

Developing Fundamental Models of
Colloid Transport and Absorption in Sand
Filters

by

Matthew R. Killeen

Thesis

for the award of

Doctor of Philosophy



The
University
Of
Sheffield.

Supervised by Dr. Karl P. Travis

Department of Materials Science and Engineering

The University of Sheffield

April 2018

Acknowledgements

First of all, I would like to express my sincerest gratitude to Dr. Karl Travis, whose vast knowledge on all things computer simulation has been invaluable over the last 4 years. His supervision and constant optimism has provided untold motivation in completing this work. He has always pointed me in the right direction, and has been a first-rate teacher. In a similar vein, these thanks are extended to Dr. Mark Bankhead, whose input and contributions have been received gratefully. He has given me insight into the nuclear industry, allowing me to spend time at the offices at NNL gaining useful industrial experience.

Of course, I am also hugely thankful for the camaraderie of my colleagues in H7, and those in the rest of the computer modelling group, who have provided years of much needed support and entertainment. In particular, Adam, with whom I shared many important bug fixing sessions with, and Kai, whose ability to blame all of life's problems on "the referee" is uncanny.

Finally, my gratitude towards my family cannot be expressed enough. My parents have provided me with every opportunity I could have asked for, and have always encouraged and nurtured an interest in science. My partner, Rose, has provided untold support throughout this process, and I will always be grateful for this. An honourable mention must also go to my writing buddy and dog, Odie, who has sat at my feet for hours on end writing the following.

Developing Fundamental Models of Colloid Transport and Absorption in Sand Filters

Project Summary

This work was undertaken as part of an Industrial Collaborative Awards in Science and Engineering (iCASE) research programme, jointly funded by the National Nuclear Laboratory (NNL) and the Engineering and Physical Sciences Research Council (EPSRC). The aim was to probe the mechanisms of clogging of sand bed filters using particle based computer simulation methods. Existing models take a top down approach, making use of an empirical clogging parameter. Simulation holds the prospect of relating this parameter to properties of the effluent and the sand bed.

The problem was approached using two computational methods: molecular dynamics, and smooth particle applied mechanics. The molecular dynamics model yielded successful results, qualitatively agreeing with existing experimental data with regards to the rate of deposition within the bed, and the associated observed pressure drop. The model was systematically explored by varying the nature of the colloid-fluid-sand forces, the geometry and packing fraction of the sand bed, and the concentration of the colloids. An investigation into the fractal nature of the deposits was also performed, suggesting that a lower fractal dimension creates greater physical hinderance to the flow. This serves as additional validation for existing theories.

The smooth particle model yielded less successful results. Substantial parameterisation of the model was undertaken, however, the model still showed signs of instability under certain conditions. Again, it produced qualitative agreement with existing literature, but showed substantial deviation from the results gained from the molecular dynamics model. Ultimately, further parameterisation of this model is required to allow for a more effective comparison of the models.

Table of Contents

1. Introduction	1
1.1 Background and thesis outline	1
1.2 SIXEP	3
1.3 Modelling techniques	6
1.3.1 Microscoping approach	7
1.3.2 Macroscopic approach	11
1.3.3 Modelling colloid transport	17
1.3.4 Modelling colloid interactions	21
1.3.5 Experimental techniques	22
1.4 Mechanics	23
1.5 Summary	24
1.6 References	26
2. Molecular dynamics	30
2.1 Initial conditions	32
2.2 Boundary conditions	33
2.2.1 Periodic boundary conditions	33
2.2.2 Minimum image convention	35
2.2.3 Other boundary conditions	36
2.3 Integration schemes	39
2.3.1 Euler algorithm	39
2.3.2 Verlet algorithm	40
2.3.3 Velocity Verlet algorithm	40
2.3.4 Runge-Kutta algorithm	41
2.3.5 Comparison of integration schemes	43
2.4 Potential energy functions	47
2.4.1 Hard sphere potential	48
2.4.2 Square well potential	49
2.4.3 Lennard-Jones potential	50
2.4.4 Short range soft repulsive potential	52

2.5	Link cells	53
2.6	Calculating thermodynamics properties	54
2.6.1	Potential energy	55
2.6.2	Kinetic energy	55
2.6.3	Total energy	56
2.6.4	Temperature	56
2.6.5	Pressure	57
2.7	Thermostats	58
2.7.1	ad hoc velocity rescaling	58
2.7.2	Andersen thermostat	59
2.7.3	Gaussian thermostat	59
2.7.4	Nosé-Hoover thermostat	60
2.8	References	62
3.	Molecular dynamics simulation of sand bed filtration	64
3.1	Overview	63
3.2	Model design	65
3.2.1	Equilibration stage	65
3.2.2	Filtration stage	71
3.2.3	Calculating properties	75
3.3	Model validation	85
3.3.1	Specific deposit	85
3.3.2	Pressure drop	91
3.3.3	Concentration with depth	94
3.3.4	Summary	96
3.4	Systematic exploration	97
3.4.1	Interatomic potentials	97
3.4.2	Lattice structure	107
3.4.3	Colloid concentration	109
3.4.4	Porosity	111
3.4.5	Fractal dimension	115
3.5	Summary	126
3.6	References	128

4. Continuum mechanics	130
4.1 Conservation equations	130
4.1.1 Conservation of mass	130
4.1.2 Conservation of momentum	132
4.1.3 Conservation of energy	133
4.1.4 Constitutive relations	135
4.2 Smooth particle applied mechanics	136
4.2.1 Overview	136
4.2.2 SPAM approximation of the continuity equations	137
4.2.3 Weight functions	139
4.2.4 Artificial viscosity	140
4.2.5 Boundary conditions	141
4.2.6 Initial conditions	142
4.2.7 Particle size	143
4.3 References	144
5. Constitutive relations	145
5.1 Overview	145
5.2 Equilibrium equation of state	146
5.2.1 $\varepsilon = 100$	146
5.2.2 $\varepsilon = 10$	157
5.3 Viscosity	162
5.3.1 Sliding wall method	163
5.3.2 Lees-Edwards method	164
5.3.3 The SLLOD method	165
5.3.4 Calculating viscosity through extrapolation	167
5.3.5 $\varepsilon = 100$	170
5.3.6 $\varepsilon = 10$	175
5.4 References	180
6. SPAM simulation of sand bed filtration	181
6.1 Overview	181
6.2 Model design	181
6.3 Fluid behaviour	185
6.3.1 Binary mixture of colloids and fluid	185

6.3.2 Fluid behaviour under gravity	186
6.4 Simulation instabilities	188
6.5 Specific deposit and pressure drop	190
6.5.1 Specific deposit	190
6.5.2 Pressure drop	191
6.6 Comparison with MD	194
6.6.1 Colloid-colloid potential	194
6.6.2 Porosity	195
6.6.3 Square lattice	197
6.6.4 Equation of state	199
6.7 Summary	200
6.8 Reference	201
7. Conclusions and further work	202
7.1 References	204

List of figures

Figure		Page
1.1	Discharge of radiation to the Irish Sea. SIXEP was opened in 1985	4
1.2	Schematic showing a simplified flow diagram for SIXEP.	5
1.3	Basic transport mechanisms, adapted from O'Melia. A is sedimentation, B is interception, C is diffusion.	8
1.4	An illustration of straining, adapted from Bradford <i>et al.</i>	9
1.5	Illustrations of deposit profiles as a function of depth. Left – exponential, middle – hyperexponential, right – non-monotonic.	12
1.6	Illustrations of deposits with a low (left) and high (right) fractal dimension, adapted from Mays ¹⁵ .	13
1.7	Head loss against time for the fitted model and experimental data. Results obtained by O'Melia and Ali.	16
1.8	Schematic of Bradford's model. Region 2 is the SWI and region 1 is the bulk aqueous phase. J is the colloid flux, S is the colloid concentration in the solid phase, C is the suspended colloid concentration, k is a rate constant and α is the colloid attachment efficiency.	19
2.1	An illustration of periodic boundary conditions in two dimensions. The red circles represent the starting positions, and the blue circles the final positions. As a particle leaves the central simulation cell it is replaced by its periodic image.	34
2.2	An illustration of the minimum image convention. Each particle from the central cell is considered in turn. It is placed at the center of a box with the same dimensions as the cell. The force between itself and all	36

other particles within this cell is calculated, shown by the coloured particles. The circle represents the potential cutoff.

- 2.3 An illustration of elastic boundary conditions. The particle's perpendicular component of velocity is reversed upon impact with the wall. The red circles represent the starting positions, and the blue circles the final positions. 37
- 2.4 An illustration of mirror boundary conditions where the dashed line is the mirror boundary. The particles interact with a mirrored version of themselves. The red circles represent the simulation positions, and the blue circles the mirror particles. 37
- 2.5 An illustration of stone wall boundary conditions. The perpendicular component of velocity is set to zero, the parallel component of velocity is set to the velocity of the wall upon collision with the wall. The red circle represents the starting positions, and the blue circle the final positions. 38
- 2.6 A comparison of phase space plots for the Euler method (top left), RK2 method (top right), RK4 method (bottom left), and the analytical solution (bottom right). Both the Euler and RK2 methods are shown to have significant integration errors compared to the RK4 method. Δt is the same in all cases. 45
- 2.7 Global error as a function of timestep size for four integration schemes described in the text., applied to a 1D simple harmonic oscillator. Lines are least squares fits to the data. 46
- 2.8 A plot of energy against interparticle distance for the hard sphere potential. When the separation is equal to or smaller than the diameter of the hard sphere the potential energy is infinite. 49
- 2.9 A plot of energy against interparticle distance for the square well potential. When the separation is smaller than the diameter of the hard 50

	sphere the potential energy is infinite, and the attractive region is approximated by a square well. $R = 1$.	
2.10	A plot of potential energy against interparticle distance for the Lennard-Jones potential.	51
2.11	A plot of potential energy against interparticle distance for the soft-sphere potential.	52
2.12	An illustration of link cells. The simulation cell is split into link cells, and only particles in the selected cell and adjacent cells are considered when calculating forces. The red particles are close enough to interact, and the dark particles are too far away.	54
3.1	A schematic showing the molecular dynamics model. The yellow particles are the sand particles, the blue particles are fluid particles and the red particles are colloid particles. The dashed line is an elastic boundary that initially separated the fluid mixture and the sand bed.	65
3.2	Illustration showing the initial fluid conditions. The blue particles are fluid and the red are colloids. The type of each particle is randomly decided based on the ratio. This example shows a 4:1 ratio of fluid:colloid.	67
3.3	Snapshots showing the computed particle locations of a binary mixture after 100 and 10,000 timesteps using periodic boundary conditions. $\epsilon_{11} = \epsilon_{22} = 20.0$ $\epsilon_{12} = 25.0$. $\Delta t = 0.01$.	68
3.4	Computed particle locations showing the initial conditions (left) and final coordinates (right) of a binary mixture containing 1024 particles (512 colloid, 512 fluid). The simulation length was 1,000,000 time periods, $\Delta t = 0.01$, $g = 0.1$, $\tau = 1 / \Delta t$.	70
3.5	Illustration showing the shifted core potential. The distance is calculated from the surface of the sand particle rather than the centre.	71

3.6	Illustration showing the sand particle coordinates. The left shows a triangular lattice (coordination number of 6) and the right shows a square lattice (coordination number of 4).	72
3.7	Image depicting the start of the second phase. The elastic boundary has been removed, and the fluid and the sand can now interact with each other.	74
3.8	Schematic showing the weighted method of planes. Each particle whose distance, r_j , from the plane of interest is less than the smoothing length, h , contributes to the property at that plane. Red particles are within the smoothing length of the plane highlighted in red, grey particles are not.	76
3.9	Plot showing the one-dimensional Lucy weight function used to weight particle influence at a plane. $h = 3$.	77
3.10	Computed particle locations for two simulations where the gravitational force, $g = 0.0$ (left) and 1.0 (right). $\Delta t = 0.01$, total timesteps = 100,000, $\rho_0 = 1.0$.	78
3.11	Density profiles without (left) and with (right) a gravitational field. Density was calculated at 100 planes in the y-directions every 10 timesteps.	79
3.12	Velocity profile across a pore.	81
3.13	Image showing the profile planes used to measure the thermodynamic pressures across the filter. The pressure drop is the difference in pressure at profile planes A and B.	82
3.14	Snapshots showing particle coordinates at 5 stages during the same simulation. The purple circles are the colloidal particles, the light blue circles are the fluid particles, the dark blue circles are the clogged colloidal particles and the orange circles are the sand particles. Sticking probability = 0.001, sticking distance = 0.75.	83

3.15	Pressure profiles for the 5 computed particle locations in figure 3.13.	84
3.16	Specific deposit against time showing the two filtration stages. Here blocking is modelled: colloids only stick to the sand particles, colloid clogging distance = 0, colloid sticking probability = 0.	87
3.17	Computed particle locations at 5 stages during the simulation shown in figure 3.15.	88
3.18	Specific deposit against time showing the two filtration stages. Here ripening is modelled: colloids only stick to both sand particles and deposited colloid particles. Colloid clogging distance = 0.5, colloid sticking probability = 0.00.	89
3.19	Computed particle locations at 5 stages during the simulation shown in figure 3.18.	90
3.20	Pressure drop against time for a clean filter, used to establish ΔH_0 .	92
3.21	Pressure drop as a function of specific deposit. The linear fit describes the relationship between pressure drop and specific deposit during the first phase. The quadratic fit shows the dependence between pressure drop and specific deposit in the second phase.	93
3.22	Illustrations of deposit profiles as a function of depth. Left – exponential, middle – hyper-exponential, right – non-monotonic.	94
3.23	Specific deposit against depth. Porosity = 0.7.	95
3.24	Specific deposit against depth. Porosity = 0.4.	96
3.25	Specific deposit (a) and pressure drop (b) against time whilst varying the strength of the colloid-colloid pairwise potential from $\varepsilon = 2.0$ to $\varepsilon = 100.0$.	98

3.26	Density profiles comparing the weakest (left) and strongest (right) colloid-colloid interactions from figure 3.1. The density includes contributions from all fluid particles.	99
3.27	Specific deposit (a) and pressure drop (b) against time whilst varying the strength of the colloid-sand pairwise potential from $\varepsilon = 2.0$ to $\varepsilon = 100.0$.	101
3.28	Specific deposit (a) and pressure drop (b) against time whilst varying the strength of the colloid-fluid pairwise potential from $\varepsilon = 2.0$ to $\varepsilon = 100.0$.	103
3.29 (part 1)	Snapshots of the particle locations after 100,000 timesteps for the simulations in figure 3.27, with increasing fluid-colloid repulsion from left to right. The red particles are colloid, the light blue fluid, and the dark blue are clogged.	104
3.29 (part 2)	Snapshots of the particle locations after 100,000 timesteps for the simulations in figure 3.27, with increasing fluid-colloid repulsion from left to right. The red particles are colloid, the light blue fluid, and the dark blue are clogged.	105
3.30	Comparison of specific deposit against time (top) and pressure drop against time (bottom) whilst varying the colloid-colloid interaction strength for a square lattice (left) and triangular lattice (right).	108
3.31	Comparison of the coordinates of clogged particles in a triangular lattice (left) and a square lattice (right). The dashed line shows a suggested flow path.	109
3.32	Specific deposit (a) and pressure drop (b) against time whilst varying the concentration of colloid particles from 10% to 50%.	110
3.33	Example results from probing porosity. The purple particles are misses, and the blue particles are hits. 1,000,000 test particles were generated.	111

3.34	Porosity as a function of number of test particles created, N .	112
3.35	Specific deposit (a) and pressure drop (b) against time whilst varying the porosity of the sand bed from 1.44 to 70.57.	113
3.36	Image showing the coordinates of the clogged particles for the simulations shown in figure 3.36, with porosity decreasing from left to right and top to bottom.	114
3.37	Pressure drop as a function of porosity. The data is taken from figure 3.35, where $t = 90,000$.	115
3.38	An illustration of the Koch Curve, a deterministic fractal created using an iterative process.	116
3.39	Illustrations of deposits with a low (left) and high (right) fractal dimension, adapted from Mays.	116
3.40	Flow diagram showing how to determine which sand particle a colloid has stuck to.	118
3.41	Snapshots of the deposited colloid particles on five sand particles at varying time throughout the simulation. Colloid particles could deposit on top of other colloid particles. The colloid clogging distance was 0.3.	120
3.42	Log-log plot of number of particles against radius of gyration. The colloid clogging distance is 0.3. The fractal dimension is taken from the gradient of the slope.	121
3.43	Snapshots of the deposited colloid particles on five sand particles at varying time throughout the simulation. Colloid particles can deposit on top of other colloid particles. The colloid clogging distance is 0.4.	122
3.44	Log-log plot of number of particles against radius of gyration. The colloid clogging distance is 0.4. The fractal dimension is taken from the gradient of the slope.	123

3.45	Fractal dimension against sticking distance.	124
3.46	Specific deposit as a function of time for two data sets, with fractal dimensions of 1.671 and 1.296.	125
3.47	Fractal dimension against sticking distance.	126
4.1	Lucy's 2-dimensional weight function (red), where $h = 3$, along with the first (blue) and second (green) derivatives.	140
4.2	Comparison of elastic (left) and mirror (right) boundary conditions. The density calculated at the particle at the boundary for the elastic case is half that of the mirrored case. When calculating the density at the dotted line, an additional contribution is included from the mirrored particles.	142
5.1	P ressure against density, comparing results from MD simulations of 1024 particles with different timesteps. The black line shows the results generated using the equation of state (equation 5.8). The length of all simulations was 500 reduced units.	148
5.2	P ressure against density, comparing results from MD simulations against the predictions from the equation of state.	149
5.3	Average pressure against energy, comparing results from MD simulations against the predictions from the equation of state.	151
5.4	K inetic temperature against density, comparing results from MD simulations against the predictions from the equation of state.	152
5.5	Snapshots of the computed particle locations for a constant density of 1.2 (left) and 1.4 (right).	153
5.6	Average kinetic temperature against energy, comparing results from MD simulations against the predictions from the equation of state. Density was fixed at unity.	154

5.7	Total energy against density, comparing results from MD simulations against the predictions from the equation of state. The kinetic temperature was fixed at unity.	155
5.8	Average total energy against temperature, comparing results from MD simulations against the predictions from the equation of state. The density is fixed at unity.	156
5.9	Difference in pressure against difference in density. The reference point, (0,0), is at a density and temperature of 1.0.	158
5.10	Difference in pressure against difference in energy. The reference point, (0,0), is at a density and temperature of 1.0.	159
5.11	Difference in temperature against difference in density. The reference point, (0,0), is at a density and temperature of 1.0.	160
5.12	Difference in temperature against difference in energy. The reference point, (0,0), is at a density and temperature of 1.0.	161
5.13	An illustration of the sliding wall method. The red particles make up the wall. The wall moves with a fixed velocity in the x-direction proportional to the shear rate. The structure of the walls is maintained using Hookean springs.	163
5.14	Lees-Edwards periodic boundary conditions. The top and bottom layer of periodic images are offset proportionally to the effective strain rate, n . The particle leaving the simulation cell is shifted when being mirrored.	164
5.15	Shear rate against viscosity for a series of SLLOD simulations, run at unit density and temperature. Each simulation was run for 7,000,000 timesteps, with a timestep of 0.001, using 1024 particles. A Lorentzian function is used to fit the data. The inset shows the same data on a log scale.	168

5.16	Shear rate against viscosity for a series of SLLOD simulations, run at unit density and temperature. Each simulation was run for 7,000,000 timesteps, with a timestep of 0.001, using 1024 particles. A four parameter Cross equation is used to fit the data.	169
5.17	Viscosity against shear rate for a series of SLLOD simulations, run at unit temperature. Density was increased from 0.2 to 1.3 gradually. Top: $\rho = 0.9 - 1.3$, bottom $\rho = 0.2 - 0.8$.	171
5.18	Zero-shear rate viscosity dependence on density. The data is taken from extrapolation of figure 5.17.	172
5.19	Viscosity against shear rate for a series of SLLOD simulations, run at unit density. Temperature was increased from 0.2 to 1.6 gradually.	173
5.20	Viscosity dependence on temperature at unit density, the viscosity data is taken from extrapolation of the SLLOD simulations in figure 5.18.	174
5.21	Results from SLLOD simulations run at constant temperature, varying density from 0.7 up to 1.4. $\epsilon = 10.0$.	175
5.21 (part B)	Results from SLLOD simulations run at constant temperature, varying density from 0.7 up to 1.4. $\epsilon = 10.0$. The bottom graph shows the same data on a logarithmic scale.	176
5.22	Viscosity against density. The data is extrapolated from SLLOD simulations, using the same scale as figure 5.18 for comparison.	177
5.23	Results from SLLOD simulations run at constant density, varying temperature from 0.6 up to 2.5. $\epsilon = 10.0$.	178
5.24	Viscosity as a function of temperature. The data was extrapolated from figure 5.22.	178
6.1	A schematic showing the smooth particle model. The yellow particles are static sand particles, the blue particles are fluid particles modelled	182

as SPAM particles, and the red particles are colloids, modelled as coarse-grained discs.

6.2	Snapshots showing the computed particle locations of a binary mixture of 1024 particles (512 fluid, 512 colloid) after 100 and 10,000 timesteps using periodic boundary conditions. $\Delta t = 0.01$. Blue particles are the SPAM fluid, and red particles are colloids.	186
6.3	Computed particle locations of the SPAM fluid under a gravitational force of 0.1. $t = 100$ (left), $t = 100,000$ (right), $\Delta t = 0.001$. The bottom boundary is elastic, and the lateral boundaries are periodic.	187
6.4	Computed particle locations, highlighting a local instability. $g = 2.0$.	189
6.5	Specific deposit as a function of time showing the two linear stages.	191
6.6	Pressure drop as a function of specific deposit, comparing the results from MD and SPAM simulations.	192
6.7	Specific deposit as a function of colloid-colloid interaction strength.	194
6.8	Specific deposit as a function of porosity, comparing the MD and SPAM simulations.	195
6.9	Pressure drop as a function of porosity, comparing the MD and SPAM simulations.	196
6.10	Specific deposit as a function of porosity, comparing the MD and SPAM simulations.	197
6.11	Specific deposit as a function of porosity, comparing the MD and SPAM simulations.	198

List of tables

Table		Page
3.1	Fractal dimension for each of the five clusters. The fractal dimension is the slope of the linear fit shown in figure 3.40.	121
3.2	Fractal dimension for each of the five clusters. The fractal dimension is the slope of the linear fit shown in figure 3.42.	124
6.1	Linear rate constants for the MD and SPAM simulations.	191

1 Introduction

1.1 Background and thesis outline

Clogging is the reduction in the permeability of a porous membrane due to the build-up of deposit within it. It is not a particularly well understood process. The aim of this work is to provide a mechanistic insight into clogging of sand bed filters using a combination of particle and continuum mechanics.

Unfortunately, it is difficult to probe the mechanisms experimentally due to difficulties with non-invasive imaging of particle deposits within a sand bed to high enough resolution. This makes computer modelling an important technique in helping to gain an insight that can be verified experimentally. Models based on colloid filtration theory exist, but are largely empirical with little mechanistic basis, particularly in relation to the dependence of the pressure head-loss on the specific deposit. It is predicted that the flow rate, the nature of the colloid-colloid and colloid-sand interactions, and the fractal dimension of the deposit will be contributing factors.

Chapter one introduces the processes of the Site Ion-Exchange Plant (SIXEP) and explains why the problem of clogging of sand bed filters is of importance to the nuclear waste cycle. It then reviews the current literature around filtration science discussing the different transport mechanisms that bring colloidal particles into contact with sand grain particles, the reasons that colloids are deposited, and the time dependent effects of this deposition. It highlights existing top down modelling work and experimental techniques that can be used to test hypotheses. It also provides a brief introduction to the modelling techniques the rest of this work will make use of.

Chapter two discusses the first modelling technique used in this work: molecular dynamics (MD). The fundamentals of molecular dynamics are explained, which provide the basis for the first model developed.

Chapter three develops the molecular dynamics model, giving details surrounding the decisions taken in creating the model. Data from the literature and from existing models is then used to verify that the model, at least qualitatively, gives realistic and reliable results. A systematic exploration of the parameter space is then undertaken, resulting in a series of mechanistic hypotheses. The effect of the pairwise

potentials, the density of the sand, the concentrations of colloids, and the fractal dimension of deposits are investigated.

Chapter four discusses the second modelling technique used in this work: smooth particle applied mechanics (SPAM). The fundamental of continuum mechanics are first discussed, then the intricacies of SPAM are examined. It also highlights the need to define constitutive relations.

In chapter five, the details on how to parameterise a continuum scale model using pseudo-experimental data from molecular simulations are discussed. An equation of state and the dependence of viscosity on density/temperature is determined for two potentials.

In chapter six, the parameterised SPAM model is used to perform a similar systematic exploration to chapter three. The continuum scale model is compared with the molecular dynamics model, allowing for a direct comparison of the modelling techniques.

1.2 SIXEP

The first 11 nuclear power plants built in the United Kingdom were Magnox plants and, as of 2015 when the final plant (Wylfa in Anglesey) was shut down, none of these remain operational.¹ The term Magnox originates from the magnesium-alloy used to clad the fuel rods within the reactor. The legacy waste from these plants accounts for the vast majority of waste remaining in the UK. Due to the low burn-up of these 1st generation plants the majority of the fuel was reprocessed to extract usable uranium, which at the time was in short supply. Reprocessing was also required to extract plutonium from Magnox stations that were built for the sole purpose of generating material for weapons programmes in the UK and US. Reprocessing operations were found to generate a large volume of radioactive effluent. By the 1970s, discharge and dispersal of this effluent into the Irish Sea was no longer an acceptable option.

The Site Ion Exchange Plant (SIXEP) at Sellafield reprocessing site is a plant designed with the aim of reducing the amount of caesium and strontium discharged to the Irish Sea from pond storage water used to store spent Magnox fuel prior to treatment. Since its opening in 1985, it has helped to vastly reduce the environmental impact of the Sellafield site (figure 1.1). Several steps are taken in order to try to limit the amount of corrosion of the fuel rods, and therefore limit the amount of strontium and caesium released to the pond water:²

- The pH is maintained at 11.5 by dosing with sodium hydroxide
- The temperature of the pool is maintained using refrigerant coolers
- The presence of non-radioactive ions is kept as low as possible




Figure 1.1: Discharge of radiation to the Irish Sea. SIXEP was opened in 1985.³

Even with these steps in place there is still a degree of corrosion that takes place, meaning subsequent measures must be taken to remove the radioactive isotopes prior to discharge to the sea. SIXEP uses three main processes to achieve this goal: sedimentation, sand bed filtration, and ion-exchange, outlined in figure 1.2.

The initial sedimentation step removes any particulates within the waste greater than approximately 10 μm , leaving an effluent waste containing other colloids and residue. The ion-exchange process must be done at a neutral pH as alkaline conditions would dissolve the zeolite catalyst. The magnesium colloids remaining in the waste are soluble in neutral conditions and would blind the ion-exchange column if present. Therefore, there is a prior sand-bed filtration process to remove any colloids suspended in the effluent waste, which at high pH are insoluble.




Figure 1.2: Schematic showing a simplified flow diagram for SIXEP.

When the plant was designed, a research programme was undertaken to optimise its efficiency. This included selecting the material to be used in the synthetic ion-exchangers and testing various polyelectrolytes to determine their effect on the process of filtration. The aim when designing the sand bed filters was not only to remove the highest percentage of solids from the waste, but also to create the longest time between backwashes; as colloids deposit within the filter bed the permeability decreases, eventually leading to a pressure drop (and reduction in performance), requiring the filter to be backwashed with water to remove the sludge before it can be used again. The research programme concluded that a low molecular weight polyelectrolyte such as Nalfloc N7607 or Magnafloc 1597 should be used to enhance the sticking of colloid particles to the sand bed.²

The plant has operated successfully and as effectively as the design intent since opening. However, due to closure of other plants on site and a change in the nature of the effluent waste being sent to SIXEP, it is becoming increasingly important to ensure that there is sufficient knowledge available to be able to accurately quantify the risks associated with changes to the waste management cycle. Future variations in the feed stream could include:⁴

- Sand and clinoptilite fines
- Corrosion products from sludge
- Miscellaneous beta gamma waste
- Skip decontamination washings
- Algae

and it is not currently known how such changes in feed stream would affect the processes at SIXEP. Morris⁴ highlighted several key areas where knowledge gaps exist associated with the SIXEP cycle, and in particular the process of sand bed filtration, and was summarised by Bridge⁵: “there has been to date a substantial need for a more detailed foundational understanding of the physics and chemistry of colloid-sand interactions within deep bed filtration, and the parameters which control these.” This works aims to help address this need.

1.3 Modelling techniques

The complex nature of how and where colloids deposit within a filter makes predicting when a filter will become clogged, and how this will vary with differing feed streams, a difficult problem. Over the last forty years there have been several mathematical, experimental, and computational investigations into the phenomena of clogging. Jegatheesan and Vigneswaran⁶ recently reviewed the existing modelling work in the area, summarising a variety approaches that have been used to tackle this problem. The techniques were split into two broad categories: microscopic and macroscopic. Macroscopic models consider the cumulative collection of deposits and the resulting time-dependent effects, whereas microscopic models consider individual particle interactions. Similarly, Tien and Payatakes⁷ further split the research into two distinct categories: a phenomenological approach and a theoretical approach. The phenomenological approach describes the behaviour of deep bed filter using a set of partial differential equations, where the model parameters are required from prior experimental work. This type of model can be used for the basis of design and scale-up, but gives very little mechanistic information; it is a pragmatic approach that requires little fundamental basis, but gives good results. The theoretical approach, however, aims to derive mathematical formulae that describe the dynamics of deep bed filtration based on the nature of the filter-fluid and filter-colloid interactions. This

section will outline some of the key theories that define the current knowledge surrounding the processes of filtration and clogging.

1.3.1 Microscopic approach

The microscopic approach investigates the interactions between a single collector (a sand grain when modelling a sand bed filter) and a single colloid and ignores any time-dependent effects. The aim is to understand the reasons that a colloid travels close enough to a collector to deposit, and the forces that control the likelihood of deposition occurring.

Colloid filtration theory

O'Melia *et al*⁸, who formed the basis of colloid filtration theory, suggested there are three mechanisms by which a suspended colloid particle can come into contact with a collector (a sand grain, in the case of a sand bed): interception, sedimentation, and diffusion, outlined in figure 1.3. It assumes steady, saturated flow conditions. A particle following the streamline may collide with a collector simply as a result of the size of the colloid, resulting in interception. If the density of the suspended particle is different to that of the water, then other phenomena, such as fluid drag and buoyancy, will affect the trajectory, resulting in sedimentation. Finally, a colloid particle is also subject to a series of random collisions with the fluid and other suspended particles, resulting in Brownian motion and diffusion. In deep bed filtration, this process is statistical; the more collectors that a colloid particle flows past, the more likely it is to be captured. The size of the colloids and the viscosity of the water are key to determining which process is dominant.




Figure 1.3: Basic transport mechanisms, adapted from O'Melia⁵. A is sedimentation, B is interception, C is diffusion.

These three processes combine to give the single-collector efficiency, η_0 , which describes the ratio between the total number of particles touching the collector surface and the total flux past the collector:

$$\eta_0 = \eta_i + \eta_s + \eta_d \quad (1.1)$$

where η_i is the number deposited by interception, η_s is the number deposited by sedimentation, and η_d is the number deposited by diffusion. Levich⁹ and Spielman¹⁰ showed that analytical expressions for each of these can be obtained in conditions where colloid-collector interactions are favourable or when they are repulsive. The efficiency of the collector, in terms of retention, would then also depend on the ratio of collisions resulting in the colloid being removed from suspension:

$$\eta = \alpha\eta_0 \quad (1.2)$$

where α is the attachment efficiency (the number of collisions resulting in deposition compared to the total number of collisions). If $\alpha = 1$ then colloid deposition through colloid filtration theory can be calculated exactly.

This model of colloid retention neglects the fact that both the porosity and sand grain-sand grain junctions also contribute to the likelihood of colloids being retained. In fact, in conditions where a collision with a collector resulting in sticking is unlikely, then the rate of deposition by these three methods is going to be low. Bradford *et al*¹¹ suggested that mechanical filtration and straining will be the dominant mechanisms for colloid deposition in such conditions. Mechanical filtration is the retention of colloids that are larger than the pores in the filter bed; the particles cannot physically pass through the bed. The deposited particles form a filter cake on top of the bed that grows with time and decreases permeability, however this is a phenomenon associated with membrane filtration rather than deep bed filtration. Straining, shown in figure 1.4, occurs at grain-grain junctions within the bed, and only occurs at a fraction of the pore space (as opposed to mechanical filtration); it is dependent on the ratio of the colloid size and the pore size.




Figure 1.4: An illustration of straining, adapted from Bradford *et al*¹¹.

Colloid-collector interactions

It is not necessarily true that removal of a colloid from a suspension is irreversible; it is a balance between the attractive and physical forces holding the colloid in place and those dragging the colloid away. Therefore the forces that control this balance must be considered. Bradford and Torkzaban¹² explain how the interaction energies between colloid and collector particles can be calculated using Derjaguin-Landau-Verwey-Overbeek (DLVO) theory, which calculates the interaction energy, ϕ_{total} , as a sum of the electrostatic, ϕ_{el} , and van der Waals forces, ϕ_{vdW} :

$$\phi_{total}(r) = \phi_{el}(r) + \phi_{vdW}(r) \quad (1.3)$$

where r is the distance between the colloid and the collector. Elimech *et al*¹³ have recorded expressions for both ϕ_{el} and ϕ_{vdW} that can be used to calculate DLVO interaction energies for specific interactions. In conditions where the collector and the colloid share the same charge then there will be a repulsive electrostatic interaction, though at an energy minimum defined by the van der Waals forces, weak colloid retention can occur. Weak retention is categorised by Bradford¹² as occurring when only a fraction of colloid collisions result in deposition, and where a fraction of those attached may detach due to an increase in kinetic energy. In the opposite conditions, where colloid-collector electrostatic interactions are favourable then it is likely that contact with the collector will result in deposition.

It is clear that, in addition to the DLVO interactions, other forces, such as hydrodynamic forces, capillary forces and those resulting from flow velocities will also play a role in the mechanism of deposition. However, Bradford also suggests that “at present, non-DLVO interactions are incompletely understood and quantitative theory has not been developed or is not generally accepted.”

1.3.2 Macroscopic approach

The theories discussed previously do not consider any time-dependent effects; it is assumed that there is always surface area available for deposition, however, this is not the case. The macroscopic investigations discuss the build up of deposit (by any mechanism) with time, and how this affects the dynamics.

Blocking/ripening

Camesano *et al*¹⁴ suggest there are two distinct mechanisms by which colloids deposit with time: blocking and ripening. Blocking occurs when deposited colloids fill up available retention sites, therefore reducing the available space for further deposition. The volume of reduction depends on both the colloid size and the interactions it has with other colloids. With time, the rate of deposition on a given colloid reduces. A dynamic blocking function can be used to describe the available surface area for deposition:

$$B(\theta) = \frac{\theta_{max} - \theta}{\theta_{max}} \quad (1.4)$$

where θ is the fractional surface coverage and θ_{max} is the maximum possible surface coverage. Ripening is in fact the opposite process; deposited colloids provide additional surface area for further deposition. This is as a result of favourable colloid-colloid electrostatic interactions. The rate of deposit for a given collector increases with time. The cluster grows continues to grow with time, causing the pore size to become vanishingly small, resulting in mechanical filtration.

