

FLORIDA STATE UNIVERSITY
COLLEGE OF ARTS & SCIENCES

A NOVEL MODEL FOR PRECIPITATION REACTIONS IN MICROFLUIDIC DEVICES
OR, A UNIFIED FRAMEWORK FOR COUPLING PHASE-CHANGE CHEMISTRY AND
FLUID-STRUCTURE INTERACTION IN LOW REYNOLDS NUMBER FLOWS

By

PATRICK S. EASTHAM

A Dissertation submitted to the
Department of Mathematics
in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

2020

ProQuest Number:28003050

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent on the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest 28003050

Published by ProQuest LLC (2020). Copyright of the Dissertation is held by the Author.

All Rights Reserved.

This work is protected against unauthorized copying under Title 17, United States Code
Microform Edition © ProQuest LLC.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 - 1346

Patrick S. Eastham defended this dissertation on July 2, 2020.
The members of the supervisory committee were:

Nick Moore
Professor Co-Directing Dissertation

Nick Cogan
Professor Co-Directing Dissertation

Wei Yang
University Representative

Richard Bertram
Committee Member

Sangyun Lee
Committee Member

The Graduate School has verified and approved the above-named committee members, and certifies that the dissertation has been approved in accordance with university requirements.

For my girls
who keep me going

ACKNOWLEDGMENTS

First and foremost, I thank my advisors, Nick Moore and Nick Cogan, who assisted in every facet of my graduate development. They were my professors while taking courses, and went on to continue mentoring me on conducting research, communicating ideas and results effectively, and navigating an academic professional environment. This dissertation would not have been possible without their consistent support and encouragement.

I thank my collaborators in the Department of Chemistry & Biochemistry. Oliver Steinbock and his graduate student, Qingpu Wang, were crucial in explaining their experimental setup, motivation, and questions that they were interested in that mathematical modeling could address.

I thank Kourosch Shoele in the Department of Mechanical Engineering for hiring me as a research assistant and collaborating with me on a project that – while not the topic of this dissertation – helped me both learn to carry out research and write papers.

I thank Bryan Quaife in the Department of Scientific Computing for being an excellent professor, teaching me modern techniques for numerically solving integral equations and linear systems, and for allowing me to use his boundary integral equation code – collaboratively build with Nick Moore – in Sec. 3.2. Without that code, the analysis in that section certainly would not have been possible.

I thank fellow graduate student Matt McCurdy who collaborated with me on the initial formulation, algorithms and coding for the sinkhole application in Chapter 5.

I would also like to acknowledge helpful discussions I have had with visiting speakers during their time in Tallahassee, including Wanda Strychalski, Michael Shelley, Stephen Morris, and Daniel Harris. Their helpful comments, critiques, and encouragements were greatly appreciated.

Finally, I gratefully acknowledge support from the National Science Foundation Graduate Research Fellowship Program under Grant 1449440.

TABLE OF CONTENTS

List of Figures	vii
List of Symbols	xi
Abstract	xii
I Framework Derivation and Validation	1
1 Motivation	2
1.1 A microfluidic experiment	3
1.2 Existing mathematical methods	5
1.3 What a new framework achieves	8
2 The Multiphase Framework	11
2.1 Volumeless ions within a multiphase fluid	12
2.1.1 Basic reaction	12
2.1.2 Complex reactions	18
2.2 Multiphase momentum equations	22
2.2.1 Stationary solid and friction	22
2.2.2 Complex discrete rheology	25
2.3 Discrete–continuum interpolation operators	28
2.3.1 Discrete-to-continuum operator	28
2.3.2 Continuum-to-discrete operator	29
2.4 Framework summary	30
2.4.1 Chemistry	31
2.4.2 Mass Balance	31
2.4.3 Momentum Equations	32
2.4.4 Interpolation Operators	32
3 Framework Validation and Analysis	33
3.1 Analysis of a reduced model	33
3.1.1 Fixed point analysis	34
3.1.2 Ricatti-Poiseuille analysis	36
3.2 Interpretation of multiphase drag	42
3.2.1 Benchmark problem	42
3.2.2 Indirect method for computing drag	43
3.2.3 Porous circle problem	45
3.3 Particles within a step channel	48
3.4 Benchmark for continuum-to-discrete operator	51
3.4.1 Gradient descent algorithm	52
3.4.2 Single particle transform	53
3.4.3 Transform of step function	53

II Applications	55
4 Precipitation within a Microfluidic Chamber	56
4.1 Model and timescale separation	57
4.1.1 Nondimensionalization	57
4.1.2 Timescale separation	60
4.2 Numerical discretization	61
4.2.1 Quasi-steady reaction	62
4.2.2 Weak formulation of PDEs	64
4.2.3 Validation	66
4.3 Results	68
4.3.1 Wavy domain geometry	68
4.3.2 Hypothesis for unidirectional growth	69
4.4 Conclusion	70
5 Sinkhole Formation	72
5.1 Geophysical model	73
5.2 Numerical discretization	74
5.3 Results	74
5.4 Conclusion	77
6 Discussion and Conclusion	79
6.1 Future applications	80
6.1.1 Ultrasonic drug delivery	80
6.1.2 Biofilms in water filtration	80
6.1.3 Framework extensions	81
Appendices	
A Solving Ricatti’s Differential Equation	83
B Tree-Based Cell Lists for Fast Interpolation	85
C Software Packages	87
References	90
Biographical Sketch	99

LIST OF FIGURES

1.1	Inactive “Black smoker” hydrothermal vent. Vents like this are integral to one theory for the “origin of life” [67]. Reproduced from [57] with permission.	4
1.2	Inorganic precipitate membrane formed in a microfluidic channel. (a) The photograph of the resulting membrane after 2 h with 0.5 M NaOH and 0.5 M NiCl ₂ solutions being injected simultaneously into the microfluidic device. The mixing part of the channel is 50 mm long, 2 mm wide, and approximately 130 μm high. (b) A magnified view of a selected area from (a). (c-g) A sequence of micrographs showing the unidirectional thickening process after (c) 1 min, (d) 15 min, (e) 30 min, (f) 60 min, and (g) 120 min. Scale bars correspond to (a) 1 cm, (b) 5 mm, and (c) 200 μm. Reproduced from [36] with permission.	5
1.3	Formation of Cu(OH) ₂ precipitate in a test tube. Photograph by José Furtado, distributed under a CC BY-SA 4.0 license [40].	6
1.4	Example of FSI application: simulation of blood flow through artery using (i–iii) IBM and (iv) OpenFOAM. Reproduced from [46] with permission, with more algorithmic details found in [58].	7
1.5	Examples from analysis and simulations conducted in this thesis: (a) Velocity profile of 1D reduced model gives no-slip behavior in reaction zone (grey) in Sec. 3.1; (b) Discrete particles implemented for moving porous media give expected porous behavior in Sec. 3.3; (c) Velocity magnitude of full simulations in wavy microfluidic domain, where dark red is maximum velocity and dark blue is zero velocity in Chapter 4.	9
2.1	Schematic of precipitate reaction in a control volume. Precipitation causes solution (white) to transform into solid (shaded) after a certain concentration threshold is reached of aqueous product C . Aqueous chemicals A , B and C are volumeless scalar fields while the solvent and solid are treated as multiphase components. The volume of solid gained is exactly equal to the volume of solvent lost. Reproduced from [36] with permission.	13
2.2	Comparison of friction terms ξ . ξ_{KC} (solid) is singular in the limit $\theta_s \rightarrow 0$, ξ_H (dash) is non-singular in the porous limit ($K = 0.5$, $n = 2$) and ξ_B (dot) yields maximum friction when both phases are present in equal amounts. All terms have been normalized by choosing h such that $\xi(\theta_s^*) = \xi^*$. Reproduced from [36] with permission.	24
3.1	Dynamical system for ψ_C and θ_p . (a) Qualitative stability of $\dot{\psi}_C(\psi_C)$ ODE. The left equilibrium ψ_C^- is stable and exists for $\chi \geq 0$. (b) Visualization of planar dynamical system approaching the fixed point $(\psi_C^-, 1)$ for $\chi > 0$. The thicker line corresponds to homogeneous initial conditions for ψ_C and θ_p . Shaded region is outside of the domain of $(\psi_C, \theta_p) \in [0, \infty) \times [0, 1]$. Reproduced from [36] with permission.	36
3.2	Developing precipitate affects fluid flow. (a) Flow profile in a 1D channel transitions from one-channel to two-channel flow. The reaction region is shaded. (b) Relevant	

variables evaluated in the reaction region at $x = L/2$, normalized for legibility; ψ_C develops first, followed by θ_p , which when large enough triggers the transition from one- to two-channel flow. The percolation threshold θ_s^* is set to $\theta_s^* = 0.3$, which is why negligible change in fluid velocity is seen until $\theta_p \approx 0.7$. The specific θ_s^* was chosen based on the maximum volume fraction of an arrangement of packed spheres in three dimensions [26]. The black diamond and square represent the system state inside the reaction region at $t = 2.5$ and $t = 5$, respectively. Reproduced from [36] with permission. 38

3.3	Comparison of different models. (a)–(c) use the stress in Eq. (2.47) with the friction term given by (a) Kozeny-Carman, (b) Hill function, and (c) the ‘biofilm’ term; (d) uses a conventional multiphase stress \mathbf{T}' in Eq. (3.13) with ξ_B . The coefficients h are scaled so as to make the three coefficients comparable in strength. The first two friction terms produce the desired no-slip on the membrane interface and the third, while affecting the fluid flow, does not generate the desired no-slip boundary condition. Additionally, the stress and friction combination used in (d), which is usually employed in multiphase models, does not yield the desired no-slip behavior on the fluid-structure interface. Reproduced from [36] with permission.	39
3.4	Effect of increasing precipitate thickness. As a percentage of the domain length, membrane width is (a) 5%, (b) 15%, and (c) 33%. The increasing maximum flow speed is due to the constant-flux constraint, and is analogous to that what would occur if the precipitate boundaries and interface conditions were prescribed <i>a priori</i> in a single-phase flow. Reproduced from [36] with permission.	40
3.5	Benchmark domain for drag analysis. The circle (shaded) with radius r will either be represented by a porosity field (present framework) or by a collection of small “filling circles” (boundary integral equations).	43
3.6	Comparison of flow fields using our model framework versus a classical no-slip boundary condition. Color represents the magnitude of the velocity field, and both images use the same color scaling. Comparison of the flow generated from (a) our multiphase Brinkman equation and (b) classical Stokes flow over a circle. In the multiphase model, the circle is represented by a volume fraction scalar field, while in (b) a circle must be explicitly built into the mesh so that a no-slip boundary condition may be applied.	46
3.7	Examples from porous circles with volume fractions equal to (a) 5% (b) 10% (c) 15%, (d) 20% (e) 25% (f) 30% (g) 35% (h) 40% (i) 45% (j) 50% (k) 55% (l) 60%. Note that the actual volume fractions (cumulative area of filling circles divided by area of outline) are not exactly these values, but only approximate, due to the filling circle radii being taken from a random distribution. See Appendix C for more details. . . .	47
3.8	Drag on porous body versus volume fraction. The “porous body” is (a) composed of 40 “filling circles” and (b) represented by a volume fraction scalar field in our domain with $\xi^* = 100$, $\theta_{fluid}^* = 0.6$. The disparate scales of the vertical axis indicates that we should not interpret our porous objects as being composed of small circles, at least in two-dimensions.	48

3.9	Particle dynamics for cohesion, adhesion and gravity, where fluid flow is not incorporated. The particles, while only being represented as points, exhibit concrete radii and allow themselves to pile up next to and on top of walls. Some particles are colored for convenience.	49
3.10	Particle dynamics for only fluid flow. In this simulation, particles behave simply as volumeless tracers without any steric interactions. They follow the fluid flow over the step.	50
3.11	Particle dynamics for (a) one-way fluid coupling, including gravity, cohesion and adhesion and (b) all forces, including a two-way fluid-particle coupling. The cohesion of the particles allows clumps to move coherently as a single (or multiple) bodies. In (a) the top of the initial clump feels the full force of the fluid, and is pushed over the step; the bottom half of the initial clump is overtaken by gravity and falls before reaching the step. In (b) the entire initial clump is overtaken by gravity, unlike Fig. 3.11, because the fluid velocity inside the clump is significantly less due to the permeability of this clump affecting the fluid velocity	51
3.12	Convergence for single particle continuum-to-discrete transform. (a) Physical space convergence and (b) cost function for increasing iterations both demonstrate that the convergence experiences a “transition region” where, before and after, convergence is slow.	53
3.13	N particles attempting to approximate a step function via the inverse transform \mathcal{G}^*	54
4.1	Microfluidic setup for precipitate reaction. Ionic species enter the tubes on the left, and exit the channel on the right. Reproduced from [109] with permission.	56
4.2	Schematic of variables and parameters, with a top-down view of the channel. Precipitate forms in the middle, which affects the surrounding fluid flow.	58
4.3	Schematic representing idea that aqueous chemical dynamics reach steady state on every reaction-diffusion time step $\Delta\tau$	61
4.4	Validation of time-dependent advection-diffusion equation. (a) Convergence study for steady advection-diffusion equation using method of manufactured solution on rectangular domain. (b) Convergence study for time stepping on diffusion equation (no advection). This convergence study was conducted in the “Y-shaped” domain, and all errors were taken after one time unit of simulation. Both errors are relative.	67
4.5	Convergence study for the variable-coefficient multiphase Brinkman equation, using the method of manufactured solutions. Results show second-order convergence in space for the velocity variable, with first-order convergence for the pressure variable, as expected.	67
4.6	Schematic of wavy domain. The generality of the FEM scheme allows us to run simulations on domains with realistic geometries.	69
4.7	Flow at 4 different time points for wavy-shaped domain. Left side represents solid volume fraction (red = 1, blue = 0) along with velocity vector field, right side shows	

velocity magnitude; rows are at equivalent times. One can see the velocity field going to no-slip as the precipitate develops. Additionally, the precipitate follows the curves of the domain, as is expected. This would be a very difficult problem to simulate using a method that required explicit knowledge of where the precipitate would appear. 71

5.1 Sinkhole schematic with boundary and initial conditions. Not to scale. 73

5.2 Sinkhole mesh after finite element discretization. 74

5.3 Zoomed-in view of initial condition for particle distribution. Successively larger chunks, radially from the center, will be removed. This will simulate sinkholes that have varying degrees of “structural integrity”. 75

5.4 Varying two sinkhole parameters. Vertical axis signifies number of particles who have passed through the hole at $y = 0$, and horizontal axis is time. (a) Varying v_{\max} leads to interesting differences in dynamics while (b) varying r_{cut} does not. 76

5.5 Sinkhole formation for $r_{\text{cut}} = 0.16$ and $v_{\max} = 0.3$. Particle and flow profiles are shown for (a) $t = 0$, (b) $t = 1.7$, (c) $t = 3.3$ and (d) $t = 5$. All particles fall simultaneously, contrary to the hypothesis that collapse will start with individual particles at the hole falling before a collapse occurs. 77

6.1 Images demonstrating bi-directional, complex membrane growth from microfluidic experiments. Reproduced from [109] with permission. 81

B.1 Example of finite elements (black solid) whose indices will be associated with the cell (red dash). An element is associated with a cell when at least one of its nodes lies within the cell. Elements can be associated with more than one cell. 85

C.1 Logo for CircGeometry.jl. 88

C.2 Examples of porous geometries that can be generated using CircGeometry.jl. The software works by filling a (a) circle, (b) rectangle, or (c) arbitrary polygon with circles of random radii so that a given volume fraction is achieved. 89

LIST OF SYMBOLS

The following short list of symbols are used throughout the document. The symbols represent quantities that I use consistently.

