of range

fredrik@Workstation:~/uio$ ./segfault
Segmentation fault (core dumped)
```

The gnu-compiler has an environment called gdb in which you can run your program. “gdb” will catch segmentation faults and provide some information about where they were encountered. Furthermore, advanced editors like “qt creator” have built-in debuggers and the ability to place breakpoints in the code, making it possible to step through code line by line. Some times though, the thing that works best is to print messages or key values at various places, signaling that a code block is finished. For example, making each function print its name when it is called will quickly identify problems.

A.1.3 Finite difference methods

First and foremost: Do the discretization by hand. It will give valuable insight, which makes implementation and debugging a lot simpler. Secondly, the numerical stability of the scheme should be investigated prior to implementation. This ensures that a conditionally stable scheme does not fail tests because of a too large discretization parameter.

There is one very important rule in programming in general: “First make it work, then make right, then make it fast”. When solving PDEs numerically this principle is extra important. A numerical discretization scheme will (to my knowledge) either result in a main loop or a linear system which solves the equation. Regardless of the overall structure of the final program, this main block will be located unchanged somewhere within the final structure. This part should therefore be implemented and tested first. Structural changes and efficiency issues can be addressed at a later point.

Visualization is an invaluable tool while implementing a numerical PDE solver. Even though the code runs through all the steps without crashing, it does not have to be correct. This is illustrated in Figure A.1.

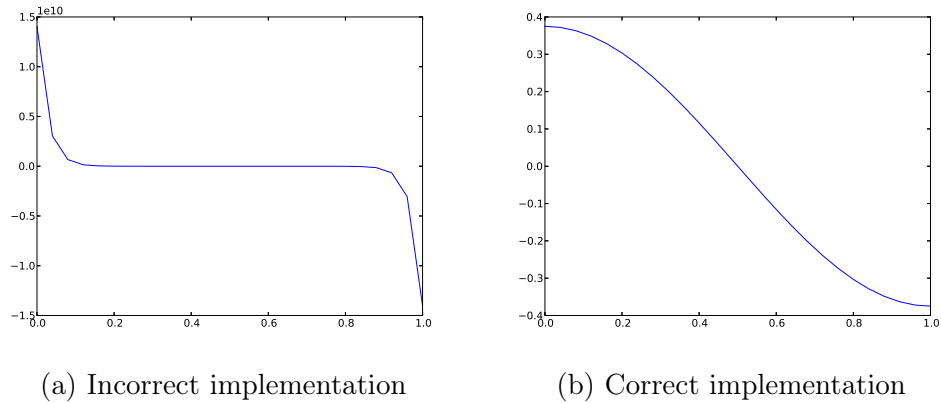


Figure A.1: Illustration of an implementation which does not crash the program, but is still clearly incorrect (a). (b) shows the correctly implemented version of the same time step.

At this point I would like to introduce rubber-duck debugging which is said to be invented by Dennis Ritchie, the developer of the C programming language. The story goes that he would keep a rubber duck at his desk and whenever he was stuck, would describe the code in detail (what each statement did and was supposed to do) to the rubber duck. Asking questions, and saying thing out loud forces you to think about things in a slightly different manner.

When the code seems to reproduce the intended results it is time to start the verification. This is where we make an error estimate and do some numerical analysis. Making sure an implementation is correct is a lot harder than it sounds, but there are a few points that should be fulfilled:

- **Manufactured solution**
As was described earlier, the source term of the equation can be used to force almost any solution upon the equation.
- **Stationary solution**
This is a special type of manufactured solution. Using the source term, and initial condition, the solution can be forced to be constant. If a scheme reproduces a stationary solution it conserves energy, which is very desirable.

- **Exact numerical solution**
Depending on the discretization, it should be possible to find an exact solution to the discretized version of the PDE. An example of this is found in chapter 3.3.3. The scheme is expected to reproduce this solution to machine precision, but there are often several factors which limit the accuracy. At the very least, the error should be significantly smaller than the largest error term.
- **Convergence test**
It is often difficult to make convergence tests work properly, but they are arguably the most powerful tests and therefore important to implement. The principle is to verify that using a smaller discretization parameter results in a smaller error, and that the error is reduced by the expected amount.

There are probably more ways to make sure that a finite difference scheme is working properly, but the ones listed will usually give a good implication.

A.1.4 Random Walk and Monte Carlo methods

While solving a PDE numerically, the intermediate calculations can be checked and verified. Since Monte Carlo methods are based on random numbers, the results of intermediate calculations will be random, making it impossible to verify them completely. Instead, some statistical properties of the random numbers can be verified. For example, using uniformly distributed random numbers will give a certain mean and standard deviation, while a Gaussian distribution will give another. The random number generator should be tested in order to verify that these properties are reproduced to a reasonable accuracy.