Concentration with depth

Considering the filter as a whole (as opposed to the growth of an individual cluster), it is expected that the concentration of deposited colloid particles will decrease exponentially with depth.⁶ However, in conditions where there are multiple types of colloid particle, or a large particle size distribution, specific colloids can be preferentially deposited at the top of the sand bed resulting in a hyper-exponential decay. Additionally, if there are only weak forces (such as those predicted under certain conditions by DLVO theory) there is the potential for deposited colloids to remobilise and deposit further down the bed resulting in a non-monotonic profile. Figure 1.5 outlines these profiles.




Figure 1.5: Illustrations of deposit profiles as a function of depth.

Left – exponential, middle – hyperexponential, right – non-monotonic. The dotted line shows the expected linear dependence.

Experimental work performed by Veerapaneni¹⁵ suggested four observations when deposition is favourable within the bed:

- For a given specific deposit, a larger head loss is observed when the fluid velocity is lower.
- Clogging does not depend primarily on porosity reduction. With a head loss increase of 2-3 times, specific deposit only accounts for less than 1% of the filter volume at a porosity of 40%.
- When clogging is low, head loss scales linearly with specific deposit.
- When clogging is high, head loss scales quadratically with specific deposit.

Clogging

As colloids deposit within the bed there is a reduction in permeability caused by clogging of the filter. This results in a pressure drop and a loss in performance. Clearly, the mechanisms by which colloids deposit within a filter vary extensively based on a wide variety of parameters, meaning that there is no overriding relationship between the amount of deposit within a filter and the pressure drop that results from it. Mays¹⁶ suggests that the rate of clogging does not depend solely on the physical and chemical effects of the deposited colloids but also on the morphology of the deposits. He proposes the fractal dimension as a means of quantifying this:

$$M \propto L^D \quad (1.5)$$

where M is the deposit mass, L is its characteristic size (or radius of gyration) and D is the fractal dimension. It is theorised that deposits with a high fractal dimension, resulting from conditions where hydrodynamic forces dominate, leave colloid-free flow paths. Where colloid-colloid or colloid-collector interactions are dominant then it is likely that the deposits will have a lower fractal dimension, causing more disruption to the flow (figure 1.6)




Figure 1.6: Illustrations of deposits with a low (left) and high (right) fractal dimension, adapted from Mays.¹⁶ The white circles represent sand particles and the red spheres represent colloid particles. The fractal dimension is illustrated by the clustering of the colloids.

O'Melia and Ali¹⁷ parameterised a model based on filtration experiments with the aim of predicting head loss from a given specific deposit (specific deposit defines the ratio of deposited colloids to the volume of the filter). The model begins with Darcy's Law (or the Kozeny-Carman equation^{18,19}):

$$u = -\frac{k\rho g \Delta H}{\mu \Delta x} \quad (1.6)$$

where u is the approach velocity, k is the permeability, ρ is the density, g is the acceleration due to gravity, μ is the viscosity, ΔH is the observed head loss (the difference in pressure between the top and bottom of the filter), and Δx is a distance in the direction of the fluid flow. The permeability depends inversely on the surface area within the filter:²⁰

$$k \propto M^{-2} \quad (1.7)$$

where M is the specific area (surface area per bed volume). The ratio of head loss to clean bed head loss is therefore:

$$\frac{\Delta H}{\Delta H_0} = \frac{k_0}{k} = \left(\frac{M}{M_0}\right)^2 \quad (1.8)$$

where k_0 and M_0 are the clean bed permeability and specific area. For a clean bed, the specific area is defined by:

$$M_0 = \frac{6(1 - \epsilon)}{d_c} \quad (1.9)$$

where ϵ is the porosity and d_c is the diameter of the spherical collectors. With deposition, the surface area of the collector increases:

$$A_c = \pi d_c^2 + \beta' N A_p^2 \quad (1.10)$$

where A_c is the surface area per collector, β' defines the fraction of retained particles contributing to the increase in area, N is the number of particles per collector, and A_p is the surface area per particle. The constant bed volume per collector is given by:

$$V_c = \frac{\left(\frac{\pi}{6}\right) d_c^3}{(1 - \epsilon)} \quad (1.11)$$

Combing equations 1.6 and 1.7 gives the specific area for the clogging filter:

$$M = 6(1 - \epsilon) \left(\frac{1}{d_c} + \frac{\beta' N A_p^2}{\pi d_c^3} \right) \quad (1.12)$$

Substituting equations 1.5 and 1.8 into 1.3 gives:

$$\frac{\Delta H}{\Delta H_0} = \left(1 + \frac{\beta' N A_p^2}{\pi d_c^3} \right)^2 \quad (1.13)$$

The number of attached atoms relates to the specific deposit through:

$$N = \frac{\pi d_c^3}{6(1 - \epsilon)V_p} \sigma \quad (1.14)$$

where V_p is the volume per particle and σ is the specific deposit. This leads to the final relationship:

$$\frac{\Delta H}{\Delta H_0} = (1 + \gamma\sigma)^2 \quad (1.15)$$

where:

$$\gamma = \frac{\beta' d_c A_p}{6(1 - \epsilon) V_p} \quad (1.16)$$

Figure 1.7 shows the results of their model compared to the experimental data used to fit it. They concluded that the size and concentration of the suspended colloids were the most important physical variables influencing the filtration process. As Wingert *et al*²¹ note, this model has very little predictive power; the empirical nature of the clogging parameter and the collector efficiency cannot be easily transposed from one case to another. This top down approach does not allow for prediction of how changes in chemistry (or indeed feed stream) would affect the filtration process.

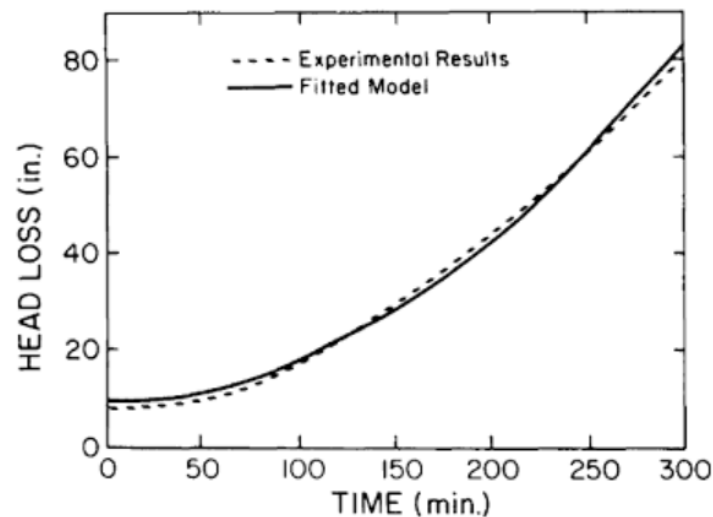


Figure 1.7: Head loss against time for the fitted model and experimental data. Results obtained by O'Melia and Ali.¹⁷ The fitted model is based on equations 1.15 and 1.1.6.

Mays and Hunt²² expanded on this work, revealing a power law correlating the clogging parameter and the Peclet number, which is a function of the flow rate:

$$\gamma = 10^6 N_{Pe}^{-0.55 \pm 0.09} \quad (1.17)$$

$$N_{Pe} = \frac{u d_c}{D_p} \quad (1.18)$$

$$D_p = \frac{kT}{3\pi\mu d_p} \quad (1.19)$$

where N_{Pe} is the Peclet number, u is the flow velocity, d_c is the diameter of the collector, d_p is the diameter of the particle D_p is the particle diffusivity, k is Boltzmann's constant, T is the temperature and μ is the dynamic viscosity. Still, the problem remains that clogging parameter is empirically determined, and is not based on any fundamental physics or chemistry.

1.3.3 Modelling colloid transport

The processes considered so far are zero-dimensional; they do not describe the movement of particles through a filter, rather they describe an instantaneous situation. An advection-dispersion equation (ADE) is typically used as an equation of motion to describe the rate of change of colloids in the mobile phase as a function of the rate of deposition, the rate of advection through the pores, and the rate of dispersion in the pore space.

One of the first models in this area was developed by Herzig *et al*²³. A series of traditional mass balance equations were used to describe the rate of deposition at a given depth in the bed with time. This mass balance is a function of the probability of particle sticking per unit depth, through the “filter coefficient,” where the coefficient itself changes with time. Both the pressure drop and filter coefficient were expressed as ratios of the instantaneous value and its initial value, where empirical data was used

to fit the equations. Walata *et al*²⁴ later linked the filter coefficient ratio and the pressure drop ratio to the amount of deposit within the bed. Although these models show the trends of the amount of deposit and the pressure drop with time, they do not provide mechanistic insight into clogging, and require prior experimental work to define the parameters used in the expressions; they are not predictive.

As an improvement on this, Tien *et al*²⁵ suggested that the deposition process should be split into two phases. The initial phase is the blocking/ripening phase, where the colloids form a layer around the collectors. This was modelled using colloid filtration theory and the sphere in a cell model.²⁶ The second phase is of cluster growth, where the pore space reduces with time, leading to an increase in pressure drop, and was modelled on flow through a constricted tube.²⁷ The model required the porosity of the deposit and the amount of deposit at which the mechanisms swap to be experimentally determined.

More recent examples include models by Bradford *et al*²⁸ and Yuan *et al*²⁹. The model consists of two separate regions: the bulk aqueous phase and the solid-water interface (SWI), and is illustrated in figure 1.8. Colloids are transported through region 1 using an ADE. First order kinetic expressions define the rate of mass transfer between the two regions, and the flow through region 2 is again defined by an ADE (although at a much lower velocity). Only a small fraction of the particles in region 2 will interact with the solid phase at a given time, and a fraction of these will deposit. Immobilised colloids in the solid phase fill up possible retention sites with time.

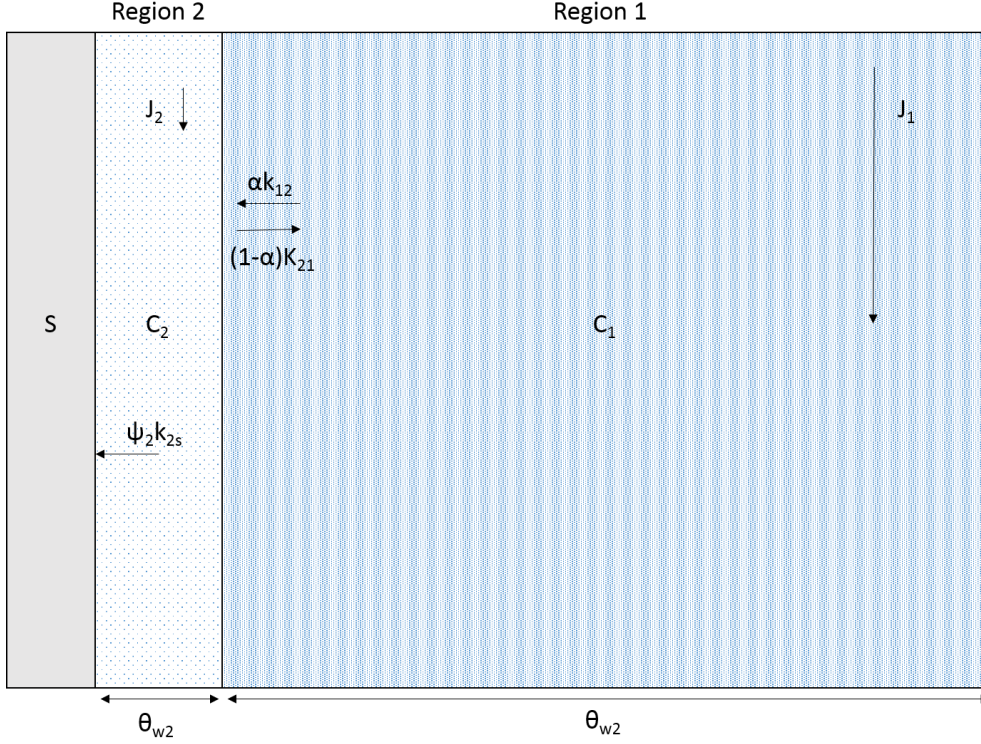


Figure 1.8: Schematic of Bradford's²⁹ model. Region 2 is the SWI and region 1 is the bulk aqueous phase. J is the colloid flux, S is the colloid concentration in the solid phase, C is the suspended colloid concentration, k is a rate constant and α is the colloid attachment efficiency.

The ADE can be written as:

$$\frac{\partial c}{\partial t} + \frac{\partial(\rho S)}{\partial t} = \frac{\partial}{\partial x} \left(D_c \frac{\partial c}{\partial x} \right) - \frac{\partial(q_c c)}{\partial x} \quad (1.20)$$

where D_c is the dispersion coefficient, q_c is the volumetric water constant and ρ is the sand bulk density. The terms on the left hand side represent the change in concentration of suspended and deposited colloids with time, and the terms on the right hand side describe the dispersive and advective transport of the colloids through the filter. DLVO theory (see earlier) is used to describe the suspended concentration:

$$\frac{\partial(\rho S)}{\partial t} = Bk_{ac}C - \rho k_{dc}S + \varphi_{str}k_{str}C \quad (1.21)$$

where B is the dynamic blocking function (eq 1.4), k_{ac} and k_{dc} are the attachment and detachment rates, φ_{str} describes straining as a function of depth and k_{str} is the first order rate of straining. The attachment rate is described through colloid filtration theory:

$$k_{ac} = \frac{3(1 - \varepsilon)}{2d_g} \alpha \eta_0 v_c \quad (1.22)$$

where ε is the porosity of the sand bed, d_g is the sand grain diameter, η_0 is the single grain collector efficiency (described by equation 1.1), and v_c is the colloid velocity. The rate constant defining the mass transfer between region 2 and the solid phase is defined based on the current filtration state:

$$k_{12} = k_r \text{ and } k_{21} = 0 \text{ when } 0 < \sigma \leq \sigma_r \quad (1.23)$$

$$k_{12} = k_{ac} \text{ and } k_{21} = k_d \text{ when } \sigma_r < \sigma \leq \sigma_u \quad (1.24)$$

$$k_{12} = 0 \text{ and } 0 = k_d \text{ when } \sigma = \sigma_u \quad (1.25)$$

where σ is the specific deposit, k_r is the attachment rate during the initial ripening stage, k_{ac} is the attachment rate during the operable stage, and k_d is the detachment rate during the operable stage. Equations 1.15-1.19 are used to define the pressure drop based on the amount of deposit. This type of model is used widely in the water, filtration, and aerosol industries as a method of scaling from laboratory experiments to operational filters. Still, the problem persists that many of the parameters are required from existing literature or from experimental work; the attachment efficiency, collector efficiency, attachment and detachment rates, threshold specific deposits and the clogging parameter must all be specified. Additionally, the model gives no

mechanistic information describing the morphology or locations of deposited particles. In order to create truly predictive models, a more fundamental approach is needed.

1.3.4 Modelling colloid interactions

There are numerous methods that have been used to model colloid interactions in suspension. Due to the fact that colloids often move on a timescale orders of magnitude slower than that of solvent particles, modelling colloid particles in suspension is a difficult task. Therefore, truly atomistic molecular dynamics simulations may not be the ideal approach.

Dissipative particle dynamics (DPD) has often been used to model colloidal behavior with success. Particles trajectories are defined by Newton's second law of motion, where a dissipative force proportional to particle velocities and a random force create the total force on each particle. E. Boek *et al*³⁰ showed that a colloid suspension can be modelled as a system of large colloidal particles in a liquid of interacting point particles, where DPD defines the interactions. They were able to detect characteristic shear-thinning behaviour when applying a steady-shear rate, and suggested this method as a useful technique in studying the rheology of particulate suspensions.

Dzwiniel *et al*³¹ showed how a combination of DPD and molecular dynamics (MD) can be used to realistically model colloidal behaviour. They used a two-level model, where colloids interacted with other colloids, and with solvent particles, through a combination of DPD and MD potentials, where each interaction type (colloid-colloid, solvent-solvent, and colloid-solvent) was treated with discrete parameters. The MD potential used was a Lennard-Jones potential. They were able to recreate 2D hexagonal colloidal arrays that can be seen in experiments, and suggested that a generalization of this combination, called the Fluid Particle Model, could be used to increase the length scales of simulation.³²

Horbach and Frenkel³³ suggested Lattice-Boltzmann as a method for modelling colloid transport. They argue that this is an improvement on modelling colloids using a Yukawa-like potential as it uses Navier-Stokes equations to model the hydrodynamics of the system, in addition to the interaction-dependent properties. Similarly, Padding and Louis used coarse-grained molecular dynamics and stochastic

rotation dynamics to successfully model the effects of Brownian motion and hydrodynamics forces in colloidal suspensions.³⁴

1.3.5 Experimental techniques

There have been several experimental techniques, both on a micro-scale and on a meso-scale, used to either try to investigate the mechanisms involved in particle deposition or to help validate or parameterise models.

Bridge *et al*³⁵ used time-lapse fluorescence imaging as a non-invasive means of imaging colloid and fluid transport to a millimeter level. In this technique, suspended tracers (particles tagged with fluorescent molecules) are pumped through a small-scale sand filter. A camera with a filter fitted to the emission wavelength is used to capture images of the fluorescent molecules, allowing the visualization of the trajectories and locations of the deposited particles. Both the feed stream and the sand bed can be treated to provide different electrostatic conditions. They were able to measure the effect of changing the pH, flow rate, and ionic strength, and even calculate an attachment efficiency based on the retained mass profiles.

In a similar vein, Zevi *et al*³⁶, have used confocal microscopy to image sand grain-grain junctions and the methods of colloid deposition within them. They were able to obtain images containing 512 square pixels, representing 374 μm^2 . They concluded that the technique is a viable method of quantifying the number of colloids passing a collector in a given time and the number of colloids deposited by different mechanisms.

Various other imaging techniques are discussed and analysed by Ochiai *et al*³⁷, and they highlight the current reliance on the porous media being “ideal” such that it does not represent natural colloids or porous media. They conclude that “a key challenge that remains is the development of tools to investigate transport of natural colloids in natural porous media.”

1.4 Mechanics

There are two broad categories that underpin the work undertaken throughout this project: atomistic simulation and continuum simulation. This brief section will introduce and outline the methods, which are discussed in detail in chapters 2 and 4.

Matter can be described using two sets of fundamentally different mechanics¹. Particle mechanics is used to describe the interactions between individual particles, whereas on a larger scale, continuum mechanics uses assumed relationships to describe the continuum nature of matter.³⁸ The methods differ significantly in their assumptions, the nature of the results obtained, and, usually, in the type of algorithm used to solve problems.

Particle mechanics, as the name suggests, evolves the trajectories of a number of mass points in order to obtain information about a system. The ordinary differential equations (ODEs) that define the classical equations of motion are solved over time. A force law defines the interactions between particles, usually, due to difficulties calculating higher order interactions, between pairs of particles, which in turn, define their motion. This means there is no need for any additional information, such as transport laws or an equation of state; this is all inherent within the dynamics. This makes it a particularly useful technique when trying to determine transport properties such as heat flux or viscosity, as they can be easily calculated as statistical averages of functions of particle momenta and coordinates. The main limitations of particle mechanics are associated with the length scales that can be readily reached and the accuracy of the force law.

Continuum mechanics, however, makes use of pre-defined relationships where particles and force laws have no role. Instead, partial differential equations (PDEs) evolve the state of the continuum, and the fluxes of energy, mass and momentum within it. It is based on the fundamental assumption that the macroscopic appearance of a system is continuous, and that a system's properties vary gradually. The combination of the assumed relationships and the conservation laws form the partial differential equations that have an infinite number of degrees of freedom, which, when simulating, need to be reduced to a finite number. The PDEs are sufficient to evolve a system when constitutive relationships and a series of boundary conditions are defined. Microscopic fluctuations are not present in continuum mechanics, and it cannot be used for determining transport

¹ When not considering quantum mechanics

properties. Instead, once the appropriate continuum descriptions have been obtained from other means, such as from experimental data or from atomistic simulation, continuum mechanics allows simulations to reach much larger time and length scales, and to make contact with experimental data.

The two methods can indeed overlap and agree when the force laws correspond to the continuum equations, and when microscopic fluctuations can be ignored (as they are absent in continuum mechanics but not particle mechanics). Furthermore, the two approaches can be related by using particles to represent a continuum; smooth particle applied mechanics (SPAM) uses a series of particles that represent the fluxes of the conserved quantities to solve problems in continuum mechanics. The similarity between the algorithms used in MD and SPAM is an area that this work aims to make use of by using MD as a means of not only validating, but also parameterising, SPAM simulations.

1.5 Summary

Filtration is a complex process that is not easily predictable. There are several transportation methods via which a colloid may end up in contact with a collector, and the chance of the colloid being retained is a delicate balance of the electrostatic and van der Waals forces holding it in place, and the hydrodynamic forces dragging it away. Regardless of the mechanism by which colloids stick, over time the buildup of deposit causes clogging, a process that has been subject to experimental work resulting in empirical relationships linking the specific deposit, head loss, and the Peclet number through a clogging parameter, γ .

Currently there is a knowledge gap that exists in describing this parameter; it is not based on any fundamental physics or chemistry. This work aims to give a mechanical and physical insight into the clogging parameter through computer simulation by linking the rate of deposit and head loss to properties of the sand bed and the colloid particles themselves. The existing work uses a top down approach; experimental data is used to parameterise models to predict future changes. Instead, this work will make use of molecular and continuum simulation to build models from the bottom up, starting with the fundamental physics of particle trajectories. Additional elements will be added to tune the model towards a filtration process. The mechanisms by which a single colloid can stick are well defined, and described by a

combination of colloid filtration theory, straining, DLVO and non- DLVO interactions, and will be vastly simplified in this work. Instead, the effect of a growing cluster of colloid particles around a collector will be the focus. This work aims to achieve the following:

- Create a coarse grained molecular dynamics model that contains the essential physics of filtration, while simplifying the process of deposition.
- Test whether this model qualitatively agrees with existing literature regarding the rate of build-up specific deposit, increase in pressure drop, and deposit concentration with depth.
- Explore this model by systematically changing the available parameters. The geometry, porosity and lattice structure of the sand particles can be varied, in addition to the nature of the fluid/colloid/sand interactions, the flow rate and the probability of deposition.
- Test whether deposits created using this model have a fractal dimension, and if so, test the effect of the fractal dimension on the system.
- Parameterise a continuum scale model of the same process using smooth particle applied mechanics. This allows for the time and length scales of simulation to be increased.

1.5 References

- [1] BBC, 2105, *Wylfa: Last day for Anglesey nuclear plant*. [ONLINE] Available at: <http://www.bbc.co.uk/news/uk-wales-north-west-wales-35158939> [Accessed September 2017]
- [2] M. Howden, *British Nuclear Fuel's site ion exchange effluent plant*, Inst. Chem. Eng. Symp. Ser. **77**, 253 (1983)
- [3] S. Owens, M. Higgins-Bos, M. Bankhead, J. Austin, NNL Science Issue 3, 4 (2015)
- [4] Y. Morris, *Review of SIXEP Particle Abatement NNL (11) 11425 Issue 1*, Sellafield, Cumbria (2011)
- [5] J. Bridge, *Technical Review of Colloid Transport, Deposition and Remobilisation Processes in Sand Bed Filter Relevant to SIXEP Treatment Processes*, 2012
- [6] V. Jegatheesan, S. Vigneswaran, *Deep bed filtration: mathematical models and observations*, Critical Reviews in Environ. Sci. and Tech., **35**, 515 (2007)
- [7] C. Tien, A. Payatakes, *Advances in deep bed filtration*, AiChE, **25**, 737 (1979)
- [8] K. Yao, M. Habibian, C. O'Melia, *Water and waste water filtration. Concepts and applications*, Env. Sci. Tech., **5**, 1105 (1971)
- [9] V. Levich, *Physical-chemical hydrodynamics*, J. Electrochemical Soc. **110**, 251C (1963)
- [10] L. Spielman, S. Friedlander, *Role of the electrical double layer in particle deposition by convective diffusion*, J. Colloid Interface Sci, **46**, 22 (1974)
- [11] S. Bradford, J. Simunek, M. Bettahar, M. Van Genuchten, *Significance of straining in colloid deposition: Evidence and implications*, Water Research J. **42**, (2006)

- [12] S. Bradford, S. Torkzaban, *Colloid transport and retention in unsaturated porous media: A review of interact-, collector-, and pore-scale processes and models*, *Vadoze Zone Journal*, **7**, 17 (2008)
- [13] M. Elimelech, J. Gregory, X. Jia, R. Williams, *Particle Deposition & Aggregation*, Elsevier, Amsterdam (1995)
- [14] T. Camesano, K. Unice, B. Logan, *Blocking and ripening of colloids in porous media and their implications for bacterial transport*, *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, **160**, 291 (1999)
- [15] S. Veerapaneni, *Formation and morphology of colloidal deposits in porous media*, Ph.D. Thesis, Rice University, Houston, TX, 1996
- [16] D. Mays, *Contrastive clogging in granular media filters, soils, and dead-end membranes*, *J. Environ. Eng.*, **136**, 475, 2010
- [17] C. O'Melia, W. Ali, *The role of retained particles in deep bed filtration*, *Progress in Water Research*, **10**, 167, 1978
- [18] J. Kozeny, *Capillary transport of water in soil*, *Sitzungsber Akad Wiss Wien*, **136**, 271 (1927)
- [19] P. Carman, *Fluid through granular beds*, *Institution of Chemical Engineers*, **15**, 532 (1937)
- [20] D. Mays, J. Hunt, *Hydrodynamics aspects of particle clogging in porous media*, *Environ. Sci. Technol.*, **39**, 577 (2005)
- [21] L. Wingert, N. Bardin-Monnier, A. Charvet, D. Bemer, D. Thomas, *Modelling of the deep granular bed clogging by nanoparticles*, *Separation and Purification Technology*, **176**, (2016)
- [22] D. Mays, J. Hunt, *Hydrodynamic and chemical factors in clogging by montmorillonite in porous media*, *Environ. Sci. Technol.*, **41**, 5666, 2007

- [23] J. Herzig, D. Leclerc, P Le Goff, *Flow of suspensions through porous media – application to deep bed filtration*, *Industrial and Engineering Chemistry*, **62**, 8 (1970)
- [24] S. Walata, T. Takahashi, C. Tien, *Effect of particle deposition on granular aerosol filtration: a comparative study of methods in evaluating and interpreting experimental data*, *Aerosol Science and Technology*, **5**, 23 (1986)
- [25] C. Tien, R. Turian, H Pendse, *Simulation of the dynamic behaviour of deep bed filters*, *AIChE*, **25**, 385 (1979)
- [26] J.Happel, *Viscous flow in multiparticle systems: slow motion of fluids relative to beds of spherical particles*, *AIChE*, **4**, 197 (1958)
- [27] A. Payatakes, *A new model for granular porous media*, *AIChE*, **19**, 58 (1973)
- [28] S. Bradford, J. Simunek, M. Bettahar, M van Genuchten, S. Yates, *Modelling colloid attachment, straining, and exclusion in saturated porous media*, *Environ. Sci & Tech.*, **37**, 2242 (2003)
- [29] H. Yuan, A. Shapiro, *A mathematical model for non-monotonic deposition profiles in deep bed filtration systems*, *Chem. Eng. J.*, **166**, 105 (2011)
- [30] E. Boek, P. Coveney, H. Lekkerkerker, P. van der Schoot, *Simulating the rheology of dense colloidal suspensions using dissipative particle dynamics*, *Phs. Rev.* **55**, 3124 (1997)
- [31] W. Dzwinel, D. Yuen. K. Boryczko, *A Two-Level, Discrete-Particle Approach for Simulating Ordered Colloidal Structures*, *J. of Coll. and Interace Sci.* **225**, 179 (2000)

- [32] W. Dzwinel, D. Yuen, K. Boryczko, *Mesoscopic dynamics of colloids simulated with dissipative particle dynamics and fluid particle model*, J. Mol. Model, **43**, 33 (2001)
- [33] J. Horbach, D. Frenkel, *Lattice-Boltzmann method for the simulation of transport phenomena in charged colloids*, Phy. Rev. E, **64**, 061057 (2001)
- [34] J. Padding, A. Louis, *Hydrodynamics interactions and Brownian forces in colloidal suspensions: Coarse-graining over time and length scales*, Phy. Rev. E, **74**, 031402 (2006)
- [35] J. Bridge, S. Banwart, L. Heathwaite, *Noninvasive quantitative measurement of colloid transport in mesoscale porous media using time lapse fluorescence imaging*, Environ. Sci. Technol, **40**, 5930 (2006)
- [36] Y. Zevi, A. Dathe, B. Gao, B. Richards, *Quantifying colloid retention in partially saturated porous media*, Water Resources Research, **42**, 2006
- [37] N. Ochiai, E. Kraft, J. Selker, *Methods for colloid transport visualization in pore networks*, Water Research Journal, **42**, 2006
- [38] Hoover W. *Smooth Particle Applied Mechanics*. World Scientific, Singapore, 2006

2 Molecular dynamics

Molecular Dynamics (MD) is a computer simulation method for studying the evolution of many body systems using the laws of classical mechanics. It has been an area of active research since the 1950s, and has nearly as long a history as modern computers. Berni Alder pioneered the method when developing techniques to study dense fluids, and it is now accepted that, alongside experiment and theory, computer simulation is a third branch of physics.³

In conventional MD simulations, a series of particles, which are modeled as point objects, move within a simulation cell. The trajectories of these particles are evolved by solving the classical equations of motion, where the force acting on each particle is a vector sum of its neighbours, and any non-conservative forces (such as friction). The pairwise force is calculated from a chosen potential energy function (or force field). Once the total force acting on a particle, and therefore the acceleration it experiences is calculated, the trajectory it would travel on can be predicted.

Once the potential energy function, boundary conditions and initial conditions have been defined, MD is essentially exact. This allows for thermodynamic properties of interest to be calculated from time averages of functions of the positions and momenta of particles, through Boltzmann's statistical mechanics. In order to obtain useful information from a computer simulation, the model must be both accurate in terms of modeling the particle interactions, but also be able to run quickly enough that results can be obtained on a reasonable timescale.

Take, for example, a system of N particles. The entire state of the system is defined by the $3N$ generalised coordinates, \mathbf{q} , and the $3N$ generalised velocities, $\dot{\mathbf{q}}$ (where \mathbf{q} is shorthand for $\{q_i\}$ etc.). The equations of motion are used to relate the current coordinates and velocities to future time steps. The Lagrangian equation of motion, based on Hamilton's principle of least action, is a second order ordinary differential equation:

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{\mathbf{q}}} \right) = \frac{\partial L}{\partial \mathbf{q}} \quad (2.1)$$

where L is the Lagrangian of the system, which is simply the difference between the kinetic energy, K , and the potential energy, U :

$$L(\mathbf{q}, \dot{\mathbf{q}}) = K - U \quad (2.2)$$

Hamilton devised a description of the mechanics of a system of particles (more general than that of Newton) where the momenta, \mathbf{p} , are used to move the system forward in time, rather than the velocity. It gives the same results as Newtonian mechanics (and Lagrangian). In Hamilton's formulation the generalised momenta are defined as:

$$\mathbf{p} = \frac{\partial L(\mathbf{q}, \dot{\mathbf{q}})}{\partial \dot{\mathbf{q}}} \quad (2.3)$$

which are on an equal footing with the generalised positions. The equations of motion are defined as two sets of coupled first order ordinary differential equations

$$\dot{\mathbf{q}} = \frac{\partial H}{\partial \mathbf{p}} \quad (2.4)$$

$$\dot{\mathbf{p}} = -\frac{\partial H}{\partial \mathbf{q}} \quad (2.5)$$

where H , the Hamiltonian, is defined as

$$H(\mathbf{q}, \mathbf{p}) = (\dot{\mathbf{q}}, \mathbf{p}) - L \quad (2.6)$$

Using these equations of motion, having specified the initial \mathbf{q} and \mathbf{p} , the entire phase space trajectory of the system is known, and the equations are themselves time reversible; meaning the history of the trajectory can also be determined from the solution going backwards in time. The use of Hamiltonian and Lagrangian mechanics and general coordinates as opposed to Newtonian mechanics allows a much larger class of problems in mechanics to be solved more easily, particularly those systems involving constraints.

2.1 Initial conditions

In order to evolve the particle trajectories, the initial state or initial conditions of the system must be specified. This means assigning each particle within the system a position and a momentum.

Any convenient set of positions can be used. Typically when simulating a liquid a lattice structure is used, such as a cubic lattice (in three dimensions), or a square lattice (in two dimensions). This is for computational simplicity only, as during the equilibrium phase all memory of the initial structure is lost. Usually, MD simulations are carried out at constant density, as the unit cell size does not change, therefore the initial particle spacing and the cell size are determined by the required density.

The momenta may be chosen so that the system starts at a particular energy or temperature. This can be achieved by sampling from a distribution, where the variance is related to the required temperature, or by running the simulation in the presence of a thermostat until the correct temperature is reached. To stop the simulation cell from drifting away from the origin it is important to rescale the velocities so that there is no initial net momentum in any direction. Finally, the momenta are rescaled to give the starting temperature, as zeroing the linear momentum will change the sum of the squared momenta, from which the temperature is computed.

2.2 Boundary conditions

Molecular dynamics problems require the use of boundary conditions in order to create realistic systems. They are used to describe the behavior of particles at the edge of the simulation cell. At the boundaries, mass, momentum and energy can enter or leave the system, and the location, velocity and temperature of the boundaries can be specified. They create the required conditions for the simulation, not only by defining where the edges of the system are, but also, in some cases, by affecting the dynamics and the flow. Several different types of boundary condition can be used in molecular dynamics, and are used throughout this work.

2.2.1 Periodic boundary conditions

Periodic boundary conditions, first introduced by Born and Von Karman in 1912⁴, are used to allow the simulation of a bulk material by creating a pseudo infinite system (usually) free from edge effects. In MD, the simulation time is proportional to the number of atoms squared, (unless special measures are taken – see later), therefore one is restricted to a moderate value for the number of atoms. Without periodic boundary conditions, the simulation cell would have a high proportion of particles at the edge, due to the restriction on the total. Consider a two-dimensional cell containing 100 atoms in a 10 x 10 grid; over one third of the particles would be at the edge. When simulating a fluid the particles at the surface are not of interest, instead, the properties of the bulk are being investigated. In most conditions, replication of the simulation cell through space removes the edge effects, however it remains possible for longer range hydrodynamic effects to persist across the boundary.




Figure 2.1: An illustration of periodic boundary conditions in two dimensions. The red circles represent the starting positions, and the blue circles the final positions. As a particle leaves the central simulation cell it is replaced by its periodic image.

Figure 2.1 illustrates two-dimensional PBCs. The simulation box is the central cell. A particle leaving the central cell and entering the cell below it is replaced by its periodic image entering the central cell from above. The length of the cell must be longer than the range of the interatomic potential being used to make sure that a particle does not interact with its own periodic image. It is important to note that PBCs impose an artificial periodicity on the system, which can have unintended consequences when calculating transport coefficients. In this example, the simulation cell is a square, though this is not strictly necessary; alternative PBCs include hexagonal, octahedral etc.

2.2.2 Minimum image convention

The force resulting from the pairwise interatomic potential on each particle is a vector sum of pairwise interactions between itself and all periodic images. As the PBCs create an infinite number of particle neighbours, the potential needs to be truncated. The minimum image convention is used to limit the number of particle interactions to a computable number. Here, each particle in the primary cell is considered in turn. The particle of interest is considered to be at the centre of a box with the same dimensions as the primary cell and the pair forces between it and all other particles that lie within this box are calculated, including those from periodic image particles. The minimum image separation, r_{ij}^m between two particles, i and j , is:

$$r_{ij}^m = r_{ij} - n L \quad (2.7)$$

where n is the is number of box lengths that particle i is from particle j (must be an integer), and L is the box length. To additionally limit the number of pair force calculations, the potential is usually truncated at some cut-off point within the unit cell. Potentials that fall off slowly with distance require larger truncation radii than those that fall off more quickly.