θ_i	volume fraction of multiphase component i
ψ_i	molar concentration of aqueous chemical species i
ϕ_i	mass concentration of aqueous chemical species i
M_i	molar mass of aqueous chemical species i
$\mathcal{M}_i(\Omega)$	total mass of substance i in domain Ω
\mathbf{q}_i	Darcy velocity of multiphase component i
\mathbf{v}_i	Darcy velocity of multiphase component i
p	common pressure shared among phases
η_i	viscosity of multiphase component i
\square_s	subscript s denotes “solvent”, e.g. water; the fluid phase
\square_p	subscript p denotes “precipitate” membrane or “porous” media; the solid phase

ABSTRACT

Chemical processes within flows are ubiquitous. There exists an important class of reactions that result in a phase change from liquid to solid: precipitation reactions. Inspired by recent microfluidic experiments, this dissertation develops a unified mathematical framework for handling such reactions occurring within a slow-moving fluid flow. A key challenge for precipitate reactions is that, in general, the location of the developed solid is unknown *a priori*. To model this situation, we use a multiphase framework with fluid and solid phases; the aqueous chemicals exist as scalar fields that react within the fluid to induce phase change. We conduct several analytic and numerical validation studies to verify that the model exhibits desired fluid–structure behaviors without requiring interface boundary conditions. To demonstrate the functionality of this framework, we conduct full-scale simulations of two scientific applications: a microfluidic reaction experiment as well as a geophysical study of sinkhole formation. The results of this dissertation are (1) a rigorous derivation of a model framework that conserves mass and incorporates fluid-structure interaction, and (2) numerical methods for solving the full PDE system, which have been implemented in open-source software. The framework can be applied to precipitate reactions where the precipitate greatly affects the surrounding flow, a situation appearing in many laboratory and geophysical contexts including the hydrothermal vent theory for the origin of life. More generally, this model can be used to address low Reynolds number fluid–structure interaction problems that feature the dynamic generation of structures.

Part I

Framework Derivation and Validation

CHAPTER 1

MOTIVATION

Before one reads a dissertation on any topic, it is helpful to first address why one should bother. This chapter answers this question of motivation, to which the answer is twofold: (1) the ability to numerically simulate precipitate reactions would be useful to the costly process of running microfluidic experiments, and (2) the framework developed in the following chapters is widely applicable to modeling flows with complex, dynamically-formed structures. Even more than the resulting system of equations, we hope that the reasoning throughout the derivations offers a rational explanation for why certain modeling choices are made. All multiphase models, to some degree, must rely on fundamental assumptions – none are “true physics”, but rather useful approximations. It is necessary to be transparent about parameter regimes where a model should – and should not – be applied. For example, are there known parameter regimes where a model fails to capture real-world behavior? This thesis derives, analyzes, and simulates a novel framework primarily aimed at modeling a precipitate reaction inside a microfluidic device, but can also be applied to a wider range of scientific phenomena.

Precipitate reactions result in the growth of nano-crystals in complex patterns. We can model how these patterns grow and form, e.g. how one end is pointy but another other side is smooth, etc. However, when one looks at the big picture in the presence of fluid motion, the result of the collective growth of all crystals is that the precipitate impedes flow. Even though crystal growth “should”, for full resolution, be described on the nano-scale, our ultimate interest is on the scale of millimeters or more. Multi-scale physics are not unique to this problem, and in fact have gained considerable interest over the past several years, including in areas of mechanics, chemistry [112, 50], and biological processes [42]. To emphasize, the goal is not to accurately model each and every crystal growth. It is not even to try to model the probabilistic nature of this or that crystal growing at a particular point in space. Instead, the goal is a system of equations that describe the chemical reaction process “accurately enough”, and the macro response of the fluid to the dynamically generating precipitate “accurately enough”. It is the inherent tension between these two objectives that makes developing a model nontrivial, and the resulting novel framework applicable to a wide class of problems.

An analogy to our goal is the objective of modeling fluid flow through porous media. Porous media, resolved in all its gory detail, involves an intricate, inter-connected web of material which allows fluid to zig-zag through voids. However, the porous media literature often provides a “homogenization” and simply treats the material as having a single scalar (or possibly tensor) quantity to describe how fluid flows within the medium. In this way, one loses the inherent complexity of the porous material, and additionally makes all porous materials – every beach on every ocean – exactly the same. But that is okay, because in the end, hydrogeologists do not really care about the exact motion of the fluid dripping through layers of rock, sand and clay as rainwater collects. What they care about is the *average* amount of fluid that passes through the grains; for example, the rate of water collected into an underground reservoir. The results provided by porous media experiments are well approximated by these continuum models which demonstrates the practical success of this approach.

This chapter continues in Sec. 1.1 with a description of the experiment that motivated this dissertation. Sec. 1.2 describes currently available mathematical methods for modeling problems similar to our own, as well as explanations for why these methods are inadequate for our particular goals. Sec. 1.3 provides a brief explanation for what our new model will accomplish and how this is done, as well as a bird’s-eye view of what one will learn by reading this dissertation.

1.1 A microfluidic experiment

One hypothesis for the “origin of life” is that the first biomolecules were formed in undersea hydrothermal vents. In this theory, passive, anisotropic diffusion across a precipitate membrane supports the transmembrane gradients necessary for the formation of the first biochemical molecules [67]. An experimental approach to study this theory examines simpler systems in microfluidic reactors which allows for the controlled study of the prebiotic chemistry in hydrothermal vent chimneys (see Fig. 1.1) [7, 67].

Microfluidic devices have become an important tool in modern chemistry and biomedical analytics [78]. One application is the “lab on a chip”, i.e. the miniaturization of chemical separation and analysis procedures onto a disposable device as small as a few square centimeters. The devices are typically made from etched glass or lithographically-processed elastomers and the fluid is usually controlled mechanically by external pumps or electrically via electro-osmotic flows [66].

Recent studies have used microfluidic methods to form inorganic membranes within Y-shaped devices [9, 111]. The membranes result from chemical reactions between two different solutions that are injected separately but later merge in a long reaction channel that brings the reactants




Figure 1.1: Inactive “Black smoker” hydrothermal vent. Vents like this are integral to one theory for the “origin of life” [67]. Reproduced from [57] with permission.

into direct contact. This merging is usually performed under low Reynolds number (Re) conditions and for miscible liquids, such as aqueous solutions of $NaOH$, and a dissolved metal salt, such as $NiCl_2$. At the reaction interface between these liquids, a precipitate, such as $Ni(OH)_2$, swiftly forms a thin porous membrane (see Fig. 1.2). This precipitate reaction typically involves the formation of microscopically small colloid particles and either their aggregation or addition to a membrane. This phenomenon is related to so-called “chemical gardens” which consist of thin cylindrical precipitate membranes separating a metal salt solution from silicate or hydroxide solutions [88, 65].

Ding et al. showed experimentally that membrane thickness increases with the square root of time, indicating diffusion-controlled growth [31]. The membrane thickening occurs only in the direction of the metal salt solution (e.g. $NiCl_2$) and not in the direction of the anionic precipitation partner (e.g. OH^-), indicating that the membrane is more permeable to anions than cations. This phenomenon has been qualitatively explained by the charged nature of the membrane that suppresses the transmembrane transport of the cation [9, 111].

The modeling challenges presented by this experiment involve a confluence of topics that have been studied before, namely ionic reactions [94, 117], precipitation [120, 1], passive diffusion through a membrane [73, 22, 31], and fluid-structure interaction [107, 20, 87, 100, 99, 70, 84]. The particular combination of these aspects provides the opportunity for a new model that captures them all. Before explaining the approach to this new model, we first detail why currently available methods are insufficient for our goals.

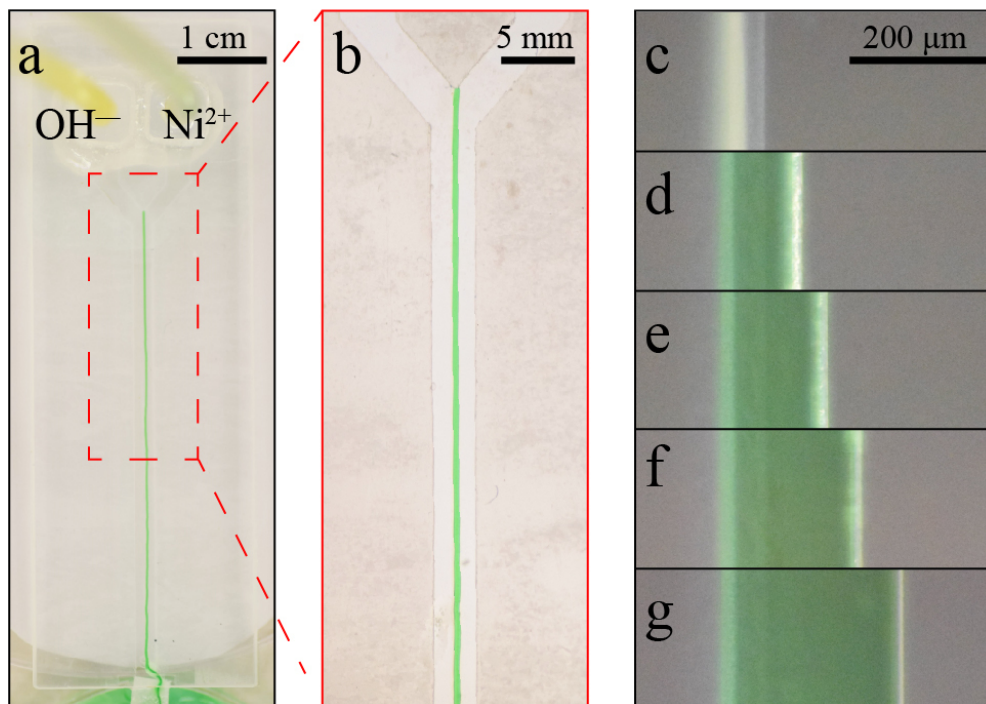



Figure 1.2: Inorganic precipitate membrane formed in a microfluidic channel. (a) The photograph of the resulting membrane after 2 h with 0.5 M NaOH and 0.5 M NiCl₂ solutions being injected simultaneously into the microfluidic device. The mixing part of the channel is 50 mm long, 2 mm wide, and approximately 130 μm high. (b) A magnified view of a selected area from (a). (c-g) A sequence of micrographs showing the unidirectional thickening process after (c) 1 min, (d) 15 min, (e) 30 min, (f) 60 min, and (g) 120 min. Scale bars correspond to (a) 1 cm, (b) 5 mm, and (c) 200 μm. Reproduced from [36] with permission.

1.2 Existing mathematical methods

We now review existing models and describe how they are insufficient for our purposes. It must be emphasized that these models have all proven extremely successful in the areas to which they have been applied. The new modeling challenges, motivated by the experiment detailed above, pose a confluence of objectives such that no existing model can be directly taken and applied. The current review will consist of two parts: chemical reactions and fluid-structure interactions. These two fields have developed independently, which is one reason why a new model is useful – it provides a unifying framework.

Equations for chemical reactions rely on knowledge of the “stoichiometry”, which is an equation for how reactants combine to form products. For example,



represents the net ionic equation for the reaction, visualized in Fig. 1.3, of copper sulfate (CuSO_4) with sodium hydroxide (NaOH) to form the precipitate copper hydroxide ($\text{Cu}(\text{OH})_2$); the sulfate and sodium ions are excluded for clarity, as they do not form a precipitate. The (aq) parenthetical indicates that the ions on the left hand side are part of an aqueous solution, e.g. dissolved in water, and (s) indicates the right hand side is a solid, i.e. the precipitate.

The advantage of stoichiometric equations is that there is no context of the environment, no mention of solubility, no indication as to whether the reaction occurs deep underground or in the far reaches of outer space, and so on. These equations are incredibly useful because they boil down the interactions of trillions and trillions of individual molecules, interacting in space and time in an assortment of complex geometric configurations, into a single line which accurately reflects the *mass transfer* from reactants to products. What Eq. (1.1) demonstrates is that exactly one Cu^{2+} ion reacts with exactly two OH^- ions to form exactly one $\text{Cu}(\text{OH})_2$ molecule, always and forever. One key detail missing in stoichiometry is the *rate* at which the reaction proceeds.