Random walkers have some additional properties which can be tested. Both a one-dimensional and a two-dimensional random walk will fill all available space, given infinitely many time steps. Letting 4-5 walkers jump around for around 10^4 time steps, and plotting their paths should give an indication of how well this property is reproduced.

Finally, it is very useful to specify a random seed. This will ensure that the random number generator produces the same sequence of numbers each simulation, which often makes testing easier.

A.1.5 The developed software

Since the developed software combines two solvers, it is of course very important to make sure that each solver is working properly. This can be verified by performing the tests mentioned earlier.

While implementing the hybrid solver, there were four main points which either made the work easier, or would have made the work easier:

- Unit testing
This is a way to test a single unit of code, often a single function, and verify that it does exactly what it should do.
- Version control
During implementation and testing, a lot of changes will be made to the code. It is surprisingly difficult to remember all the changes, which often means that functioning code will be lost. Using version control software eliminates this problem. Version control software allows a developer to save snapshots of code with a comment, as well as some more advanced options. This allows for tracking of progress, committing working code before changes are made, and much more. Should some changes be made to a working piece of code, it can be reverted to the last working version.
- Make independent modules
This applies to most code development. If the code is divided into independent modules, it is not only easier to test the code, it will also be much easier to reuse the code.
As an example, the block tridiagonal solver has been implemented as a separate module, which since then has been used by Torbjørn Sæland in his masters thesis.
- Do not be nostalgic
Because you are now using version control, there is no need to comment out sections of code in case it is needed later. This makes the rest of the code less messy, and reduces the risk of forgetting to remove something.

A.1.6 Last resorts

If nothing else seems to work, I like to try these techniques in order to get a fresh view on the problem.

First off all, doing some form of verification of intermediate steps by hand is often very useful. If everything seems to be implemented correctly, but the results are still wrong, it is time to really verify as much as possible. Try verifying indexing, counting, etc.

Changing parameters is often useful. As a rule of thumb, it is unwise to use “simple” parameter values, like zero or one, because values like these are often special cases which give simple effects. Changing a constant from one to a random value between zero and ten can often illustrate errors in parts of the code which has already been tested. Similarly, any matrix calculations done by hand should be done with matrices larger than 3×3 . It is much harder to recognize patterns in smaller matrices.

If neither of these work, try starting from scratch in a different programming language (preferably an interpreted language). The slight difference in syntax often reveals a lot of problems.

Finally, it often helps to get some distance from the problem. Work with something else for a couple of days, and come back with a fresh, and more critical view.

Bibliography

- [1] Milton Abramowitz and Irene A Stegun. *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*. Courier Dover Publications, 2012.
- [2] Jon I Arellano et al. “Ultrastructure of dendritic spines: correlation between synaptic and spine morphologies”. In: *Frontiers in neuroscience* 1.1 (2007), p. 131.
- [3] Mary L Boas. *Mathematical Methods in the Physical Sciences*. John Wiley & Sons., Inc, 2006. Chap. 8.
- [4] Madeleine L Craske et al. “Spines and neurite branches function as geometric attractors that enhance protein kinase C action”. In: *The Journal of cell biology* 7.170 (2005), pp. 1147–1158.
- [5] L Farnell and WG Gibson. “Monte Carlo simulation of diffusion in a spatially nonhomogeneous medium: A biased random walk on an asymmetrical lattice”. In: *Journal of Computational Physics* 208.1 (2005), pp. 253–265.
- [6] EG Flekkøy, J Feder, and G Wagner. “Coupling particles and fields in a diffusive hybrid model”. In: *Physical Review E* 64.6 (2001), p. 066302.
- [7] Bruce Graham, Andrew Gillies, and David Willshaw. *Principles of computational modelling in neuroscience*. Cambridge, 2011.
- [8] Anders Hafreager. “Flow of dilute gasses in complex nanoporous media”. MA thesis. University of Oslo, 2014.
- [9] George Marsaglia. “Xorshift RNGs”. In: *Journal of Statistical Software* 8.14 (2003), pp. 1–6.
- [10] Sean T O’Connell and Peter A Thompson. “Molecular dynamics–continuum hybrid computations: a tool for studying complex fluid flows”. In: *Physical Review E* 52.6 (1995), R5792.

- [11] Mathis Plapp and Alain Karma. “Multiscale finite-difference-diffusion-Monte-Carlo method for simulating dendritic solidification”. In: *Journal of Computational Physics* 165.2 (2000), pp. 592–619.
- [12] Naoaki Saito and Yasuhito Shirai. “Protein kinase C γ (PKC γ): function of neuron specific isotype”. In: *Journal of biochemistry* 132.5 (2002), pp. 683–687.
- [13] James M Varah. “On the solution of block-tridiagonal systems arising from certain finite-difference equations”. In: *Mathematics of Computation* 26.120 (1972), pp. 859–868.
- [14] Patrick B Warren. “Dissipative particle dynamics”. In: *Current opinion in colloid & interface science* 3.6 (1998), pp. 620–624.