Figure 2.2: An illustration of the minimum image convention. Each particle from the central cell is considered in turn. It is placed at the center of a box with the same dimensions as the cell. The force between itself and all other particles within this cell is calculated, shown by the coloured particles. The circle represents the potential cutoff.

2.2.3 Other boundary conditions

Other types of useful boundary conditions include elastic boundaries, stone wall boundaries, and mirror boundaries. Elastic (or rigid) boundaries, shown in figure 2.3, are used when the boundary is a physical boundary fixed in space. They are implemented by either having a rigid wall with which particles elastically collide, where kinetic energy is conserved (their velocity in the appropriate direction is reversed upon impact), or by having a series of particles at the boundary with a steep repulsive potential. Both of these methods prevent particles from passing through the boundary, and are useful when modelling systems with a physical boundary, such as the flow of a fluid through a porespace.

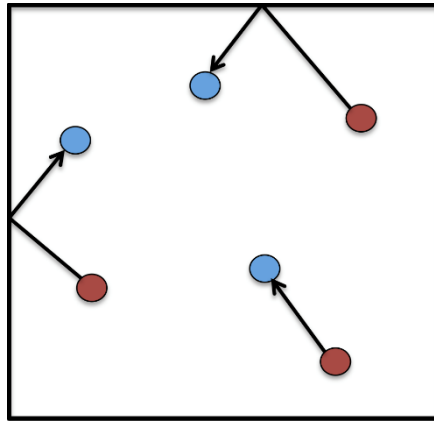


Figure 2.3: An illustration of elastic boundary conditions. The particle's perpendicular component of velocity is reversed upon impact with the wall. The red circles represent the starting positions, and the blue circles the final positions.




Figure 2.4: An illustration of mirror boundary conditions where the dashed line is the mirror boundary. The particles interact with a mirrored version of themselves. The red circles represent the simulation positions, and the blue circles the mirror particles.

Mirror boundaries, shown in figure 2.4, work by having a mirror image of particles near the boundary on the other side with which they can interact. The mirrored particles can move or be static. This allows properties to be continuous across the boundary, giving rise to such phenomena as heat flow through the fluid properties. Mirror boundaries are useful in inhomogeneous models, such as modelling the interaction between a liquid and a heated surface.⁵

In confined fluid problems, the most common boundary condition is the stonewall (or non-slip) boundary. When a particle comes into contact with the boundary its normal velocity is set to zero, and its tangential velocity is set to the velocity of the wall. This is illustrated in figure 2.5, and is used when modelling Poiseuille flow.




Figure 2.5: An illustration of stone wall boundary conditions. The perpendicular component of velocity is set to zero, the parallel component of velocity is set to the velocity of the wall upon collision with the wall. The red circle represents the starting positions, and the blue circle the final positions.

There are, of course, many other types of boundary condition. It is sometimes necessary to use boundary conditions that, rather than representing bulk properties in a confined space, create a certain set of conditions. The Lees-Edwards boundary conditions (see section 5.3.2) are an example of this, where they are used to create a shear flow.

2.3 Integration schemes

In all but a very few trivial examples, numerical methods are required to integrate the equations of motion of a many body system. Numerical integration schemes are used to reduce a derivative to a finite difference quotient; formally it is a truncated Taylor Series. The aim, given a set of variables at the current time, t , is to obtain the same set of variables at a future time, $t + \delta t$. A Taylor series takes the form⁶:

$$f(t + \delta t) = f(t) + f'(t)\delta t + f''(t)\frac{\delta t^2}{2!} + f'''(t)\frac{\delta t^3}{3!} + \dots + f^n(t)\frac{\delta t^n}{n!} \quad (2.8)$$

There are a multitude of integration schemes available for solving problems in particle mechanics, with varying accuracy and computing power required. Methods are either self-starting, where trajectories can be evolved with only information from the current time step, or not self-starting, where data from previous or other time steps is required. The timestep, δt , will vary with the algorithm used; generally, a smaller timestep will increase accuracy but also increase simulation time. The timestep should be much shorter than the time it takes a particle to travel its own length.

2.3.1 Euler algorithm

The Euler method is the simplest method of numerical integration. It is a self-starting method that assumes that the velocity is constant over the timestep⁷:

$$\mathbf{r}(t + \delta t) = \mathbf{r}(t) + \mathbf{v}(t)\delta t \quad (2.9)$$

The Taylor series is truncated at the first derivative. This results in poor numerical accuracy. It is a first order method, meaning that the error per time step is proportional to δt^2 .

2.3.2 Verlet algorithm

A more widely used integration scheme, often used in computer games, and based on the centred difference method of Störmer, is the Verlet algorithm⁷. It is an example of an integration scheme that is not self-starting, rather, it is based on the current positions and accelerations, $\mathbf{r}(t)$ and $\ddot{\mathbf{r}}(t)$, and the positions from the last time step, $\mathbf{r}(t-\delta t)$. The addition of two Taylor series, one for $\mathbf{r}(t+\delta t)$ and one for $\mathbf{r}(t-\delta t)$ leads to the following expression for $\mathbf{r}(t+\delta t)$:

$$\mathbf{r}(t + \delta t) = 2\mathbf{r}(t) - \mathbf{r}(t - \delta t) + \dot{\mathbf{v}}(t)\delta t^2 \quad (2.10)$$

where odd powered δt terms have cancelled, and the series has been truncated at the quadratic term.

One clear disadvantage of this method is that it does not explicitly handle velocities, which, for instance, are required for calculation of temperature. They may be calculated from the predicted positions:

$$\mathbf{v} = \frac{\mathbf{r}(t + \delta t) - \mathbf{r}(t - \delta t)}{2\delta t} \quad (2.11)$$

This method leads to improved numerical accuracy. It is a second order method where the local error per time step is proportional to δt^4 .

2.3.3 Velocity Verlet algorithm

The velocity Verlet algorithm makes improvements on the original Verlet algorithm by storing the positions, velocities, and accelerations at once. A truncated Taylor series expansion for $\mathbf{r}(t+\delta t)$ is used to advance the positions:

$$\mathbf{r}(t + t) = \mathbf{r}(t) + \mathbf{v}(t)\delta t + \dot{\mathbf{v}}(t)\frac{\delta t^2}{2} \quad (2.12)$$

From the new positions (and the potential), the pair forces can be calculated, and hence the accelerations. From the current velocities and the accelerations, the new velocities can be calculated:

$$\mathbf{v}(t + \delta t) = \mathbf{v}(t) + [\dot{\mathbf{v}}(t) + \dot{\mathbf{v}}(t + \delta t)] \frac{\delta t}{2} \quad (2.13)$$

This then allows for the calculation of energies, and it has the distinct advantage that it is self-starting. This method has the same error associated with it as the original Verlet algorithm.

2.3.4 Runge-Kutta algorithm

The Runge Kutta (RK) methods overcome the deficiency of the Euler method by using a weighted average of a number of time derivatives. They were developed by Carl Runge, and are based on Martin Kutta's root finding technique, and use information about the slope at more than one point to extrapolate to the next time step.^{9,10} The two methods outlined below, RK2 and RK4, use two and four time derivatives respectively, designed to give the first n terms of the Taylor series expansion.

RK2

This method is the simpler and less accurate of the two. It advances the coordinates and velocities in two stages:

$$\begin{aligned} \{ \mathbf{r}'(t + \delta t) &= \mathbf{r}(t) + \mathbf{v}(t)\delta t ; \\ \mathbf{v}'(t + \delta t) &= \mathbf{v}(t) + \dot{\mathbf{v}}(t)\delta t \} \end{aligned} \quad (2.14)$$

$$\{\mathbf{r}(t + \delta t) = \mathbf{r}(t) + \frac{\delta t}{2} [\mathbf{v}(t) + \mathbf{v}'(t + \delta t)] ; \quad (2.15)$$

$$\mathbf{v}(t + \delta t) = \mathbf{v} + \frac{\delta t}{2} [\dot{\mathbf{v}}(t) + \dot{\mathbf{v}}'(t + \delta t)] \}$$

where ' denotes the first guess. The first guess accelerations $\ddot{\mathbf{r}}'$ are calculated from the pair forces generated by the first guess coordinates, \mathbf{r}' . The method is second order with regards to the timestep δt , meaning that the errors in the resulting velocities and coordinates at time $t+\delta t$ are proportional to δt^3 .

RK4

The four-stage Runge-Kutta scheme is arguably the most useful of the Runge Kutta schemes due to the large increase in accuracy. The first stage is the same as in the RK2 method:

$$\{\mathbf{r}'_1(t + \delta t) = \mathbf{r}(t) + \mathbf{v}(t)\delta t ; \quad (2.16)$$

$$\mathbf{v}'_1(t + \delta t) = \mathbf{v}(t) + \dot{\mathbf{v}}(t)\delta t\}$$

The remaining three stages are calculated iteratively based on the coordinates and velocities, and therefore the accelerations, calculated in the previous step:

$$\{\mathbf{r}'_2\left(t + \frac{\delta t}{2}\right) = \mathbf{r}(t) + \frac{\delta t}{2} \mathbf{v}_1(t) ; \quad (2.17)$$

$$\mathbf{v}'_2\left(t + \frac{\delta t}{2}\right) = \mathbf{v}(t) + \frac{\delta t}{2} \dot{\mathbf{v}}_1(t)\}$$

$$\{\mathbf{r}'_3\left(t + \frac{\delta t}{2}\right) = \mathbf{r}(t) + \frac{\delta t}{2} \mathbf{v}'_2\left(t + \frac{\delta t}{2}\right) ; \quad (2.18)$$

$$\mathbf{v}'_3\left(t + \frac{\delta t}{2}\right) = \mathbf{v}(t) + \frac{\delta t}{2} \dot{\mathbf{v}}'_2\left(t + \frac{\delta t}{2}\right)\}$$

$$\{\mathbf{r}'_4(t + \delta t) = \mathbf{r}(t) + \delta t \mathbf{v}'_3\left(t + \frac{\delta t}{2}\right) ; \quad (2.19)$$

$$\mathbf{v}'_4(t + \delta t) = \mathbf{v}(t) + \delta t \mathbf{v}'_3\left(t + \frac{\delta t}{2}\right)\}$$

The last step is to average the four derivatives:

$$\{\mathbf{r}(t + \delta t) = \mathbf{r}(t) + \frac{\delta t}{6} \left[\mathbf{v}(t) + 2\mathbf{v}'_2\left(t + \frac{\delta t}{2}\right) + 2\mathbf{v}'_3\left(t + \frac{\delta t}{2}\right) + \mathbf{v}'_4(t + \delta t) \right] ; \quad (2.20)$$

$$\mathbf{v}(t + \delta t) = \mathbf{v}(t) + \frac{\delta t}{6} \left[\dot{\mathbf{v}}(t) + 2\dot{\mathbf{v}}'_2\left(t + \frac{\delta t}{2}\right) + 2\dot{\mathbf{v}}'_3\left(t + \frac{\delta t}{2}\right) + \dot{\mathbf{v}}'_4(t + \delta t) \right]\}$$

The four stages are weighted $\{1/6, 1/3, 1/3, 1/6\}$. This method is fourth order with respect to the timestep δt , giving a local error per timestep proportional to δt^5 . It is therefore clear that this will give vast improvements in accuracy. The increase in evaluations of the right hand side associated with an increasing number of stages causes a rise in the amount of computing power required. The key is to find the right balance between accuracy and calculation time. For higher order methods, such as RK5, the extra accuracy is outweighed by the increased numerical work.

2.3.5 Comparison of integration schemes

Simple simulations can be used to illustrate the differences between the methods. Take, for example, the problem of a one-dimensional harmonic oscillator. Here, a particle oscillates around a fixed equilibrium point, obeying the equation:¹¹

$$F = -k(q - q_{eq}) \quad (2.21)$$

where F is the restorative force, $q - q_{eq}$ is the displacement and k is the restorative force. q_{eq} is typically taken to be the origin for simplicity. In the absence of friction the system undergoes sinusoidal oscillation with a constant amplitude and frequency (simple harmonic motion). If there is a dampening force present the system's amplitude decreases with time towards the equilibrium point. Hamilton's equations of motion for this example are:

$$\dot{q} = p \quad (2.22)$$

$$\dot{p} = -kq \quad (2.23)$$

where q is the coordinate and p is the momentum. It is easy to show that equations 2.22 and 2.23 have an analytical solution:

$$q(t) = q_0 \cos \omega t + \frac{p_0}{\omega m} \sin \omega t \quad (2.24)$$

$$p(t) = -q_0 m \omega \sin \omega t + p_0 \cos \omega t \quad (2.25)$$

where $\omega = \sqrt{\frac{k}{m}}$, and p_0 and q_0 are the momentum and coordinate at $t=0$ respectively. Squaring each of the above equations and adding them to one another shows that the phase space trajectory is an ellipse:

$$p^2 + m^2 \omega^2 q^2 = q_0^2 m^2 \omega^2 + p_0^2 \quad (2.26)$$

Setting $p_0 = 0$, $m = 1$, and $k = 1$ creates a special case in which the trajectory is a circle of radius q_0 . Three different numerical algorithms have been used to solve equations 2.21 and 2.22 approximately: Euler, RK2 and RK4 respectively. Figure 2.6

shows the resulting trajectories together with the analytical solution. The same timestep, $\delta t = 0.1$ (reduced units) was used in all three numerical schemes, and each simulation was run for 1000 time periods. The Euler method is shown to be the most inaccurate as both the amplitude and frequency are increasing with time, forming a spiral phase space trajectory. The RK2 method is also shown to be inaccurate, although to a lesser extent: the line thickens with time rather than forming an ellipse. The RK4 forms the expected phase space, showing how comparatively accurate the method is, and is the integration scheme used throughout this work.




Figure 2.6: A comparison of phase space plots for the Euler method (top left), RK2 method (top right), RK4 method (bottom left), and the analytical solution (bottom right). Both the Euler and RK2 methods are shown to have significant integration errors compared to the RK4 method. δt is the same in all cases.

A more enlightening analysis can be obtained by examining the maximum error produced by each integrator as a function of timestep size.

Taking $k = 1$, $m = 1$, and an initial condition $q_0 = 1$, $p_0 = 0$, the analytical solution for the coordinate is $q = \cos(t)$. Defining the coordinate error as $q - \cos(t)$, the trajectory was followed for 2 periods ($\tau = 4\pi$) using all 4 integration schemes, across a series of 10 timestep sizes, starting at $\Delta t = 2\pi/50$, decreasing by a factor of 2 in each case. The maximum error was recorded over the cycle. Figure 2.7 shows a plot of the natural logarithm of this maximum error as a function of the natural logarithm of the time step size. Straight line fits to the 4 data sets reveal slopes of 1, 2, 2 and 4, for the Euler, RK2, Verlet, and RK4 integration schemes respectively, confirming their global errors as δt , δt^2 , δt^2 and δt^4 . The *local* single step errors in each case are δt , δt^3 , δt^4 and δt^5 .




Figure 2.7: Global error as a function of timestep size for four integration schemes described in the text, applied to a 1D simple harmonic oscillator. Lines are least squares fits to the data.

2.4 Potential energy functions

When using conservative forces, MD simulations require the use of a potential energy function (sometimes referred to as a force field) to describe the way in which the atoms interact. They can be written in many forms, ranging from a fully quantum description to the readily used pair potentials, which are an embodiment of classical mechanics. The potential energy of a system can be written as a sum of one-body, two-body, three-body (...etc.) terms, with the terms representing the contribution of single atoms, pairs of atoms and triplets of atoms (...etc.):

$$\begin{aligned}\Phi = \sum_i \phi_1(\mathbf{r}_i) + \sum_i \sum_{j>i} \phi_2(\mathbf{r}_i, \mathbf{r}_j) \\ + \sum_i \sum_{j>i} \sum_{k>j>i} \phi_3(\mathbf{r}_i, \mathbf{r}_j, \mathbf{r}_k) + \dots\end{aligned}\tag{2.27}$$

where ϕ is the potential energy. The first term represents the potential energy from any external field acting on the system, and each successive term represents two-body, three-body (...etc.) interactions. Due to the difficulties with including three-body (and higher) interactions in potential energy calculations, a two-body potential with adjusted parameters is often used:

$$\Phi = \sum_i \sum_{j>i} \phi_2^{eff}(r)\tag{2.28}$$

where ϕ_2^{eff} is an effective pair potential, accounting for higher body interactions, and r is the interparticle separation $|\mathbf{r}_{ij}|$. Once the potential energy for all pairs of particles within a system is known, the force acting on each particle can be calculated using the following equation:

$$\mathbf{F}_i = -\nabla_i \phi \equiv -\phi'(r)r_{ij}; \quad (2.29)$$

where r_{ij} is the distance between particles i and j . The total force on particle i is calculated by summing the force from all pairwise interactions.

$$\mathbf{F}_i = \sum_{j=1}^N \mathbf{F}_{ij} \quad (2.30)$$

where \mathbf{F}_{ij} is the force between particles i and j . Newton's third law means that only half the number of calculations are required ($\mathbf{F}_{ij} = \mathbf{F}_{ji}$). There are many different potential energy functions that are widely used, and this section will discuss those relevant to this work.

2.4.1 Hard sphere potential

The simplest potential energy function is the hard sphere potential:¹²

$$\begin{aligned} \phi_{ij}^{hs} &= \infty \quad r < \sigma \\ &= 0 \quad r \geq \sigma \end{aligned} \quad (2.31)$$

where ϕ^{hs} is the potential energy, r is the distance between the centers of the particles, and σ is the diameter of the hard sphere. This is a discontinuous potential, where particles undergo elastic collisions. Figure 2.8 describes a hard sphere potential. This model finds use in developing statistical mechanic theories of liquids, particularly in perturbation theory.




Figure 2.8: A plot of energy against interparticle distance for the hard sphere potential. When the separation is equal to or smaller than the diameter of the hard sphere the potential energy is infinite.

2.4.2 Square well potential

An improvement on the hard-sphere potential is the square-well potential.¹²³ This potential has an attractive region in addition to the hard repulsive core. As, before, the potential goes to infinity when the separation is less than σ , but it also approximates the attractive region using a rectangular well of depth ε . The width of the well is dependent on a dimensionless parameter, $R/\sigma > 1$.

$$\begin{aligned}
 \phi_{ij}^{sw} &= \infty & r \leq \sigma \\
 &= -\varepsilon & \sigma < r < R\sigma \\
 &= 0 & R\sigma \leq r
 \end{aligned}
 \tag{2.32}$$

The width of the well is the product of σ and R . Figure 2.9 illustrated the potential.




Figure 2.9: A plot of energy against interparticle distance for the square well potential. When the separation is smaller than the diameter of the hard sphere the potential energy is infinite, and the attractive region is approximated by a square well. $R = 1$.

2.4.3 Lennard-Jones potential

One of the most commonly used potential energy functions is the Lennard-Jones potential, given by:¹⁴

$$\phi_{ij}^{LJ} = 4\epsilon \left[\left(\frac{\sigma}{r_{ij}} \right)^{12} - \left(\frac{\sigma}{r_{ij}} \right)^6 \right] \quad (2.33)$$

It is a pair-additive function used to describe the interaction between a pair of neutral atoms, and is illustrated in figure 2.10. When the distance between the atoms is

small they repel each other. The distance, σ , is the van der Waals radius. This is an approximation of Pauli repulsion, and is represented by the r^{-12} term. The r^{-6} term is the attractive term, and describes the long-range attractive forces. At an infinite separation, the attraction is considered to be zero. As the particles get closer the potential energy becomes increasingly negative, until a potential energy minimum is reached. With a judicious choice of parameters this potential can accurately simulate Argon.




Figure 2.10: A plot of potential energy against interparticle distance for the Lennard-Jones potential.

2.4.4 Short range, soft repulsive potential

Using a pair core potential that is only active at short separations is an effective way to maintain particle separation.²

$$\phi_{ij} = \varepsilon \left[1 - \frac{r_{ij}^2}{\sigma^2} \right]^4 ; |r_{ij}| < \sigma \quad (2.34)$$

This potential produces a force that discourages overlaps, and can be combined with other potentials, such as the embedded atom density dependent potential energy function. Without a pair core potential, the embedded atom function has a tendency to make string structures, causing a lack of uniformity of density within a relaxed system. The pair core potential fixes this problem, allowing the system to relax towards a potential energy minimum. Figure 2.11 describes a pair core potential.




Figure 2.11: A plot of potential energy against interparticle distance for the soft-sphere potential.

2.5 Link cells

Performing loops over pairs of particles is computationally expensive, and when using the RK4 integration scheme, the pairwise forces are calculated 4 times per timestep. To further optimise the algorithm used to calculate the pair forces (and energies), link cells can be used to remove unnecessary calculations, shown in figure 2.3. The central simulation cell is split into a series of link cells, whose width and height are at least equal to the cut-off distance of the interatomic pair potential. At the start of each time step each particle is assigned to the appropriate link cell. When calculating the pairwise forces, instead of looping over all pairs of particles, each link cell is considered in turn. Then each particle inside the link cell is selected, and the usual force calculation is performed, however, only particles in the current link cell, and any adjacent cells (including any boundary conditions) are considered. In figure 2.12, the current link cell is highlighted in red, and all particles considered for the force calculation of particles within this cell are red. The dark particles are more than 2 link cells away from the selected cell, and therefore the particles within cannot be close enough to interact (and are not considered in the double loop). The initial process of sorting particles into cells is computationally inexpensive, so the time benefit from using link cells is vast.




Figure 2.12: An illustration of link cells. The simulation cell is split into link cells, and only particles in the selected cell and adjacent cells are considered when calculating forces. The red particles are close enough to interact, and the dark particles are too far away.

2.6 Calculating thermodynamic properties

In order to obtain meaningful results from simulations it is important to be able to calculate properties of interest, such as kinetic energy (hence also temperature), and pressure.

In molecular dynamics, thermodynamic properties can be calculated instantaneously from functions of the readily available coordinates and momenta of all of the atoms within the system. In a system at thermodynamic equilibrium the instantaneous values fluctuate around the average. This means that taking a statistical average of properties over time will result in accurate expressions of these properties.

Equilibrium molecular dynamics generates the microcanonical ensemble,¹⁴ where N , V , and E are constant. Here, the temperature and pressure are not explicitly conserved; rather, they fluctuate around their average values. Temperature and pressure can be calculated by taking time averages of the instantaneous values.

2.6.1 Potential energy

The potential energy may be calculated by simply averaging the instantaneous values of potential energy, Φ . This is done at the same time as the force is calculated, and although it is not necessary to perform this calculation, is useful in verifying that energy is being conserved (although energy is not conserved when using a thermostat).

$$\Phi(t) = \sum_i \sum_{j>i} \phi(\mathbf{r}_{ij}) \quad (2.35)$$

This gives the instantaneous total potential energy. This is simply summed at each timestep then divided by the number of timesteps.

2.6.2 Kinetic energy

Instantaneous kinetic energy, K , is given by:

$$K(t) = \frac{1}{2} \sum_i m_i \dot{\mathbf{r}}_i^2 = \sum_i \frac{\mathbf{p}_i^2}{2m_i} \quad (2.36)$$

where m_i is the mass. Again, the average kinetic energy is calculated by time averaging the instantaneous values.

2.6.3 Total energy

Total energy, E , is simply the sum of the potential and kinetic energies:

$$E = K + \Phi \quad (2.37)$$

In the microcanonical ensemble generated in equilibrium molecular dynamics the total energy should remain constant (to machine accuracy if a smooth potential is used in combination with the RK4 integration scheme, though there will be some drift in other conditions), throughout a simulation. If the timestep used is too large, or the forces are discontinuous at the cut-off, then integration errors can take effect, causing slight fluctuations. A drift in energy will diverge exponentially from the expected value. If the total energy drifts with time then the thermodynamic state of the system is also changing so equilibrium properties cannot be measured.

2.6.4 Temperature

In the canonical ensemble the temperature is constant, however, in the microcanonical ensemble the temperature fluctuates around the average total. The equipartition formula states that the average kinetic energy of a system of N particles is equal to $k_B T/2$ per degree of freedom:

$$K = \frac{D}{2} N k_B T \quad (2.38)$$

where DN is the number of degrees of freedom. If each particle has three degrees of freedom, the kinetic energy would equal $3Nk_B T/2$. In a molecular dynamics system the total linear momentum is often conserved, meaning that the number of degrees of freedom becomes $D(N-1)$.

2.6.5 Pressure

The most common method for calculating pressure in a molecular dynamics simulation is based on the Clausius virial function. This virial is the sum of the product of the coordinates of all the particles in the system and the forces acting on them:

$$W = \sum_{i=1}^N \mathbf{r}_i \cdot \dot{\mathbf{p}}_i \quad (2.39)$$

where \mathbf{r} is the coordinate of particle i and $\dot{\mathbf{p}}$ is the force. The virial theorem gives the pressure:

$$PV = Nk_B T + \frac{1}{D} \left\langle \sum_{i=1}^N \mathbf{r}_i \cdot \dot{\mathbf{p}}_i \right\rangle \quad (2.40)$$

which, when using periodic boundary conditions, becomes:¹⁶

$$PV = Nk_B T + \frac{1}{D} \left\langle \sum_{i=1}^N \mathbf{r}_{ij} \cdot \dot{\mathbf{p}}_{ij} \right\rangle \quad (2.41)$$

The pressure (and temperature) fluctuate throughout a simulation, so the average over many timesteps (usually the whole duration of the production simulation) should be taken.

2.7 Thermostats

Being able to monitor and control the temperature is desirable in molecular dynamics. In standard molecular dynamics the system is driven by only the interatomic forces; there are no external forces used to control the dynamics, and energy is conserved. In many real experiments, however, the temperature is controlled instead of the energy, making the ability to control temperature in molecular dynamics simulations desirable. This is achieved by working in the Canonical ensemble.

Thermostats are used in NEMD in order to control the temperature of the system. There are many different techniques with the aim of achieving this, some of which will be discussed here, and one must consider whether both the thermodynamics and the dynamics are preserved.

2.7.1 *ad hoc* velocity rescaling

The simplest method of applying a thermostat is to multiply all of the velocities of the particles in a system by the same factor, α :

$$\alpha = \sqrt{\frac{T_{req}}{T(t)}} \quad (2.42)$$

where T_{req} is the target temperature, and $T(t)$ is the instantaneous microscopic temperature. Because the same factor is applied to all of the velocities there is no effect on the centre of mass. This rescaling is usually applied after a pre-determined number of timesteps, or when the calculated kinetic energy goes outside a set limit around the target value. This method can produce useful results when the time averaged properties do not depend on the ensemble chosen, however it falls down when the properties of interest are dependent on fluctuations (as opposed to averages) as it creates discontinuities in phase space, and it does not generate the properties of

the canonical ensemble; if the rescaling were to be applied every timestep the kinetic energy would remain constant, not allowing for any energy fluctuations.

2.7.2 Andersen Thermostat

The Andersen thermostat is shown to generate the correct canonical ensemble over infinitely long trajectories. It works by introducing random collisions of the particles in the system with an imaginary heat bath, sampling the new velocity from a Maxwell-Boltzmann distribution at the desired temperature.¹⁷

$$\rho(v_{x,i}) = \left(\frac{m_i}{2\pi k_B T}\right)^{1/2} \exp\left(-\frac{m_i v_{x,i}^2}{2k_B T}\right) \quad (2.43)$$

Particles are either chosen randomly and a collision performed, or all particles have all components of their velocities reassigned simultaneously. After this event, the centre of mass motion needs to be removed. As with velocity rescaling, this is not done every time step, but is performed on a collision frequency.

Although this method accurately generates the correct kinetic and potential energies, due to the stochastic nature of the collisions, correct molecular kinetics is not maintained. This means that particle trajectories de-correlate from previous timesteps more quickly than in the canonical ensemble, creating erroneous results when measuring properties such as diffusion coefficients.

2.7.3 Gaussian thermostat

As mentioned previously, the main drawback of the *ad hoc* velocity rescaling method is that it produces discontinuities in the momentum phase space, due to the rescaling mechanism. The Gaussian thermostat applies Gauss' principle of least constraint to add a constraint force term to the equations of motion ensuring the kinetic energy (and hence temperature) is a constant of the motion:¹⁸

$$\dot{\mathbf{q}} = \frac{\mathbf{p}}{m} \quad (2.44)$$

$$\dot{\mathbf{p}} = -\mathbf{F} - \alpha\mathbf{p} \quad (2.45)$$

$$\alpha = \frac{\sum \mathbf{F} \cdot \mathbf{p}}{\sum \mathbf{p} \cdot \mathbf{p}} \quad (2.46)$$

where $\alpha\mathbf{p}$ is the constraint force, with α being calculated instantaneously equation 2.46. This generates the isokinetic ensemble, but not the canonical ensemble. The equations of motion are deterministic avoiding the problems caused by the stochastic nature of the interactions in other thermostats

2.7.4 Nosé-Hoover thermostat

The Nosé-Hoover thermostat provides a method to simulate systems, asymptotically, in the NVT ensemble by introducing a fictitious force that guides the total temperature of the system towards the target temperature.^{19,20} The coefficient, ζ , has a frictional effect that either speeds up or slows down particles. It obeys a feedback equation based on the ratio of the current kinetic energy to the target kinetic energy:

$$\dot{\mathbf{r}}_i = \frac{\mathbf{p}_i}{m} \quad (2.47)$$

$$\dot{\mathbf{p}}_i = \mathbf{F}_i - \zeta\mathbf{p}_i \quad (2.48)$$

$$\dot{\zeta} = \left[\left(\frac{\sum \mathbf{p}^2}{m} / \sum k_B T_0 \right) - 1 \right] / \tau^2 \quad (2.49)$$

where, T_0 is the temperature corresponding to the kinetic energy K_0 , and τ is the damping force (typically in the region of 500 reduced units). This frictional force is contained within the equations of motion, and equation 2.49 is treated as another ordinary differential equation to be integrated giving ζ at each timestep.

This thermostat generates the canonical ensemble. It is commonly used in the molecular simulation community.

2.8 References

- [2] Hoover W. *Smooth Particle Applied Mechanics*. World Scientific, Singapore, 2006
- [3] Alder B.J, Wainwright T.E. *Studies in Molecular Dynamics. I. General Method* . The Journal of Chemical Physics, **31**, 459 (1959)
- [4] M. Born and T. Von Karman, *On fluctuations in spatial grids*, Physik. Z., **13**, 297 (1912)
- [5] T. Dong and S. Jiang, *International Conference on Computational and Information Sciences*, Brooks/Cole, California, 1212 (2010)
- [6] W. Cheney and D. Kincaid, *Numerical Mathematics and Computing*, Brooks/Cole, California, 2008
- [7] L. Euler, *Institutionum Calculi Integralis*, Impensis Academiae Imperialis Scientiarum, Ghent, 1768
- [8] L. Verlet, *Computer “experiments” on classical fluids. I. Thermodynamical properties of Lennard-Jones molecules*, Physical Review, **159**, 98 (1967)
- [9] M. Kutta, *Beitrag zur näherungsweise Integration totaler Differentialgleichungen*, für Math. u. Phys, **46**, 435 (1901)
- [10] C. Runge, *Ober die numerische Aufltising von Differentialgleichungen*, Math Ann. **46**, 167 (1895)
- [11] S.L.Garret, *Understanding Acoustics*, Springer International Publishing, Zurich (2017)
- [12] J. Jover, A. Haslam, A. Galindo, G. Jackson, E. Muller, *Pseudo hard-sphere potential for use in continuous molecular-dynamics simulation of spherical and chain molecules*, J. Chem. Phys. **137**, 144505 (2012)

- [13] H. Davis, K. Luks, *Transport properties of a dense fluid of molecules interacting with a square-well potential*, J. Phys. Chem. **69**, 869 (1965)
- [14] L. Verlet, Physical Review, *Computer" experiments" on classical fluids. I. Thermodynamical properties of Lennard-Jones molecules*, **159**, 98 (1967)
- [15] M. Allen, D. Tildesley, *Computer Simulation of Liquids*. Oxford Science Publications, Oxford, 1989
- [16] H. Andersen, *Molecular dynamics simulations at constant pressure and/or temperature*, J. Chem. Phys, **72**, 2384 (1980)
- [17] D. Evans, W. Hoover, B. Failor, B Moran, A. Ladd, *Nonequilibrium molecular dynamics via Gauss's principle of least constraint*, Physical Review, **28**, 1016 (1983)
- [18] S. Nosé, *A molecular dynamics method for simulations in the canonical ensemble*, Mol. Phys., **50**, 225 (1984)

3 Molecular dynamics simulation of sand bed filtration

3.1 Overview

In this work, molecular dynamics was used with the twin aims of probing the mechanisms involved in the process of clogging and of providing data to help validate the continuum scale model. The idea was not to create an atomistic scale model that accurately modelled the interactions of sand particles with individual water molecules and specific colloidal particles, rather, it was to create a simplified model of sand bed filtration containing the essential physics. This allows for physical insight free from unnecessary complication. It was therefore of great importance to maintain simplicity in the design, using existing literature as validation and guidance throughout the process. This chapter will discuss the design of the model, the techniques used to validate it, and highlight the key results. It is worth noting that all units discussed in this and further chapters are reduced units.

3.2 Model design

The model aims to predict the flow of a binary mixture of fluid and colloid particles through a sand bed, where the colloids were capable of sticking to the sand particles, causing clogging. It was a two-dimensional coarse grained molecular dynamics model, where all particles (fluid, colloid, and sand) are represented by soft discs, of the same size. The sand bed consisted of a lattice of static discs. All inter-particle interactions were defined by variations of the same short-range smooth pair core potential, which is entirely repulsive. An applied field acts in the y-direction to represent gravity, which entices the fluid to flow through the sand bed. To model chemical adsorption of colloid particles on the sand grains, colloids approaching within a capture distance of a sand grain were stripped of their velocity if a random number sampled on $[0,1]$ was less than a specified sticking probability. Additional particles were introduced at the top of the system at regular intervals creating a continuous flow through the filter. Over time, colloids deposited within the sand bed causing clogging.

To analyse the simulation, the thermodynamic properties of interest (density, pressure, temperature) were calculated both globally and locally. The pressure drop and specific deposit, which are two key parameters used to analyse the problem of clogging at SIXEP, were also monitored, which gave a means of comparing this model to existing models and literature. The model is split into two stages: the equilibration stage and the filtration stage, which are depicted in figure 3.1.




Figure 3.1: A schematic showing the molecular dynamics filtration model. The yellow particles are the sand particles, the blue particles are fluid particles and the red particles are colloid particles. The dashed line is an elastic boundary that is present in the first stage.

3.2.1 Equilibration stage

The left-hand image in figure 3.1 depicts the initial stage, where a binary mixture of fluid and colloid particles was suspended over the lattice of sand particles. A gravitational field was applied and the fluid was allowed to relax. The binary fluid

mixture cannot interact with the sand particles at this point. When equilibrium was reached the elastic boundary separating the fluid and the sand particles was removed, ending the first stage.

Initial conditions

Each particle was assigned an initial coordinate and momentum. The coordinates of the fluid were chosen to create a square lattice. This was a convenient choice as it is computationally simple to assign the coordinates in this way. Within a few timesteps, all memory of this lattice will be lost as the system becomes a fluid.

The density and number of particles were defined as input parameters for each simulation, which therefore defined the size of the column of fluid and the initial spacing between particles. The type of particle (fluid or colloid) was decided randomly to create a system where the colloids were initially scattered throughout the fluid, and is illustrated in figure 3.2. The desired ratio of colloid/fluid particles was defined in the input file, and a random number generator was used to sequentially determine the type of each particle.