The simplest chemical reaction is given by “ n^{th} -order kinetics” which models the reaction of n species. These reaction equations use the law of mass action, which says that the rate of a reaction is proportional to the product of all the reactant’s concentrations [11]. Therefore, an ordinary differential equation (ODE) modeling the precipitation reaction reaction in Eq. (1.1) could be written as

$$\frac{d[\text{Cu}(\text{OH})_2]}{dt} = k [\text{Cu}^{2+}]^\alpha [\text{OH}^-]^\beta \quad (1.2)$$

where $[\cdot]$ means the concentration of the chemical species, k is the reaction rate constant and α and β determine the order of the reaction. In general, k , α , and β must be determined experimentally and, in particular, α and β *cannot* simply be taken from the stoichiometry (e.g. $\alpha = 1$, $\beta = 2$). Assuming that k , α , and β are obtained from running costly experiments, the above equation is a good model for obtaining the time-dynamics of $[\text{Cu}(\text{OH})_2]$. However, there is one subtle point



Figure 1.3: Formation of $\text{Cu}(\text{OH})_2$ precipitate in a test tube. Photograph by José Furtado, distributed under a CC BY-SA 4.0 license [40].

that needs to be addressed; what is meant by “concentration”? There are different ways to define concentration, all of which are equivalent in the case of single-phase chemistry, but differ when phase-change occurs. This dissertation will demonstrate a conclusive way to incorporate classical reaction kinetics (like the equation described above) into precipitate reactions occurring within a fluid flow while preserving important physical principles such as conservation of mass.

The second component of our model is the fluid-structure interaction (FSI). The models, techniques and results of FSI have been successfully applied to many applications from vibration of aircraft wings to pumping through human vasculature [32]. FSIs occur in multi-physics problems where separate components of the domain have different mechanical properties. “Fluid” flow causes deformation of a boundary “structure”, which in turn affects the fluid flow, and so on. For example, a researcher might be interested in the mechanical stress occurring on a blood vessel wall as varying amounts of blood are pumped (see Fig. 1.4); inside the vessel is viscous liquid, governed by the Navier-Stokes equations, and the vessel itself is elastic which is governed by its own, separate, system of equations. The efficient and accurate coupling of these two equations comprises the main difficulty in FSI problems. There have been several numerical approaches developed to address FSI problems [51], which we will briefly review.




Figure 1.4: Example of FSI application: simulation of blood flow through artery using (i–iii) IBM and (iv) OpenFOAM. Reproduced from [46] with permission, with more algorithmic details found in [58].

A key difficulty in the numerical solutions of FSI problems is that the different governing equations partition the domain, such that they are separated by an interface. This interface moves dynamically, and therefore one approach to FSI focuses on the interface-tracking problem [101, 106]. The obvious issue with using an interface-tracking method for the present precipitation problem is: where does one define the interface of a dynamically generated structure? Precipitate will be generated continuously, according to solutions to differential equations, so there is a need to avoid any method that requires the singular creation of an interface from nothing.

An alternative approach to FSI is called the “immersed boundary method” (IBM), in which a single domain partitioning is used that is non-conforming

with the fluid-structure interface. The “trick” employed by the IBM is to treat the effect of the structure on the fluid as a force in the fluid governing equations. This allows one to avoid remeshing, but creates its own issues related to the accurate approximation of a “structure” that lies off the fluid grid; one must always interpolate the structure onto the fluid mesh, which necessarily introduces its own errors. Furthermore, there is still the same problem of how to define the forces from a solid that is being continuously generated via a reaction, e.g. what is the effect of a solid that is only 20% developed?.

Finally, one lesser-known technique, developed in the late 1990s, is similar to the IBM but instead of representing the structure as a force within the fluid momentum equation, it is represented by a scalar field [5]. This technique is reminiscent of representing the porous media in Darcy’s equation by a scalar (or tensor) “permeability” field. The model introduced in this dissertation is similar to this technique, with the additional property of a “fuzzy” interface between fluid and structure and a mechanism for the the dynamic generation of the underlying structure. Additionally, an independent derivation of the underlying governing system from a classical multiphase framework is provided.

To summarize, a key assumption of these FSI techniques is that the existence of the structure in question is binary – it either does or does not exist. There is no mechanism to smoothly introduce a structure, which is exactly what happens in a precipitate reaction: in the beginning, there is only fluid and no structure. As the chemicals react, a structure appears and the fluid responds to the presence of the structure. Of course, one could simply ignore the fact that at one moment the structure is not there, and at the next the structure is, but that also calls into question how one maintains conservation of mass – the mass of the precipitate structure is necessarily derived from the aqueous chemical species that react inside the liquid phase. The framework derived in this dissertation allows the seamless coupling of phase-change chemistry with fluid-structure interaction. The next section describes the goals of this new model and how they follow the philisophy of historical literature while constructing a new framework that suits our purpose.

1.3 What a new framework achieves

The multiphase framework developed here builds on previous ones [23, 120, 61], but with some keys differences that are guided by a combination of physical principles, model simplicity, and the microfluidic experiments described in Sec. 1.1. First, our formulation conserves the total mass of the components — solvent, dissolved species, and precipitate — throughout the reaction. In

particular, the model accounts for changes in solute concentrations that result from the formation of new precipitate and the associated exclusion of solvent volume. This effect is essential for overall mass conservation, but is neglected in previous models that treat reaction chemicals as scalar fields distinct from the multiphase material. The treatment of reaction chemicals as additional components of a multiphase material, i.e. with their own volume fractions, has been successfully modeled by many [79, 117] but greatly complicates the analysis, interpretation, and simulation of the governing equations. Second, by making certain choices in the averaging procedure for the multicomponent stress, our formulation becomes equivalent to an incompressible Brinkman system.

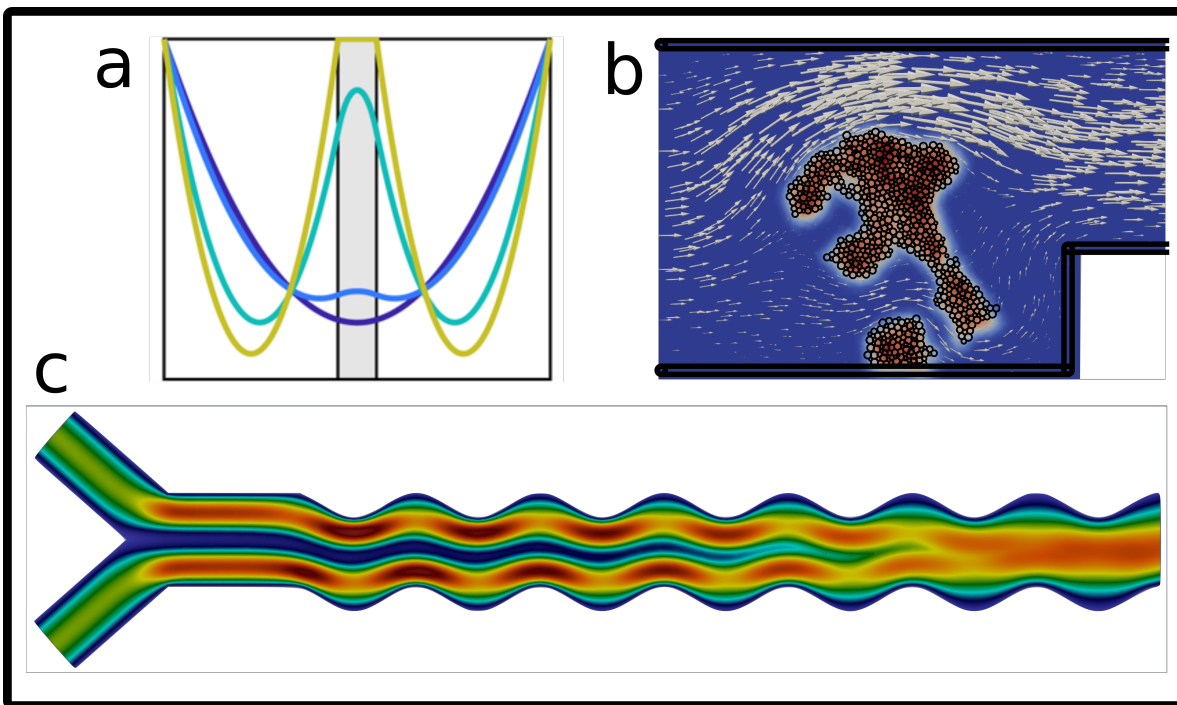


Figure 1.5: Examples from analysis and simulations conducted in this thesis: (a) Velocity profile of 1D reduced model gives no-slip behavior in reaction zone (grey) in Sec. 3.1; (b) Discrete particles implemented for moving porous media give expected porous behavior in Sec. 3.3; (c) Velocity magnitude of full simulations in wavy microfluidic domain, where dark red is maximum velocity and dark blue is zero velocity in Chapter 4.

This equivalence is important for two reasons. First, it guarantees that the model reduces to the Stokes equations in the fluid limit and to Darcy’s equation in the porous-medium limit [15, 35, 49]. In particular, it guarantees that when the precipitate is fully developed, the interface behaves as an impermeable surface without requiring any interface boundary conditions [10, 91]. As demonstrated in Section 3.1.2, many existing multiphase models fail to exhibit this behavior [14, 25, 23, 96], as they were developed primarily for highly permeable systems. Second, the equivalence to Brinkman

significantly simplifies the overall structure of the partial differential equation (PDE) system by eliminating certain cross-terms in the stress divergence that arise in other models.

This dissertation is divided into two parts. Part 1 introduces a model framework that unifies reaction chemistry, which allows a structure to be dynamically generated, and fluid-structure interaction effects, where a structure *smoothly* affects the surrounding fluid flow in a manner that approximates classical fluid-structure interaction problems. Specifically, Chapter 2 develops the model framework, both in a simple case designed to capture the microfluidic experiment of Sec. 1.1 while also providing an extension in order to capture more complex phenomenon, such as geophysical and biological dynamics. The model framework is validated in Chapter 3 using both theoretical analysis and numerical benchmarks – see Figs. 1.5(a,b).

Part 2 showcases the model framework for two applications and ends on a discussion. Chapter 4 focuses on the motivating microfluidic experiment – see Fig. 1.5(c) – while Chapter 5 simulates sinkhole formation. Finally, Chapter 6 summarizes the dissertation and discusses the implications of the framework to future problems. The modeling framework presented in this dissertation, coupled with freely available open-source packages for implementing simulations, will be of use to the wider scientific community.

CHAPTER 2

THE MULTIPHASE FRAMEWORK

We require a framework that accurately describes several aspects of a microfluidic experiment, including the advective transport of ionic species, their reaction to form a product, the precipitation of the product out of solution, and finally the response of the bulk fluid motion to the dynamically-generated precipitate membrane. Advection-diffusion-reaction (ADR) equations are derived for the aqueous chemical concentrations, while the fluid and membrane dynamics are described by multiphase mass and momentum balance equations. In many multiphase models, either constituent can be viscous, viscoelastic, or otherwise. Since the membrane adheres to the substrate, we assume the membrane is an immobile solid, leading to considerable simplifications. Later, a method will be presented to handle solids with complex rheology and the ability to move throughout a domain.

We assume that aqueous reactants and products contribute mass, but not volume, to the fluid phase. The solvent and precipitate each have their own distinct, constant mass densities, and any arbitrary control volume can be divided into solvent and precipitate volume fractions. The formation of new membrane involves the precipitation of product out of solution and the sequestration of solvent. A key modeling assumption is that the volume of fluid sequestered equals the volume of the resulting membrane. As we will show, this assumption ensures incompressibility of the phase-averaged velocity field, i.e. the so-called Darcy velocity.

The derivation of the model framework proceeds as follows. Sec. 2.1 derives evolution equations for the reaction of aqueous ionic species and mass balances for all chemical species as well as fluid and solid phases. Sec. 2.2 describes the momentum equation for the fluid and solid phases, including a discrete description for complex, mobile solids. Sec. 2.3 details a scheme to switch from discrete to continuum descriptions of the solid phase via invertible operators. Regardless of the specific choices, using the following framework will result in a closed, coupled PDE system governing the chemistry and physics of the system, where total mass is conserved throughout aqueous reactions and phase transitions. Sec. 2.4 summarizes the model framework for the interested reader.

2.1 Volumeless ions within a multiphase fluid

In this section we derive equations for the chemical reactions. We follow the “nucleation and growth” model of precipitation [69] and separate the reaction into two sequential parts: in the first, reactants come together to form an aqueous product, and the second describes the aggregation of this aqueous product into a solid precipitate. While many aqueous chemical reactions do not alter the solution volume significantly, the formation of a precipitate excludes fluid volume and therefore can alter the local concentration of the dissolved species. Accordingly, our model neglects the volume occupied by the aqueous species but does account for changes in species concentration that are due to the precipitated solid excluding fluid volume. This effect introduces additional terms in the aqueous reaction equations that are required for mass conservation. To our knowledge these additional terms are not accounted for in the literature that treats aqueous chemicals as scalar fields distinct from the multiphase material.

2.1.1 Basic reaction



We take the pedagogical approach of first deriving our model for the simple case of two reactants combining to form a single product. In later sections, the possible reactions that the framework can handle are expanded, although many of the technical details follow straightforwardly from this basic case. The aqueous reaction is written as a generic net ionic equation



where $A(\text{aq})$, $B(\text{aq})$, and $C(\text{aq})$ are chemicals in the aqueous phase and a , b , and c are their respective stoichiometric coefficients; the precipitation reaction is written simply as



As a concrete example consider the reaction described in Sec. 1.1,



Then $A = \text{Ni}^{2+}$, $B = \text{OH}^{-}$ and $C = \text{Ni}(\text{OH})_2$ and $a = c = 1$, $b = 2$.

The aqueous chemicals will be measured with a variable for the *number* of chemicals per unit *solvent* volume, i.e. molarity, which we will call ψ_i for chemical species $i \in \{A, B, C\}$. Reaction rates depend on a reactants’ molarity, and molarity can change due to two independent factors: either

the number of molecules changes due to the aqueous reaction, or the solvent volume changes due to precipitation. Because either one can occur in a precipitation reaction, these two competing effects must be carefully considered when formulating the reaction equations.