Figure 3.2: Illustration showing the initial fluid conditions. The blue particles are fluid and the red are colloids. The type of each particle is randomly decided based on the ratio. This example shows a 4:1 ratio of fluid:colloid.

The momenta were calculated so that the average kinetic energy per particle equalled the value that the simulator chose (typically 1.0). First, a random number generator assigned each particle a velocity of between -1.0 and 1.0. The centre of mass velocity was then removed to stop the simulation box from drifting:

$$\dot{\mathbf{q}}_i' = \dot{\mathbf{q}}_i - \frac{1}{N} \sum_{i=1}^N \dot{\mathbf{q}} \quad (3.1)$$

Finally, the momentum was rescaled so that the average kinetic energy equalled the input value.

It is worth noting some details about the random number generator used to define the initial velocities. A random number generator requires a new seed each time a series of random numbers is generated, otherwise the exact same series of pseudo “random numbers” will be created. In order to create a new series of random numbers each time a simulation was run the system clock time was used as a seed, resulting in a unique seed, hence, unique velocities. If repeated simulations with the *same* initial velocities were required, the simulator could choose not to use a seed to generate the numbers. See appendix A for a segment of FORTRAN code used to generate a random seed.

Equations of motion

The equations of motion solved for the colloid and fluid particles were:

$$\dot{\mathbf{q}}_i = \frac{\mathbf{p}_i}{m_i} \quad (3.2)$$

$$\dot{\mathbf{p}}_i = \sum_{j=1}^N \mathbf{F}_{ij} - m_i g \mathbf{k} \quad (3.3)$$

where \mathbf{F} is the force acting on each particle as a result of the inter-particle potential, Φ , g is the gravitational field strength and \mathbf{k} is a unit vector in the positive y -direction. The mass, m , of all particles in the simulations was 1. The particles that comprised the fluid were allowed to relax under gravity, with a viscous damping force,

τ , being used to remove excess kinetic energy until equilibrium was reached. The damped equation of motion becomes:

$$\dot{\mathbf{p}}_i = \sum_{j=1}^N \mathbf{F}_{ij} - m_i g \mathbf{k} - \frac{m_i |\dot{\mathbf{q}}}{\tau} \quad (3.4)$$

Interatomic potential

All of the inter-particle interactions were based on variations of the same short-range smooth pair core potential:

$$\Phi(r < \sigma) = \varepsilon \left(1 - \left(\frac{r}{\sigma}\right)^2\right)^4 \quad (3.5)$$

where σ and ε define the strength and effective diameter of the particle respectively, and r is the inter-particle distance ($r = |\mathbf{r}_{ij}|$; $\mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j$). This potential was chosen as it is computationally simple and short-ranged, but still capable of producing complex results. Related potentials have been used extensively in dissipative particle dynamics (DPD) to successfully model colloidal behaviour.^{1,2,3} The potential is entirely repulsive, however, in the presence of multiple particle a depletion force can be created by adjusting ε . To demonstrate this, a simulation was run modelling the behaviour of a binary mixture of 400 particles (200 particles of type fluid, and 200 of type colloid). The ε value for fluid-fluid and colloid-colloid interactions was 20.0, but the ε value for fluid-colloid interactions was 25.0. Figure 3.3 shows the computed particle coordinates after 100 and 10,000 timesteps. Using this simple distinction, phase separation was shown. This serves as good evidence that although the potential is simple, it can still produce complex and realistic results.




Figure 3.3: Snapshots showing the computed particle locations of a binary mixture after 100 and 10,000 timesteps using periodic boundary conditions. $\epsilon_{11} = \epsilon_{22} = 20.0$ $\epsilon_{12} = 25.0$. $\Delta t = 0.01$.

With 2 types of particle, there are 4 possible ϵ values, however, due to symmetry, only 3 of these are distinct: colloid/colloid, fluid/fluid, and colloid/fluid. Each of these ϵ values was defined in the input file.

Boundary conditions

Figure 3.4 illustrates the boundary conditions employed in the first stage. An elastic boundary was used to keep the fluid above the sand particles and the lateral boundaries were periodic. The dimensions of the simulation cell were chosen based on the input values of the number of fluid/colloid particles (and their density). An arbitrary choice was made to start the column of fluid in an aspect ratio of 2:1. This ratio created a column with a sufficient width ensuring that particles would not interact with their own image. The image on the right hand side of figure 3.6 shows the coordinates at the end of stage 1.

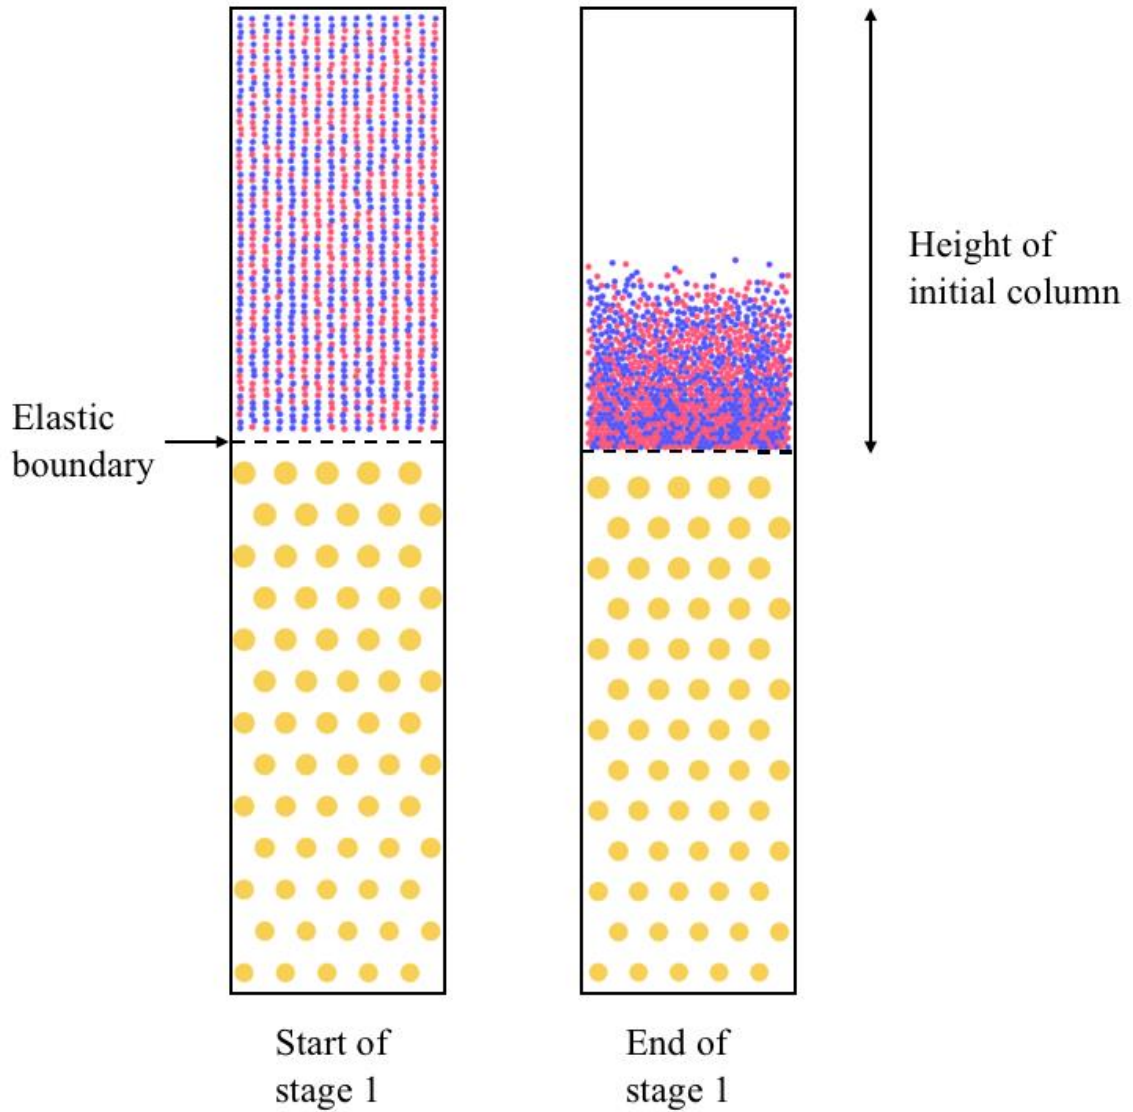


Figure 3.4: Computed particle locations showing the initial conditions (left) and final coordinates (right) of a binary mixture containing 1024 particles (512 colloid, 512 fluid). The simulation length was 1,000,000 time periods, $\Delta t = 0.01$, $g = 0.1$, $\tau = 1/ \Delta t$.

3.2.1 Filtration stage

The second stage of the simulation was the filtration stage. During this stage the fluid could flow past and interact with the sand particles, and the thermodynamic properties of interest were monitored. The sand particles were static, so they did not have any equations of motion associated with them. The equations of motion for the fluid mixture remained the same.

Interatomic potential

In order to probe the effect that the porosity of the filter has on the flow rate, it was important that the sand particles had a size. A potential with a slightly different form was used to achieve this; a shifted core potential allows the sand particles to be given a finite size, r_s :

$$\Phi(\delta r < \sigma) = \varepsilon \left(1 - \frac{\delta r^2}{\sigma}\right)^4 \quad (3.6)$$

$$\delta r = r - r_s \quad (3.7)$$

Instead of the potential energy (and force) being calculated from the spacing between two mass points, it is calculated from the nearest point on the outer radius of the sand particle, as shown in figure 3.5




Figure 3.5: Illustration showing the shifted core potential. The distance is calculated from the surface of the sand particle rather than the centre.

This potential shifts the interaction between the sand particles and the colloids by the size of the sand particle. With 3 types of particle, there are 9 possible ϵ values, however, due to symmetry, only 6 of these are distinct: colloid/colloid, fluid/fluid, sand/sand, colloid/fluid, colloid/sand, and fluid/sand. Each of these ϵ values was defined in the input file.

Initial conditions

The sand particles were fixed in space, therefore they did not have any momentum associated with them. They either formed a square lattice, a triangular lattice, both of which are illustrated in figure 3.6. The algorithms to define the lattice coordinates were simple, and just required the simulator to define the initial packing fraction.




Figure 3.6: Illustration showing the sand particle coordinates. The left shows a triangular lattice (coordination number of 6) and the right shows a square lattice (coordination number of 4).

Boundary conditions

The elastic boundary was removed to start the second stage. The lateral boundaries remained periodic. The original fluid phase would eventually work its way through the filter resulting in termination of the simulation. To prolong the simulation, new fluid particles were introduced at regular intervals. The number of particles, N_{new} , and how often they were created, were defined in the input file. They were created at

the top of the simulation cell (as shown in figure 3.7), with an even spacing, where the spacing, r_{new} , is obtained by:

$$r_{new} = \frac{w}{N_{new}} \quad (3.8)$$

where w is the width of the simulation cell. The new particles had a velocity in the y-direction of 0, as the gravitational force soon generated this. The x-velocity was chosen randomly. Similarly, the type of each new particle (colloid or fluid) was also selected at random, where the chance of each type of particle being created was based on the starting ratio.

Because the integration scheme loops over pairs of particles, additional particles vastly increase simulation time. It was therefore important to keep the number of particles to a minimum. Fluid particles that travelled beyond the final layer of sand particles were removed. The simulation continued for either a set number of timesteps, or until a certain pressure drop was reached. In the latter case, once the filter was clogged, the addition of new particles would cause the simulation to become unreasonably lengthy due to their slower transport through the membrane before they were ultimately removed.




Figure 3.7: Image depicting the start of the second phase. The elastic boundary has been removed, and the fluid and the sand can now interact with each other.

Particle sticking

In a similar method to the models discussed in section 1.3.3, a sticking probability was used to determine the chance of a colloid depositing within the sand bed. Two variables were used in order to achieve this. The first d , was the clogging distance. This was the maximum distance between the surface of the sand particle and the colloid at which there was a chance of sticking, which was akin to the solid-water interface described in NNL's model. The second, p , was the probability of deposition,

and was a value between 0 and 1; $p = 0$ represents zero sticking, while particles would always stick if $p = 1$. When the separation between the colloid and the sand particle was less than d , a random number was generated. If this number was less than the sticking probability then the particle was taken to be deposited, and in this state its velocity was set to be zero. This was an irreversible process, and the immobile deposited particle continued to interact with all other particles.

This method of sticking modelled a situation where similar to the process of physisorption. There was a finite space around each sand particle for colloids to deposit, and as time passed the possible retention sites were filled, resulting in a maximum amount of deposit per sand particle. This was representative of blocking, which was modelled in previous work using a dynamic blocking function.

The opposite effect to blocking is ripening, which is likely to result from chemisorption. This occurs when colloid-colloid interactions are attractive, where deposited colloids actually provide additional surface area for further colloid deposition. This can lead to improved colloid retention, though if the attractive forces are too strong, it will result in large deposits that cause complete clogging. To model this, an additional sticking probability was used. The same algorithm was implemented, though the distance calculated was the colloid-colloid separation. This allowed colloids to deposit on top of other colloids, continually growing the clusters.

3.2.3 Calculating properties

To probe this model effectively it was necessary to monitor how properties varied across the filter. There are several techniques that allow local thermodynamic properties to be calculated, such as the histogram approximation (bins method)⁴ and the method of planes.⁵

An inherently simple approach is to employ smooth particle averages, where a weight function is used to calculate each particle's contribution to the property at that plane.⁶ At a given plane, the total value of the property of interest is a sum of the weighted contributions of all particles within the cut-off distance of the weight function (the smoothing length), h . This method is illustrated in figure 3.8.




Figure 3.8: Schematic showing the weighted method of planes. Each particle whose distance, r_j , from the plane of interest is less than the smoothing length, h , contributes to the property at that plane. Red particles are within the smoothing length of the plane highlighted in red, grey particles are not.

An appropriate choice of weight function is Lucy's one-dimensional weight function (which is used in SPAM). The weight function is illustrated in figure 3.9, and takes the form:

$$w(r < h) = \frac{5}{4\pi} \left[\left(1 + \left(3 \frac{r}{h} \right) \right) \right] \left[1 - \left(\frac{r}{h} \right) \right]^3 \quad (3.9)$$




Figure 3.9: Plot showing the one-dimensional Lucy weight function used to weight particle influence at a plane.

A series of tests were performed in order to illustrate how this method can be used to monitor various properties.

Mass density

The mass density at a plane is given by:

$$\rho(r, t) = \sum_i m_i w_{y_\alpha i} \quad (3.10)$$

$$w_{y_\alpha i} = w(|y_\alpha - y_i|) \quad (3.11)$$

where $w_{y_\alpha i}$ is the weight function calculated for the distance between the plane of interest, y_α , and each particle, i . Two simulations were run to illustrate this method using the short-range pair core potential: the first with no gravitational force, and the second with a gravitational force of 1.0. 400 identical particles were simulated for 100,000 timesteps, where $\Delta t = 0.01$. The side boundaries were periodic, and the bottom and top boundaries were elastic. Figure 3.10 shows the computed particle coordinates at the start and the end of the simulation.




Figure 3.10: Computed particle locations for two simulations where the gravitational force, $g = 0.0$ (left) and 1.0 (right). $\Delta t = 0.01$, total timesteps = 100,000, $\rho_0 = 1.0$.

The density in the y -dimension was calculated every 10 timesteps at 100 planes equally spaced in the y -direction using equation 3.9 and the Lucy weight function (equation 3.8). It was then averaged by the number of snapshots taken to give the density profiles shown in figure 3.11. In the first case, where there was no gravitational field; the density was evenly spread across the system. The density profile shows a uniform value of 1.0 across all coordinates. When there was a gravitational field applied, there was a density gradient produced. A higher density was found at the bottom of the system compared to the top, and this was correctly captured using this profiling technique.




Figure 3.11: Density profiles without (left) and with (right) a gravitational field. Density was calculated at 100 planes in the y -directions every 10 timesteps.

Momentum density

The momentum density at a given plane is calculated using the following expression:

$$J_y(y_\alpha, t) = \sum_i m_i \mathbf{v}_i w(|y_\alpha - y_i|) \quad (3.12)$$

where \mathbf{v}_i is the lab-frame velocity of particle i , consisting of the thermal (microscopic) and streaming (macroscopic) components. The streaming velocity is calculated from the ratio of the momentum and mass density, J/ρ .

A suitable test to validate the reliability of the momentum density calculation is to model Poiseuille flow. This is the flow of a fluid through two parallel walls, where the flow is generated by an external field, and is well described in literature; Poiseuille flow should produce a quadratic velocity profile across the pore. This process was modelled by simulating the flow of 2048 particles through a pore space for 1,000,000 timesteps, with a timestep of 0.001. The flow was generated using a force in the y -direction of 0.1. 100 profile planes evenly spaced in the x -direction were used to calculate velocity across the pore. The walls were elastic boundaries, where the excess energy generated by the external force was removed using velocity rescaling upon collisions with the wall, following the example of Ziaran and Mohamad.⁷ The resulting velocity profile is shown in figure 3.12. The least squares fit of the velocity profile does indeed show the expected quadratic dependence, with the results agreeing excellently across most of the pore, showing slight deviations towards the walls.




Figure 3.12: Velocity profile across a pore

Pressure

Pressure at a plane is calculated using the same approach, however there is an important addition. The first term on the right-hand side is the kinetic contribution, which is straightforward to calculate. The potential contribution (the second term) requires an arbitrary choice to be made on how to distribute the force from each particle. This work follows previous examples by Hoover in assigning half of the $y_{ij}F_{ij}^y$ contribution to each member of the interaction pair.

$$\begin{aligned}
 P_{yy}(y_\alpha, t) = & \sum_i m_i \dot{y}_i^2 w(|y_\alpha - y_j|) \\
 & + \sum_i \sum_j y_{ij} F_{ij}^y w(|y_\alpha - y_j|)
 \end{aligned}
 \tag{3.13}$$

The pressure drop (or head loss) is used at SIXEP as an indicator of clogging, and was therefore an important parameter in this model. The pressure drop was calculated as the difference between the pressure at the top of the filter (which was the plane with the highest pressure) and the bottom of the filter (the plane with the lowest pressure). Figure 3.13 shows how the planes were used to measure thermodynamic properties across the filter, and highlights the two planes from which the pressure drop was calculated.




Figure 3.13: Image showing the profile planes used to measure the thermodynamic pressures across the filter. The pressure drop is the difference in pressure at profile planes A and B.

To illustrate this, a simulation was run with a sufficiently high sticking probability and sticking distance for the colloids to quickly deposit at the top of the bed, causing complete clogging of the filter. Figure 3.14 shows the computed particle locations at 5 stages during the simulation. As can be seen, as soon as the colloids were allowed to interact with the sand they began to stick, clogging the filter. The

pressure was calculated at 100 evenly spaced planes in the y-direction every 10 timesteps.




Figure 3.14: Snapshots showing particle coordinates at 5 stages during the same simulation. The purple circles are the colloidal particles, the light blue circles are the fluid particles, the dark blue circles are the clogged colloidal particles and the orange circles are the sand particles. Sticking probability = 0.001, sticking distance = 0.75.

Figure 3.15 shows the obtained pressure profiles. The filter bed begins at a y-coordinate of 0. The black line shows the pressure profile at $t = 0$, before the particles were allowed to interact with the sand. As expected, there was no pressure observed

below a y-coordinate of 0 during the initial phase. As the colloids started to stick, the pressure at the top of the filter became increasingly high; as more particles entered the simulation from the top, and with very few particles passing through the top layer of clogged colloids, the pressure continued to build. This, of course, was an exaggerated situation, where the purpose was to show that the techniques being used to analyse the simulation were sufficient to monitor the pressure drop.




Figure 3.15: Pressure profiles for the 5 computed particle locations in figure 3.14.

3.3 Model validation

Before systematically exploring this model to probe the mechanism of clogging, it was important to verify that the results, at least qualitatively, showed the same trends as those in other models and experiments. Two main sources were used to validate the model: existing literature, and a one-dimensional model created by the National Nuclear Laboratory (NNL). The model was based on an advection-dispersion equation and empirical rates of deposition and clogging, and was discussed in section (1.3.3). There are 3 main areas for comparison: the rate of specific deposit, the rate of increase in pressure drop, and the deposit concentration with depth. It is worth noting that all simulations discussed in this section, unless otherwise stated, used the input parameters detailed in appendix B.

3.3.1 Specific deposit

Specific deposit, σ , is defined as the mass of deposit per unit of filter volume. There are two phases during the filtration lifecycle: blocking/ripening and operation. According to Camesano *et al*⁸, deposition should increase linearly with time, with the rate of deposit being dependent on the current phase. This theory has been used to parameterise existing models in the field, with the NNL model using the following equations to determine the rate of deposit:

$$k_{12} = k_r \text{ and } k_{21} = 0 \text{ when } 0 < \sigma \leq \sigma_r \quad (3.14)$$

$$k_{12} = k_{ac} \text{ and } k_{21} = k_d \text{ when } \sigma_r < \sigma \leq \sigma_u \quad (3.15)$$

$$k_{12} = 0 \text{ and } 0 = k_d \text{ when } \sigma = \sigma_u \quad (3.16)$$

where k_{12} is the rate of attachment, k_{21} is the rate of detachment, k_r is the rate during the ripening/blocking phase, k_{ac} and k_d are the rates during the operation stage, and σ_r and σ_u are the threshold values of specific deposit.

MD model prediction of specific deposit

All particles in the molecular dynamic simulations had a mass of 1.0, therefore the specific deposit was simply calculated as the number of deposited particles divided by the volume of the filter. The volume of the filter was defined as the width of the simulation cell multiplied by the depth of the sand bed (see figure 3.6). There were two distinct sticking mechanisms that were used in the molecular dynamics model. The first was where colloids could only stick to the sand particles (blocking), the second was where colloids could stick to both sand particles and already deposited colloids (ripening).

Figure 3.16 shows the specific deposit build up for a simulation using the *blocking* mechanism. This result qualitatively agrees with the predictions of Camesano; there were two phases shown, which equate to the two phases of filtration (blocking and operation). As the first layer of colloids surrounded the sand bed particles there was a steep increase in the specific deposit. At a threshold value of around $\sigma_r = 2.7$, the rate of deposition decreased to a second linear rate. Taking the gradient of the slopes of the two regimes results in values of 2.0×10^{-4} and 9.2×10^{-5} for k_r and k_{ac} respectively.




Figure 3.16: Specific deposit against time showing the two filtration stages. Here blocking is modelled: colloids only stick to the sand particles, colloid clogging distance = 0, colloid sticking probability = 0.

Figure 3.17 shows the growth of the clusters throughout this simulation. Before 1000 timesteps, where the change in rate occurred, there was little coverage of the sand particles. It was easy for colloids to get close enough to deposit. At 1000 timesteps there was approximately one layer of colloids surrounding each sand particle, causing the change in deposition rate; it was now more difficult for colloids to deposit on the sand particles. Throughout the rest of the simulation, the clusters gained density, but were *not* growing. The process being modelled was blocking; the deposited colloids did not provide additional surface area for deposit, rather, they hindered further deposition. This serves as evidence that the molecular dynamics model both qualitatively shows the correct linear trends when modelling deposition that causes blocking, but also generates two phases with a threshold specific deposit, agreeing with the existing literature.

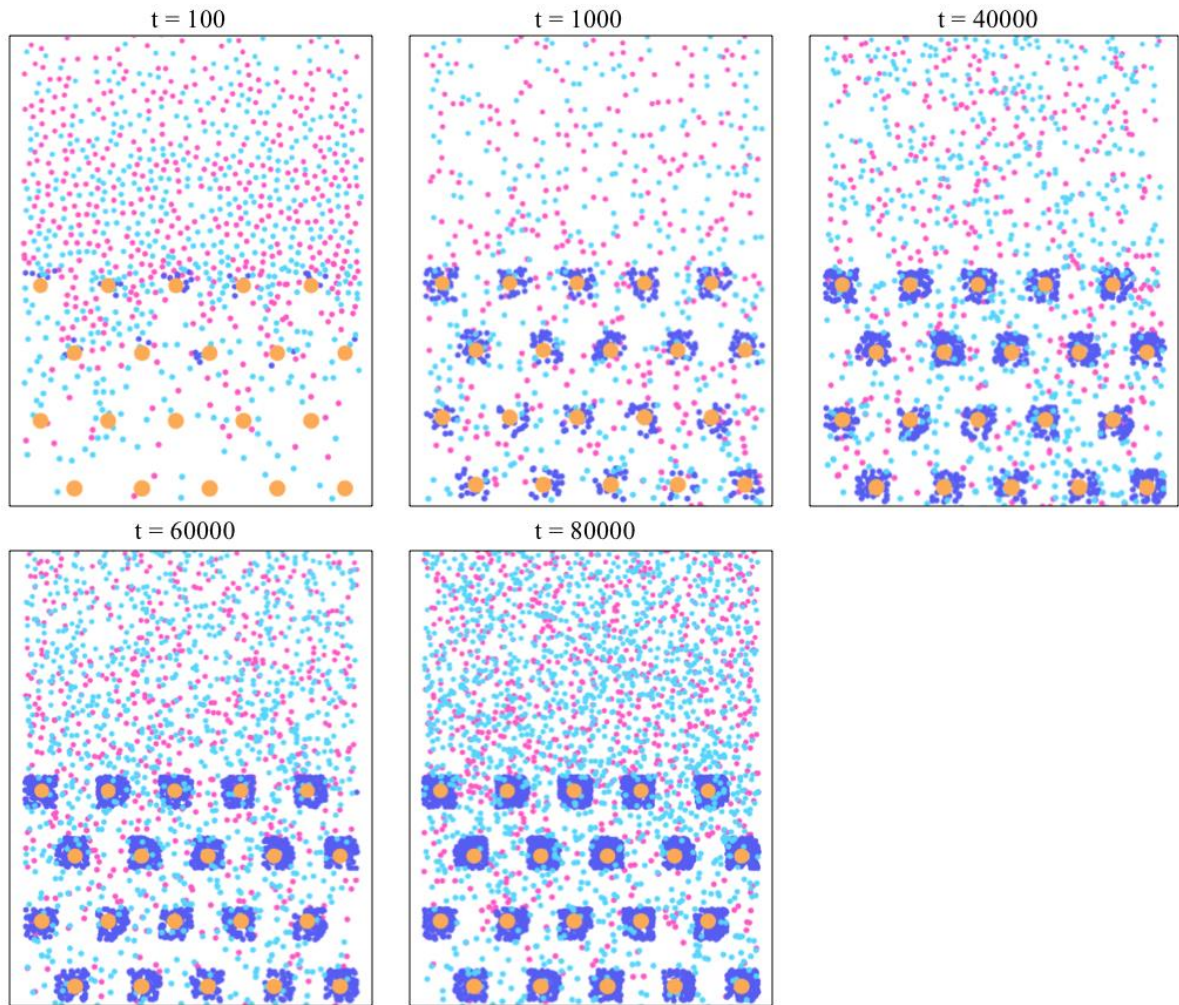


Figure 3.17: Computed particle locations at 5 stages during the simulation shown in figure 3.15.

Figure 3.18 shows the results from a simulation where *ripening* was modelled; colloid particles could stick to already deposited colloids, causing continual cluster growth. Again, there were two distinct regimes, though this time the rate of deposition *increased* after a threshold value of specific deposit. Taking the gradient of the slopes of the two regimes results in values of 1.2×10^{-4} and 2.2×10^{-4} for k_r and k_{ac} respectively.




Figure 3.18: Specific deposit against time showing the two filtration stages. Here ripening is modelled: colloids stick to both sand particles and deposited colloid particles. Colloid clogging distance = 0.5, colloid sticking probability = 0.00.

Figure 3.19 shows the particle locations at varying stages throughout this simulation. At $t = 1000$, there was a slight coverage of each sand particle, which is when the mechanism began to change. Instead of having to deposit on top of the sand particles, the deposited colloids provided additional surface area for deposition, resulting in an increase in the rate. The clusters continued to grow throughout the simulation, decreasing the pore space. At $t = 100,000$ the pore spaces are nearly completely blocked, which would result in mechanical filtration.

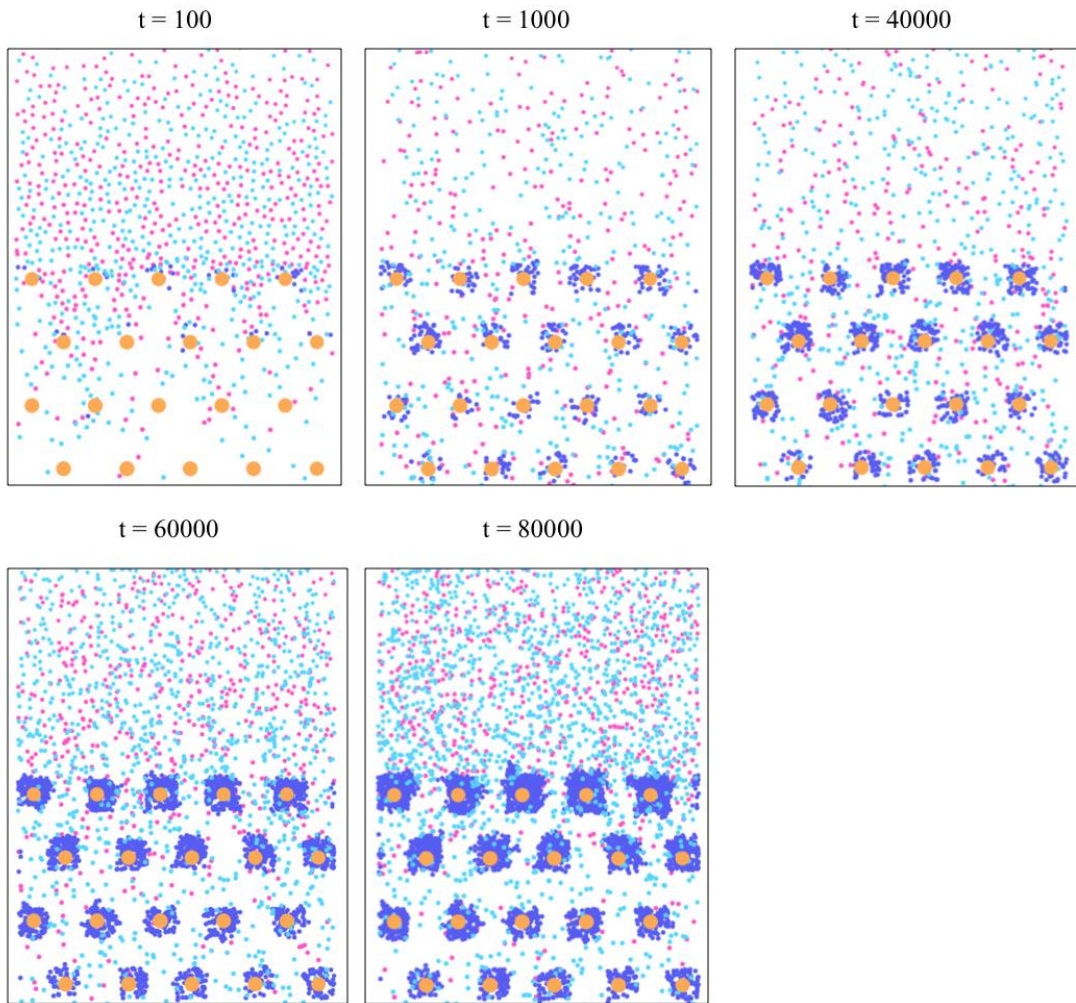


Figure 3.19: Computed particle locations at 5 stages during the simulation shown in figure 3.18.

Again, the model qualitatively agrees with existing literature regarding the build-up of specific deposit: it successfully models the two stages of filtration, where the rate of deposit changes at a threshold value. When blocking was modelled, the rate of deposition *decreased* in the second phase, and when ripening was modelled, the rate of deposition *increased* in the second phase. The gradient of the slopes equate to the rate constants used in NNL's model, and the value of specific deposit at which the rates change equates to σ_r .

3.3.2 Pressure drop

The second area for validation was the pressure drop. Experimental work done by Veerapaneni⁹ suggested that the pressure drop should scale linearly with specific deposit during the initial phase, and scale quadratically with specific deposit during the second stage. Mays and Hunt¹⁰ determined the following relationship between specific deposit and normalised pressure drop:

$$\frac{\Delta H}{\Delta H_0} - 1 = 2\gamma\sigma + (\gamma\sigma)^2 \quad (3.17)$$

where ΔH is the instantaneous pressure drop, ΔH_0 is the pressure drop across a clean filter, and γ is the clogging parameter. This relationship suggests a transition from a linear to a quadratic dependence based on the specific deposit, and is the equation that the NNL model uses to predict the pressure drop.

MD model calculation ΔH_0

The method used to calculate the pressure drop in the molecular dynamics simulations is outlined in section 3.2.5. Pressure drop is often defined as a ratio of the instantaneous pressure drop and the pressure drop across a clean filter. Therefore, a simulation was run to determine the pressure drop of a clean filter using the molecular dynamics model in order to allow future results to be normalised. Instead of the fluid containing a mixture of colloid and water particles, it contained only water particles. This meant that there was no deposition within the filter bed, and therefore, no clogging. Figure 3.20 shows the pressure drop as a function of time. Initially the pressure drop was overestimated, creating the transient peak. This was due to the fact that when the fluid first started to flow through the filter, pressure was *only* created on the top of the filter – this is an artefact of the method of planes as opposed to a physical property of the system. Once the first fluid particles reach the bottom of the filter a more reasonable trend was seen. Once a steady state was reached the value for the pressure drop of this clean filter settled at around $\Delta H_0 = 800$. This value was used to normalise future simulations.




Figure 3.20: Pressure drop against time for a clean filter, used to establish ΔH_0 .

MD Model dependence of ΔH with σ

Figure 3.21 shows the pressure drop for the same sample molecular dynamics simulation as figure 3.16, where the pressure drop has been normalised. The results are in excellent agreement with the trends predicted by Veerapaneni and Mays and Hunt. During the first phase, there was indeed a linear increase in the pressure drop with specific deposit. After 10000 timesteps, where the mechanism changed, the dependence became quadratic, again agreeing with equation 3.11. This change occurred at the *same* threshold value of specific deposit seen previously in section 3.3.1. The model not only successfully shows the two regimes expected of a filtration process, it also shows a transition from a linear to a quadratic relationship between pressure drop and specific deposit. Based on the linear fit in the first phase, this result

suggests a clogging parameter of 1.4, and based on the quadratic fit for the second phase, a clogging parameter of 4.0.




Figure 3.21: Pressure drop as a function of specific deposit. The linear fit describes the relationship between pressure drop and specific deposit during the first phase. The quadratic fit shows the dependence between pressure drop and specific deposit in the second phase.

3.3.3 Concentration with depth

A constant first order rate of deposition predicts that the concentration of deposited colloids will decrease exponentially with depth¹¹. However, in conditions where there are multiple types of colloid particle, or a large particle size distribution, specific colloids can be preferentially deposited at the top of the sand bed resulting in a hyper-exponential decay. Additionally, if there are only weak forces (such as those predicted under certain conditions by DLVO theory) there is the potential for deposited colloids to remobilise and deposit further down the bed resulting in a nonmonotonic profile (see figure 3.22).




Figure 3.22: Illustrations of deposit profiles as a function of depth.

Left: exponential, middle: hyper-exponential, right: non-monotonic.

Adapted from Bridge.¹²

The NNL model did not provide any information regarding the concentration of deposited colloids with depth, therefore the only means of available validation in this area was with the existing literature. This highlights the key advantage of a particle-based approach: additional mechanistic information can be easily obtained.

MD model prediction of concentration with depth

Figure 3.23 shows the concentration of deposited colloids as a function of depth at a porosity of 0.7 obtained using the molecular dynamics model. The specific deposit decreased exponentially with depth, which is in agreement with the experimental predictions of Jegetheesan.¹¹




Figure 3.23: Specific deposit against depth obtained using the MD filtration model. Porosity = 0.7. Error bars are negligibly small. The solid line is a linear fit.

When the porosity of the filter was decreased (by increasing the packing fraction of the sand) the trend became hyper-exponential. Figure 3.24 shows the concentration profile where the porosity was decreased to 0.4. This decrease in porosity caused colloids to be deposited favourably at the top of the bed, changing the deposit profile.

A non-monotonic profile cannot be observed using the current molecular dynamics model as this occurs when colloids remobilise and redeposit further down the bed; sticking is permanent in this model.




Figure 3.24: Specific deposit against depth obtained using the MD filtration model. Porosity = 0.4. Error bars are negligibly small. The solid line is an illustration of an expected linear trend.

3.2.4 Summary

A series of tests were performed using the molecular dynamics model to validate it against existing literature and modelling work. The model showed the two expected phases of filtration: ripening/blocking and operation. Additionally, the results showed that the specific deposit increased linearly with time, with the rate of increase changing at a threshold value of specific deposit. Similarly, the results showed a change from a linear increase in pressure drop to a quadratic dependence at the same threshold value. These trends agree with both experimental work and with the assumption inherent in the NNL model. Furthermore, the simulations show an exponential decrease in deposited colloid concentration with depth, that becomes hyper-exponential under certain conditions. This also agrees with experimental results.

3.4 Systematic exploration

Having used the data from available literature to try to validate the results and trends from the MD model, the next stage was to systematically explore this model to further probe the mechanism of clogging. The following sections will show the results gained from varying parameters in isolation.

3.4.1 Interatomic Potentials

One of the primary areas for investigation of the molecular dynamics model was to examine how the nature of the interatomic potentials effects the dynamics. It has already been shown that the soft core repulsive potential being used is capable of causing phase separation, in a similar way to the depletion force. The effect of the strengths of the interatomic potentials was investigated by simply varying each interatomic pair potential in isolation.

Colloid-colloid potential

The first pair potential investigated was the colloid-colloid potential, which defined both the interactions between colloids suspended in the fluid, and the interactions between deposited colloids and those suspended in the fluid. ϵ_{cc} was varied between 2.0 and 100.0, and the resulting rates of build-up of specific deposit and rate of increase in pressure drop were measured. Figure 3.25 (a) shows the effect on specific deposit, and figure 3.25 (b) shows the effect on the pressure drop.

It can be seen that the initial ripening phase showed little dependence on the colloid interaction strength, with the rate of deposition being constant throughout. This was expected as the mechanism for deposition depends solely on the strength of the interaction between the sand and the colloid particles. However, once this phase had completed and there was a layer of colloidal particles surrounding the sand particles, further deposition was prohibited by a strong repulsive colloid/colloid interaction; the rate of deposition *decreased* with an *increasing* ϵ . The stronger repulsive forces kept the separation between the colloids larger, allowing fewer of them to get close enough to the sand particles to interact and deposit.

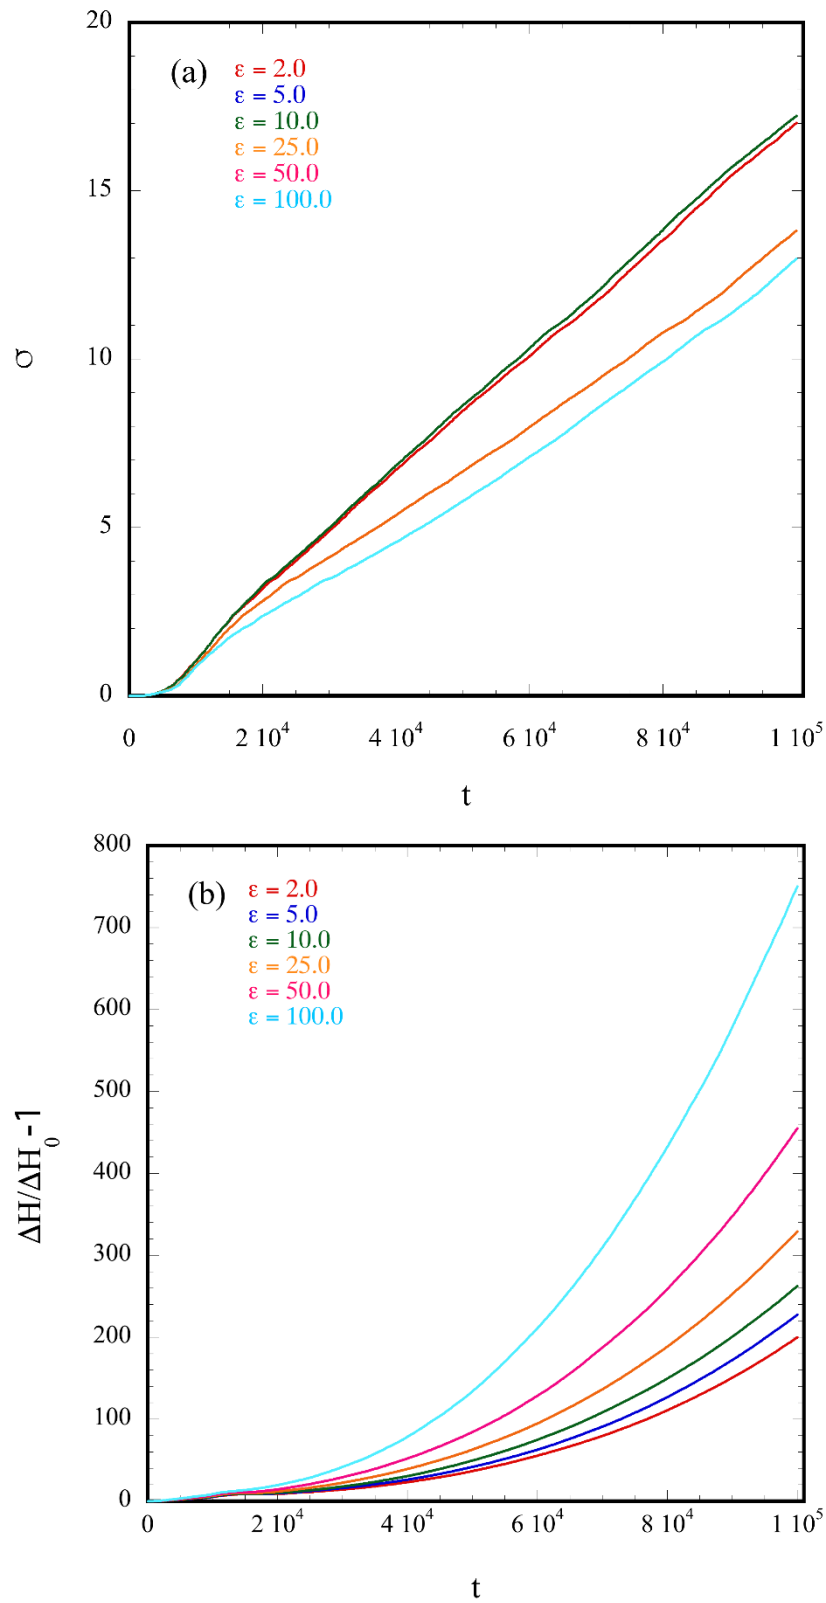


Figure 3.25: Specific deposit (a) and pressure drop (b) against time whilst varying the strength of the colloid-colloid pairwise potential from $\varepsilon = 2.0$ to $\varepsilon = 100.0$.

Even with a lower value of specific deposit, the pressure drop increased with increasing colloid repulsion. The strong repulsive forces not only hindered further deposition, but also stopped colloid particles from passing through the bed, causing an increase in pressure due to the increased number of particles above the bed.

Figure 3.26 shows the density profiles obtained from the simulations of the highest and the lowest interaction strength. It is clear that the hyper-exponential nature of the deposit increased with increasing colloid repulsion: the strong repulsive interactions caused a sharper drop-off in the density profile. This means that colloids were being deposited favourably towards the top of the filter compared to the bottom; the additional repulsion between colloids hindered them from penetrating further into the bed, resulting in a higher density of deposit at the top of the bed. Even though, overall, there was less deposit within the filter with higher colloid-colloid repulsion, the pressure drop was higher due to the favourable deposition at the top of the bed.




Figure 3.26: Density profiles comparing the weakest (left) and strongest (right) colloid-colloid interactions from figure 3.1. The density includes contributions from all fluid particles.

To maximise efficiency in the industrial process, it is therefore ideal to create a situation where the colloid-colloid interaction is not so unfavourable that the deposited

colloid particles not only block other colloids from depositing, but also stop them penetrating further into the bed causing a premature head loss. Having strong colloid-colloid repulsion both hinders deposition, increases the pressure drop, and causes favourable deposition at the top of the bed.

Colloid-sand potential

The second potential investigated was the colloid-sand potential, which affected how close the colloid particles could get to the sand particles. The strength of the potential was increased from 2.0 to 100.0, and the rates of build up of specific deposit and the increase in pressure drop were measured.

Figure 3.27 shows the results of this exploration. It is clear that there was no significant deviation in either the rate of deposition (a) or the increase in pressure drop (b). Once the sand particles had a layer of colloid particles surrounding them, and the second phase began, the rate-limiting interaction for the rate of deposit was the strength of the colloid-colloid interaction, as seen earlier.

However, what is of more interest is that the initial ripening phase also showed no dependence on the strength of the colloid-sand interaction. It may seem natural to assume that a stronger repulsive force between the sand and the colloid would cause the ripening phase to lengthen. This is likely to be due to the fact that the sticking probability (which is akin to a chemical bond) dominates the weaker van der Waals forces between sand and the colloids.

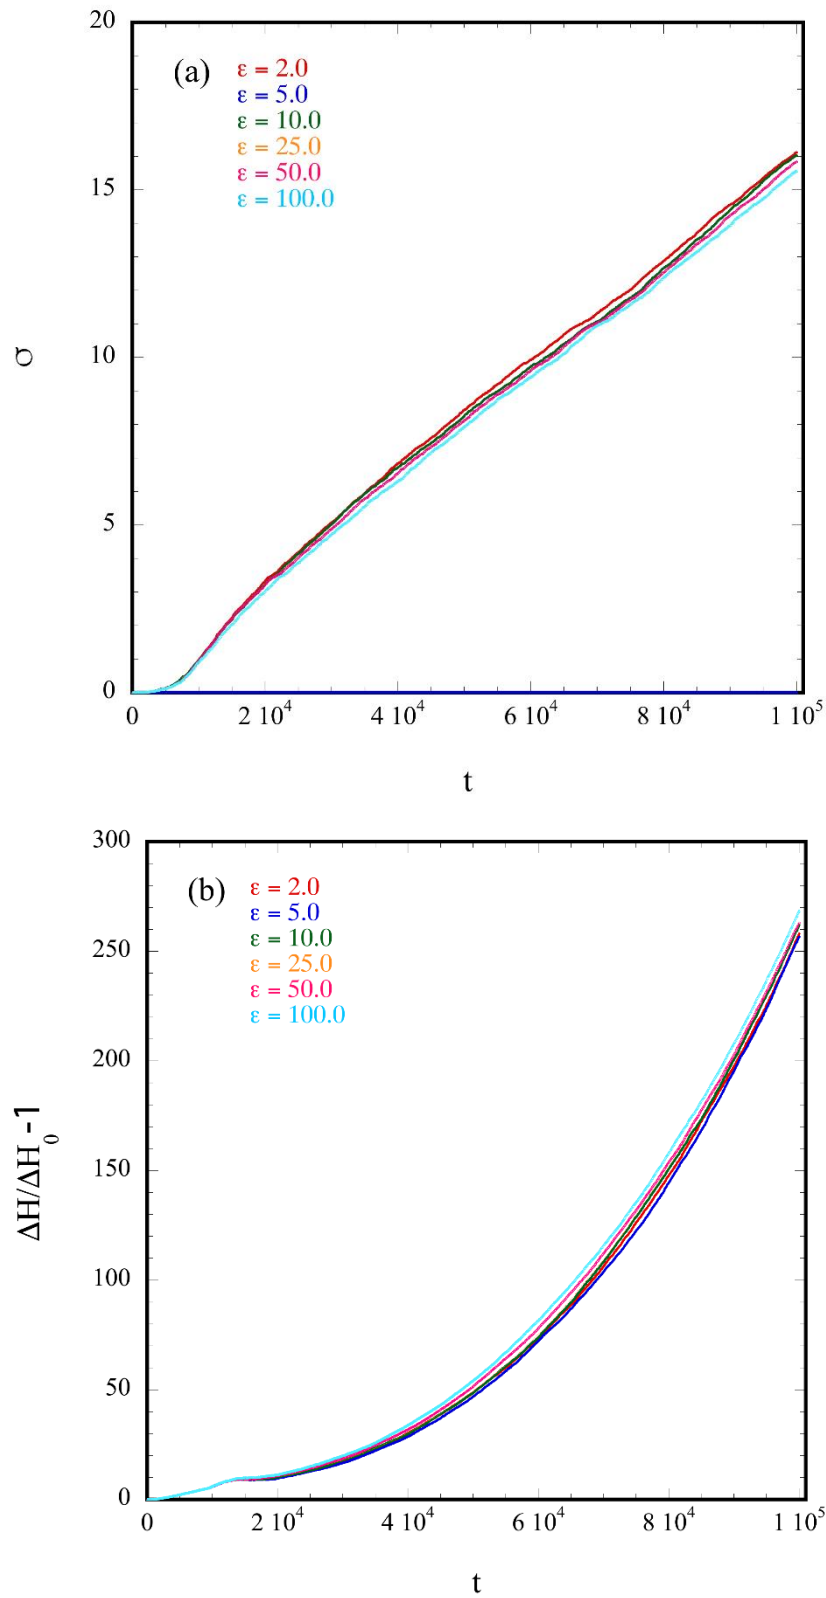


Figure 3.27: Specific deposit (a) and pressure drop (b) against time whilst varying the strength of the colloid-sand pairwise potential from $\epsilon = 2.0$ to $\epsilon = 100.0$.

Fluid-colloid potential

The fluid-colloid interaction was investigated in a similar manner to the previous examples. The strength was varied from 2.0 to 100.0, and the results are shown in figures 3.28 (a) and 3.28 (b).

As the repulsion between the colloid and fluid increased, the pressure drop also increased; the colloids blocked the flow paths, and the stronger repulsive forces stopped the fluid from passing. However, what is of more interest in this case is the specific deposit data; there was an increase in the rate of deposit from an epsilon value of 2.0 to 25.0, then a sharp decrease in the rate of deposit from 25.0 to 100.0.

Figure 3.29 shows snapshots of these simulations. In image *a*, where the interaction strength between the fluid and colloid ($\epsilon = 2.0$) was actually weaker than that colloid-colloid interaction ($\epsilon = 10.0$), it was favourable for the fluid to pass down the column more quickly than the colloids. The pressure build up comes from the colloids, as opposed to the fluid, struggling to penetrate, which can be seen from the dominance of the red particles in the image. As the fluid-colloid interaction strength increased, the fluid particles contributed more to the pressure build up as they too struggled to pass through the filter. At $\epsilon = 50.0$ (f), and subsequently $\epsilon = 100.0$ (g), the fluid blocked all of the available flow paths towards the top of the sand bed, meaning that not only was there an increase in the pressure drop, there was also a decrease in the rate of deposit. The filter was effectively clogged, not by a layer of filtrate causing mechanical clogging as seen previously, rather, the repulsion between the clogged particles and the fluid stopped the fluid from penetrating into the sand bed, blocking all available flow paths.





Figure 3.28: Specific deposit (a) and pressure drop (b) against time whilst varying the strength of the colloid-fluid pairwise potential from $\epsilon = 2.0$ to $\epsilon = 100.0$.




Figure 3.29 (part 1): Snapshots of the particle locations after 100,000 timesteps for the simulations in figure 3.27, with increasing fluid-colloid repulsion from left to right. The red particles are colloid, the light blue fluid, and the dark blue are clogged.




Figure 3.29 (part 2): Snapshots of the particle locations after 100,000 timesteps for the simulations in figure 3.27, with increasing fluid-colloid repulsion from left to right. The red particles are colloid, the light blue fluid, and the dark blue are clogged.

Pairwise potential - summary

The results of the investigation surrounding the pairwise potentials used to model the processes involved in filtration have yielded some interesting developments.

Firstly, it has been shown that the rate of deposit depends greatly on the interactions between the fluid and the colloid particles, showing little dependence on any interaction with the sand particles. Once a layer of colloid particles surrounds the sand particles (which is a relatively rapid process), the mechanics are dominated by the interactions of those colloid particles with the fluid passing them. If there is a strong repulsion hindering the trajectories of the colloid particles then the fluid passes easily down the bed and the colloids struggle to penetrate. This results in a preference for the particles to deposit at the top of the bed, leading to an increased pressure drop and, ultimately, clogging. A hyper-exponential deposit profile is an indicator of this. If the opposite is true, and the fluid particles suffer from a larger repulsion, colloid particles travel deeper into the bed resulting in a linear deposit profile, and a pressure build up from the fluid particles.

In terms of the mechanisms involved in complete clogging, this model suggests that there are two distinct mechanisms by which this occurs. The first is mechanical clogging, which occurs when deposited colloids completely block the flow paths between the sand particles. The other is not caused directly by colloid particles clogging the pore spaces, but is caused by a backlog of fluid particles that are unable to travel down the bed, and is shown in figure 3.29. A small amount of deposited colloid particles that strongly repel the fluid cause the pore spaces to become blocked with fluid particles unable to flow down the bed, even though the pore space still exists to travel through. This results in the same indicators of clogging; there is a large increase in pressure drop, and also there is a decrease in the rate of deposit. The clogging could be further investigated using a model that had *different* sizes for the colloid and fluid particles; a fluid made up of many fluid particles per colloid would allow further probing of the diffusion of fluid particles through a clogged pore space, though such a simulation (with differing timescales of movement of colloid and fluid) would require considerably more computational power.

3.4.2 Lattice structure

The lattice structure of the sand particles is a variable that can be easily adjusted. Two sets of simulations were run, where the only difference was the structure of the lattice. A square lattice, where each sand particle has 4 nearest neighbours, and a triangular lattice, where each sand particle has 6 nearest neighbours, were investigated.

The results are shown in figure 3.30. To give a range of simulations for comparison, the lattice structure was compared for a series of simulations where the strength of the colloid-colloid interactions was varied from 2.0 to 100.0. This choice was motivated by the fact that this parameter has a large influence on both the specific deposit and pressure drop. Both sets of simulations were run with the same packing fraction, meaning that *only* difference between the two series was the coordinates of the sand particles. The results show that in the case of a square lattice both the rate of deposit and pressure drop varied remarkably little with interaction strength. This is in stark contrast to the results when the sand bed has a triangular lattice.

Figure 3.31 shows the coordinates of the clogged particles of the highest interaction strength for the two lattice structures, where the red dashed line is a hypothesised flow path for the fluid. This schematic can help to explain the observations above: as the deposits build up around the sand particles, a much simpler flow path is created when the sand particles are in a square lattice. In a triangular lattice, there is more interference from the clogged particles causing the increase in pressure drop, whereas in a square lattice, the fluid can flow through the channels relatively easily.

In terms of the lower rate of deposit build up in a square lattice, it can be seen that a colloid that travels around a sand particle is less likely to come close enough to a sand particle in the next layer to interact. In a triangular lattice, it is more difficult for the particles to flow through the layers without interacting with the sand, and will therefore deposit at a faster rate.

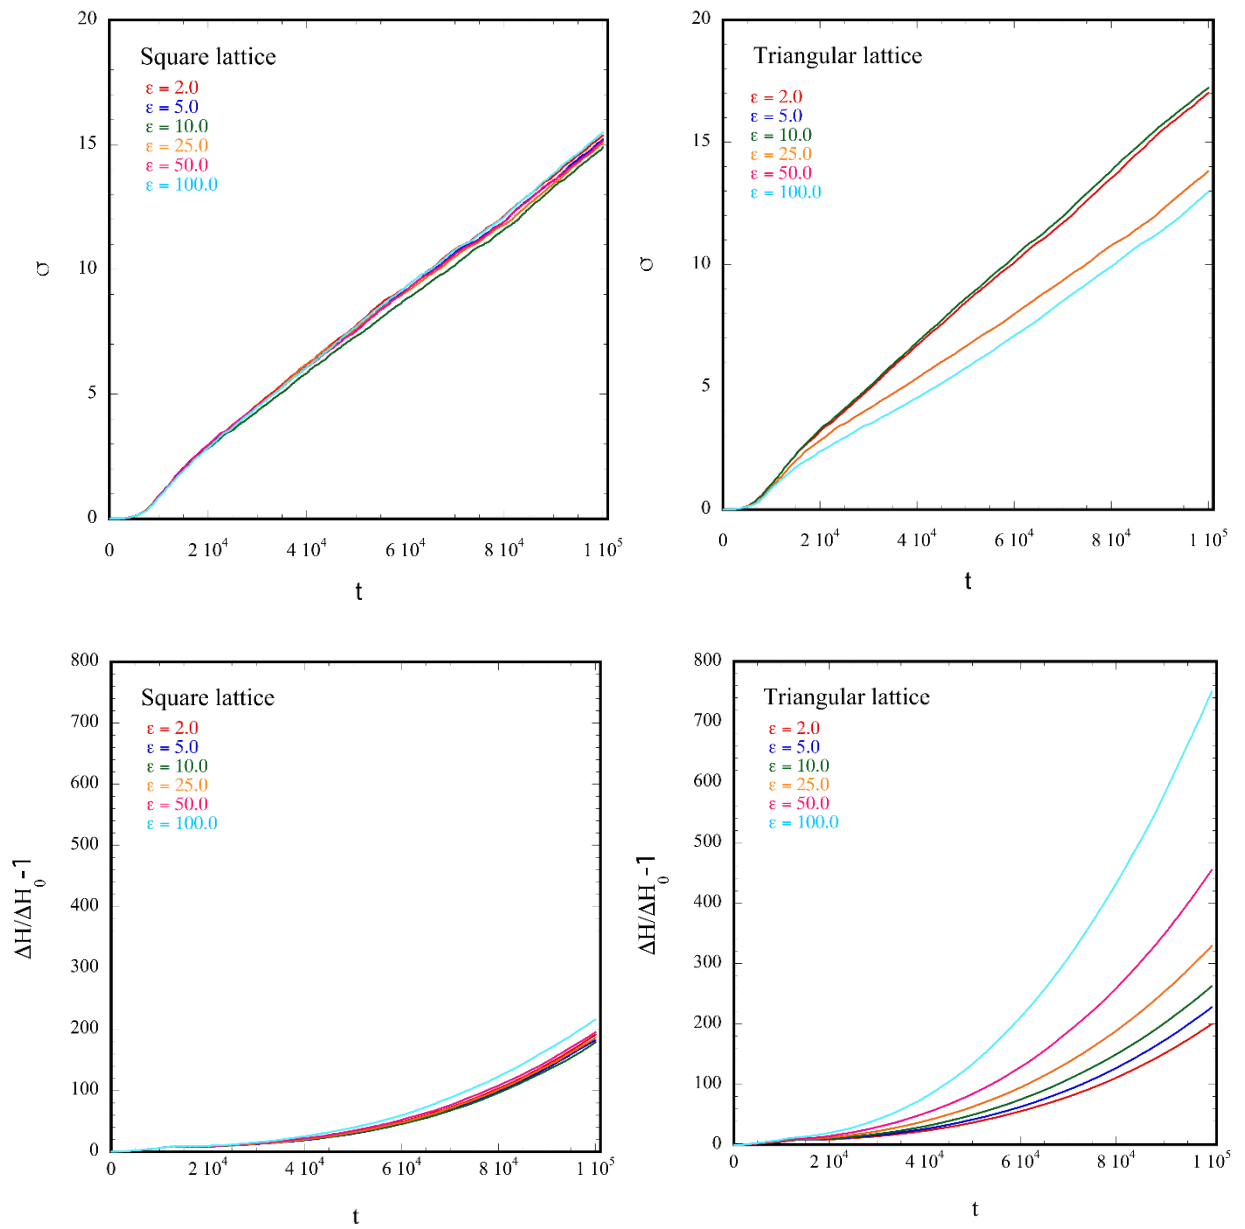


Figure 3.30: Comparison of specific deposit against time (top) and pressure drop against time (bottom) whilst varying the colloid-colloid interaction strength for a square lattice (left) and triangular lattice (right)




Figure 3.31: Comparison of the coordinates of clogged particles in a triangular lattice (left) and a square lattice (right). The dashed line shows a suggested flow path.

3.4.3 Colloid concentration

The ratio of colloid/fluid particles is an important parameter, yet it is difficult to control on plant. This is, however, easy to vary within a molecular dynamics simulation. On plant, the concentration of colloids in the effluent waste is in the region of parts per million. Due to time restrictions it is not possible to model a situation where the concentration of colloids is as low as this; a 4% colloid solution, simulating 2500 particles (total of fluid + colloid particles), would run for over 24 hours without producing significant clogging. A series of simulations was conducted in which the colloid concentration was increased from a minimum value of 10%.

Figure 3.32 show the results. As expected, both the specific deposit (a) and pressure drop (b) increased with increasing colloid concentration, showing the same trends as previously. The concentration of the colloid particles had little mechanical effect; having a higher concentration simply speeds up the clogging process.

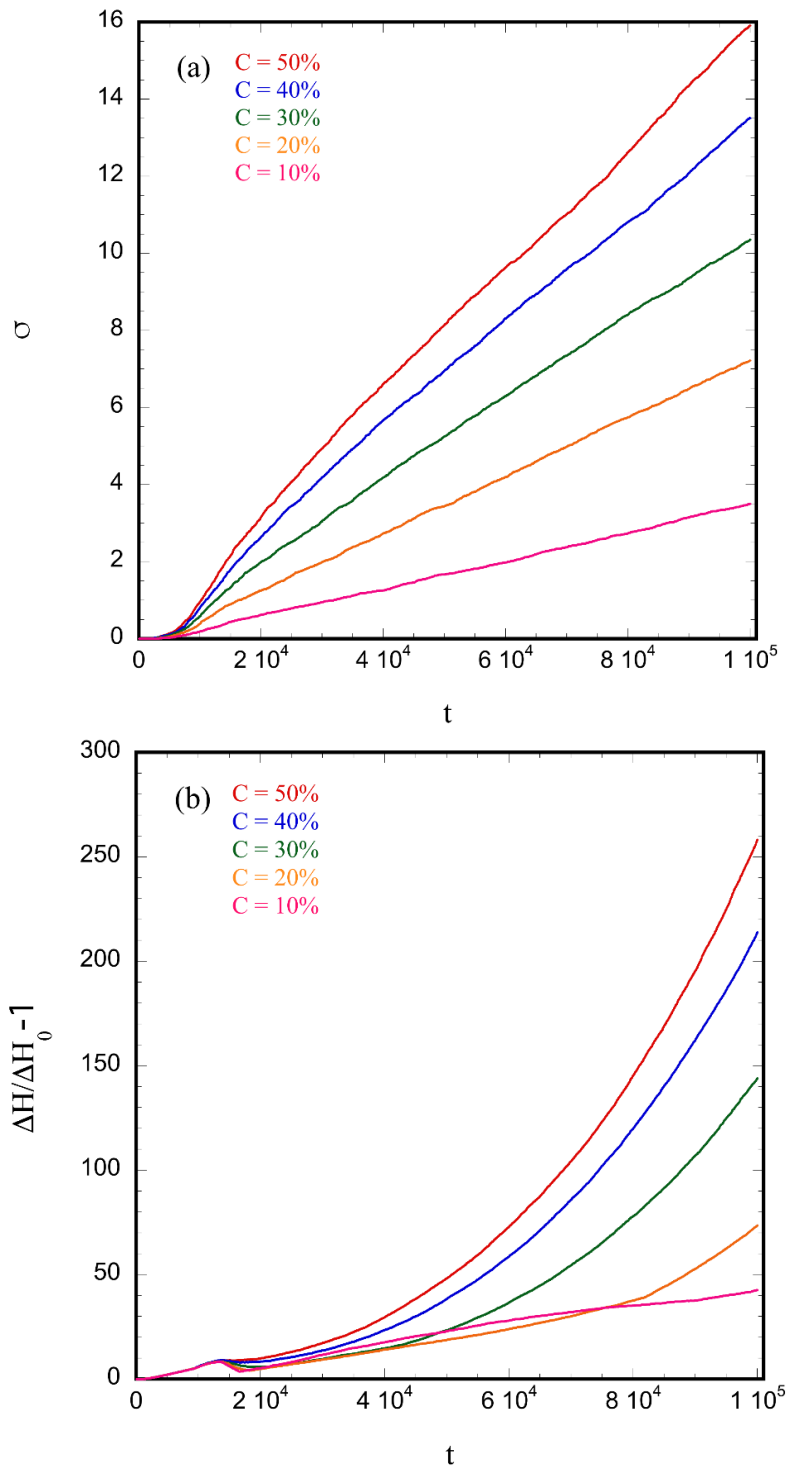


Figure 3.32: Specific deposit (a) and pressure drop (b) against time whilst varying the concentration of colloid particles from 10% to 50%. The apparent decrease in pressure drop is a transient artefact of the normalisation process, emphasised by lowering the concentration.

3.4.4 Porosity

Porosity, ϕ , is a measure of the empty space within a material and is expressed as a fraction between 0 and 1 or as a percentage between 0 and 100. Being able to measure the porosity of the filter is important in determining the effect that the packing fraction of the sand particles has on the filtration process.

A quick and simple way to measure porosity is to probe the coordinate space using random sampling. A test particle with randomly generated coordinates is placed within the simulation cell. If the coordinates of the particle lies within a circle of radius σ_{sand} centred on a sand particle, then it is considered to be a hit, otherwise it is a miss. This process is repeated to cover the coordinate space. Figure 3.33 show an example of the results obtained from this method.




Figure 3.33: Example results from probing porosity. The purple dots are misses, and the blue dots are hits. 1,000,000 trial insertions were performed.

In order to gain an accurate value for the porosity it is important the enough samples were made to extensively cover the coordinate space. Figure 3.34 shows the results of a series of simulations where the number of trial insertions was increased. It is clear that the results start to converge between 10,000 and 100,000 particles. Therefore, it was decided that 100,000 trial insertions was more than sufficient to obtain an accurate value.




Figure 3.34: Porosity (as a percentage) as a function of number of test particle trial insertions created, N .

A series of simulations were run where the porosity was varied from 1.44% to 70.57% and all other variables were kept constant, the results of which are shown in figure 3.35. As expected, the pressure drop increased with decreasing porosity. However, what is of more interest is the rate of deposition. Instead of there being a constant increase in the rate of deposit with decreasing porosity, there was actually a point at which the opposite effect was seen.





Figure 3.35: Specific deposit (a) and pressure drop (b) against time whilst varying the porosity of the sand bed from 1.44 to 70.57.

At a porosity of less than 42.24, the rate of deposition decreased with increasing porosity, even though the rate of pressure drop continued to increase. By examining the coordinates of the clogged particles (figure 3.36) it can be seen that the mechanism of deposition changed at this porosity. Rather than forming clusters around the sand particles, the entire filter became mechanically clogged. Colloids could not penetrate past the top layer of sand easily, causing a decrease in the rate of deposition, though the rate of pressure drop still increased.

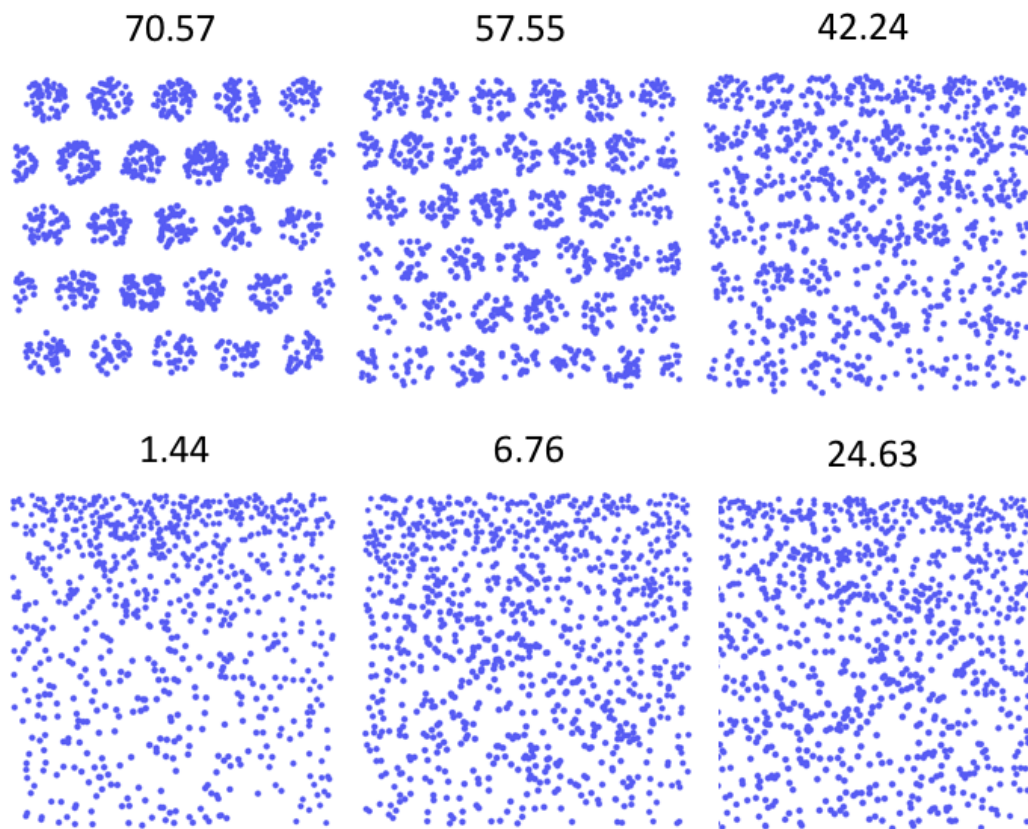


Figure 3.36: Image showing the coordinates of the clogged particles for the simulations shown in figure 3.36, with porosity decreasing in a clockwise direction.

Figure 3.37 shows the pressure drop as a function of porosity at $t = 90,000$. As can be seen, this model predicted a quadratic decay, with the data converging much more quickly at lower porosity. Once the deposit mechanism changed from forming clusters to forming a solid bed (see figure 3.36), the dependence between porosity and pressure drop decayed.




Figure 3.37: Pressure drop as a function of porosity. The data is taken from figure 3.35, where $t = 90,000$.

3.4.5 Fractal dimension

Non-fractal objects have a definite topological dimension; a line is one-dimensional, a circle is two-dimensional, and a sphere is three-dimensional. Fractals have a dimension that lies between the usual topological dimensions. For example, a fractal object embedded in a two-dimensional space has a fractal dimension between one and two. Fractals result from one of two methods. The first method produces self-similar fractals, which are fractals built up using a deterministic set of rules. The Koch curve is an example of a self-similar fractal (see figure 3.38). It results from an iterative process where, initially, a single line has its middle third removed and replaced with two equidistant lines forming an equilateral triangle with no base. This process repeats. After each iteration, the scale is reduced by a factor of three, and four new objects that each resemble the whole at the start of the current iteration are created.

This means that fractal dimension is $\ln(4)/\ln(3) = 1.26$, and the shape looks identical at any scale.




Figure 3.38: An illustration of the Koch Curve, a deterministic fractal created using an iterative process.