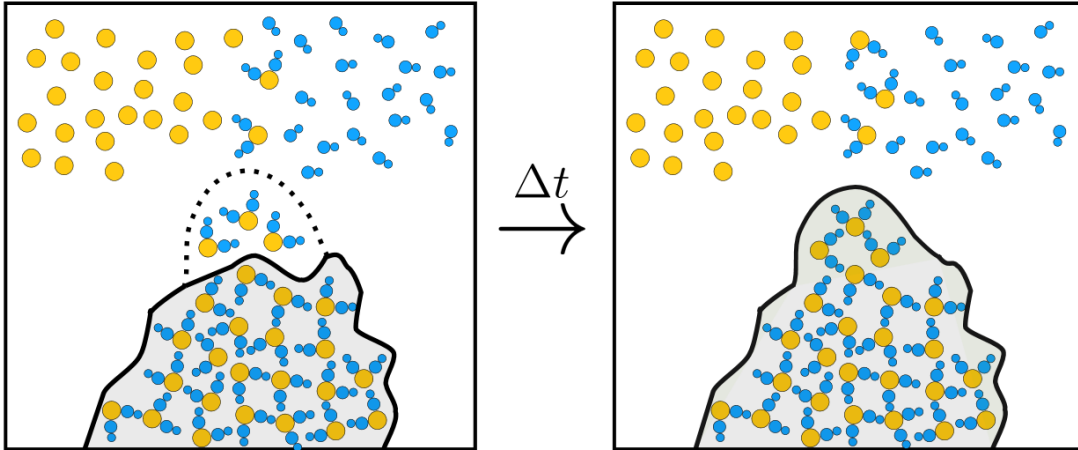


Figure 2.1: Schematic of precipitate reaction in a control volume. Precipitation causes solution (white) to transform into solid (shaded) after a certain concentration threshold is reached of aqueous product C . Aqueous chemicals A , B and C are volumeless scalar fields while the solvent and solid are treated as multiphase components. The volume of solid gained is exactly equal to the volume of solvent lost. Reproduced from [36] with permission.

We begin by deriving equations for how the aqueous reaction proceeds in a spatially homogeneous environment; later the effects of advective and diffusive spatial fluxes will be added. Suppose the chemicals exist in some aqueous solution of fixed control volume V_0 . The chemicals undergo both the aqueous and precipitation reactions which results in fluid mass and solvent volume being converted to precipitate mass and volume (see Figure 2.1). The fluid component has mass

$$\mathcal{M}_f = \left(\rho_s + \sum M_i \psi_i \right) \theta_s V_0 \quad (2.4)$$

where ρ_s is the constant solvent mass density (without any reactants or products present), θ_s is the solvent volume fraction, and M_i is the molar mass of chemical species i . The summation represents the contribution of the chemical species to fluid mass, so that the *fluid* mass density is not constant. The precipitate component has mass $\mathcal{M}_p = \rho_p \theta_p V_0$ where ρ_p is the constant precipitate mass density and θ_p is the precipitate volume fraction. Physically, precipitate mass is composed of both precipitated chemical C and sequestered solvent mass.

The change in ψ_i *purely due to aqueous reaction*, i.e. no precipitation, can be modeled as a second-order kinetics reaction

$$\dot{\psi}_A^{(\text{aq})} = -ar\psi_A\psi_B, \quad \dot{\psi}_B^{(\text{aq})} = -br\psi_A\psi_B, \quad \dot{\psi}_C^{(\text{aq})} = cr\psi_A\psi_B. \quad (2.5)$$

where the dot indicates a derivative with respect to time, r is the rate of aqueous reaction per chemical concentration, and the (aq) superscript refers to the fact that Eq. (2.5) models the change in ψ_i *purely due to aqueous reaction*, i.e. Eq. (2.1). More general power laws are sometimes used to model chemical kinetics, but here we use purely second-order kinetics for simplicity [see 17, pp. 573 - 575]. None of the analysis, however, depends specifically on this choice and the results could be carried forward for other kinetics.

To derive equations for the change in ψ_i *purely due to precipitation* we appeal to ideas from continuum mechanics. The concentration of ions A in the control volume is written as $\psi_A = n_A/(\theta_s V_0)$ where n_A is the number of A ions in V_0 . Note that this formulation makes explicit the dependence of ψ_A on both n_A and θ_s . Consider the change in a small increment of time Δt . Then the time-dependent variables are updated so that

$$\psi_A + \Delta\psi_A = \frac{n_A}{(\theta_s + \Delta\theta_s)V_0}. \quad (2.6)$$

Recall that n_A is constant during precipitation as only $C(\text{aq})$ precipitates. Approximating for small $\Delta\theta_s$ and neglecting higher-order terms gives

$$\psi_A + \Delta\psi_A = \frac{n_A}{\theta_s V_0} \left(1 - \frac{\Delta\theta_s}{\theta_s}\right) = \psi_A \left(1 - \frac{\Delta\theta_s}{\theta_s}\right). \quad (2.7)$$

Then, cancelling the ψ_A , dividing both sides by Δt , and letting $\Delta t \rightarrow 0$ gives the change in ψ_A purely due to precipitate reaction as $\dot{\psi}_A^{(\text{p})} = -\psi_A \dot{\theta}_s / \theta_s$. By symmetry, a similar formula holds for $\dot{\psi}_B^{(\text{p})}$. Note that both of these are essentially applications of the product rule for $\partial_t(\psi_i \theta_s) = 0$, which physically means that the total number of ions of $i \in \{A, B\}$ in the control volume does not change in time *due to precipitation*.

A similar procedure can be followed for ψ_C , except now the number of aqueous chemicals n_C changes as $C(\text{aq})$ precipitates,

$$\psi_C + \Delta\psi_C = \frac{n_C + \Delta n_C}{(\theta_s + \Delta\theta_s)V_0}. \quad (2.8)$$

Above, both n_C and θ_s change in time. Expanding both expressions while linearizing for small $\Delta\theta_s$, dividing by Δt , and taking the limit as $\Delta t \rightarrow 0$ one obtains $\dot{\psi}_C^{(\text{p})} = -\psi_C \dot{\theta}_s / \theta_s + \dot{n}_C / \theta_s$. The first term in this expression is analogous to those obtained for reactants A and B , and simply describes

the effect on concentration when solvent volume is changing. The second term, however, is new and describes the effect on ψ_C as aqueous C molecules are converted into precipitate membrane. We write $\dot{n}_C = \alpha \dot{\theta}_s$ where the specific value of α will be found shortly to guarantee conservation of mass throughout the entire reaction. The expressions for the rate of change of aqueous chemical concentrations due to precipitation are thus:

$$\dot{\psi}_A^{(p)} = -\psi_A \dot{\theta}_s / \theta_s, \quad \dot{\psi}_B^{(p)} = -\psi_B \dot{\theta}_s / \theta_s, \quad \dot{\psi}_C^{(p)} = -\psi_C \dot{\theta}_s / \theta_s + \alpha \dot{\theta}_s / \theta_s, \quad (2.9)$$

where the (p) superscript refers to the fact that Eq. (2.9) models the change in ψ_i *purely due to precipitation of the membrane out of solution*, i.e. Eq. (2.2). Assuming that the aqueous and precipitate reactions act independently, $\dot{\psi}_i = \dot{\psi}_i^{(aq)} + \dot{\psi}_i^{(p)}$, gives

$$\dot{\psi}_A = -ar\psi_A\psi_B - \psi_A \dot{\theta}_s / \theta_s \quad (2.10a)$$

$$\dot{\psi}_B = -br\psi_A\psi_B - \psi_B \dot{\theta}_s / \theta_s \quad (2.10b)$$

$$\dot{\psi}_C = cr\psi_A\psi_B - \psi_C \dot{\theta}_s / \theta_s + \alpha \dot{\theta}_s / \theta_s \quad (2.10c)$$

These equations describe the dynamics of aqueous species molarity in the absence of spatial fluxes.

To obtain the value of α that guarantees conservation of mass, we again apply a continuum mechanics argument. The change in precipitate mass after a small time step is $\Delta \mathcal{M}_p = \rho_p \Delta \theta_p V_0$. To simplify the expression for change in fluid mass, we expand $\Delta \mathcal{M}_f$ while neglecting second order terms to get

$$\Delta \mathcal{M}_f = V_0 \rho_s \Delta \theta_s + V_0 \theta_s \sum_i M_i \Delta \psi_i + V_0 \Delta \theta_s \sum_i M_i \psi_i. \quad (2.11)$$

Replacing the $\Delta \psi_i$ with their respective differential terms in Eqs. (2.10) and performing some algebraic manipulation produces

$$\Delta \mathcal{M}_f = V_0 \rho_s \Delta \theta_s + V_0 \theta_s r \psi_A \psi_B (cM_C - aM_A - bM_B) + V_0 M_C \Delta n_C. \quad (2.12)$$

Conservation of mass during the aqueous reaction (2.1) implies

$$aM_A + bM_B = cM_C. \quad (2.13)$$

Thus, the term in parenthesis in (2.12) vanishes. Meanwhile, conservation of mass of the entire system implies $\Delta \mathcal{M}_p = -\Delta \mathcal{M}_f$, i.e. the mass lost by the fluid equals the mass gained by precipitate. Additionally, the assumption that fluid volume is converted perfectly to precipitate volume implies $\Delta \theta_p = -\Delta \theta_s$. Using the respective definitions of M_i and solving for Δn_C gives

$$\Delta n_C = \left(\frac{\rho_p - \rho_s}{M_C} \right) \Delta \theta_s. \quad (2.14)$$

Dividing by Δt and taking the limit $\Delta t \rightarrow 0$ gives $\dot{n}_C = \alpha \dot{\theta}_s$ where

$$\alpha = \frac{\rho_p - \rho_s}{M_C}. \quad (2.15)$$

Physically, this value of α corresponds to the concentration of $C(\text{aq})$ that must leave the fluid phase during precipitation in order for mass to be conserved. This completes our description of temporal dynamics; now we introduce spatial fluxes.

In experiments, the aqueous reaction occurs within the flow of a microfluidic device and therefore spatial fluxes must be considered. To describe these fluxes, consider the general conservation law for the *chemical mass per unit control volume* $\phi = M_i \psi_i \theta_s$,

$$\frac{\partial \phi}{\partial t} + \nabla \cdot \mathbf{J} = \Gamma \quad (2.16)$$

where \mathbf{J} is the flux of ϕ and Γ is a transfer term for the rate that ϕ enters the system. We choose \mathbf{J} to account for advection and diffusion of the chemical concentrations,

$$\frac{\partial(M_A \psi_A \theta_s)}{\partial t} + \nabla \cdot (M_A \psi_A \theta_s \mathbf{v}_s - \kappa_A \nabla(M_A \psi_A)) = \Gamma_A \quad (2.17a)$$

$$\frac{\partial(M_B \psi_B \theta_s)}{\partial t} + \nabla \cdot (M_B \psi_B \theta_s \mathbf{v}_s - \kappa_B \nabla(M_B \psi_B)) = \Gamma_B \quad (2.17b)$$

$$\frac{\partial(M_C \psi_C \theta_s)}{\partial t} + \nabla \cdot (M_C \psi_C \theta_s \mathbf{v}_s - \kappa_C \nabla(M_C \psi_C)) = \Gamma_C \quad (2.17c)$$

where \mathbf{v}_s is the (tracer) velocity of the solvent and κ_i are diffusion coefficients which possibly depend on the solvent volume fraction. Note that the diffusive flux used above transports mass according to gradients in *molarity* ψ_i , not gradients in ϕ_i . This choice produces the physically realistic steady state of uniform molarity in a quiescent, non-reacting fluid that has inhomogeneous volume fraction.

Assuming that the reactions and spatial fluxes act independently, the Γ_i correspond to the rates given in Eqs. (2.10). Rearranging and multiplying each equation by its respective molar mass M_i gives

$$\Gamma_A = -ar M_A \theta_s \psi_A \psi_B \quad (2.18a)$$

$$\Gamma_B = -br M_B \theta_s \psi_A \psi_B \quad (2.18b)$$

$$\Gamma_C = cr M_C \theta_s \psi_A \psi_B + \alpha M_C \dot{\theta}_s \quad (2.18c)$$

where $\alpha = (\rho_p - \rho_s)/M_C$. Now that mass balance equations for the chemistry are established, mass balance equations for the multiphase solvent-precipitate system are needed.

A simple but necessary assumption is that our volume is occupied by only solvent and membrane, i.e. there are no “voids”. This no-void assumption implies

$$\theta_s + \theta_p = 1 \quad (2.19)$$

everywhere. Mass balances for the solvent and membrane phases provide

$$\frac{\partial(\rho_s\theta_s)}{\partial t} + \nabla \cdot (\rho_s\theta_s\mathbf{v}_s) = R_s \quad (2.20a)$$

$$\frac{\partial(\rho_p\theta_p)}{\partial t} = R_p \quad (2.20b)$$

where R_i denotes the rate of mass added to phase i . Eq. (2.20b) has no advective term since the precipitate membrane is assumed to be immobile.

To ensure conservation of total mass, the rates R_s and R_p must be related. To derive this relationship, let V_0 be an arbitrary control volume. The total mass (of all components) in V_0 is

$$\mathcal{M}(V_0) = \int_{V_0} \rho_s\theta_s + \rho_p\theta_p + \sum M_i\psi_i\theta_s \, dV. \quad (2.21)$$

Summing the five mass balance equations, (2.17a)–(2.17c) and (2.20a)–(2.20b), integrating over V_0 , and applying the divergence theorem gives

$$\frac{d}{dt}\mathcal{M}(V_0) + \int_{\partial V_0} \underbrace{(\rho_s\theta_s\mathbf{v}_s + \sum \mathbf{J}_i) \cdot \hat{\mathbf{n}}}_{\text{boundary flux}} \, dS = \int_{V_0} \underbrace{R_s + R_m + \sum \Gamma_i}_{\text{transfer \& reaction}} \, dV. \quad (2.22)$$

where $\hat{\mathbf{n}}$ is the outward unit normal vector. Summing Eqs. (2.18) and applying (2.13) gives $\sum \Gamma_i = \alpha M_C \dot{\theta}_s$. For the sake of obtaining a relationship between R_s and R_p , briefly consider the case of zero boundary flux. Then, to conserve total mass within any control volume, we must have $R_s + R_p + \alpha M_C \dot{\theta}_s = 0$ holding point-wise. Substituting the value of α in Eq. (2.15), using the no-void assumption (2.19) and the definition of R_m in (2.20b) gives

$$R_s = -\frac{\rho_s}{\rho_p} R_p. \quad (2.23)$$

This relationship is required for overall mass conservation.