The second method results in disordered fractals, where the self-similarity only arises when looking at the average properties. Processes that are random or chaotic can result in disordered fractals, the coastline of Britain being one such example.

It is hypothesised that the fractal dimension of the deposits within a filter will have an impact on the flow rate¹²; a high fractal dimension, where the deposits tend towards a two-dimensional shape, will result in less physical and chemical interruption to the flow compared to low fractal dimension deposits. This is because a high fractal dimension will divert the flow through colloid-free pores, as opposed to being more distributed throughout space, and is illustrated in figure 3.39.




Figure 3.39: Illustrations of deposits with a low (left) and high (right) fractal dimension, adapted from Mays¹⁵.

Measuring fractal dimension

The fractal dimension arises from the fact that fractals are shown to follow scaling laws, whereby the number of particles within a cluster scales as a power law in the radius of the object.

$$N \sim R^D \quad (3.18)$$

where N is the number of particles, R is the radius of the cluster and D is the fractal dimension. This relation can be linearised by taking logarithms, giving:

$$\ln N = D \ln(R) + \ln(C) \quad (3.19)$$

where C is a constant. The fractal dimension may then be obtained from the slope of a linear least squares fit to this logarithmically transformed data.

The radius of gyration is used to measure the radius of the cluster as a function of the number of particles. This refers to the distribution of the particles around the centre of mass of the cluster. It is defined for a cluster with N particles as:¹³

$$R_{ij}^2 = \frac{1}{N} \sum_{\alpha=1}^N (x_i^\alpha - x_i)(x_j^\alpha - x_j) \quad (3.20)$$

where x_i is the i th coordinate of the centre of mass of the cluster (and ranges from 1 to 2 in two-dimensions), and $\alpha = 1, \dots, N$ denotes the particular particle within the cluster. The tensor can be diagonalised to obtain its eigenvalues, λ_i , $i = 1, \dots, d$. For $d = 2$, which is the case in this model, the asphericity parameter, which is a measure of the shape of the cluster, A_2 , is defined as:

$$A_2 = \frac{\langle(\lambda_1 - \lambda_2)^2\rangle}{\langle(\lambda_1 + \lambda_2)^2\rangle} \quad (3.21)$$

where the angled brackets represent an average over multiple clusters. Calculating these values gives important information about the shape of the cluster. For a circular cluster, $A_2 = 0$ as $\lambda_1 = \lambda_2$, whereas $A_2 = 1$ for a straight rod ($\lambda_1 = \lambda, \lambda_2 = 0$). If $\lambda_1 \neq \lambda_2$, the ratio of the eigenvalues gives the lengths of the axes the ellipsoid.

Before attempting to determine whether the molecular dynamics model shows a relationship between the fractal dimension of deposits and the rate of deposition or pressure drop, it was first necessary to determine whether the deposits do indeed have a fractal dimension. As has been discussed previously, there are two mechanisms by which particles can stick within the filter. The first mechanism only allows colloids to deposit directly onto the sand particles, and the second allows colloids to additionally deposit onto already deposited colloids. In order to determine the radius of gyration and the number of particles in each cluster the simulation must record which sand particle each colloid has deposited onto. In the case of the first mechanism this is simple. However, in the case of the second mechanism this requires a more complex recursive process, illustrated in figure 3.40. This algorithm occurs after the distance and sticking probability calculation, and therefore only happens once a particles is definitely going to stick.

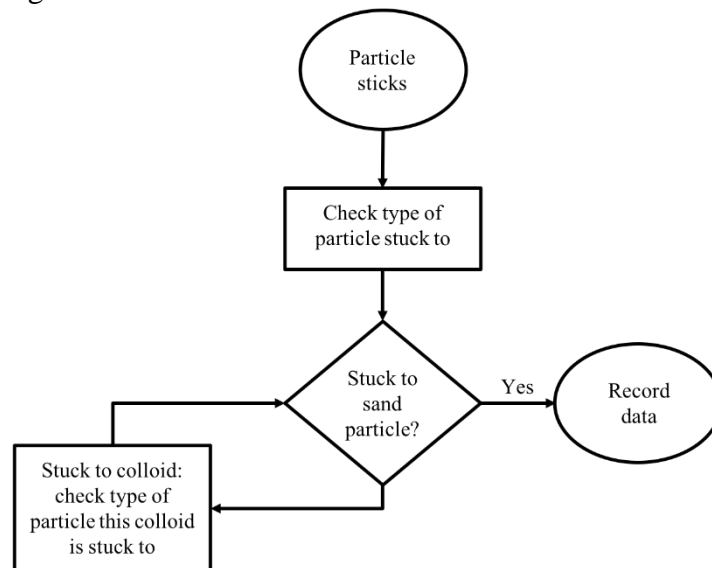


Figure 3.40: Flow diagram showing how to determine which sand particle a colloid has stuck to.

Fractal dimension in molecular dynamics simulations

In order to determine whether the deposits have a fractal dimensions a series of simulations were run in both the blocking and ripening regime. Figure 3.40 shows the deposited particle locations in the ripening regime where the colloid clogging distance was 0.3. As can be seen, the clusters continue to grow with time, and form patterns that show the expected features of fractal shapes.

Figure 3.42 shows the log-log plot of the number of particles in each cluster against the radius of gyration. It is clear that all of the clusters showed the linear relationship expected of a fractal shape. The gradient of each fit is shown in table 3.1. For the first four clusters, the fractal dimension is significantly far from both 1.0 and 2.0, giving good evidence that the clusters *do* have a fractal dimension. Cluster 5 has a dimension of 1.044, suggesting that this cluster does not have a fractal dimension.

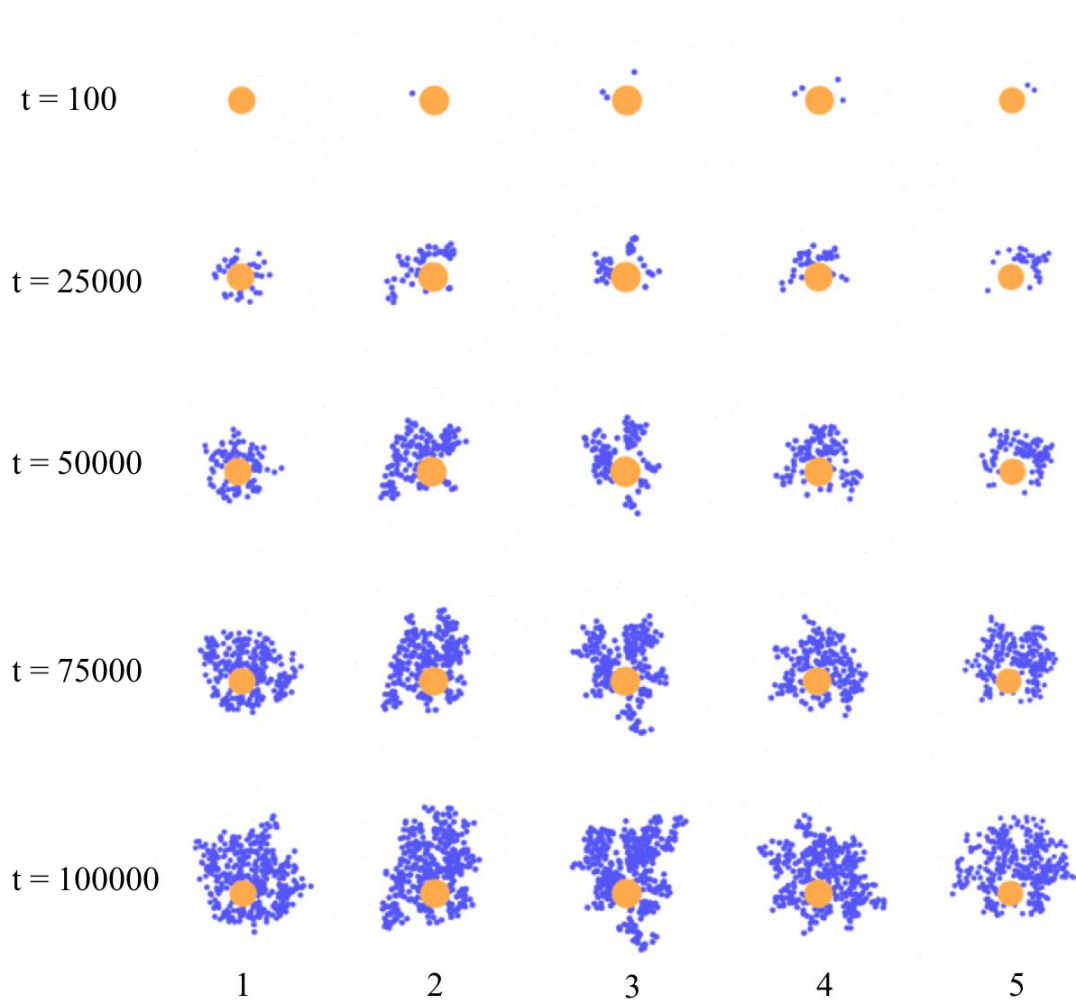


Figure 3.41: Snapshots of the deposited colloid particles on five sand particles at varying time throughout the simulation. Colloid particles could deposit on top of other colloid particles. The colloid clogging distance was 0.3.




Figure 3.42: Log-log plot of number of particles against radius of gyration. The colloid clogging distance is 0.3. The fractal dimension is taken from the gradient of the slope.

Cluster	Fractal dimension
Cluster 1	1.564
Cluster 2	1.724
Cluster 3	1.784
Cluster 4	1.614
Cluster 5	1.044

Table 3.1: Fractal dimension for each of the five clusters. The fractal dimension is the slope of the linear fit shown in figure 3.40.

Figure 3.43 shows the cluster growth when the colloid clogging distance was increased from 0.3 to 0.4. The colloids can now stick to each other at a slightly longer distance causing a faster growth in the cluster size.

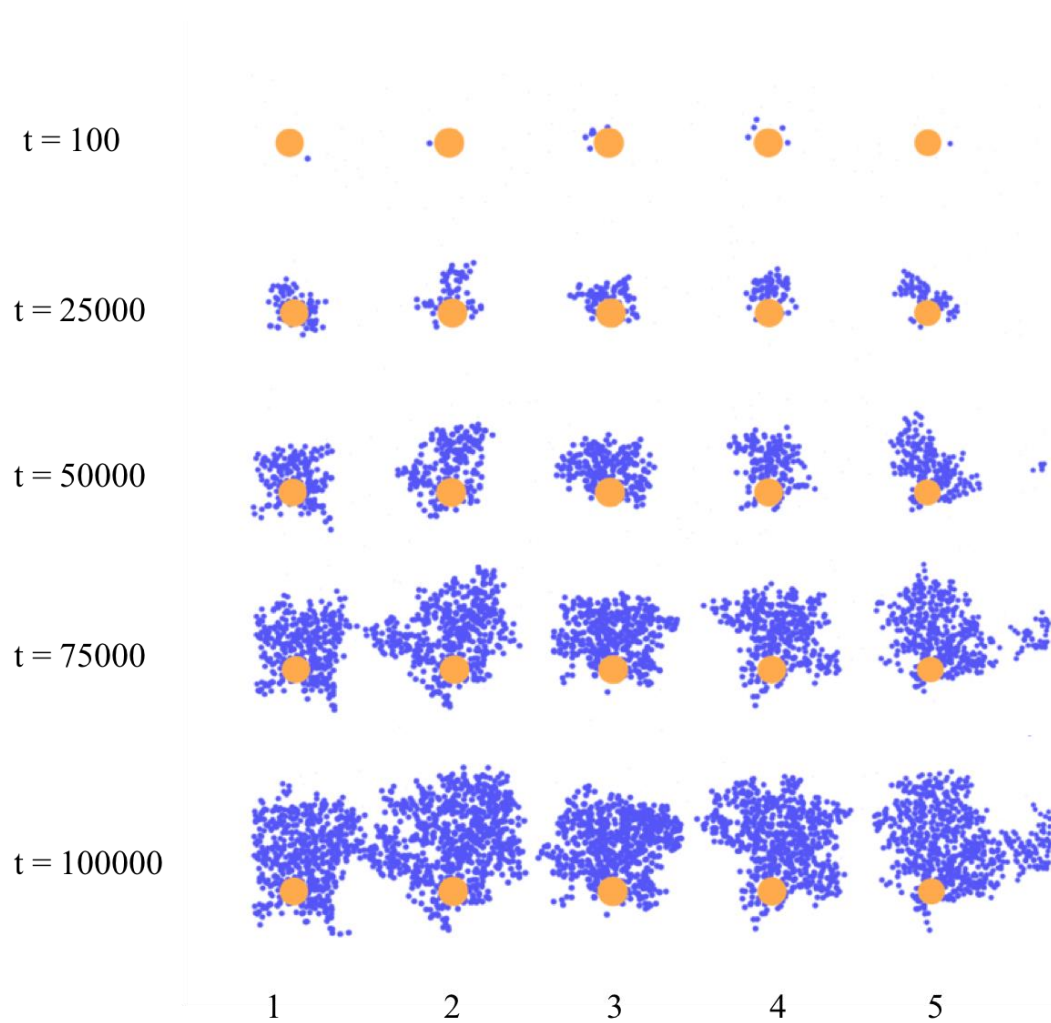


Figure 3.43: Snapshots of the deposited colloid particles on five sand particles at varying time throughout the simulation. Colloid particles can deposit on top of other colloid particles. The colloid clogging distance is 0.4.

As figure 3.44 shows, the clusters again show a linear relationship between number of particles and the radius of gyration, though there was a significant reduction in the average fractal dimension (approximately 20 %): an increase in the separation distance at which the colloids can stick to each other causes a decrease in the fractal dimension. Importantly, this gives a direct means of testing the theory proposed by Mays (see figure 3.37).




Figure 3.44: Log-log plot of number of particles against radius of gyration. The colloid clogging distance is 0.4. The fractal dimension is taken from the gradient of the slope.

Cluster	Fractal dimension
Cluster 1	1.618
Cluster 2	1.367
Cluster 3	1.267
Cluster 4	1.179
Cluster 5	1.374

Table 3.2: Fractal dimension for each of the five clusters. The fractal dimension is the slope of the linear fit shown in figure 3.42.

A series of simulations were run where the colloid sticking distance was reduced from 4.0 to 3.0. The average fractal dimension was measured in addition to the rate of deposition and the rate of pressure drop. Figure 3.45 shows the dependence of the fractal dimension of the clusters against the sticking distance. There is a general decrease in the fractal dimension with increasing sticking distance.




Figure 3.45: Fractal dimension against sticking distance.

The two data sets with the highest (1.671) and lowest (1.296) fractal dimensions were chosen for comparison. Figure 3.46 shows the difference in rate of deposition for the two simulations. In the case of a lower fractal dimension, a faster rate of deposit was observed. This in itself is not surprising; the longer sticking distance means colloids are more likely to stick.




Figure 3.46: Specific deposit as a function of time for two data sets, with fractal dimensions of 1.671 and 1.296.

A better means of comparison is to compare the number of fluid particles that pass through the filter for a given specific deposit. The fluid particles are not affected by the change in sticking distance, therefore any observed differences cannot be attributed to this change. Figure 3.47 shows this comparison. It is clear that for the *same* value of specific deposit, there is a significant difference between the number of fluid particles that are able to pass through the filter. The deposits with a *lower* fractal deposit cause more interference, allowing *fewer* fluid particles to pass through. This agrees with the theory proposed by Mays.




Figure 3.47: Concentration of fluid passing through the filter against specific deposit.

3.5 Summary

A coarse grained molecular dynamics model was used to investigate the process of clogging of sand bed filters. The model was validated against existing literature, experimental work, and simulations, and qualitatively agreed in several key areas. Firstly, the amount of specific deposit was shown to increase linearly with time, showing two distinct regimes. The linear rate changed at a threshold value of deposit, to either a faster rate when modelling ripening, or a slower rate when modelling blocking. Secondly, the dependence of the pressure drop with specific deposit was shown to change from a linear relationship to a quadratic relationship at the same threshold value. Furthermore, the concentration of deposited colloids was seen to decrease exponentially with depth. Significantly, these results agree with the experimental observations of Veerapaneni.

A systematic exploration of this model was then performed, allowing for a mechanistic insight into process of clogging. The strength of the colloid-colloid and

the fluid-colloid potentials were shown to be factors that significantly affected the rate of clogging. The model showed two mechanisms by which clogging occurs: the first where deposited colloids completely block the pore spaces, and the second where a small amount of deposited colloids that strongly repel the fluid particles cause the pore spaces to be filled with fluid that cannot penetrate through the filter. The effects of the lattice structure, the porosity of the filter, and the colloid concentration were also examined.

A final investigation into the fractal dimension of the deposits was performed. It was shown that the colloid-colloid sticking distance was a parameter that allowed for control of the fractal dimension; a longer sticking distance resulted in a lower fractal dimension. Deposits with a lower fractal dimension were seen to increase the rate of clogging (for the same value of specific deposit). This agrees with the predictions of Mays, and is a significant result of the investigation.

3.6 References

- [1] W. Pan, A. Tartakovsky, *Dissipative particle dynamics model for colloid transport in porous media*, *Advances in Water Resources*, **58**, 41, 2013
- [2] E. Boek, P. Coveney, H. Lekkerkerker, P. van der Schoot, *Simulating the rheology of dense colloidal suspensions using dissipative particle dynamics*, *Phs. Rev.* **55**, 3124 (1997)
- [3] W. Dzwinel, D. Yuen. K. Boryczko, *A Two-Level, Discrete-Particle Approach for Simulating Ordered Colloidal Structures*, *J. of Coll. and Interace Sci.* **225**, **179** (2000)
- [4] A. Jabbarzadeh, *Molecular dynamics simulations of flow over a nanocylinder in a rectangular nano-channel*, *TechConnect Briefs*, **2**, 733 (2010)
- [5] P. Daivis, K. Travis, B. Todd, *A technique for the calculation of mass, energy, and momentum densities at planes in molecular dynamics simulations*, *J. Chem. Phys.*, **104**, 9651 (1996)
- [6] W. Hoover, *Smooth Particle Applied Mechanics*, World Scientific, Singapore, 2006
- [7] A. Ziarani, A. Mohamad, *A molecular dynamics study of perturbed Poiseuille flow in a nanochannel*, *Microfluidics and Nanofluidics*, **2**, 12 (2006)
- [8] T. Camesano, K. Unice, B. Logan, *Blocking and ripening of colloids in porous media and their implications for bacterial transport*, *Colloids and Surfaces A*, **160**, 291 (1999)
- [9] S. Veerapaneni, *Formation and morphology of colloidal deposits in porous media*, Ph.D. Thesis, Rice University, Houston, TX, 1996

- [10] D. Mays, J. Hunt, *Hydrodynamic aspects of particle clogging in porous media*, Environ. Sci. Technol., **39**, 577 (2005)
- [11] V. Jegetheesan, S.Vigneswaran, *Deep bed filtration: mathematical and observations*, Critical Reviews in Environ. Sci. and Tech., **35**, 515 (2007)
- [12] J. Bridge, *Technical Review of Colloid Transport, Deposition and Remobilisation Processes in Sand Bed Filter Relevant to SIXEP Treatment Processes*, 2012
- [13] S. Quandt, A P Young, *The shape of two-dimensional percolation and Ising clusters*, J.Phys, **20**, 851-856 (1987)

4 Continuum mechanics

Continuum mechanics describes the motion of matter using pre-defined relationships that define the fluxes of energy, mass and momentum contained within it. The densities of these conserved quantities can only change via redistribution, and this process is observable on a macroscopic scale.¹ It is assumed that the matter completely fills the space it occupies, and continuum mechanics therefore ignores the fact that matter is made of atoms. On length scales much larger than inter-atomic distances, this assumption allows for accurate modelling. Fundamental laws, such as the conservation of mass, momentum, and energy, may be applied to derive differential equations that describe the behaviour of a fluid in a continuum, and constitutive relationships describe information particular to the fluid being studied.

4.1 Conservation equations

The equations that define the fluxes can be obtained by considering an infinitesimal volume, V , enclosed by an arbitrary surface, S .

4.1.1 Conservation of mass

The total mass, M , contained within the volume element is given by:

$$M = \int_V \rho(\mathbf{r}, t) \cdot d\mathbf{r} \quad (4.1)$$

where $\rho(\mathbf{r}, t)$, is the mass density at time t and position \mathbf{r} . As mass is conserved, it follows that the mass will only change with flow through the enclosed surface:

$$\frac{dM}{dt} = - \int_S \rho(\mathbf{r}, t) \mathbf{u}(\mathbf{r}, t) \cdot d\mathbf{S} \quad (4.2)$$

where the mass flux is the product of the mass density and the streaming velocity, $\rho(\mathbf{r}, t)\mathbf{u}(\mathbf{r}, t)$. Using the divergence theorem:

$$\frac{dM}{dt} = - \int_V \nabla \cdot [\rho(\mathbf{r}, t)\mathbf{u}(\mathbf{r}, t)] d\mathbf{r} \quad (4.3)$$

where ∇ is the spatial gradient operator with components, $\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z}$. The rate of change of mass can also be written in terms of the change in mass density:

$$\frac{dM}{dt} = \int_V \frac{\partial \rho(\mathbf{r}, t)}{\partial t} d\mathbf{r} \quad (4.4)$$

Equating 2.52 and 2.53 gives the mass continuity equation in an Eulerian frame:

$$\frac{\partial \rho(\mathbf{r}, t)}{\partial t} = -\nabla \cdot [\rho(\mathbf{r}, t)\mathbf{u}(\mathbf{r}, t)] \quad (4.5)$$

An alternative form can be obtained by using a co-moving frame of reference:

$$\frac{d\rho(\mathbf{r}, t)}{dt} = -\rho(\mathbf{r}, t)\nabla \cdot \mathbf{u}(\mathbf{r}, t) \quad (4.6)$$

4.1.2 Conservation of momentum

A similar method can be used to obtain the momentum continuity equation. If $\mathbf{G}(t)$ is the total momentum contained within the volume element, then the rate of change of momentum is:

$$\frac{d\mathbf{G}(t)}{dt} = \int_V \frac{\partial[\rho(\mathbf{r}, t)\mathbf{u}(\mathbf{r}, t)]}{\partial t} d\mathbf{r} \quad (4.7)$$

The total momentum can change in two ways. Firstly, momentum can flow through the surface resulting in convection. This convective contribution is given by:

$$\frac{d\mathbf{G}_{conv}}{dt} = - \int_S \rho(\mathbf{r}, t)\mathbf{u}(\mathbf{r}, t)\mathbf{u}(\mathbf{r}, t) \cdot d\mathbf{S} \quad (4.8)$$

The second contribution is from the pressure exerted on the volume by the surrounding fluid, called the stress contribution. The force, $d\mathbf{F}$, exerted is linearly proportional to the surface area, $d\mathbf{S}$:

$$d\mathbf{F} = -d\mathbf{S} \cdot \mathbf{P} \quad (4.9)$$

where \mathbf{P} is the pressure tensor. The stress contribution is therefore given by:

$$\frac{d\mathbf{G}_{stress}}{dt} = - \int_S d\mathbf{S} \cdot \mathbf{P} \quad (4.10)$$

Adding equations 4.10 and 4.7 results in the following:

$$\int_V \frac{\partial[\rho(\mathbf{r}, t)\mathbf{u}(\mathbf{r}, t)]}{\partial t} d\mathbf{r} = - \int_S d\mathbf{S} \cdot [\rho(\mathbf{r}, t)\mathbf{u}(\mathbf{r}, t)\mathbf{u}(\mathbf{r}, t) + \mathbf{P}] \quad (4.11)$$

Again, using divergence theorem to convert surface integrals to volume integrals gives the momentum continuity equation in an Eulerian frame:

$$\frac{\partial[\rho(\mathbf{r}, t)\mathbf{u}(\mathbf{r}, t)]}{\partial t} = -\nabla \cdot [\rho(\mathbf{r}, t)\mathbf{u}(\mathbf{r}, t)\mathbf{u}(\mathbf{r}, t) + \mathbf{P}] \quad (4.12)$$

It can be shown that an alternative form of the momentum continuity equations, in the co-moving frame, is:

$$\rho(\mathbf{r}, t) \frac{d\mathbf{u}(\mathbf{r}, t)}{dt} = -\nabla \cdot \mathbf{P} \quad (4.13)$$

4.1.3 Conservation of energy

If the total energy per unit mass is $e(\mathbf{r}, t)$ the total energy density is $\rho(\mathbf{r}, t)e(\mathbf{r}, t)$. The total energy consists of two contributions: a kinetic (convective) component and an internal energy density:

$$\rho(\mathbf{r}, t)e(\mathbf{r}, t) = \frac{1}{2} \rho(\mathbf{r}, t)\mathbf{u}(\mathbf{r}, t)^2 + \rho(\mathbf{r}, t)U(\mathbf{r}, t) \quad (4.14)$$

The total energy inside the volume, E , will therefore change according to:

$$\frac{dE}{dt} = \int_V \frac{\partial[\rho(\mathbf{r}, t)e(\mathbf{r}, t)]}{\partial t} d\mathbf{r} \quad (4.15)$$

The total energy can change by convection through the surface, diffusion across the surface, and by work done by the surface stresses. These three mechanisms result in:

$$\begin{aligned} \frac{dE}{dt} = \int_S d\mathbf{S} \cdot [\rho(\mathbf{r}, t)e(\mathbf{r}, t)\mathbf{u}(\mathbf{r}, t)] + \int_S d\mathbf{S} \cdot \mathbf{J}_Q \\ + \int_S (d\mathbf{S} \cdot \mathbf{P}(\mathbf{r}, t)) \cdot \mathbf{u}(\mathbf{r}, t) \end{aligned} \quad (4.16)$$

where \mathbf{J}_Q is the heat flux vector. Again, using divergence theorem equation 4.16 give the energy continuity equation:

$$\frac{dE}{dt} = -\nabla \cdot \left[\begin{array}{c} \rho(\mathbf{r}, t)e(\mathbf{r}, t)\mathbf{u}(\mathbf{r}, t) + \mathbf{J}_Q(\mathbf{r}, t) \\ \mathbf{P}(\mathbf{r}, t) \cdot \mathbf{u}(\mathbf{r}, t) \end{array} \right] \quad (4.17)$$

This can be expressed in a co-moving frame as²⁰:

$$\frac{dE}{dt} = -\nabla \cdot \mathbf{J}_Q(\mathbf{r}, t) - \mathbf{P}(\mathbf{r}, t)^T : \nabla \mathbf{u}(\mathbf{r}, t) \quad (4.18)$$

4.1.4 Constitutive relations

The conservation equations described in section 2.3.1 express the fluxes of mass, momentum and energy. In order to obtain a closed system of equations constitutive relations are required that relate the forces and fluxes. A combination of the constitutive relations and the continuity equations leads to the Navier-Stokes equations; once the boundary conditions are applied this leads to a complete description of a fluid close to equilibrium.

Newton's law of viscosity

Newton's law of viscosity is an approximation based on experimental evidence that holds true in certain situations. It states that the shear stress between adjacent fluid layers is proportional to the negative value of the velocity gradient between the two layers:

$$\tau_{yx} = -\mu \frac{du_x}{dy} \quad (4.19)$$

where μ is the shear viscosity and τ is the stress. This law holds true for Newtonian fluids, but has its limitations. Materials that display shear thinning or shear thickening do not obey this law, and are non-Newtonian fluids.

Fourier's law

Fourier's law is an empirical relationship relating the heat flux vector, \mathbf{J}_Q , to the temperature gradient:

$$\mathbf{J}_Q = -\lambda \nabla T \quad (4.20)$$

where λ is the thermal conductivity and T is the thermodynamic temperature.

Fick's law of diffusion

Fick's law of diffusion relates the mass flux vector to the mass gradient:

$$\mathbf{j}_A = -p_A D_{AB} \nabla \omega_A \quad (4.21)$$

where \mathbf{j}_A is the mass flux vector, p_A is the concentration of species A, ω_A is the mass fraction of species A, and D_{AB} is the diffusion coefficient of species A in a mixture of A and B.

4.2 Smooth particle applied mechanics

4.2.1 Overview

Conventional continuum mechanics is a field theory, meaning the field variables (such as density, velocity etc.) vary continuously in space and time. The equations to be solved are the conservation equations of mass, momentum, and energy discussed earlier. Combined with the constitutive relations and boundary conditions, they form a closed system which can be solved. However, due to the spatial variation of the field variables, the equations naturally have an infinite number of degrees of freedom; in order to be able to simulate systems, the number of degrees of freedom needs to be reduced to a computable number. Techniques such as the finite element method and finite difference method exist that split the continuum into a grid, with each grid segment having a finite number of degrees of freedom. These techniques fall down when the structure under simulation undergoes extreme changes in shape, suffers from mechanical failure, or involves a combination of solids and liquids. Instead, more flexible techniques are needed.

SPAM overcomes these problems by using a set of smooth particles whose coordinates define the grid for the interpolation of field variables. This makes it a

flexible technique that can inherently handle mixtures, mechanical failure, and chaotic flows such as turbulence without any meshing problems. The technique was created by Gingold, Lucy and Monaghan in 1977 as a means of overcoming the problems associated with using a fixed grid, and has a wide range of applications since its conception.²

4.2.2 SPAM approximation of the continuity equations

The partial differential equations to be solved in Lagrangian continuum mechanics are:

$$\dot{\rho} = -\rho \nabla \cdot \mathbf{v} \quad (4.22)$$

$$\dot{\mathbf{v}} = -\frac{1}{\rho} \rho \nabla \cdot \mathbf{P} \quad (4.23)$$

$$\dot{e} = -\frac{1}{\rho} \mathbf{P} : \nabla - \frac{1}{\rho} \nabla \cdot \mathbf{Q} \quad (4.24)$$

Smooth particle applied mechanics interpolate the field variables $(\rho, \mathbf{v}, e, \mathbf{P}, \mathbf{Q})$ using a weight function. For any of these variables, the local average value is calculated as a weighted average of all particles whose distance, $r = |r_{ij}|$ is less than the smoothing length, h :

$$f(r_i) = \sum_j m_j \frac{f_j}{\rho_j} w(r) \quad (4.25)$$

where f is the value of the field variable, m is the particle mass, and w is the weight. The smooth particle approximation for the local density is obtained by setting $f_j = \rho_i \equiv \rho(r_i)$:

$$\rho_i = \sum_j m_j w(r) \quad (4.26)$$

The density at a given point or particle location is calculated by summing the weighted contribution of all particles within the smoothing length: there is therefore no need to solve the mass continuity equation.

Approximations for the conservation laws of momentum and energy can also be derived, resulting in the following two expressions:

$$m\dot{v} = - \sum_j m_i m_j \left(\frac{\mathbf{P}_i}{\rho_i^2} + \frac{\mathbf{P}_j}{\rho_j^2} \right) \cdot \nabla_j w(r) \quad (4.27)$$

$$m\dot{e} = - \frac{1}{2} \sum_j m_i m_j \left(\frac{\mathbf{P}_i}{\rho_i^2} + \frac{\mathbf{P}_j}{\rho_j^2} \right) : \mathbf{v}_{ij} \nabla_j w(r) \quad (4.28)$$

$$- \sum_j m_i m_j \left(\frac{\mathbf{Q}_i}{\rho_i^2} + \frac{\mathbf{Q}_j}{\rho_j^2} \right) \cdot \nabla_j w(r)$$

These ordinary differential equations can be treated with an algorithm that is very similar to that used in molecular dynamics, though the right-hand side takes on a different form. An appropriate numerical integrator such as RK4 (see section 2.2.3) can be used to evolve these equations; given a set of initial conditions and boundary conditions, the particle trajectories and the flux of energy between these particles can be calculated, along with the field variables at particle locations. The constitutive relations must be defined by the simulator, and are discussed in detail in chapter 5.

4.2.3 Weight functions

The weight function spreads each particle's influence through space, making the choice of the smoothing length, h , of great importance. If h is too short, the discrete nature of particles is given too much influence meaning there is poor variable interpolation, and if h is too large then the computation time required vastly increases. It is up to the simulator to find a suitable balance, and it is suggested that the smoothing length should be high enough to incorporate contributions from approximately 20 particles.³

The function must satisfy several conditions: $w'(0) = 0$, $w(h) = 0$, $w''(h) = 0$ and $w'''(h) = 0$ (where ' denotes the first derivative). A suitable choice was introduced by Lucy to be used in two dimensions (shown in figure 4.1):

$$w(r < h) = \left(\frac{5}{\pi h^2}\right) \left(1 + 3\frac{r}{h}\right) \left(1 - \frac{r}{h}\right)^3 \quad (4.29)$$

Figure 2.10 shows the Lucy weight function. The weight function, and both the first and second derivative, cut off at the smoothing length, and both the derivatives are continuous below this. This means that the spatial derivatives of the field variables can have no discontinuities. This function has been chosen for this work as it has been used extensively in previous work, making for easier comparison with existing results.




Figure 4.1: Lucy's 2-dimensional weight function (red), where $h = 3$, along with the first (blue) and second (green) derivatives.

4.2.4 Artificial viscosity

The numerical integration of the Navier-Stokes equations used in SPAM can result in instabilities; shock waves can arise from local pressure waves, which, when too steep, cannot be resolved by the mesh. As a result, it is necessary to introduce an artificial viscosity as a device to spread the shockwaves out over several particles. One of many ways of achieving this is to add an additional viscous component to the force expression:

$$m\dot{\mathbf{v}} = - \sum_j m_i m_j \left(\frac{P_i}{\rho_i^2} + \frac{P_j}{\rho_j^2} + \Pi_{ij} \right) \nabla_j w(r) \quad (4.30)$$

where Π_{ij} is the viscous component, defined by:

$$\Pi_{ij} = -\alpha h \frac{c_i + c_j}{\rho_i + \rho_j} \frac{\mathbf{v}_{ij} \cdot \mathbf{r}_{ij}}{r_{ij}^2 + \epsilon h^2} \quad (4.31)$$

where c is the speed of sound of the fluid particle, α is a dimensionless factor used to control the strength of the dissipation and $\epsilon \sim 0.01$ and avoids singularities in the case that particles are too close to each other. In order for energy to conserve, the artificial viscosity must also be added to the energy expression:

$$m\dot{e} = -\frac{1}{2} \sum_j m_i m_j \left(\frac{P_i}{\rho_i^2} + \frac{P_j}{\rho_j^2} + \Pi_{ij} \right) : \mathbf{v}_{ij} \nabla_j w(r) \quad (4.32)$$

$$- \sum_j m_i m_j \left(\frac{Q_i}{\rho_i^2} + \frac{Q_j}{\rho_j^2} \right) \cdot \nabla_j w(r)$$

A more pragmatic approach is to maintain a minimum particle separation using an interparticle potential (as used in molecular dynamics). A hard sphere potential (equation 2.31) would stop the formation of shockwaves by means of elastic collisions, and a short range repulsive potential (equation 2.34) would achieve this by discouraging overlaps. Hoover has previously used this technique, showing that the properties of the system do not depend on the artificial viscosity providing the interaction distance is significantly smaller than the smoothing length.

4.2.5 Boundary conditions

The boundary conditions used in SPAM are the same as those used in MD (see section 2.2.2). There is, however, one important addition to note. If non-periodic boundaries are used it would be the case that the calculated particle density (and other field variables) at the simulation edge would be half the value of the bulk density; the weight function would only receive a contribution from half the number of particles compared to the bulk⁴ (figure 4.2). This can cause irregularities such as a layer of particles that appear to be stuck to the edge. This problem can be overcome by using mirrored particles on the other side of the boundary to replace the “missing” particles.




Figure 4.2: Comparison of elastic (left) and mirror (right) boundary conditions. The density calculated at the particle at the boundary for the elastic case is half that of the mirrored case. When calculating the density at the dotted line, an additional contribution is included from the mirrored particles.