Now consider the so-called Darcy velocity field $\mathbf{q}_s = \theta_s\mathbf{v}_s$. Substituting the value of R_s from Eq. (2.23) into Eq. (2.20a), using a consequence of the no-void assumption ($\dot{\theta}_s = -\dot{\theta}_p$) and substituting R_p by its value in Eq. (2.20a) implies that the fluid Darcy velocity is divergence-free:

$$\nabla \cdot \mathbf{q}_s = 0 \quad (2.24)$$

and therefore the multiphase material is incompressible. Now return to Eq.(2.22) and consider the entire domain Ω with total mass $\mathcal{M} = \mathcal{M}(\Omega)$. From the above argument, the right-hand-side of this equation vanishes as a necessary condition on mass conservation. Then applying incompressibility via Eq. (2.24), gives the total mass balance:

$$\frac{d\mathcal{M}}{dt} = - \int_{\partial\Omega} \sum M_i (\mathbf{q}_s \psi_i - \kappa_i \nabla \psi_i) \cdot \hat{\mathbf{n}} \, dS. \quad (2.25)$$

As expressed in this equation, the total mass of the system is conserved as long as the chemical flux at the boundary vanishes. More generally, the total mass of the system can change according to how much chemical mass is being injected or removed via the boundary flux terms. This completes the proof of conservation of mass for our system.

We now specify our choice for the form for the precipitation term R_p . Although complicated models of precipitation exist [69, 81], we employ a simple model in which the rate of membrane mass growth is proportional to the amount of product, provided that the product concentration exceeds some precipitation threshold, i.e.

$$R_p = \beta \psi_C \theta_s \mathcal{H}(\psi_C - \psi_C^*) \quad (2.26)$$

where β is a rate constant, \mathcal{H} is the standard Heaviside function and ψ_C^* is the concentration threshold for precipitation to occur.


2.1.2 Complex reactions

Chemical reactions in the real world exist as single components of complex reaction systems; reactants themselves come from previous reactions and products go onto form their own reactions. Notably this is true in biochemical reactions, where hundreds of cascading reactions form a system necessary for an organism's survival. In this section we extend the simple reaction system described previously to a more general form that can handle the vast majority of reaction stoichiometries that one can encounter.

The reaction involving arbitrary numbers of reactants and products generalizes easily from the case outlined in the previous section. To be specific, we consider the general aqueous reaction of n reactants into m products



and, of these m products, one will precipitate from the aqueous to solid phase, which is given by



The case of multiple components of the product precipitating will be handled later. The fluid component in a control volume V_0 now has mass equal to

$$\mathcal{M}_f = \left(\rho_s + \sum_{i=1}^n M_{A_i} \psi_{A_i} + \sum_{i=1}^m M_{B_i} \psi_{B_i} \right) \theta_s V_0. \quad (2.29)$$

Analogous to the derivation in the previous section, the change in ψ_i purely due to aqueous reaction can be modeled by the equations

$$\dot{\psi}_{A_j}^{(\text{aq})} = -a_j r \prod_{i=1}^n \psi_{A_i}^{\gamma_i} \quad j \in \{1, \dots, n\} \quad (2.30a)$$

$$\dot{\psi}_{B_j}^{(\text{aq})} = b_j r \prod_{i=1}^n \psi_{A_i}^{\gamma_i} \quad j \in \{1, \dots, m\} \quad (2.30b)$$

where r is still the aqueous reaction rate and γ_i is the reaction order of species i , which will have to be determined experimentally; note that if $\gamma_i \equiv 1$, then the above is an n^{th} -order reaction. To be clear, in order for the above equation to be well-defined, the $n+1$ parameters γ_i and r will need to be determined, likely using costly experiments. The reaction purely due to precipitation is

$$\dot{\psi}_{A_j}^{(\text{p})} = -\psi_{A_j} \dot{\theta}_s / \theta_s, \quad j \in \{1, \dots, n\} \quad (2.31a)$$

$$\dot{\psi}_{B_j}^{(\text{p})} = -\psi_{B_j} \dot{\theta}_s / \theta_s + \alpha \delta_{jk} \dot{\theta}_s / \theta_s, \quad j \in \{1, \dots, m\} \quad (2.31b)$$

where δ_{jk} is the classic Kronecker delta function. Analogous to the previous chapter, α is defined to be

$$\alpha = \frac{\rho_p - \rho_s}{M_{B_k}}. \quad (2.32)$$

We still assume that the aqueous reaction (2.27) and precipitate reaction (2.28) occur independently, such that the total rate of change of ψ_i is simply the sum of these reactions. Due to conservation of mass throughout these reactions, the following relationship holds:

$$\sum_{i=1}^n a_i M_{A_i} = \sum_{i=1}^m b_i M_{B_i} \quad (2.33)$$

where M_j is the molar mass of species j . The system of ADR equations is

$$\frac{\partial(M_i \psi_i \theta_s)}{\partial t} + \nabla \cdot (M_i \psi_i \theta_s - \kappa_i \nabla(M_i \psi_i)) = \Gamma_i, \quad i \in \{A_1, \dots, A_n, B_1, \dots, B_m\} \quad (2.34)$$

where we need to derive all of the mass transfer terms Γ_i . These are given by

$$\Gamma_{A_j} = -a_j r M_{A_j} \theta_s \prod_{i=1}^n \psi_{A_i}^{\gamma_i}, \quad j \in \{1, \dots, n\} \quad (2.35a)$$

$$\Gamma_{B_j} = b_j r M_{B_j} \theta_s \prod_{i=1}^n \psi_{A_i}^{\gamma_i} + \delta_{jk} \alpha M_{B_k} \dot{\theta}_s, \quad j \in \{1, \dots, m\} \quad (2.35b)$$

where an explanation for how these are derived is given in the previous section. We now have extended our model to handle aqueous reactions of arbitrary complexity.

We briefly note that reversible reactions can also fit into our model framework – one would simply need to add reaction rates corresponding to the product concentrations, in accordance with the law of mass action and traditional reversible reaction kinetics. To reverse the precipitation reaction, one simply changes the sign of α , i.e. the mass flows from precipitate to fluid, as opposed to the other way around. We now address the slightly more complicated case of multiple precipitates in the same reaction.

If multiple (aqueous) products undergo precipitation, then one must necessarily also provide more volume fraction variables, i.e. add more phases. Therefore, if K aqueous products form precipitates ($K \leq m$) then we must have in the model $K + 1$ distinct volume fraction variables. The no-void assumption then becomes

$$\theta_s + \sum_k \theta_p^{(k)} = 1, \quad (2.36)$$

where $\theta_p^{(k)}$ denotes the k^{th} precipitate phase. We then need to obtain the reaction equations and, in particular, need to define the analogous rate “ α ” for each precipitate such that total mass is conserved throughout the reaction. To be clear, we wish to extend the results of the above section where we allow *multiple* precipitate membranes to form from the aqueous products:



where \mathcal{I}_K is the set of K indices of the aqueous products that precipitate. Most of the equations of the previous section still hold, so we will not repeat the entire derivation. We only note that Eq. (2.35b) is changed to

$$\Gamma_{B_j} = b_j r M_{B_j} \theta_s \prod_{i=1}^n \psi_{A_i} + \sum_{k \in \mathcal{I}_K} \alpha_k \delta_{jk} M_{B_k} \dot{\theta}_s, \quad j \in \{1, \dots, m\} \quad (2.38)$$

in order to account for the (possibly) different α_k values, which are only defined for $k \in \mathcal{I}_K$. The mass of the precipitate phase k is

$$\mathcal{M}_p^{(k)} = \rho_p^{(k)} \theta_p^{(k)} V_0 \quad (2.39)$$

where $\rho_p^{(k)}$ is the mass density of the k^{th} precipitate. We also have the change in fluid mass as

$$\Delta \mathcal{M}_f = V_0 \rho_s \Delta \theta_s + V_0 \theta_s r \left[\sum_{i=1}^m b_i M_{B_i} - \sum_{i=1}^n a_i \psi_{A_i} \right] \prod_{j=1}^n \psi_{A_j} + V_0 \sum_{k \in \mathcal{I}_K} \alpha_k M_{B_k} \Delta \theta_s. \quad (2.40)$$

Like in the previous section, one can use the stoicheometry of the reaction to set the term in brackets to zero. We are now in a position to derive the α_k terms. Remember that mass is conserved during precipitation, i.e. $\Delta M_f = -\sum_{k \in \mathcal{I}_K} \Delta \mathcal{M}_m^{(k)}$; writing this out using Eqs. (2.39) and (2.40) gives

$$V_0 \rho_s \Delta \theta_s + V_0 \sum_{k \in \mathcal{I}_K} \alpha_k M_{B_k} \Delta \theta_s = - \sum_{k \in \mathcal{I}_K} \rho_p^{(k)} \Delta \theta_p^{(k)} V_0. \quad (2.41)$$

Volume is also conserved but, because there are multiple precipitates, there is a question as to how to divide the volume among the newly formed precipitates.

Cancelling out the V_0 and plugging in conservation of volume one obtains

$$\sum_{k \in \mathcal{I}_K} \alpha_k M_{B_k} = -\rho_s + \sum_{k \in \mathcal{I}_K} \rho_p^{(k)} \quad (2.42)$$

which is a single equation for K unknowns α_k . So how do we determine the specific α_k ? The most straightforward method is to simply select the α_k , based on intuition from the case of a single precipitate, and then use the above equation to verify that this is an admissible solution. Therefore, we let each α_k obey

$$\alpha_k = \frac{K \rho_p^{(k)} - \rho_s}{K M_{B_k}} \quad (2.43)$$

which one can verify satisfies Eq. (2.42). Therefore, while these choices of α_k do conserve mass, they are *not unique*, and other choices could give the same conservation of mass scheme. Additionally, if $K = 1$, the single-precipitate scheme is recovered exactly, as expected.

The new $K + 1$ mass balance equations are then

$$\frac{\partial(\rho_s \theta_s)}{\partial t} + \nabla \cdot (\rho_s \theta_s \mathbf{v}_s) = R_s \quad (2.44a)$$

$$\frac{\partial(\rho_p^{(k)} \theta_p^{(k)})}{\partial t} = R_p^{(k)}, \quad k \in \mathcal{I}_K \quad (2.44b)$$

where we need to find an analogous way to find a relationship between R_s and all $R_p^{(k)}$. By inspection of Eq. (2.23), the mass reaction terms R_i are required to obey

$$R_s = -\rho_s \left(\sum_{k \in \mathcal{I}_K} \rho_p^{(k)} \right)^{-1} \sum_{k \in \mathcal{I}_K} R_p^{(k)}. \quad (2.45)$$

which one can verify does indeed conserve mass in a manner similar to the single-precipitate equations of Sec. 2.1.1. This section has established a complete description of chemistry which can occur in the vast majority of precipitation reactions. The next section pertains to the second half of the framework: the fluid-structure interaction.

2.2 Multiphase momentum equations

A significant advantage of our model is that the chemistry and fluid dynamics are described separately – chemicals exist as volumeless scalar fields which obey traditional reaction kinetics, while the volume fractions of the multiphase material obey mass and momentum transfer equations. Until this point, we have focused exclusively on the chemical reactions and mass movements. In this section, we focus on momentum transfer between the phases.

The nondimensional Reynolds number is defined to be $\text{Re} = \rho UL/\eta$, where ρ is a characteristic mass density, U is a characteristic velocity, L is a characteristic length, and η is a characteristic viscosity. The Reynolds number represents the ratio of inertial to viscous forces in a problem. We are concerned with applications in microfluidic devices (small L) or with slow moving fluid (small U), so it is safe to assume that the Reynolds number is negligible ($\text{Re} \ll 1$). Since inertial effects are assumed negligible, the solvent momentum balance can be written as

$$\nabla \cdot \mathbf{T} - \xi \mathbf{v}_s = \theta_s \nabla p \quad (2.46)$$

where \mathbf{T} is the multiphase stress tensor, p is a so-called common pressure that is shared by the fluid and precipitate phases [23], and ξ is a friction coefficient which depends on volume fraction. Deviating slightly from the predominant multiphase literature, we chose the form of the stress tensor as

$$\mathbf{T} = \eta(\nabla \mathbf{q}_s + \nabla \mathbf{q}_s^T) \quad (2.47)$$

where η is the fluid viscosity. In particular, since $\mathbf{q}_s = \theta_s \mathbf{v}_s$, we have placed the fluid volume fraction θ_s *inside* the gradient, whereas many multiphase models place the θ_s outside of the gradient but inside the divergence [16, 25, 23]. Such a choice must be made for model closure, and neither is fully justified by first principles. We make the choice above to obtain equivalence to the Brinkman system [15] when the membrane is immobile, as discussed below.

2.2.1 Stationary solid and friction

Indeed, applying incompressibility of the Darcy velocity, as derived in Eq. (2.24), to Eq. (2.46) simplifies it to a Brinkman equation with variable coefficients:

$$\eta \nabla^2 \mathbf{q}_s - \frac{\xi}{\theta_s} \mathbf{q}_s = \theta_s \nabla p. \quad (2.48)$$

Note that Eq. (2.48) does not have any cross-derivative terms that would appear if the traditional multiphase stress tensor were used [23, 56]. Because the membrane is assumed immobile ($\mathbf{v}_m \equiv \mathbf{0}$), no momentum equation is needed for it.

We emphasize that our choice of stress tensor in Eq. (2.47), which is slightly unconventional in the multiphase literature, is responsible for producing equivalence to the Brinkman system. The Brinkman equations were originally formulated in 1949 [15] and, since then, have undergone fundamental theoretical development [19] in close cooperation with experimental measurements [77], and, in recent decades, numerical simulation [80, 74]. We will demonstrate in section 3.1.2 that this system produces the expected no-slip behavior on a fully formed membrane. It is interesting that, with the above modification, the modern multiphase averaging framework can be made to produce the Brinkman system, and therefore is consistent with 70 years of research on porous media systems.

The friction coefficient ξ should be chosen in such a way that, at high membrane volume fraction, friction becomes the dominant effect in Eq. (2.48). The choice made here, and mentioned briefly in Leiderman and Fogelson [62], is to use the Kozeny-Carman (KC) formula for permeability as it depends on porosity [34]. In the present notation, the KC relationship gives the friction coefficient as

$$\xi_{KC}(\theta_s) = h \frac{(1 - \theta_s)^2}{\theta_s} \quad (2.49)$$

where h is an arbitrary constant. This friction coefficient will provide the desired no-slip behavior in the precipitate limit $\theta_p \rightarrow 1$. Angot [5] discusses the implications of a similar singular friction term in a Brinkman system, although their model does not include the solvent volume fraction term in front of the pressure gradient and only applies to domains with spatially discontinuous volume fractions; the current framework generalizes this notion by being able to account for smooth spatial and temporal gradients in the volume fraction.