4.2.6 Initial conditions

As with molecular dynamics, both the initial coordinates and momenta of all particles in the simulation must be defined. This is achieved using the same methods

as described in section 2.1. Additionally, SPAM simulations require the starting energies of all particles to be defined, to allow for the change in energy according to equation 4.32. This can be calculated from the equation of state; energy is a function of density and temperature.

4.2.7 Particle size

It is important to note that the size of the particles used in a SPAM simulation is in fact arbitrary; the same sample can be modelled using different simulations with differing particle sizes. Consider replacing each particle in a simulation with two particles each with half of the mass. Calculating the density at a given point through the equation:

$$\rho_i = \sum_j m_j w(r - r_j) \quad (4.33)$$

would result in the same value; there is a contribution from twice as many particles, each with half of the mass. The same applies for calculating the trajectories. This means that SPAM simulations can be altered to give more detail in areas of interest; one large particle can be split into multiple small particles to increase resolution.

4.3 References

- [1] W. Hoover, *Phys. Rev.*, **31**, 1695 (1985)
- [2] D. Evans, G. Morriss, *Statistical Mechanics of Nonequilibrium Liquids*, Elsevier, Amsterdam, 1990
- [3] W. Hoover, *Smooth Particle Applied Mechanics: The state of the art*, World Scientific, Singapore, 2006
- [4] O. Kum, W. Hoover, *Smooth particle boundary conditions*, C. Hoover, *Physical Review E*, **68**, (2003)

5 Constitutive relations

5.1 Overview

In order to model a continuum the simulator must first define a series of constitutive relations that predict how the material will behave in both equilibrium and non-equilibrium conditions. These relations replace the microscopic particle interactions that naturally define the nature of the system in particle dynamics as a method of transporting both momentum and energy. There are two methods that can be used to gather the required information to define these relations: performing real experiments, and using molecular simulation to gather pseudo-experimental data. This work will use pseudo-experimental data from molecular dynamics simulations to derive the constitutive relations. This gives the distinct benefit of being able to use the coarse grained molecular dynamics filtration model as a means of validating the SPAM model, as the fluid being modelled in the continuum is parameterised from the *same* interatomic pair potential.

This chapter will outline the techniques used to define an equilibrium equation of state, and to gather the data to create the constitutive relation defining the viscosity of the fluid away from equilibrium. The interatomic pair potential being investigated is the same as that used throughout the previous molecular dynamics work (chapter 3):

$$\Phi(r) = \varepsilon \left[1 - \left(\frac{r}{\sigma} \right)^2 \right]^4 \quad (5.1)$$

where r is the inter-particle distance, ε is the interaction strength, and σ is the effective diameter of the particle. Two sets of constituent relations will be defined: the first where $\varepsilon = 10.0$ and $\sigma = 1.0$, the second where $\varepsilon = 100.0$ and $\sigma = 1.0$. Although this requires additional simulation time to gather the data required, it gives the advantage of being able to compare not only the agreement between the molecular dynamic filtration model and the continuum model with different fluid parameters, it also allows the continuum model to be further explored by examining how changing the fluid interaction strength effects the flow.

5.2 Equilibrium Equation of State

An equation of state relates pressure, P , and energy, E , to density, ρ , and temperature, T , at thermodynamic equilibrium. This allows the pressure and energy to be calculated for any given scenario for the system and is necessary when modelling a continuum. It follows that, in order to create an equilibrium equation of state, the change in both pressure and energy of the fluid as functions of density and temperature must be defined. This can be done through relatively simple molecular dynamics or Monte Carlo simulations. It is worth noting that an equation of state is usually only valid in one phase. If a system that requires multiple phases is being modelled then the equation of state either has to incorporate both phases, or a separate equation of state is needed for each phase. This work focuses only on a dense fluid, therefore the equation of state needs to hold for the fluid phase.

5.2.1 $\varepsilon = 100$

Hoover¹ investigated the potential of equation 5.1 using $\varepsilon = 100$ and unit cutoff ($\sigma = 1$). He chose a reference state of unit density and unit temperature, resulting in values for $PV/N\varepsilon$ and $E/N\varepsilon$ of 5.040 and 1.443 respectively. For small deviations from this reference state, the pressure, temperature, and energy can be expanded as a double Taylor series in the derivatives $\delta\varepsilon$, δT , and $\delta\rho$. A series of simulations run at constant energy (NVE) and constant temperature (NVT) were then conducted in order to determine the expansion coefficients, resulting in the following equations:

$$\frac{PV}{N\varepsilon} = 5 + 8\delta\rho + 2.5\delta\varepsilon + 9\delta\rho^2 + 2\delta\rho\delta\varepsilon \quad (5.2)$$

$$\frac{kT}{\varepsilon} = 1 - \delta\rho + 0.7\delta\varepsilon - 0.86\delta\rho^2 - 0.5\delta\rho\delta\varepsilon \quad (5.3)$$

$$\frac{E}{N\varepsilon} = 1.443 + 1.5\delta\rho + 1.5\delta\varepsilon + 2.4\delta\rho^2 + 1.26\delta\rho\delta\varepsilon \quad (5.4)$$

$$\delta\rho = \left(\frac{N\sigma^2}{V} \right) - 1.0 \quad (5.5)$$

$$\delta\varepsilon = \left(\frac{E}{N\varepsilon} - 1.443 \right) \quad (5.6)$$

$$\delta\tau = \left(\frac{kT}{\varepsilon} - 1.0 \right) \quad (5.7)$$

In SPAM, the first two expansions are used to provide the particle local equilibrium pressures and temperatures, which contribute to the particle accelerations and heat flux respectively. The third expansion is needed when there is heat flow.

The timestep and integration scheme used in the molecular dynamics were not reported in the literature, therefore, in order to investigate the range over which this equation of state holds, it was important to first examine the sensitivity to the timestep. To achieve this, a series of NVE simulations were run at several state points, including the reference state, spanning to a significant deviation above and below, with varying timesteps. At constant energy, equation 5.2 becomes:

$$\frac{PV}{N\varepsilon} = 5 + 8\delta\rho + 9\delta\rho^2 \quad (5.8)$$

The RK4 integration scheme was used for all simulations in this chapter. Figure 5.1 shows the results of a series of simulations, for timesteps ranging from 0.05 to 0.0001. The calculated pressure from the MD simulation is plotted against the density. The black line shows the results predicted by the equation of state. It is clear that the results with a larger timestep (0.05 and 0.03) deviate significantly from the expected values. As the timestep decreases, the results converge to those obtained from the equation of state, with there being no significant increase in accuracy when reducing the timestep below 0.01. This highlights the importance of the timestep when performing molecular dynamics simulations; using a large timestep can result in vast integration errors that change behaviour of the fluid.




Figure 5.1: Pressure against density, comparing results from MD simulations of 1024 particles with different timesteps. The black line shows the results generated using the equation of state (equation 5.8). The length of all simulations was 500 reduced units.

To test how far away from the reference point the full set of equations (4.2 – 4.4) are valid, a series of molecular dynamics simulations were conducted. For each simulation in the series, the variable of interest (density, temperature, or energy) was increased slightly, and the average of this property was measured throughout the simulation. This was repeated across the required range to develop the dependence of pressure and energy on density and temperature. Logically, from figure 5.1, it follows that any discrepancies between the equation of state predictions and the molecular dynamics simulations cannot be attributed to the timestep, providing the timestep is below 0.01. All simulations run throughout this section were run with a timestep of 0.001, using the RK4 integration scheme, simulating 1024 particles for 1,000,000 timesteps.

Pressure

In order to validate the first expansion (equation 5.2), two series of simulations were run: one at constant energy whilst varying the density of the system, and one at constant density whilst varying the total energy of the system.

Firstly, to test the dependency of pressure on density, the system was kept at a constant energy, $\frac{E}{N\varepsilon} = 1.443$, such that $\delta\varepsilon = 0$ and the density was varied between 0.5 and 1.5. The pressure energy dependence is therefore defined by equation 5.8. This is included for the sake of completeness, even though the same simulation was performed to produce figure 5.1.




Figure 5.2: Pressure against density, comparing results from MD simulations against the predictions from the equation of state.

Figure 5.2 shows the results of these simulations. At the reference state, the agreement between the equation of state and the molecular dynamics data is, as expected,

exact. The results obtained away from the reference state also show excellent agreement; even up to a density of 1.4 there was only slight difference between the prediction and the simulated data. As the density decreased below 0.7, the results started to diverge, suggesting the equation of state was not optimised below this value.

The second set of simulations were run to determine the dependency of pressure on energy. The energy, rather than being fixed at 1.443, was varied around the reference state, with the density being kept constant at unity, so that $\delta\rho = 0$. At unit density, equation 5.2 becomes:

$$\frac{PV}{N\varepsilon} = 5 + 2.5\delta\varepsilon \quad (5.9)$$

The results (figure 5.3) show the linear trend expected from examination of equation 5.9, and the agreement between the simulations and the equation of state was excellent to a significant deviation from the reference state. As the energy decreased below 1.1 the results started to diverge.

Figure 5.2 and 5.3 show that the parameters used in the equation of state for predicting pressure are highly accurate to *at least* a density of 1.4 and an energy of 1.8. The accuracy begins to diverge below values of 0.7 and 1.1 for density and energy respectively.




Figure 5.3: Average pressure against energy, comparing results from MD simulations against the predictions from the equation of state.

Temperature

To investigate the range over which the temperature expansion (equation 5.3) is accurate a similar series of simulations were conducted. Rather than measuring the pressure throughout the simulation, the average kinetic temperature was calculated. The first series of simulations were run at a constant energy of 1.443, and density was varied around the reference state. At this energy, equation 5.3 becomes:

$$\frac{kT}{\varepsilon} = 1 - \delta\rho - 0.86\delta\rho^2 \quad (5.10)$$

Figure 5.4 shows the results of this series of simulations. At the reference point the agreement is exact, and remains accurate to a large deviation at higher and lower

densities. However, at a density of 1.3, significant deviation was observed. Figure 5.5 shows the particle locations after 1,000,000 timesteps at densities of 1.2 and 1.4. It is clear that above a density of 1.2 the systems started to change phase, explaining the deviations observed in figure 5.4. The equation of state has only been parameterised using the fluid phase.




Figure 5.4: Kinetic temperature against density, comparing results from MD simulations against the predictions from the equation of state.




Figure 5.5: Snapshots of the computed particle locations for a constant density of 1.2 (left) and 1.4 (right).

The second series of simulations were run in order to test the dependence of temperature on energy. The total energy of the system was fixed for each simulation, and was varied around the reference state. The density of the system was fixed at unity, so that $\delta\rho = 0$. Equation 5.3 becomes:

$$\frac{kT}{\varepsilon} = 1 + 0.7\delta\varepsilon \quad (5.11)$$

Similarly to the dependence of pressure on energy, the results show excellent agreement not only at the reference point, but also to a substantial deviation. The linear trend predicted by the equation of state matches the results closely, showing that the equation of state is accurate to a sufficient variation around the reference point.




Figure 5.6: Average kinetic temperature against energy, comparing results from MD simulations against the predictions from the equation of state. Density was fixed at unity.

Energy

Finally, equation 5.4 was validated by testing the dependence of the total energy of the system on both density and temperature. Again, the variable of interest was varied around the reference point, and the average total energy of the system was calculated. Here, however, the simulations were conducted in the NVT ensemble; a Gaussian isokinetic thermostat was used to control the kinetic temperature.

The first set of simulations were run at a constant kinetic temperature, $\frac{kT}{\epsilon} = 1$, such that $\delta\tau = 0$. The density of the system was then varied around unity, and the total energy was measured. Equation 5.4 becomes:

$$\frac{E}{N\varepsilon} = 1.443 + 1.5\delta\rho + 2.4\delta\rho^2 \quad (5.12)$$

The results are shown in figure 5.7. Again, the results show excellent agreement in and around the reference state. Similarly to figure 5.5, the results begin to diverge as the density was increased past 1.4. This is due to the system solidifying at this point (see figure 5.6). Additionally, the results begin to diverge as density goes below 0.6, suggesting the equation of state does not hold true below this value.




Figure 5.7: Total energy against density, comparing results from MD simulations against the predictions from the equation of state. The kinetic temperature was fixed at unity.

The final set of simulations conducted investigate energy as a function of temperature. Pressure was kept constant at $\rho = 1.0$, so that $\delta\rho = 0$, and the temperature was varied around the reference state. Equation 5.4 becomes:

$$\frac{E}{N\varepsilon} = 1.443 + 1.5\delta\rho + 1.5\delta\tau + 2.4\delta\rho^2 + 1.26\rho\delta\tau \quad (5.13)$$

Again, the agreement between the equation of state and the simulation data is almost exact, showing that the equation of state can be used to accurately predict the total energy to a high temperature.




Figure 5.8: Average total energy against temperature, comparing results from MD simulations against the predictions from the equation of state. The density is fixed at unity.

Summary

The range at which equations 5.2 - 5.4 are accurate was investigated by running a series of NVE and NVT molecular dynamic simulations. The only significant deviations between the predictions of the equation of state and the calculated data occurred at high density where a phase transition was observed. The equation of state

will be used in the continuum scale model, and will give an accurate prediction of P , T , and E , providing the density of the fluid stays within the validated range.

5.2.2 $\varepsilon = 10$

Additionally, a macroscopic dense fluid equation of state with an interaction strength, ε , of 10 was investigated. This interatomic pair potential will have weaker interactions between individual particles when compared to the previous parameters, resulting in a softer fluid that will most likely not solidify until a higher density.

To develop an equation of state, as opposed to simply validating it, a similar technique was used. NVE simulations were conducted to determine the dependence of pressure and temperature on density and energy individually. Once the dependence across the required range was determined, a plot of the difference in pressure/temperature against energy/density was used to determine the coefficients used in the expansion. At the reference state, $\rho = 1$ and $\frac{kT}{\varepsilon} = 1$, the values of $\frac{PV}{N\varepsilon}$ and $\frac{E}{N\varepsilon}$ are 2.47 and 1.44 respectively.

Pressure

To determine the relationship between pressure and density, a series of NVE simulations were conducted at various densities, and the pressure was calculated. Figure 5.9 shows the results, where the differences of the pressure and density from the reference point are plotted against each other. As expected, using the previous Taylor series expansions as a basis, it forms a 2nd order polynomial, where the first and second terms are the parameters used in the expansion:

$$\frac{PV}{N\varepsilon} = 2.47 + 4.2\delta\rho + 2.3\delta\rho^2 \quad (5.14)$$




Figure 5.9: Difference in pressure against difference in density. The reference point, (0,0), is at a density and temperature of 1.0.

Similarly, to determine the relationship between pressure and energy, simulations were conducted where the energy was varied around the reference point, and the pressure was calculated. Figure 5.10 shows the results; a linear dependence as expected. This results in the following equation:

$$\frac{PV}{N\varepsilon} = 2.47 + 1.2\delta\varepsilon \quad (5.15)$$




Figure 5.10: Difference in pressure against difference in energy. The reference point, (0,0), is at a density and temperature of 1.0.

The coefficients in figures 5.9 and 5.10 combine to give the following expansion:

$$\frac{PV}{N\varepsilon} = 2.47 + 4.2\delta\rho + 1.2\delta\varepsilon + 2.3\delta\rho^2 + \chi\delta\rho\delta\varepsilon \quad (5.16)$$

where χ remains to be determined. The final coefficient was calculated by simply running several simulations at different fixed densities and energies and calculating the pressure; pressure, $\delta\rho$ and $\delta\varepsilon$ are both known, meaning χ is the only remaining unknown. 10 simulations were run at random densities and energies, and the total pressure measured, resulting in an average χ of 1.8, giving the final expansion for pressure:

$$\frac{PV}{N\varepsilon} = 2.47 + 4.2\delta\rho + 1.2\delta\varepsilon + 2.3\delta\rho^2 + 1.8\delta\rho\delta\varepsilon \quad (5.17)$$

Temperature

The dependence of temperature on density and energy was determined by conducting a series of simulations, varying each parameter in isolation.

Figure 5.11 shows the results of the density variation. As expected, temperature decreases with density, following a second order polynomial:

$$\frac{kT}{\varepsilon} = 1 - 0.6\delta\rho - 0.19\delta\rho^2 \quad (5.18)$$




Figure 5.11: Difference in temperature against difference in density. The reference point, (0,0), is at a density and temperature of 1.0.

Figure 5.12 shows the dependence of temperature on energy, revealing a linear relationship:

$$\frac{kT}{\varepsilon} = 1 - 0.7\delta\varepsilon \quad (5.19)$$




Figure 5.12: Difference in temperature against difference in energy. The reference point, (0,0), is at a density and temperature of 1.0.

Equations 5.18 and 5.19 combine to give:

$$\frac{kT}{\varepsilon} = 1 - 0.6\delta\rho + 0.7\delta\varepsilon - 0.19\delta\rho^2 - \chi\delta\rho\delta\varepsilon \quad (5.20)$$

where χ is a constant that is yet to be determined. Again, the final coefficient was determined by running a series of simulations at different densities and temperature,

where , $\delta\rho$ and $\delta\varepsilon$ are known, and the temperature was calculated. This resulted in an average χ of 0.8, giving the final expansion for temperature:

$$\frac{kT}{\varepsilon} = 1 - 0.6\delta\rho + 0.7\delta\varepsilon - 0.19\delta\rho^2 - 0.8\delta\rho\delta\varepsilon \quad (5.21)$$

Summary

An equilibrium equation of state was defined for the soft sphere potential, where $\varepsilon = 10$. This was achieved using an analogous method to that of Hoover.¹ Pressure was shown to depend quadratically on density and linearly energy. Temperature was also shown to depend quadratically on density and linearly on temperature. Equations 5.17 and 5.20 were proposed as the equation of state.

5.3 Viscosity

Shear viscosity is a measure of a system's resistance to flow. A shear force is created when two parallel plates surrounding a flow move in opposite directions, dragging some of the fluid along with them. The rate at which the plates move is the shear rate, and the force applied to the liquid by the plates is proportional to their velocity and area:

$$\tau = \frac{F}{A} \quad (5.22)$$

where τ is the shear stress, F is the force, and A is the area. This shear stress creates an anisotropic flow fastest nearest the plates, with layers of fluid decreasing in momentum towards the center:

$$\tau_{xy} = \mu\gamma \quad (5.23)$$

where μ is the viscosity, and γ is the shear strain rate. There are several methods that have been developed to obtain shear viscosity by observing a fluid's response to an induced shear flow, either by using boundaries or a fictitious force to drive the flow.^{2,3,4,5}

5.3.1 The sliding wall method

Naturally, due to the similarities with experimental methods of calculating viscosity, using a sliding wall is a sensible starting point when looking to determine viscosity through simulation; in a physical system, heat energy would be removed through contact with the wall, rather than through a thermostat. Ashurt and Hoover introduced the sliding wall method², where the fluid is surrounded by two walls made of equivalent particles to the fluid, allowing for energy to transfer between them. The walls are forced to translate in opposite directions with a velocity proportional to the shear rate, creating a sliding wall that drags fluid along with it. This creates the steady state velocity profile expected. The particles that make the wall are attached to Hookean springs, and their momenta are rescaled using the *ad hoc* method.




Figure 5.13: An illustration of the sliding wall method. The red particles make up the wall. The wall moves with a fixed velocity in the x-direction proportional to the shear rate. The structure of the walls is maintained using Hookean springs.

However, the use of a physical wall in molecular dynamics simulations can cause issues associated with the interactions of the particles in the fluid and the wall. The wall can begin to impose its structure on fluid atoms in close proximity resulting in an inhomogeneous fluid, making interpretation of results difficult. Additionally, at large shear rates this method falls down; the viscous heat generated is produced faster than it can be removed by interactions with the wall, causing disintegration of the walls.³

5.3.2 Lees-Edwards method

Lees and Edwards⁴ introduced a technique that uses adapted periodic boundary conditions to model planar Couette flow, and is a method that overcomes the difficulties that arise from having a physical boundary. Rather than implementing a sliding wall, the boundary conditions are used to drive the flow, which ensures the system remains spatially homogenous.




Figure 5.14: Lees-Edwards periodic boundary conditions. The top and bottom layer of periodic images are offset proportionally to the effective strain rate, n . The particle leaving the simulation cell is shifted when being mirrored.

Figure 5.14 is an illustration of the Lees-Edward boundary conditions. The image cell origins in the planes above and below the central cell move relative to it:

$$\mathbf{u}(\mathbf{r}, t) = \mathbf{i}\dot{\gamma}y \quad (5.24)$$

where $\mathbf{u}(\mathbf{r}, t)$ is the velocity of the cell origin, \mathbf{i} is a unit vector in the x-direction, and $\dot{\gamma}$ is the shear rate. As the distance from the central cell increases the cells therefore move proportionally faster. When a particle crosses the periodic boundary in the y-direction, instead of being directly mirrored to its new location (as is the case with standard periodic boundary conditions), the periodic image is repositioned proportionally to the shear rate in the x-direction:

$$r_x^{new} = r_x^{old} \pm \left(n + \frac{L}{2} \right) \quad (5.25)$$

where L is the length of the cell, and n is proportional to the shear rate. This lateral movement is akin to the velocity gained by interaction with a sliding wall. The repetition of particles crossing the periodic boundary over time generates the linear velocity profile expected, without the inherent problems of modelling surface interactions.

5.3.3 The SLLOD method

The main drawback with the Lees-Edwards method is that it takes a period of time for the steady state linear velocity profile to be set up; this means that the method cannot be used to study time-dependent phenomena. Because of this, further work has been done to incorporate the strain field within the equations of motion themselves (as oppose to at the boundaries). The first attempt at this was the DOLLS tensor method in which the modified equations of motion include a mechanical flow induced by the shear stress⁵:

$$\dot{\mathbf{q}}_i = \frac{\mathbf{p}_i}{m} + \mathbf{q}_i \cdot \nabla \mathbf{u} \quad (5.26)$$

$$\dot{\mathbf{p}}_i = \mathbf{F}_i - \nabla \mathbf{u} \cdot \mathbf{p}_i \quad (5.27)$$

where m is the particle mass, \mathbf{p} and \mathbf{q} are the generalised coordinates and momentum \mathbf{F} is the total force acting on the particle, and $\nabla \mathbf{u}$ is the strain rate tensor. Although the DOLLS tensor equations simulate the correct velocity profile in the linear regime, outside of the linear shear regime they produce incorrect results.⁶ By transposing the Cartesian components coupled to the strain rate tensor the SLLOD equations of motion are obtained, the name being the inverse of DOLLS, which are used to model planar Couette flow in this work:⁷

$$\dot{\mathbf{q}}_i = \frac{\mathbf{p}_i}{m} + \mathbf{q}_i \cdot \nabla \mathbf{u} \quad (5.28)$$

$$\dot{\mathbf{p}}_i = \mathbf{F}_i - \mathbf{p}_i \cdot \nabla \mathbf{u} \quad (5.29)$$

The SLLOD equations of motion have the distinct advantage of both producing the correct shear in the linear regime and the non-linear regime, whilst avoiding the issues associated with interaction with a surface.

The SLLOD method was chosen for this work. It was hoped that Green Kubo correlation functions would be used to validate the results gained, though this fell beyond the scope of this work due to the additional time required to develop an additional code to perform these simulations.

5.3.4 Calculating viscosity through extrapolation

The viscosity at a given shear rate is calculated using the relation:

$$\mu(\gamma) = -\frac{\langle P_{xy} \rangle}{\gamma} \quad (5.30)$$

The aim of the technique is to calculate the viscosity when the shear rate, γ , is 0, however, it is clear that this cannot be achieved directly as it would involve a division by zero. Instead, a series of viscosities are calculated at different shear rates, and the viscosity in the absence of shear can then be estimated through extrapolation. The type of fit that correctly predicts the viscosity has been the subject of some debate.

Alder and Wainwright⁸ were the first to attempt to use computer simulation to determine transport coefficients. They came to the conclusion that the Navier-Stokes transport coefficients diverge in two-dimensions. This was not to suggest that a two-dimensional fluid is infinitely resistant to shear flow, rather that the constitutive relation (equation 5.28) is not a suitable definition. However, more recent work has shown this assumption to be wrong.

Gravina *et al*⁹ compared NEMD results with Green-Kubo simulation data and suggested that a Lorentzian function accurately predicts this relationship in two dimensions. The fitting function takes the form:

$$\mu(\gamma) = a + \frac{b}{(1 + c\gamma^2)} \quad (5.31)$$

Figure 5.15 shows the example results from a series of SLLOD simulation. The shear rate was gradually increased from 0.0001 to 0.5. At lower shear rates the error in the results is larger, meaning it is important to have sufficient results at low shear to get an accurate estimate of viscosity. A Lorentzian fit was used to predict the value of viscosity at zero shear rate. As can be seen, the functional fit is excellent, suggesting it is a viable method of predicting viscosity, and shows that the transport coefficient *does* converge in two-dimensions.




Figure 5.15: Shear rate against viscosity for a series of SLLOD simulations, run at unit density and temperature. Each simulation was run for 7,000,000 timesteps, with a timestep of 0.001, using 1024 particles. A Lorentzian function is used to fit the data. The inset shows the same data on a log scale. Error bars are present but negligibly small in this instance.

Alternatively, Travis *et al*¹⁰ suggested that a four parameter Cross equation can be used to describe viscosity as a function shear rate (which is formally the same as equation 5.29 in the case that $p = 2$):

$$\frac{\mu(\gamma) - \mu_\infty}{\mu_0 - \mu_\infty} = \frac{1}{(1 + (K\gamma)^p)} \quad (5.32)$$

where η_0 and η_∞ are values of viscosity at very low and very high strain rates, and K and p are constants. Figure 5.16 shows the same data fitted with a Cross equation. Again, the functional fit is excellent, showing that both a Lorentzian and a Cross equation are acceptable methods of fitting two-dimensional SLLOD data.




Figure 5.16: Shear rate against viscosity for a series of SLLOD simulations, run at unit density and temperature. Each simulation was run for 7,000,000 timesteps, with a timestep of 0.001, using 1024 particles. A four parameter Cross equation is used to fit the data.

Hoover *et al*¹¹ conducted an investigation into the dependence of the extrapolated estimation of shear viscosity and the number of particles simulated. They varied the number of particles from 64 to 264196. They observed that the calculated

value of viscosity increased with number of particles significantly up to 1024 particles. Beyond this, there was little further increase in accuracy of the estimated value. However, further increasing the number of particles caused a large increase in computation time. Using this investigation as a basis, it was therefore decided to simulate 1024 particles in all SLLOD simulations in this work.

The dependence of shear viscosity with both density and temperature was determined by conducting a series of SLLOD simulations, analogous to those shown in figures 5.15 and 5.16. This was conducted for both $\epsilon = 100$ and $\epsilon = 10$.

5.3.5 $\epsilon = 100.0$

Each series of SLLOD simulations was run for 7,000,000 timesteps using a timestep of 0.005 and simulating 1024 particles. The first set of simulations were conducted to determine the dependency of viscosity on density. Temperature was fixed at 1.0 throughout using a Gaussian isokinetic thermostat, and density was varied from 0.2 to 1.3.

Figure 5.17 shows the results from these simulations. Each series was fitted with a Lorentzian function (weighted by the error per data point). As the density approached 1.3 the trend changed and the viscosity at low shear rates became proportionally much higher than the viscosity at high shear rates, as shown by the steepness of the curve. Figure 5.5 shows that the fluid begins to solidify at this density. As the continuum model will only model the fluid phase, this is the limit at which the viscosity data was required.

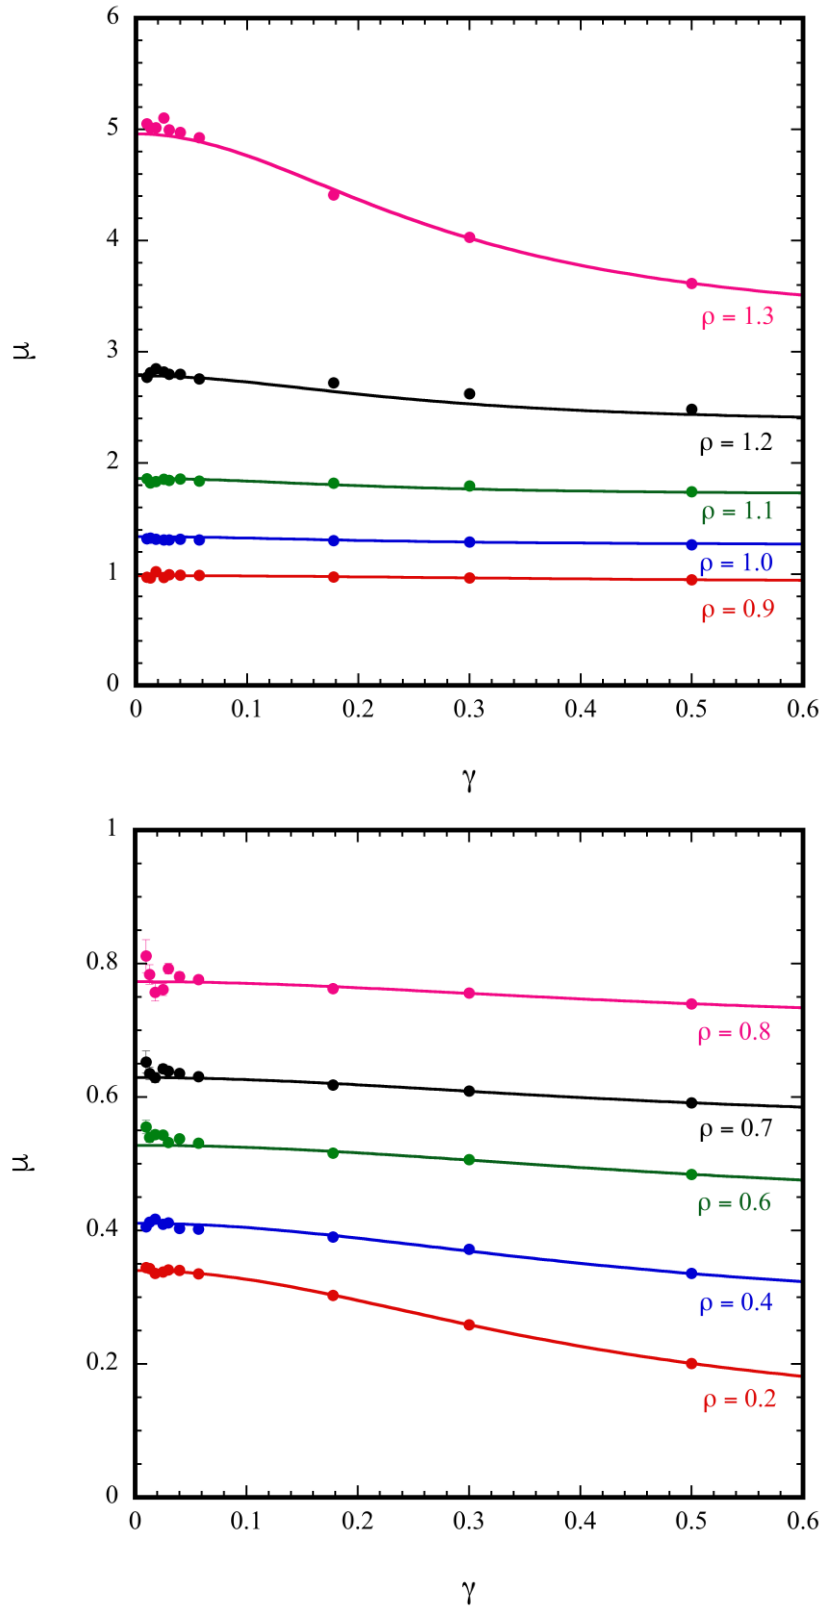


Figure 5.17: Viscosity against shear rate for a series of SLLOD simulations, run at unit temperature. Density was increased from 0.2 to 1.3 gradually. Top: $\rho = 0.9 - 1.3$, bottom $\rho = 0.2 - 0.8$.

To determine the density dependence of the zero shear rate viscosity, each set of results in figure 5.17 had to be extrapolated to zero shear rate. This was obtained from the Lorentzian model. Figure 5.18 shows the results of these extrapolations. This data can be conveniently described by the following functional fit across this range:

$$\mu = 0.443 + 0.003e^{5.68\rho} \quad (5.33)$$

The continuum mechanics simulations will use equation 5.31 to determine the viscosity at a given density. These results agree well with trends observed in experimental investigation into the dependence of viscosity on density.¹² At low density the viscosity dependence is only slight, which increases as the phase change approaches. It is expected that the curve would again level off if simulations were performed at higher densities.




Figure 5.18: Zero-shear rate viscosity dependence on density. The data is taken from extrapolation of figure 5.17.

The second set of simulations were run to calculate the relationship between viscosity and temperature; a series of SLLOD simulations were run where the temperature was gradually increased and the density was fixed at unity. Figure 5.19 shows the results from these simulations.




Figure 5.19: Viscosity against shear rate for a series of SLLOD simulations, run at unit density. Temperature was increased from 0.2 to 1.6 gradually.

Again, the viscosity data has been extrapolated to estimate the viscosity at each temperature at zero shear rate. The results are shown in figure 5.20. The potential only exhibits a weak viscosity dependence on temperature, and fits well with results obtained using similar soft-sphere potentials.¹³ The continuum simulation will use this relationship to predict the local temperature at each particle's coordinates.

$$\mu = 1.245 + 0.067T \quad (5.34)$$




Figure 5.20: Viscosity dependence on temperature at unit density, the viscosity data is taken from extrapolation of the SLLOD simulations in figure 5.18. The full lines is a linear fit used to predict viscosity for a given temperature.

5.3.6 $\varepsilon = 10.0$

The same method was used to define the relationship between density/temperature with viscosity where the effective strength of the potential was reduced to $\varepsilon = 10.0$. The softer nature of the potential resulted in significantly larger errors. In an attempt to reduce these errors, the total number of particles was increased from 1024 to 2048, and the total number of timesteps increased to 10,000,000; increasing both the number of particles and the number of timesteps gives a larger statistical sample and therefore reduces the errors.

Figure 5.21 shows the results of a series of SLLOD simulations where the density was varied from 0.7 to 1.4. It is clear that there is still a degree of noise in the data, which could be reduced by further increasing both the number of particles and timesteps. However, this lies beyond the scope of this work.




Figure 5.21: Results from SLLOD simulations run at constant temperature, varying density from 0.7 up to 1.4. $\varepsilon = 10.0$. The full lines are Lorentzian fits used to extrapolate to zero shear rate.




Figure 5.21 (part B): Results from SLLOD simulations run at constant temperature, varying density from 0.7 up to 1.4. $\varepsilon = 10.0$, shown on a logarithmic scale. The full lines are Lorentzian fits used to extrapolate to zero shear rate.

The data was extrapolated to estimate the value for viscosity at zero shear rate, the results of which are shown in figure 5.22. In the range explored the dependence of viscosity on density is linear; the softer fluid did not show the increase in viscosity associated with a phase change. The following relationship will be used to estimate the viscosity at a given density in the SPAM simulations:

$$\mu = 0.184 + 0.360\rho \quad (5.35)$$




Figure 5.22: Viscosity against density. The data is extrapolated from SLLOD simulations, using the same scale as figure 5.18 for comparison.

The final set of simulations conducted determined the relationship between temperature and density. Figure 5.23 shows the results of the SLLOD simulations. The trend is similar to that seen when $\varepsilon = 100$, showing a linear dependence within the range explored. The results were extrapolated to predict the zero shear rate viscosity, which is shown in figure 5.24. This data can be described by the following equation:

$$\eta = 0.380 + 0.157T \quad (5.36)$$




Figure 5.23: Results from SLLOD simulations run at constant density, varying temperature from 0.6 up to 2.5. $\epsilon = 10.0$.