As an alternative to the singular ξ_{KC} , the friction coefficient can be chosen to be a (non-singular) Hill function, as used in Leiderman and Fogelson [61, 62],

$$\xi_H(\theta_s) = h \frac{(1 - \theta_s)^n}{K^n + (1 - \theta_s)^n}. \quad (2.50)$$

The use of Hill functions is largely empirical, although it has significant advantages in that it is finite in the precipitate limit and therefore more numerically stable. Additionally, K determines the half-saturation point and n indicates the qualitative manner at which this saturation is achieved. These parameters allow for fine-tuning to specific experimental observations.

Finally, a friction term employed in many biofilm multiphase models is

$$\xi_B(\theta_s) = h\theta_s(1 - \theta_s). \quad (2.51)$$

This choice has become popular in the literature [16, 24, 25, 23, 96], and is often justified by the idea that friction should vanish if either phase, θ_s or θ_p , is absent. While it is an intuitive notion, we will show in Section 3.1.2 that this friction coefficient does *not* produce the physically realistic behavior of no-slip velocity on fully developed solid surfaces. To produce this behavior, it is necessary that friction *dominates*, not vanishes, in the limit $\theta_p \rightarrow 1$.




Figure 2.2: Comparison of friction terms ξ . ξ_{KC} (solid) is singular in the limit $\theta_s \rightarrow 0$, ξ_H (dash) is non-singular in the porous limit ($K = 0.5$, $n = 2$) and ξ_B (dot) yields maximum friction when both phases are present in equal amounts. All terms have been normalized by choosing h such that $\xi(\theta_s^*) = \xi^*$. Reproduced from [36] with permission.

A visual comparison of these three friction coefficients is shown in Figure 2.2. For the sake of comparison, we have chosen the constant h so that the three curves intersect at a reference porosity θ_s^* , i.e. $\xi(\theta_s^*) = \xi^*$. For $\xi^* = 3$, $\theta_s^* = 0.3$, this condition generates the constants $h_{KC} \approx 1.8$, $h_H \approx 4.5$ ($K = 0.5$, $n = 2$), and $h_B \approx 14.3$. For ξ_{KC} and ξ_H , the value θ_s^* can be loosely interpreted as the percolation threshold, i.e. the critical porosity below which the medium essentially behaves as impermeable to flow [43]; however, we will show in Sec. 3.2 that this analogy should not be taken literally, at least in two spatial dimensions.

We now have a complete multiphase model sufficient to simulate the immobile precipitates in the microfluidic experiments. But what if the membrane is not immobile? This could be important if, instead of a precipitate, we are modeling more complex media such as geophysical or biological flows that contain solids, formed out of reactions, which themselves move and are subject to complicated rheological laws. The multiphase literature is rife with continuum descriptions of complex media, and we are confident that such a continuum extension to our momentum equations is possible. However, in this dissertation, we will take a different approach. In the next section we will discuss

an extension to our model to be able to handle complex media using a discrete, as opposed to continuum, description of their interaction forces.

2.2.2 Complex discrete rheology

While simple fluids are governed by the Navier-Stokes equations – equations with considerable experimental validation – a similar continuum governing equation is not agreed upon for materials with complex rheological properties. Therefore, in this dissertation, we resort to a discrete approximation that will allow us to numerically resolve the velocity field of the porous phase and, more importantly, allow the medium to move. To accomplish this, the porous media volume fraction field θ_p is represented by N discrete point particles. Each particle occupies no space, but the individual forces prescribed on each particle will cause the collection to qualitatively behave like media with complex rheology.

To jump between this particle-description of the porous phase and its field-description – which is necessary in order to interact with our fluid governing equations described in the previous section – we define a simple Gaussian-smoother operator to “smooth out” each particle, along with an inverse operator. The aggregate of these smoothed out particles generates a field. In this manner, we can handle arbitrarily complex “solid” phases, with the only requirement that we can describe individual solid-solid forces – the only solid-fluid interaction is via a Stokes drag. These discrete particle approaches to modeling porous materials is motivated by the “Discrete Element Method” [55], although it is *not* the same as the “Combined Finite-Discrete Element Method” [95, 75], which discretizes a structure using finite elements and, among other things, can account for rotation, which our present model cannot. That being said, our discrete model, as you will see, sufficiently handles macro properties of porous media with complex rheology, and fits nicely within the model framework.

Because the problems under consideration are at a scale where inertial forces are very small compared to viscous forces, Newton’s second law can be replaced by the “Stokes force balance” approximation

$$\mathbf{F}^i = \mathbf{0}, \quad i = 1, \dots, N, \quad (2.52)$$

where \mathbf{F}^i is the sum of forces on particle i . This fundamental assumption means that the model is *only* valid where the media in question moves so slowly as to make inertial effects insignificant. This assumption also means that, from the fluid’s perspective, the structure is moving so slowly such that it is practically immobile – making our fluid continuum description of the previous section

still applicable. The practical effect of this approximation is that the movement of the particles at each moment in time depends only on the *geometry* of the configuration of particles at that instant of time. To repeat: there is no inertia, or memory, in our system.

In particular, the net force on particle i , \mathbf{F}^i , will be decomposed into a sum of the following forces: gravity \mathbf{F}_g^i , particle-particle cohesion \mathbf{F}_c^i , particle-wall adhesion \mathbf{F}_a^i , and seepage drag \mathbf{F}_d^i such that

$$\mathbf{F}^i = \mathbf{F}_g^i + \mathbf{F}_c^i + \mathbf{F}_a^i + \mathbf{F}_d^i, \quad i = 1, \dots, N. \quad (2.53)$$

Each of the component forces is now described. A description of how these forces are computed in practice, for a large number of particles ($N \gg 1$), is reserved for Chapter 5.

Gravity Force. The gravitation force is the simplest to define; because the particles are immersed in a fluid, the buoyant effect of the density differences is considered:

$$\mathbf{F}_g^i = -(\rho_p - \rho_s)gV_i\hat{\mathbf{e}}_z \quad (2.54)$$

where g is a physical gravitation density constant with units L/T^2 , V_i is the volume of particle i with units L^3 , and $\hat{\mathbf{e}}_z$ is the unit vector pointed vertically “up”. Because this force does not depend on position or velocity, it is a constant in time for every particle. This force is zero in the case that the fluid and porous medium densities to be equal. It is primarily for this reason that we do *not* make an assumption of equal densities, as buoyancy effects are non-negligible in many applications, including the sinkhole dynamics of Chapter 5.

Cohesion Force. A cohesion force exists between two particles \mathbf{p}_i and \mathbf{p}_j , $i \neq j$. For the cohesion force we use a Lennard-Jones-like potential L [85]:

$$L(r; s, d, \epsilon) = \begin{cases} s \left[\left(\frac{d}{r} \right)^2 - 2 \left(\frac{d}{r} \right) \right] & r < (1 + \epsilon)d \\ 0 & r \geq (1 + \epsilon)d \end{cases} \quad (2.55)$$

with equilibrium distance d (the sum of the two particle’s radii) and “strength” parameter s with units M/T^2 . The independent variable r is the Euclidean distance between particle centers. The interface point $(1 + \epsilon)d$ delineates the radial distance in which the cohesion force is active; ϵ is consequently taken to be a small positive number (e.g. 0.1). The force acting on particle i as a result of its interaction with j , \mathbf{F}^{ij} , is derived by taking the negative gradient of this potential to obtain

$$\mathbf{F}_L^{ij}(r; s, d, \epsilon, \hat{\mathbf{u}}) = \begin{cases} 2s \left(\frac{d}{r^2}\right) \left[1 - \frac{d}{r}\right] \hat{\mathbf{u}}_{ij} & r < (1 + \epsilon)d \\ 0 & r \geq (1 + \epsilon)d \end{cases} \quad (2.56)$$

where $\hat{\mathbf{u}}_{ij}$ is the unit vector from \mathbf{p}_i to \mathbf{p}_j . If $(1 + \epsilon)^{-1} < d/r < 1$, then the particles are separated and the force is attracting. If $d/r > 1$, then the particles are overlapping, and the force is repulsive. This force replicates cohesion by creating an equilibrium distance between the particles ($r = d$) while also preventing “overlapping” particles.

The above model diverges from the usual 12-6 Lennard-Jones potential in two ways. First, instead of 12-6, a 2-1 exponential configuration is used to minimize the stiffness of the system as much as possible. Second, the cohesion force exists only when particles are nearly touching; this reflects the physical phenomenon that particles, once in contact, stay in contact, but will not spontaneously “attract” towards one another if they are sufficiently far apart. Due to the piecewise definition of the potential, the force is also discontinuous in r . This could be remedied in a fairly straightforward extension by choosing a potential that is differentiable and has compact support for $r < (1 + \epsilon)d$.

Because the particles exist in a subset of \mathbb{R}^n , $n \geq 2$, the cohesive force for the i^{th} particle will be written as

$$\mathbf{F}_c^i = \sum_{\substack{j=1 \\ j \neq i}}^N \mathbf{F}_L^{ij} \left(\|\mathbf{x}_i - \mathbf{x}_j\|; s, r_i + r_j, \epsilon, \frac{\mathbf{x}_i - \mathbf{x}_j}{\|\mathbf{x}_i - \mathbf{x}_j\|} \right) \quad (2.57)$$

where \mathbf{x}_j is the location of particle j , and r_i, r_j are the radii of particle i and j , respectively. Due to the interaction distance $(1 + \epsilon)d$, the vast majority of the \mathbf{F}_L^{ij} quantities will be zero. This leaves the possibility for some clever algorithm to avoid computing all of those unnecessary zeros at each time step – this will be addressed in Chapter 5.

Adhesion Force. An adhesion force is used to address particle-wall interactions. Similar to the cohesion force, the adhesion force is defined using a Lennard-Jones potential by projecting the location of a particle onto a pre-defined wall and, temporarily, imagining that there is another particle there. In this manner, particles adhere to – and, more importantly, do not pass through – walls.

Fluid Drag Force. For an inert sphere in a surrounding flow, the drag on this sphere is given classically by Stokes drag $\mathbf{D} = 6\pi\eta R\mathbf{U}$, where $R = d/2$ is the radius of the particle “spheres”, η is the fluid viscosity, and $\mathbf{U} = \mathbf{q}_f - \mathbf{q}_p^i$ is the difference in Darcy velocities between the flow and

particle i . This force is in the direction of flow and, importantly, is the only force that involves the *velocity* of the particle i . So in our notation this force becomes

$$\mathbf{F}_d^i = 6\pi\eta r_i(\mathbf{q}_s(\mathbf{x}_i) - \mathbf{q}_p^i). \quad (2.58)$$

Algebraic Equation for Particle Velocity. The Stokes force balance means that there is only a single algebraic term involving the velocity of the particle. Therefore we have:

$$\mathbf{F}_g^i + \mathbf{F}_c^i + \mathbf{F}_a^i + \mathbf{F}_d^i = \mathbf{0}, \quad (2.59)$$

which is a linear, algebraic system of equations for \mathbf{q}_p^i , the velocity of particle i . Solving for \mathbf{q}_p^i is straightforward, giving

$$\mathbf{q}_p^i = \mathbf{q}_s(\mathbf{x}_i) + \frac{1}{6\pi\eta r_i}(\mathbf{F}_g^i + \mathbf{F}_c^i + \mathbf{F}_a^i), \quad (2.60)$$

where fluid velocity and gravity, cohesion, and adhesion forces are known at a given time because the particle positions are known. Knowing \mathbf{q}_p^i , a simple forward-difference discretization can be used to update the position of the particles.

2.3 Discrete–continuum interpolation operators

In order to complete this model framework, we need a way to interpolate from a continuum to a discrete representation, and back again. The discrete-to-continuum component is straightforward; we use a smoothing operator to “smooth out” the discrete particles into a continuum field. The continuum-to-discrete operator poses a challenge, however, because it is an inverse problem that is ill-defined without applying constraints. We now detail the discrete-to-continuum operator, then describe the inverse continuum-to-discrete operator.

2.3.1 Discrete-to-continuum operator

In order to generate a porosity field θ_p , we need to be able to interpolate a set of discrete particles into a continuous field. There are several legitimate ways one can choose to do this; the method described presently is not unique. We have chosen to treat our porosity field as the result of a convolution of the the discrete particles, represented as point sources, with the Gaussian kernel G defined by:

$$G(r) = G(r; h) = \exp\left(-\frac{r^2}{h^2}\right), \quad (2.61)$$

where h is the bandwidth. The bandwidth will dictate how “large” the effect of one particle is. This has the effect of “smoothing out” the point sources. Formally, for N particles in a domain

Ω , the discrete-to-continuum operator \mathcal{G} applied to a collection of particles \mathbf{p} evaluated at a target point \mathbf{x} is defined as the convolution:

$$\mathcal{G}[\mathbf{p}](\mathbf{x}) = \theta_p(\mathbf{x}) = \int_{\Omega} G(\|\mathbf{x} - \mathbf{y}\|) \left[\sum_{j=1}^N s_j \delta(\mathbf{x}_j - \mathbf{y}) \right] d\mathbf{y}, \quad (2.62)$$

where s_j is the “strength” of particle j , \mathbf{x}_j is the source point of the j th particle, and δ is the usual Dirac delta distribution. By properties of the Dirac δ , this expression can be simplified to

$$\theta_p(\mathbf{x}) = \sum_{j=1}^N s_j G(\|\mathbf{x} - \mathbf{x}_j\|). \quad (2.63)$$

Note that, for a single evaluation, this is an $\mathcal{O}(N)$ algorithm. The true cost is realized by evaluating this sum at each of M spatial nodes where the velocity is defined. Therefore, to evaluate the entire porous media volume fraction field, one must evaluate

$$\theta_p(\mathbf{x}_i) = \sum_{j=1}^N s_j G(\|\mathbf{x}_i - \mathbf{x}_j\|) \quad \text{for } i = 1, 2, \dots, M. \quad (2.64)$$

This makes this field interpolation an $\mathcal{O}(MN)$ algorithm. Note that this is essentially a matrix-vector multiplication which can be written as $\vec{\theta}_p = \mathbf{G}\vec{s}$, where \mathbf{G} is a rectangular $M \times N$ system, \vec{s} is an $N \times 1$ vector and $\vec{\theta}_p$ is a $M \times 1$ vector. However, as both N and M will number in the tens of thousands even for moderate problems, this asymptotic cost is much too high. Luckily, there exist significantly better algorithms in the literature, including using Tree search algorithms (see Appendix B) as well as a method known as the “Improved Fast Gauss Transform” algorithm [45, 116, 97]. Both of these techniques are implemented in an excellent C++ package FIGTree¹; see the associated paper for more details [72]. At the end of the day, we can conduct this discrete-to-continuous transform with cost $\mathcal{O}(N + M)$.

2.3.2 Continuum-to-discrete operator

The more difficult operation is going from continuum to discrete. While it would be nice if an inverse gauss operator existed in the literature, that is not the case; the inversion of the above transform is, in general, an ill-defined problem. What we can do, however, is define an approximate, or pseudo, inverse as the operator that minimizes some cost function. While costly to perform, this representation of an inverse gives good results in test cases (see Sec. 3.4).

We wish to define our operator \mathcal{G}^* , which acts on a field F in a domain Ω . The fundamental property that we desire of this operator is that it “undoes” the Gauss Transform \mathcal{G} . As a reminder,

¹Homepage for FIGTree: <http://users.umiacs.umd.edu/~morariu/figtree/>

the Gauss transform \mathcal{G} takes N points, defined at N spatial locations \mathbf{x}_i with sources s_i and a single “spread” parameter h . In general, a new h could be defined for each point, but because we make use of the fast gauss transform [45], we will use a single h for all particles. We now define a cost function J , with input as a $N(d + 1) + 1$ parameter vector $\boldsymbol{\theta} = \{\mathbf{x}_i, s_i, h\}$, for a given $F(\mathbf{y})$ as:

$$J(\boldsymbol{\theta}) = \int_{\Omega} \|\mathcal{G}[\boldsymbol{\theta}](\mathbf{y}) - F(\mathbf{y})\|^2 d\mathbf{y}. \quad (2.65)$$

Note that the parameter vector $\boldsymbol{\theta}$ is unrelated to the porosity field θ_p ; this notation is simply to be more consistent with the machine learning literature [44]. Now, \mathcal{G}^* can be defined by an optimization problem:

$$\mathcal{G}^*[F] = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} J(\boldsymbol{\theta}). \quad (2.66)$$

An important constraint on this operator is that N , the number of discrete particles that will represent our field F , must be chosen beforehand. Including N as a parameter in $\boldsymbol{\theta}$ makes this problem ill-defined and computationally intractable, as approximating a given field by more and more particles will naturally make the error go down, similar to transforms such as Fourier.

This definition of \mathcal{G}^* makes sense because, from $\boldsymbol{\theta}$, we can reconstruct a set of discrete particles \mathbf{p} by using \mathbf{x}_i as the particle positions and h and s_i proportional to particle radius. In this sense, \mathcal{G} and \mathcal{G}^* are inverses of one another. The method to solve this optimization problem is independent of the definition, but one example using gradient descent is given in Sec. 3.4.

2.4 Framework summary

We include here a summary of our systems of equations for the interested reader. The fact that this framework is modular should be considered an advantage, as one can “swap out” different pieces depending on one’s own specific problem. The model for chemistry assumes that one has n reactants forming m aqueous products, K of which go onto form precipitates. With this setup, one will have $n + m + K + 1 + 1$ scalar unknowns – corresponding to ψ_i ($n + m$), θ_i ($K + 1$), and p (1) the shared pressure term – and $K + 1$ vector unknowns – corresponding to the Darcy velocities \mathbf{q}_i . The specific equations that need to be chosen in order to make this a closed system are detailed in this section. Please see the previous sections of this chapter for a detailed derivation and motivation for each set of equations; this section simply serves as a convenient summary.

2.4.1 Chemistry

The main equations for determining the temporal and spatial dynamics of the aqueous chemical species are given by ADR PDEs. The $n + m$ equations are

$$\frac{\partial(M_j\psi_j\theta_s)}{\partial t} + \nabla \cdot (M_i\psi_i\theta_s\mathbf{v}_s - \kappa_i\nabla(M_i\psi_i)) = -a_jrM_{A_j}\theta_s \prod_{i=1}^n \psi_{A_i}, \quad j \in \{A_1, \dots, A_n\} \quad (2.67a)$$

$$\frac{\partial(M_i\psi_i\theta_s)}{\partial t} + \nabla \cdot (M_i\psi_i\theta_s\mathbf{v}_s - \kappa_i\nabla(M_i\psi_i)) = b_jrM_{B_j}\theta_s \prod_{i=1}^n \psi_{A_i} + \sum_{k \in \mathcal{I}_K} \alpha_k \delta_{jk} M_{B_k} \dot{\theta}_s, \\ j \in \{B_1, \dots, B_m\} \quad (2.67b)$$

where α_k is defined by

$$\alpha_k = \frac{K\rho_p^{(k)} - \rho_s}{KM_{B_k}}. \quad (2.68)$$

If one has a multi-stage reaction, one simply has to couple together the different stages, including the necessary equations for each stage. Additionally, if any of the above reactions are reversible, one would need to add the traditional reversible terms for the law of mass action on the products. See Sec. 2.1 for a detailed description of all parameters and variables.

The ADR equations above form a nonlinear, coupled PDE system that likely contains multiple disparate timescales (e.g. diffusion, advection, and reaction timescales). In Sec. 4.1.2 we provide a timescale separation method for decoupling and linearizing these PDEs to form a system which is asymptotically similar to the above.

2.4.2 Mass Balance

The no-void assumption for K solid phases is

$$\theta_s + \sum_{k \in \mathcal{I}_K} \theta_p^{(k)} = 1 \quad (2.69)$$

and the mass balance for the $K + 1$ volume fractions is

$$\partial_t(\rho_s\theta_s) + \nabla \cdot (\rho_s\mathbf{q}_s) = R_s \quad (2.70a)$$

$$\partial_t(\rho_p^{(k)}\theta_p^{(k)}) + \nabla \cdot (\rho_p^{(k)}\mathbf{q}_p^{(k)}) = R_p^{(k)}, \quad k \in \mathcal{I}_K \quad (2.70b)$$

where we require the relationship

$$R_s = -\rho_s \left(\sum_{k \in \mathcal{I}_K} \rho_p^{(k)} \right)^{-1} \sum_{k \in \mathcal{I}_K} R_p^{(k)}. \quad (2.71)$$

between the mass transfer terms for conservation of mass. If the solid is assumed immobile, then $\mathbf{q}_p^{(k)} \equiv \mathbf{0}$ in Eq. (2.70b). The $R_p^{(k)}$ terms need to be chosen – this depends on the specific context of one’s experiment, but a simple choice is the following:

$$R_p^{(k)} = \beta_k \psi_{B_k} \theta_s \mathcal{H}(\psi_{B_k} - \psi_{B_k}^*), \quad k \in \mathcal{I}_K \quad (2.72)$$

where β_k is a rate constant that is usually on the same order as r . In an ideal world it can be determined by experiment, although the more likely practical scenario is that one chooses an order-of-magnitude estimate based on the context of the specific problem. See Sec. 2.1 for a detailed description of all parameters and variables.

2.4.3 Momentum Equations

The fluid phase momentum is described by

$$\eta \nabla^2 \mathbf{q}_s - \frac{\xi}{\theta_s} \mathbf{q}_s = \theta_s \nabla p, \quad (2.73a)$$

$$\nabla \cdot \mathbf{q}_s = \mathbf{0}. \quad (2.73b)$$

If the K precipitate phases are immobile, then $\mathbf{q}_p^{(k)} \equiv \mathbf{0}$, and you are done. If not, and you are using a discrete description of the solid phases, then the porous media particle velocities are given by:

$$\mathbf{q}_p^i = \mathbf{q}_s(\mathbf{x}_i) + \frac{1}{6\pi\eta r_i} (\mathbf{F}_g^i + \mathbf{F}_c^i + \mathbf{F}_a^i). \quad (2.74)$$

See Sec. 2.2 for a detailed description of all parameters, variables, and forces.

2.4.4 Interpolation Operators

The interpolation operators are given by

$$\mathcal{G}[\mathbf{p}] = \sum_{j=1}^N s_j G(\|\mathbf{x}_i - \mathbf{x}_j\|) \quad (2.75a)$$

$$\mathcal{G}^*[F] = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} J(\boldsymbol{\theta}). \quad (2.75b)$$

where

$$G(r) = \exp\left(-\frac{r^2}{h^2}\right), \quad J(\boldsymbol{\theta}) = \int_{\Omega} \|\mathcal{G}[\boldsymbol{\theta}](\mathbf{y}) - F(\mathbf{y})\|^2 d\mathbf{y}. \quad (2.76)$$

See Sec. 2.3 for a detailed description of all parameters, variables, and operators.

CHAPTER 3

FRAMEWORK VALIDATION AND ANALYSIS

The previous chapter was aimed exclusively at theoretical derivations that focused on internal coherence of the framework – but how applicable is the framework to the physical situations it purports to model? This chapter contains four benchmarks that demonstrate the capabilities – and limitations – of the model framework. The chapter proceeds as follows: Sec. 3.1 derives a reduced model of the stationary precipitate; this reduced model is simple enough to derive largely analytic solutions and is the first evidence that the framework works as expected, while also demonstrating deficiencies in existing models. Sec. 3.2 focuses exclusively on how our multiphase Brinkman equation represents drag on two-dimensional bodies. Integral equation code is utilized to provide an “exact” benchmark case, which allows us to discuss the form of a two-dimensional multiphase “porous media” should be interpreted. Sec. 3.3 demonstrates how the continuum fluid description can be coupled with discrete particles to represent slow-moving porous media. Finally, Sec. 3.4 gives a one-dimensional test of the discrete-continuum operators.

3.1 Analysis of a reduced model

We consider a simplified system – corresponding to the reaction in Eqs. (2.1) and (2.2) – in which incoming reactant concentrations are held fixed via chemostat [89]. By assuming parallel flow and neglecting solute diffusion, the governing equations reduce to a planar system of ODEs. This nonlinear system can be linearized around a fixed point and eigenvalue analysis provides an estimate for the rate at which membrane forms. Moreover, we find that the equation for the aqueous product is a second-order nonlinear ODE known as the Riccati equation [86, 104]. Exact solutions to the Riccati equation give explicit formulas for the time dependence of the chemical product and, consequently, the formation of new membrane. Once the membrane dynamics are known exactly, the flow profile can be obtained through the numerical solution of a simple boundary value problem (BVP). Visualization of the resulting flow profile allows a comparison between variants of the multiphase framework. In particular, we demonstrate that the framework developed here properly captures the transition from one-channel to two-channel flow as membrane develops.

Analysis of any complicated system is aided by reduction into a form that is analytically tractable. Inspired by microfluidic experiments, we assume that some chemostat controls the influx of reactants' molarity far upstream. All variables are kept constant along the longitudinal axis by neglecting diffusion and assuming parallel flow. This requirement of parallel flow also means that the reaction takes place everywhere along the longitudinal axis simultaneously. Finally, we assume that the precipitation threshold is negligible. These assumptions approximately model the fast-timescale, *initial* membrane growth observed in the experiments shown in Fig. 1.2 for a single transverse cross-section of the domain; the slow-timescale membrane thickening is diffusion-controlled, and therefore we do not seek to capture it in this analysis.

Applying these assumptions to the governing equations reduces the system considerably so that it becomes a Poiseuille analysis; these assumptions generate the following reduced system

$$\dot{\psi}_C = cr\psi_A\psi_B - (\psi_C - \alpha)\dot{\theta}_s/\theta_s \quad (3.1a)$$

$$\theta_s + \theta_p = 1 \quad (3.1b)$$

$$\dot{\theta}_s = -\beta\psi_C\theta_s/\rho_p \quad (3.1c)$$

$$\eta\frac{\partial^2 q_y}{\partial x^2} - h\frac{\theta_p^2}{\theta_s^2}q_y = \theta_s\frac{\partial P}{\partial y}. \quad (3.1d)$$

where only ψ_C , θ_s , θ_p , q_y and $\partial_y P$ are unknown. The pressure gradient $\partial_y P$ is determined by requiring a constant total flux for all time – see Eq. (3.12). Note that the longitudinal axis is chosen to be y such that only this component of the Darcy velocity $\mathbf{q}_s = q_x\hat{\mathbf{x}} + q_y\hat{\mathbf{y}}$ remains. In Section 3.1.1 we obtain an analytic estimate on the rate of membrane formation by treating ψ_C and θ_p as a planar dynamical system. Then in Section 3.1.2 we solve Eqs. (3.1) with a combination of analytic and numerical methods to visualize how the membrane affects the flow profile in time.

3.1.1 Fixed point analysis

Our approach is to linearize the reduced model system, then perform an eigenvalue analysis about a steady state fixed point. The benefit of this is the eigenvalue gives an approximate estimate to the rate that membrane develops, a quantity that is possible to measure experimentally.

Before doing a fixed point analysis, it is helpful to understand the conditions on which the existence and stability of fixed points depend. To do so, eliminate the explicit dependence of $\dot{\psi}_C$ on volume fraction and replace all $\dot{\theta}_s$ in Eq. (3.1a) with Eq. (3.1c) to obtain a quadratic ODE of the form

$$\dot{\psi}_C = \frac{\beta}{\rho_m}\psi_C^2 - \frac{\alpha\beta}{\rho_m}\psi_C + cr\psi_A\psi_B. \quad (3.2)$$

which is an ODE in time alone, as the x -dependence of ψ_A and ψ_B are determined by the initial conditions. Examining the qualitative behavior of this ODE by considering $\dot{\psi}_C = \dot{\psi}_C(\psi_C)$, it is quadratic in ψ_C , intercepts the $\dot{\psi}_C$ axis at $cr\psi_A\psi_B \geq 0$, is concave up, and has equilibria at

$$\psi_C^\pm = \frac{1}{2} \left(\alpha \pm \frac{1}{\beta} \sqrt{\alpha^2 \beta^2 - 4cr\rho_p\beta\psi_A\psi_B} \right) \quad (3.3)$$

whose existence depends on the sign of

$$\chi = \alpha^2 \beta^2 - 4cr\rho_p\beta\psi_A\psi_B. \quad (3.4)$$

If $\chi > 0$, Eq. (3.2) will have two fixed points, for $\chi = 0$ these fixed points coalesce, and for $\chi < 0$ there are no fixed points and $\dot{\psi}_C$ will grow without bound; see Fig. 3.1(a).