Figure 5.24: Viscosity as a function of temperature. The data was extrapolated from figure 5.22.

Summary

The temperature and density dependence of zero shear rate viscosity was investigated for two soft-sphere potentials: $\varepsilon = 100$, $\varepsilon = 10$. These relationships can be used by continuum scale models to predict the behaviour of the soft sphere fluid within the range investigated. The first potential, $\varepsilon = 100$, was shown to depend linearly with density at low density, and exponentially with density as the phase transition approached. It was also shown to increase linearly with temperature. The second potential, $\varepsilon = 10$, was shown to linearly with both temperature and density within the range explored, with the shift in phase transition being a result of the softer nature of the potential.

5.4 References

- [1] W. Hoover, T. Pierce, C. Hoover, J. Shugart, M. Stein, A. Edwards, *Molecular dynamics, smoothed-particle applied mechanics, and irreversibility*, Computers Math, **28**, 155, 1994
- [2] T Ashurst, W. Hoover, *Dense-fluid shear viscosity via nonequilibrium molecular dynamcs*, Physical Review, **11**, 658 (1975)
- [3] S. Liem, D. Brown, J. Clarke, *Investigation of the homogeneous-shear nonequilibrium-molecular-dynamics method*, Phys. Rev. A, **45**, 3706 (1992)
- [4] A. Lees, S. Edwards, *The computer study of transport processes under extreme conditions*, J. Phys, **C5**, 1921, 1972
- [5] W.Hoover, D.Evans, R.Hickman, A. Ladd, W. Ashurst, B.Moran, *Lennard-Jones triple-point bulk and shear viscosities. Green-Kubo theory, Hamiltonian mechanics, and nonequilibrium molecular dynamics*, Physical Review A, **22**, 1690 (1980)
- [6] D. Evans, G. Morriss, *Statistical mechanics of nonequilibrium liquids*, ANU Press, Canberra, 1990
- [7] D. Evans, G. Morriss, *Nonlinear-response theory for steady planar Couette flow*, Physical Review A, **30**, 1528 (1984)
- [8] B. Alder, T. Wainwright, *Molecular dynamics by electronic computers*, Proc. of the Int. Symp. on Stat. Mech. Theory of Transport Processes, 97 (1956)
- [9] D. Gravina, G. Giccotti, B. Holian, *Linear and nonlinear viscous flow in two-dimensional fluids*, Physical Review E, **52**, 6123 (1995)
- [10] K. P. Travis, D. Searles, D. Evans, *On the wavevector dependent shear viscosity of a simple fluid*, Molecular Physics, **97**, 415 (1999)
- [11] W. Hoover, H. Posch, O Kum, *Steady-state shear flows via nonequilibrium molecular dynamics and smooth-particle applied mechanics*, Physical Review E, **51**, 273 (1995)
- [12] B. Wang, K. Chu, A. Yu, A. Vince, *Modeling the multiphase flow in a dense medium cyclone*, Ind. Eng. Chem. Res, **48**, 3628 (2009)
- [13] W. Ashurst, W. Hoover, *Dense-fluid viscosity via nonequilibrium molecular dynamics*, Phy. Rev. A, **11**, 658 (1975)

6 SPAM simulation of sand bed filtration

6.1 Overview

Smooth particle applied mechanics was used in order to create a continuum scale simulation of sand bed filtration, where the previous data collected from the molecular dynamics model was used as validation. The equations of state and calculated shear viscosity data from chapter 5 were used to construct this model. An introduction to the methodology is given in chapter 4. The aim was to test the validity of a continuum scale model parameterised from pseudo-experimental data by comparing its agreement with existing literature, and also directly with the molecular dynamics model.

6.2 Model design

The smooth particle model was designed to be as similar to the molecular dynamics model as possible: the boundary conditions, initial conditions, sticking mechanism and property averaging techniques were all implemented using the same methods described in chapter 3. Using this analogous approach allows for direct comparison between the two methods (SPAM/MD), giving validation to the SPAM simulation in a similar method to that discussed by Travis and Hiddlestone.¹

The model aims to predict the flow of a binary mixture of fluid and colloid particles through a bed of static sand particles. The key difference between the two models being that the fluid particles are now modelled as smooth particles obeying the SPAM equations of motion (equations 4.26 – 4.27). The fluid is assumed to be incompressible, leaving the equation of state and the shear viscosity as the constitutive relations remaining to be defined. Chapter 5 outlined the techniques used to define these relations. Figure 6.1 shows a schematic of the simulation, highlighting the two types of particle used. The colloid particles are still treated as soft, repulsive discs similar to how they were treated in MD, interacting through a soft-disc potential (equation 6.1). This avoided the need to develop a more complex equation of state for a binary fluid mixture.

$$\phi_{ij} = \varepsilon \left[1 - \frac{r_{ij}^2}{\sigma^2} \right]^4 ; |r_{ij}| < \sigma \quad (6.1)$$




Figure 6.1: A schematic showing the smooth particle model. The yellow particles are static sand particles, the blue particles are fluid particles modelled as SPAM particles, and the red particles are colloids, modelled as coarse-grained discs.

Two sets of constitutive relations have been calculated for the soft-disc potential, one with $\varepsilon = 100$, and one with $\varepsilon = 10$. $\sigma = 1$.

$\varepsilon = 100$

Hoover calculated the equilibrium equation of state for this force law:² The full set of equations giving the dependence of pressure and temperature on density and energy are summarised below.

$$\frac{PV}{N\varepsilon} = 5 + 8\delta\rho + 2.5\delta\varepsilon + 9\delta\rho^2 + 2\delta\rho\delta\varepsilon \quad (6.2)$$

$$\frac{kT}{\varepsilon} = 1 - \delta\rho + 0.7\delta\varepsilon - 0.86\delta\rho^2 - 0.5\delta\rho\delta\varepsilon \quad (6.3)$$

$$\delta\rho = \left(\frac{N\sigma^2}{V} \right) - 1.0 \quad (6.4)$$

$$\delta\varepsilon = \left(\frac{E}{N\varepsilon} - 1.443 \right) \quad (6.5)$$

The relationship between shear viscosity, temperature and pressure was established in chapter 5. The results were:

$$\eta = 0.443 + 0.003e^{5.68\rho} \quad (6.6)$$

$$\eta = 1.245 + 0.067T \quad (6.7)$$

$\varepsilon = 10$

Chapter 5 also details the method used to define a new equilibrium equation of state when $\varepsilon = 10$:

$$\frac{PV}{N\varepsilon} = 2.47 + 4.2\delta\rho + 1.2\delta\varepsilon + 2.3\delta\rho^2 + 1.8\delta\rho\delta\varepsilon \quad (6.8)$$

$$\frac{kT}{\varepsilon} = 1 - 0.6\delta\rho + 0.7\delta\varepsilon - 0.19\delta\rho^2 - 0.8\delta\rho\delta\varepsilon \quad (6.9)$$

$$\delta\rho = \left(\frac{N\sigma^2}{V} \right) - 1.0 \quad (6.10)$$

$$\delta\varepsilon = \left(\frac{E}{N\varepsilon} - 1.443 \right) \quad (6.11)$$

Finally, chapter 5 also outlines the method used to define the relationship between pressure/temperature and viscosity for this potential:

$$\eta = 0.184 + 0.360\rho \quad (6.12)$$

$$\eta = 0.286 + 0.2213T \quad (6.13)$$

The constitutive equations (equations 6.2 – 6.13) feed into the SPAM conservation equations that define the density, motion and energy of the SPAM particles:

$$\rho_i = \sum_j m_j w(r) \quad (6.14)$$

$$m\dot{v} = - \sum_j m_i m_j \left(\frac{\mathbf{P}_i}{\rho_i^2} + \frac{\mathbf{P}_j}{\rho_j^2} \right) \cdot \nabla_j w(r) \quad (6.15)$$

$$m\dot{e} = -\frac{1}{2} \sum_j m_i m_j \left(\frac{\mathbf{P}_i}{\rho_i^2} + \frac{\mathbf{P}_j}{\rho_j^2} \right) : \mathbf{v}_{ij} \nabla_j w(r) \quad (6.16)$$

$$- \sum_j m_i m_j \left(\frac{\mathbf{Q}_i}{\rho_i^2} + \frac{\mathbf{Q}_j}{\rho_j^2} \right) \cdot \nabla_j w(r)$$

where w is the weight calculated from Lucy's weight function:

$$w(r < h) = \left(\frac{5}{\pi h^2}\right) \left(1 + 3\frac{r}{h}\right) \left(1 - \frac{r}{h}\right)^3 \quad (6.17)$$

All simulations run in this chapter were performed with the coefficients for $\varepsilon = 100$, besides the section detailing the comparison between the two different constitutive relations. Unless otherwise stated, all simulations used the parameters defined in appendix C, using a RK4 integration scheme, and a timestep of 0.001. The constitutive equations combine with the conservation equations to give a closed system where the motion of particles can be calculated.

6.3 Fluid behaviour

Before modelling the process of filtration, a series of simulations were run to test the behaviour of the parameterised fluid under different conditions. These tests aimed to prove that the SPAM fluid behaved in a similar way to the MD fluid from which it was parameterised.

6.3.1 Binary mixture of colloids and fluid

The behaviour of a binary mixture of fluid and colloid particles was tested by simulating 1024 particles (512 colloid and 512 fluid) for 10,000 timesteps, with a timestep of 0.01. All boundaries were periodic, and there was gravitational force. The trajectories of the fluid particles were defined by the SPAM equations of motion, and the trajectories of the colloid particles by the soft disc potential with the *same* parameters from which the SPAM fluid was parameterised ($\varepsilon = 100$, $\sigma = 1$). The cross-particle interaction was also described by the soft disc potential, though with a *different* interaction strength ($\varepsilon = 25.0$). It would therefore be expected that with time the mixture would separate into two distinct phases, each with the same properties. Figure 6.2 shows snapshots at the start (left) and end (right) of the simulation. The binary mixture did indeed show signs of separating, showing similar behaviour to the mixture discussed in section 3.2.2. This serves as evidence that the SPAM fluid experiences a depletion force similar to that observed in the equivalent MD simulations.




Figure 6.2: Snapshots showing the computed particle locations of a binary mixture of 1024 particles (512 fluid, 512 colloid) after 100 and 10,000 timesteps using periodic boundary conditions. $\Delta t = 0.01$. Blue particles are the SPAM fluid, and red particles are colloids.

6.3.2 Fluid behaviour under gravity

The behaviour of the parameterised fluid under gravity was tested by simulating 1024 particles under a gravitational force, $g = 1$, for 100,000 time periods, with a timestep of 0.01. The lateral boundaries were periodic and the bottom boundary was elastic. Figure 6.3 shows the particle coordinates at the start and the end of the simulation for both equations of state. In general, the fluid behaved as expected; the fluid density increased towards the bottom of the simulation, maintaining particle separation. However, the structured order of the particles obtained when $\varepsilon = 100.0$ suggests that the equation of state may struggle to correctly model the fluid behaviour at increased density. Figure 3.10 shows an MD simulation of the same process; it is clear that the MD more closely resembles the behaviour observed when $\varepsilon = 10.0$.

Additionally, an important phenomenon was observed at the elastic boundary. A layer of isolated particles appeared to be “stuck” to the boundary, which remained there throughout the length of the simulation. This is an inherent problem with the implementation of an elastic boundary when using weighted particle averages, and has been observed previously by Hoover.³ This problem could be overcome by

implementing mirrored boundaries (discussed in section 4.2.5). In the interests of pragmatism, mirrored boundaries were not implemented. When the elastic boundary was removed, to allow the fluid to flow through the filter, all memory of this structured layer is lost within a few timesteps, and it did not cause any functional issues.




Figure 6.3: Computed particle locations of the SPAM fluid under a gravitational force of 0.1. $t = 100$ (top), $t = 100,000$ (bottom), $\Delta t = 0.001$. The images on the left are for $\varepsilon = 100.0$, and the images are the right for $\varepsilon = 10.0$. The bottom boundary is elastic, and the lateral boundaries are periodic.

6.4 Simulation instabilities

A known issue when modelling a continuum with smooth particles is the tendency for local pressure waves to turn into shockwaves, causing instabilities in the simulations.⁴ Instabilities were indeed observed when investigating the SPAM fluid and appeared to be as a result of increased local pressure or density. As discussed in chapter 5, the constitutive relations are only considered to be valid within a certain range of density, therefore, modelling a system outside of this range will cause errors; the continuum equations of motion are unable to resolve the pressure waves at increased density.

Figure 6.4 highlights an example of such instabilities. In this case the gravitational force, $g = 2.0$, caused particles to become too densely packed, resulting in a local instability. Shortly after this snapshot, the simulation suffered from complete failure.

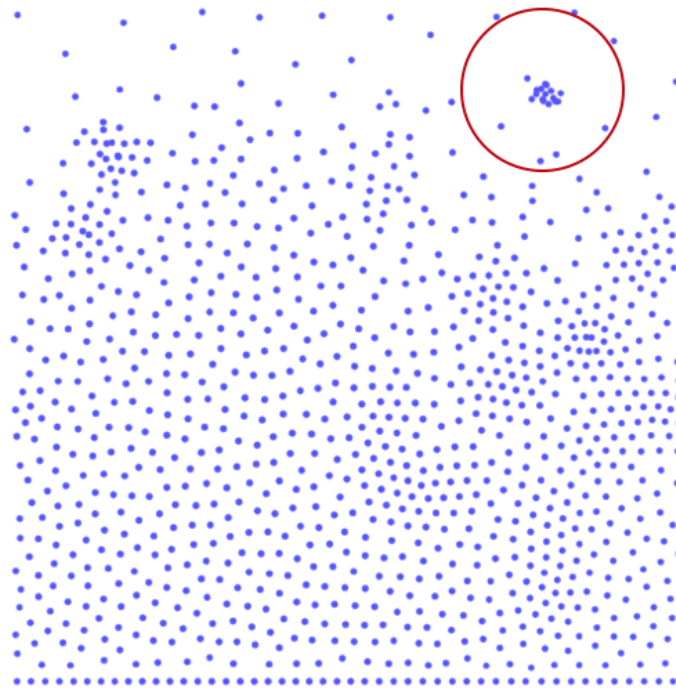


Figure 6.4: Computed particle locations, highlighting a local instability. $g = 2.0$.

There are several known techniques that can be employed to try to alleviate this problem. Introducing an artificial viscosity into the SPAM equations of conservation is one such method, where an additional viscous force contributes to the pressure at any coordinates where the flow is contracting. This is a purely numerical device (that is to say it is not based on any physics) used to stabilise a spreading instability over several particle diameters. Section 4.2.4 discusses one method of achieving this, where the SPAM equations of conservation of momentum and energy are altered to include an additional force.

This work followed the example of Hoover⁵ by employing an additional force to particle pairs where the relative velocity of the pair was negative. This additional force, therefore, was only applied when the particles in the pair were *approaching* each other. A pragmatic choice was to use the same smooth core potential that defined interactions in the MD simulations (equation 6.1). In order to minimise the influence on the continuum dynamics, it was important to keep its effect to a minimum. Values of $\varepsilon = 1.0$, and $\sigma = 0.75$ were used; this provided a strong enough repulsion to separate particles that clustered, but was also only active at distances significantly *less* than the smoothing length ($h = 3$). The addition of this force removed the instabilities observed at a constant number of particles.

However, further instabilities occurred when new particles were added to the system at regular intervals. If the new particles were introduced before the previous set of particles had traversed a distance equal to or greater than the smoothing length down the filter, significant problems arose. The frequency at which new particles could be added to the system was therefore greatly limited; it was found introducing new particles every 1000 timesteps was sufficiently slow to avoid this problem. This frequency is 1/5 of the frequency employed in the MD simulations. The SPAM simulations were therefore run for 500,000 timesteps (5 times longer than the MD simulations) to give the same total number of particles.

6.5 Specific deposit and pressure drop

To validate that the SPAM model qualitatively agreed with the MD model in terms of the trend of rate of deposition and the trend of rate of increase in pressure drop, several simulations were run. It is unlikely that the results would quantitatively agree due to the small changes in simulation conditions; the SPAM flow rate was 1/5 that of the MD flow rate, and the SPAM fluid employed an additional artificial viscosity.

6.5.1 Specific deposit

Specific deposit, σ , was defined as the number of deposited colloids divided by the volume of the filter. A more detailed discussion of the expected relationship with time is discussed in section 3.3.1. The key points are:

- the amount of deposit should increase linearly with time
- the rate of deposit should change at a threshold value of σ

The dependence on specific deposit with time for the SPAM model was determined using a similar method to section 3.3.1. Figure 6.5 shows the results. The results do indeed show two linear regimes, agreeing with both the MD model and the experimental predictions of Camesano.⁶ Taking the gradient of the slopes of the two phases gives rate constants of $k_r = 4.1 \times 10^{-5}$ and $k_{ac} = 1.7 \times 10^{-5}$ for the blocking and operation phases, respectively. Table 6.1 compares these values with those from the equivalent MD simulations. Although the rates cannot be directly compared, the relative rates can; both the MD and SPAM show a decrease in rate of deposit of just over half from stage one to stage two. This result is encouraging, suggesting the SPAM model is capturing similar dynamics to the MD model.




Figure 6.5: Specific deposit as a function of time showing the two linear stages.

	MD	SPAM
k_r	2.000×10^{-04}	4.100×10^{-05}
k_{ac}	9.200×10^{-05}	1.700×10^{-05}

Table 6.1: Linear rate constants for the MD and SPAM simulations.

6.5.2 Pressure drop

As discussed in section 3.3.2, the pressure drop is expected to scale linearly with time during the first phase, and quadratically with time during the second phase.

The pressure drop was calculated using the same method outlined in section 3.2.6. Figure 6.6 shows the pressure drop as a function of specific deposit, comparing the SPAM simulation to the MD simulation. The general trend of the results is the

same, with the SPAM model showing a quadratic relationship between the normalised pressure drop and specific deposit. This agrees with the experimental results of Veerapeneni⁷, and the predictions of Mays and Hunt⁸. However, there are obvious differences in the quantitative data between the two methods: at low values of specific deposit the SPAM simulation slightly overestimates the pressure drop, whereas at high values of specific deposit the SPAM simulation *greatly underestimates* it.




Figure 6.6: Pressure drop as a function of specific deposit, comparing the results from MD and SPAM simulations.

The parameterisation of the equation of state is likely to be the cause of this underestimation. As the simulation progresses, the clogging of particles causes an increased density (and therefore increased pressure) at the top of the filter bed. Without the implementation of an artificial viscosity (in this case, through a soft sphere potential), this would most likely lead to instabilities through shockwaves. The additional potential prevents this, but it does *not* correct the values of pressure

obtained from the equation of state; pressure remains underestimated at high densities. This discrepancy widens with time. To fix this further parameterisation of the soft-sphere potential at higher densities is required, which is beyond the scope of this work.

6.6 Comparison with MD

To further compare the SPAM model with the MD simulation results, a systematic exploration of the parameter space was performed. This involved varying the porosity of the filter, the colloid interaction strength, and the lattice structure. These parameters were chosen for two reasons: analogous simulations had already been run using the MD model allowing for direct comparison, and these parameters were shown to have distinct trends when varied in the MD model. The MD data for comparison is taken from chapter 3. All simulations in this section were run using the Runge-Kutta 4 integration scheme, with $\Delta t = 0.001$.

6.6.1 Colloid-colloid potential

The comparison with the MD model was performed by varying the strength of the colloid-colloid interaction from 2 to 100. Figure 6.7 shows the dependence of σ on ε for the two methods. It is clear that the results differ in several areas. The SPAM model consistently predicts a total deposit of approximately 30 % less than the MD model at high interaction strength. The trend above an interaction strength of $\varepsilon = 20$ was similar, with there being no significant dependence on deposit with interaction strength. However, at lower interaction strength the MD predicted an increase in specific deposit, whilst the SPAM predicted a decrease. This is likely to be a result of the mechanics being dominated by the SPAM particles, rather than the colloid particles at low interaction strength.




Figure 6.7: Specific deposit as a function of colloid-colloid interaction strength.

6.6.2 Porosity

The porosity of the filter was varied from 70.57 to 1.44. Figure 6.8 shows the dependence on specific deposit with porosity for the two models. Again, the SPAM simulations underestimated the overall deposit in the filter. SPAM also showed an increase in deposit with decreasing porosity. Conversely, the MD results showed a maximum at a porosity of 42.24, where a further decrease in porosity did not equate to an increase in deposition. This was shown to be because the filter became mechanically clogged at this point (section 3.4.4). The fact that the SPAM simulations did *not* show this maximum suggests that there was not enough deposit in the sand bed to cause complete clogging.




Figure 6.8: Specific deposit as a function of porosity, comparing the MD and SPAM simulations.

Figure 6.9 shows the calculated pressure drop as a function of porosity for the same simulations. The increase observed with decreasing porosity was only slight for the SPAM model, and is constant with the underestimation of pressure seen previously. The trend seen is the same, suggesting further parameterisation of the pressure equation of state would help to reduce the difference in results.




Figure 6.9: Pressure drop as a function of porosity, comparing the MD and SPAM simulations.

6.6.3 Square lattice

The effect of changing the lattice structure from a triangular lattice to a square lattice was examined by running an analogous set of simulations to section 3.4.2. The results are shown in figure 6.10. The MD simulations showed a slight increase in specific deposit, along with a slight increase in pressure drop with increasing colloid repulsion. As discussed previously, the square lattice creates a much simpler flow path, meaning an increase in repulsion does not necessarily hinder the trajectories of the particles. It is clear that the results from the SPAM simulations are distinctly different. A much larger increase in pressure drop is observed with increasing repulsion, suggesting that the SPAM particles suffer from increased hindrance when travelling through pore spaces than the MD particles. The reasons for this are, again, likely to be as a result of the parameterisation. The increased density in the pore spaces takes the simulation away from the accurately modelled

area of the equation of state. This could be further explored by separating the fluid particles in the pore spaces into multiple smaller particles, giving higher resolution, highlighting one of the important advantages of the SPAM method. This would require additional code changes, and would be an area explored further in future work.




Figure 6.10: Specific deposit as a function of porosity, comparing the MD and SPAM simulations.

6.6.4 Equation of state

The effect of changing the continuum potential from $\varepsilon = 100$ to $\varepsilon = 10$ was investigated by running a series of simulations with varying colloid interaction strength (analogous to section 6.4.1), with the second parameterised potential. Figure 6.11 shows the results. The consistent underestimation of deposit remains, however, the trend now matches that seen in the MD simulations, with a higher rate of deposition observed at lower colloid interaction strength. This suggests that the $\varepsilon = 10$ equation of state is better parameterised than the $\varepsilon = 100$.




Figure 6.11: Specific deposit as a function of porosity, comparing the MD and SPAM simulations.

The softer fluid is capable of withstanding a much higher density before giving rise to significant deviations in either pressure or viscosity (compare figure 4.9 to 4.2, and 4.21 to 4.18). Indeed the predicted relationship between viscosity and density observed was linear up to a density of 1.6 when $\varepsilon = 10$, and exponential when $\varepsilon = 100$. The observed difference in pressure from a density of 1 to 1.6 was 3 when $\varepsilon = 10$, and 8 when $\varepsilon = 100$. The softer continuum fluid is capable of more

accurately predicting the behaviour of the soft-sphere potential at the densities created by the filtration simulations. At extreme densities, it is likely that significant deviations between the two methods would again be observed.

6.7 Summary

Smooth particle applied mechanics was used to create a continuum scale model of a filtration process. Instabilities were initially observed at increased densities, ultimately causing simulation failure. A type of artificial viscosity that discouraged particle clustering was used to alleviate this. This was a pragmatic choice, and was made with the intent of causing as little effect on the dynamics as possible. Introducing surface tension, bulk viscosity, or an artificial viscosity incorporated into the SPAM equations of motion would perhaps be an improvement on this, but the implemented method was ultimately successful.

An attempt to validate the model against the results gained from MD and trends reported in literature was undertaken. The dependence on specific deposit with time showed the characteristic two phases associated with filtration, with both phases being linear. In addition to this, the pressure drop was shown to depend quadratically with the amount of deposit. These are encouraging results, suggesting the smooth particle approach can be used to model filtration.

However, the SPAM model was shown to consistently underestimate both the total deposit within a filter, and the associated pressure drop. It is likely that this was a result of the range under which the equation of state was valid. In order to accurately model a filtration process, where the continuum model can manage sufficiently high densities to give accurate predictions of the pressure drop, further parameterisation is required.

6.8 References

- [1] K. Travis, T. Hiddleston, *Multiscale modelling of material failure using particles*, *Molecular Simulation*, **40**, 141 (2014)
- [2] W. Hoover, T. Pierce, C. Hoover, J. Shugart, M. Stein, A. Edwards, *Computers Math*, **28**, 155, 1994
- [3] O. Kum, W. Hoover, C. Hoover, *Physical Review E*, **68** (2003)
- [4] LAMMPS, *The implementation of Smooth Particle Hydrodynamics in LAMMPS: A guide to the SPH-USER package*. [ONLINE] Available at: <https://www.cs.cornell.edu/~bindel/class/cs5220-f11/code/sph.pdf> [Accessed September 2016]
- [5] W. Hoover, *Smooth Particle Applied Mechanics*. World Scientific, Singapore, 2006
- [6] T. Camesano, K. Unice, B. Logan, *Colloids and Surfaces A*, **160**, 291 (1999)
- [7] S. Veerapaneni, *Formation and morphology of colloidal deposits in porous media*, Ph.D. Thesis, Rice University, Houston, TX, 1996
- [8] D. Mays, J. Hunt, *Environ. Sci. Technol.*, **39**, 577 (2005)

7. Conclusions and further work

The overriding aim of this work was to probe the mechanism of clogging of sand bed filters using computer simulation. The existing work in this field makes use of either small scale experimental techniques or top down empirical modelling techniques. These methods provide useful information about the specific deposit and pressure drop observed in filtration processes, allowing for prediction of these parameters under certain conditions. However, they do *not* give truly predictive results; prior experimental work is required to define parameters used in the resulting expressions. Computer simulation allows for a mechanistic insight into the process of clogging, giving the prospect of relating the empirical parameters seen in previous work to properties of the colloid, fluid, and sand particles that make up a filter.

The scientific case and technical scope provided by NNL (who part funded this work) provided several key aims for the particle-based simulation:

- predict the relationship of specific deposit with time.
- predict the dependence of pressure drop on specific deposit.
- use existing literature to validate these dependencies.
- test the mechanical hypothesis proposed by Mays and Hunt:¹ “The fractal dimension of colloidal deposits has been speculated to be an important parameter controlling the amount of clogging. Deposits of high and low fractal dimension might arise as a result of the competition between the colloid-colloid interactions and hydrodynamic forces.”
- perform a systematic exploration of the system by varying properties of the sand, colloid and fluid particles, and their associated interactions.

Two particle-based simulation techniques were used to achieve these aims: molecular dynamics and smooth particle applied mechanics. The first model developed made use of coarse-grained molecular dynamics, modelling the system as a series of soft-discs interacting through a smooth, soft core potential. The simulations yielded several significant results. The build-up of specific deposit was shown to depend linearly with time, with the rate of deposit changing at a threshold value of σ . This result agreed with previous experimental work, and was directly comparable to the equations used to predict deposition in existing empirical models. Furthermore, ΔH , was shown to depend linearly

on deposit at low deposition, and quadratically with deposit at high deposition. Importantly, this result also agreed with previously developed expressions for pressure dependence on deposit. Further validation to the dynamics was given when examining the deposition concentration with depth; the concentration of deposited colloids decreased exponentially with depth in most conditions. In conditions where colloids were favourably deposited at the top of the filter, a characteristic hyper-exponential profile was observed. These results showed that the molecular dynamics model, at least qualitatively, agreed with other methods. This was a key success of this research.

An investigation into the fractal dimension of the colloidal deposits was then performed. This showed that the deposits generated through molecular dynamics did indeed have a fractal dimension. It was then suggested that varying the sticking distance parameter was a means of controlling the fractal dimension of the deposits, with a longer sticking distance resulting in a lower fractal dimension. This allowed for the comparison of the flow through filters with deposits of both high and low dimension. The results agreed with the hypothesis proposed by Mays and Hunt. This is a significant mechanistic insight, and highlights the main advantage in particle based methods.

A systematic exploration of this model was also undertaken. The strength of the colloid-colloid, colloid-fluid, and colloid-sand interactions were varied in isolation. The results suggest that both the nature of the colloid-colloid and the colloid-fluid interaction have an important impact on the dynamics. A strong colloid-colloid repulsion leads to an increased pressure drop and a decrease in deposit rate; the deposited colloids both hinder further deposition, and block colloids from flowing through the filter. Interestingly, a strong colloid-fluid repulsion has a similar effect, though the mechanism observed is different. Rather than the pore spaces being clogged with deposited colloids, pore spaces become clogged with fluid particles unable to penetrate further into the filter. The effect of the changing lattice structure, packing fraction, and concentration of colloids was also observed.

The results obtained from this model achieved all of the aims of the project from the outset, however, the work done with it is by no means at a maximum. An additional area for exploration lies in the packing fraction and size of the sand particles. A real sand bed is likely to suffer from a sedimentation process where larger particles settle toward the bottom of the bed. The model assumes a uniform size and packing fraction of sand particles. Using a particle size distribution to vary the size of the sand particles, alongside a random lattice structure would allow for further exploration of the dynamics of the

system. In addition, the model assumes a simple, irreversible, sticking mechanism, which could be improved upon. Using a short range potential with an attractive well would perhaps achieve more realistic results. This would allow deposited colloids to move and dislodge as a result of hydrodynamics, and could potentially lead to a characteristic nonmonotonic deposit profile being observed. Furthermore, quantitative, rather than qualitative, agreement with experimental data could be achieved through using interatomic potentials parameterised on a realistic system of water, sand and colloid particles. This is far beyond the scope of this project, but would provide an extra layer of validation to the results proposed.

A more novel approach was used to model the same system using continuum mechanics. Smooth particle applied mechanics is a particle based solver of continuum mechanics that holds the potential to increase the time and length scales accessible by computer simulation. This required substantial parameterisation to describe the behaviour of the fluid in various conditions. To allow for direct comparison between the two models, the continuum model was parameterised using pseudo-experimental data from molecular dynamics. The equation of state used to describe the fluid was obtained using equilibrium molecular dynamics simulations of the smooth, soft core potential. The shear viscosity dependence on pressure and temperature for the same potential was obtained using non-equilibrium molecular dynamics.

An equilibrium equation of state developed by Hoover was investigated, showing it to be valid to a substantial deviation from the reference point. A similar equation of state was developed using the same method for a much softer interaction strength ($\varepsilon = 10$). The dependence on viscosity with shear rate was probed through a series of SLLOD simulations, yielding relationships between pressure and temperature with viscosity for both potentials ($\varepsilon = 10$ and $\varepsilon = 100$). The results of this parameterisation were used to develop a continuum scale model of filtration. Additionally, future simulators can use these results to create continuum scale models other processes, giving use to the collected data far beyond this work.

The SPAM model showed signs of instability under certain conditions. An artificial viscosity was used to alleviate this problem, giving rise to a stable continuum scale model of filtration. The model showed the same characteristic linear dependence of specific deposit with time, and quadratic dependence of pressure drop with deposit, which is a positive result. However, the quantitative agreement between the models was far from exact. Further parameterisation of the fluid would lead to a better agreement

between the two methods. The SPAM model consistently underestimated the pressure drop at high deposits, suggesting that further characterisation of the fluid at high density is required. Additionally, the model did not include surface tension, a constitutive relation that could also be obtained from molecular dynamics simulation.

7.1 References

- [1] D. Mays, J. Hunt, *Environ. Sci. Technol.*, **39**, 577 (2005)

Appendix A: A useful random number generator. Creates a random seed based on the system clock time.

```
SUBROUTINE init_random_seed()
  INTEGER :: i, n, clock
  INTEGER, DIMENSION(:), ALLOCATABLE :: seed

  CALL RANDOM_SEED(size = n)
  ALLOCATE(seed(n))

  CALL SYSTEM_CLOCK(COUNT=clock)

  seed = clock + 37 * (/ (i - 1, i = 1, n) /)
  CALL RANDOM_SEED(PUT = seed)

  DEALLOCATE(seed)
END SUBROUTINE
```


Appendix B: Input parameters used to in MD simulation of filtration

Parameter	Type	Value	Notes
Number of timesteps	Integer	100000	
Timestep	Double	0.001	
Initial density	Double	1.000	
Load coordinates	Integer	1	0 = start from square lattice, 1 = load coordinates
Side boundary	Integer	1	0 = none, 1 = periodic, 2 = elastic
Top boundary	Integer	2	0 = none, 1 = periodic, 2 = elastic
Shape of simulation cell	Integer	1	0 = square, 1 = rectanlge
Gravity	Double	1.000	
Damp	Double	0.000	
Seed	Integer	1	0 = no input seed, 1 = use system clock as seed
Initial kinetic energy	Double	1.000	
Number of fluid particles	Integer	512	
Number of colloid particles	Integer	512	
Sigma	Double	1.000	
Epsilon fluid	Double	10.000	
Epsilon colloid	Double	10.000	
Epsilon sand	Double	10.000	
Epsilon fluid/colloid	Double	10.000	
Epsilon fluid/sand	Double	10.000	
Epsilon colloid/sand	Double	10.000	
Number of planes	Integer	100	
Average planes frequency	Integer	100	

Filter	Integer	1	0 = do not filter, 1 = filter
Filter height	Double	30.000	
Number of sand particles per row	Integer	5	
Number of new particles	Integer	30	
Frequency of new particles	Integer	200	
Sticking probability	Double	0.001	
Clogging distance	Double	0.500	
Colloid sticking probability	Double	0.001	
Colloid clogging distance	Double	0.500	
Lattice type	Integer	1	0 = square, 1 = triangular
Smoothing length	Double	3.000	

Appendix C: Input parameters used in SPAM simulations

Parameter	Type	Value	Notes
Number of timesteps	Integer	100000	
Timestep	Double	0.001	
Initial density	Double	1.000	
Load coordinates	Integer	1	0 = start from square lattice, 1 = load coordinates
Side boundary	Integer	1	0 = none, 1 = periodic, 2 = elastic
Top boundary	Integer	2	0 = none, 1 = periodic, 2 = elastic
Shape of simulation cell	Integer	1	0 = square, 1 = rectanlge
Gravity	Double	1.000	
Damp	Double	0.000	
Seed	Integer	1	0 = no input seed, 1 = use system clock as seed
Initial kinetic energy	Double	1.000	
Number of fluid particles	Integer	512	
Number of colloid particles	Integer	512	
Sigma	Double	1.000	
Epsilon colloid	Double	10.000	
Epsilon sand	Double	10.000	
Epsilon fluid/colloid	Double	10.000	
Epsilon fluid/sand	Double	10.000	
Epsilon colloid/sand	Double	10.000	
Number of planes	Integer	100	
Average planes frequency	Integer	100	
Filter	Integer	1	0 = do not filter, 1 = filter

Filter height	Double	30.000	
Number of sand particles per row	Integer	5	
Number of new particles	Integer	30	
Frequency of new particles	Integer	200	
Sticking probability	Double	0.001	
Clogging distance	Double	0.500	
Colloid sticking probability	Double	0.001	
Colloid clogging distance	Double	0.500	
Lattice type	Integer	1	0 = square, 1 = triangular
Smoothing length	Double	3.000	
MD separation force	Integer	1	0 = no separation force, 1 = use separation force
MD separation sigma	Double	1.000	
MD separation epsilon	Double	0.750	
Equation of state	Integer	2.000	1 = eps(10), 2 = eps(100)