We now ask the question of whether fixed points exist in the reduced system, i.e. $\chi \geq 0$ or $\alpha^2 \beta \geq 4cr\rho_p\psi_A\psi_B$? We interpret this condition based on the physical meaning of the parameters: $\alpha = (\rho_p - \rho_s)/M_C$ has dimension of molarity and is $\mathcal{O}(10 \text{ M})$ where M refers to molar units, 1 M = 1 mol/liter; for example, using the reaction system in the introduction gives $\alpha \approx 30 \text{ M}$. We note that a similar analysis also justifies neglecting the precipitation threshold ψ_C^* , as $\psi_C^* \approx 0.001 \text{ M}$. Although both r and β scale the rates of the aqueous and precipitate reactions, respectively, they have different units. r has units of volume per time while β has units of mass per time. Because experimental values of r and β are expensive to acquire, for the sake of this simplified analysis we will assume that $\beta \approx r\rho_p$ such that their effects do not impact the sign χ . The stoichiometric coefficient c for $C(\text{aq})$ can be assumed $\mathcal{O}(1)$. Finally, examine the reactants ψ_A and ψ_B . Most experiments in microfluidic chambers use molar concentrations with an upper bound of $\mathcal{O}(1 \text{ M})$; for example, in Ding et al. [31] the maximum concentration of reactants was 0.5 M. Therefore, using parameter values taken from experiments, $\chi > 0$ and fixed points exist for the reduced system.

We now consider the planar dynamical system in phase space $(\psi_C, \theta_p) \in [0, \infty) \times [0, 1]$ with fixed point $(\psi_C^-, 1)$. The dynamical system is

$$\dot{\psi}_C = f(\psi_C, \theta_p) = \frac{\beta}{\rho_p} \psi_C^2 - \frac{\alpha\beta}{\rho_p} \psi_C + cr\psi_A\psi_B \quad (3.5a)$$

$$\dot{\theta}_p = g(\psi_C, \theta_p) = \beta\psi_C(1 - \theta_p)/\rho_p \quad (3.5b)$$

where r , c , ρ_p , α , β , ψ_A , and ψ_B are assumed to be known and constant. The eigenvalues of the Jacobian generated by equations (3.5) evaluated at the fixed points provide information about the rate of growth of ψ_C and θ_p . This particular eigen-system is simple to interpret because the eigenvectors align with the coordinate axes and therefore the eigenvalues correspond to the rates

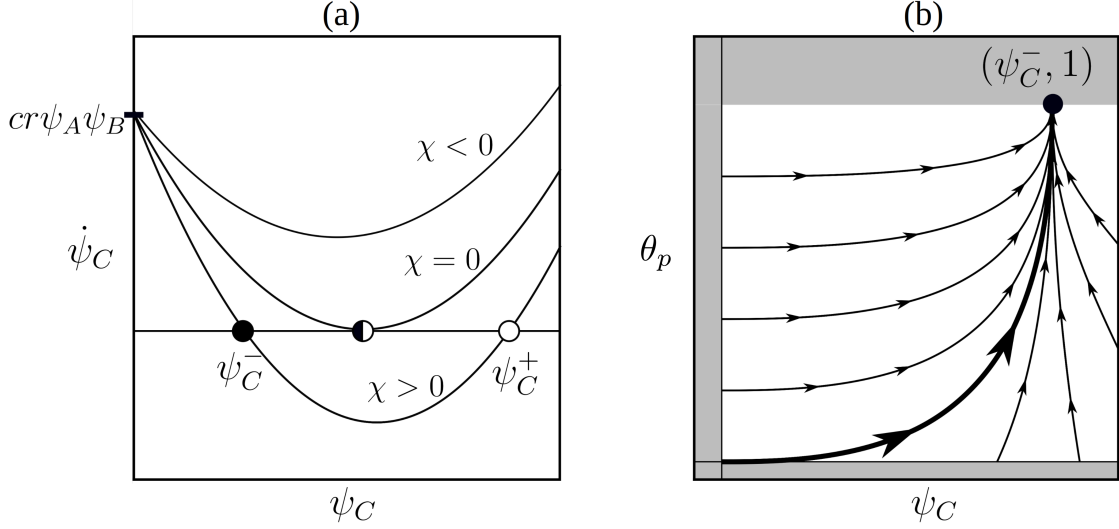


Figure 3.1: Dynamical system for ψ_C and θ_p . (a) Qualitative stability of $\dot{\psi}_C(\psi_C)$ ODE. The left equilibrium ψ_C^- is stable and exists for $\chi \geq 0$. (b) Visualization of planar dynamical system approaching the fixed point $(\psi_C^-, 1)$ for $\chi > 0$. The thicker line corresponds to homogeneous initial conditions for ψ_C and θ_p . Shaded region is outside of the domain of $(\psi_C, \theta_p) \in [0, \infty) \times [0, 1]$. Reproduced from [36] with permission.

that the physical variables (ψ_C, θ_p) approach their equilibria when close to the steady state. These rates are given by:

$$\lambda_{\psi_C} = -\frac{1}{\rho_p} \sqrt{\chi}, \quad \lambda_{\theta_p} = -\frac{1}{2} \left(\frac{\alpha\beta}{\rho_p} + \lambda_{\psi_C} \right). \quad (3.6)$$

Both eigenvalues are negative, and therefore the fixed point is stable, because

$$\alpha\beta/\rho_p + \lambda_{\psi_C} = \alpha\beta/\rho_p - \sqrt{\chi}/\rho_p > 0 \quad (3.7)$$

always. For this same reason, it is true that $|\lambda_{\theta_p}| < |\lambda_{\psi_C}|$, meaning the membrane volume fraction approaches its fixed point at a slower rate than the aqueous product. This agrees with our intuition, as the conversion of $C(\text{aq})$ to $C(\text{s})$ means we would expect θ_p production to lag behind ψ_C .

3.1.2 Ricatti-Poiseuille analysis

We now solve system (3.1) to visualize the transition in solvent velocity from one- to two-channel flow. To summarize the approach, each equation in (3.1) is solved sequentially; Eqs. (3.1)(a-c) admit exact solutions while the fluid velocity in Eq. (3.1)(d) requires numerical solution. Although in this reduced case the location of the membrane is determined through the initial conditions ψ_A^0 and ψ_B^0 , and therefore known *a priori*, the fact that no interface boundary conditions are required demonstrates a considerable advantage of this model over more traditional approaches.

For initial conditions, let $\psi_C(x, t) = 0$ and $\theta_s(x, t) = 1$. Fix the initial reagents by the piecewise-constant values

$$\psi_A^0(x) = \begin{cases} 0.4 & x \leq (L+w)/2 \\ 0 & x > (L+w)/2 \end{cases}, \quad \psi_B^0(x) = \begin{cases} 0 & x < (L-w)/2 \\ 0.3 & x \geq (L-w)/2 \end{cases} \quad (3.8)$$

such that there is only a middle region of width w in $x \in (0, L)$ in which the reactants A and B are simultaneously present. We note that the following results hold for any choice of ψ_i^0 , provided that $\chi(x) > 0$.

Substituting Eq. (3.1)(c) into Eq. (3.1)(a) produces an ODE with quadratic nonlinearity known as the Riccati equation. Despite being nonlinear, the Riccati equation admits an exact solution for $\psi_C(x, t)$; see Appendix A for a detailed solution. Given $\psi_A^0(x)$ and $\psi_B^0(x)$, the solution to this Riccati equation is:

$$\psi_C(x, t) = \gamma_1 \gamma_2 \frac{\rho_p}{\beta} \left(\frac{e^{\gamma_2 t} - e^{\gamma_1 t}}{\gamma_2 e^{\gamma_1 t} - \gamma_1 e^{\gamma_2 t}} \right) \quad (3.9)$$

where

$$\gamma_{1,2}(x) = \frac{1}{2} \left(-\alpha\beta \pm \frac{\sqrt{\chi(x)}}{\rho_p} \right). \quad (3.10)$$

In order to find the exact solution for $\theta_s(x, t)$, we must use the fact that Eq. (3.1c) is separable. Then, because the antiderivative of ψ_C can be given in terms of elementary functions, we can obtain an explicit formula for solvent volume fraction $\theta_s(x, t)$:

$$\theta_s(x, t) = \frac{\gamma_1 e^{\gamma_2 t} - \gamma_2 e^{\gamma_1 t}}{\gamma_1 - \gamma_2} \quad (3.11)$$

where we have implemented the initial condition $\theta_s(x, 0) = 1$. Then, the membrane volume fraction $\theta_p(x, t)$ can be computed easily using the no-void assumption.

Until this point, all solutions in space have been treated independently. The effect of variations in space is taken into account when solving for the longitudinal component of the Darcy velocity, q_y . We numerically solve Eq. (3.1d) for q_y using a finite difference method. The x domain is discretized into N intervals of equal width $\Delta x = 1/N$ such that $x_j = j\Delta x$, $j = 1, \dots, N-1$, and use centered-difference approximations to the derivatives. Note that $q_y(x_0) = q_y(x_N) = 0$ due to the no-slip boundary conditions. This discretization results in a tridiagonal linear system which can be solved in $\mathcal{O}(N)$ complexity by using the Thomas algorithm [see 98, pp. 78-79]. The velocity is constrained to satisfy constant flux in accordance with experiments, which mathematically is represented by

$$\int_0^L q_y dx = \text{constant}. \quad (3.12)$$

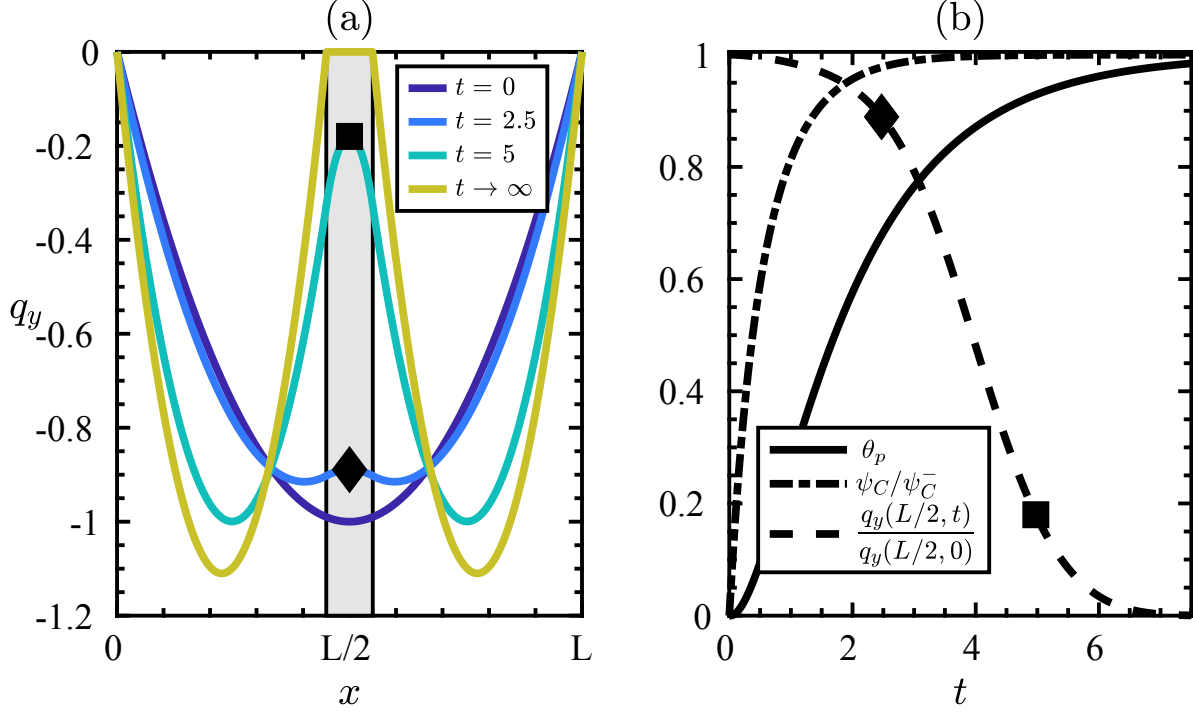


Figure 3.2: Developing precipitate affects fluid flow. (a) Flow profile in a 1D channel transitions from one-channel to two-channel flow. The reaction region is shaded. (b) Relevant variables evaluated in the reaction region at $x = L/2$, normalized for legibility; ψ_C develops first, followed by θ_p , which when large enough triggers the transition from one- to two-channel flow. The percolation threshold θ_s^* is set to $\theta_s^* = 0.3$, which is why negligible change in fluid velocity is seen until $\theta_p \approx 0.7$. The specific θ_s^* was chosen based on the maximum volume fraction of an arrangement of packed spheres in three dimensions [26]. The black diamond and square represent the system state inside the reaction region at $t = 2.5$ and $t = 5$, respectively. Reproduced from [36] with permission.

This constant-flux condition allows the computation of the required pressure gradient at each time step. We use the Julia programming language to solve the BVP [12].

The developing membrane for the 1D reduced model geometry is shown in Fig. 3.2(a). Membrane develops within the shaded region, which in this example is 10% of the domain. Because the membrane has finite width, the constant-flux condition causes the pressure gradient to increase with the developing membrane, causing the maximum speed for the two-channel flow to be slightly higher than the maximum speed for the one-channel flow. In this sense, the developing membrane splits the domain into two symmetric one-channel flows. This transition occurs *without* the need for boundary conditions at the fluid-membrane interface, indicated by the boundary of the shaded region.

Fig. 3.2(b) shows the three main variables of the reduced model as functions of time, evaluated in the middle of the reaction region ($x = L/2$). The ψ_C variable increases immediately due to the

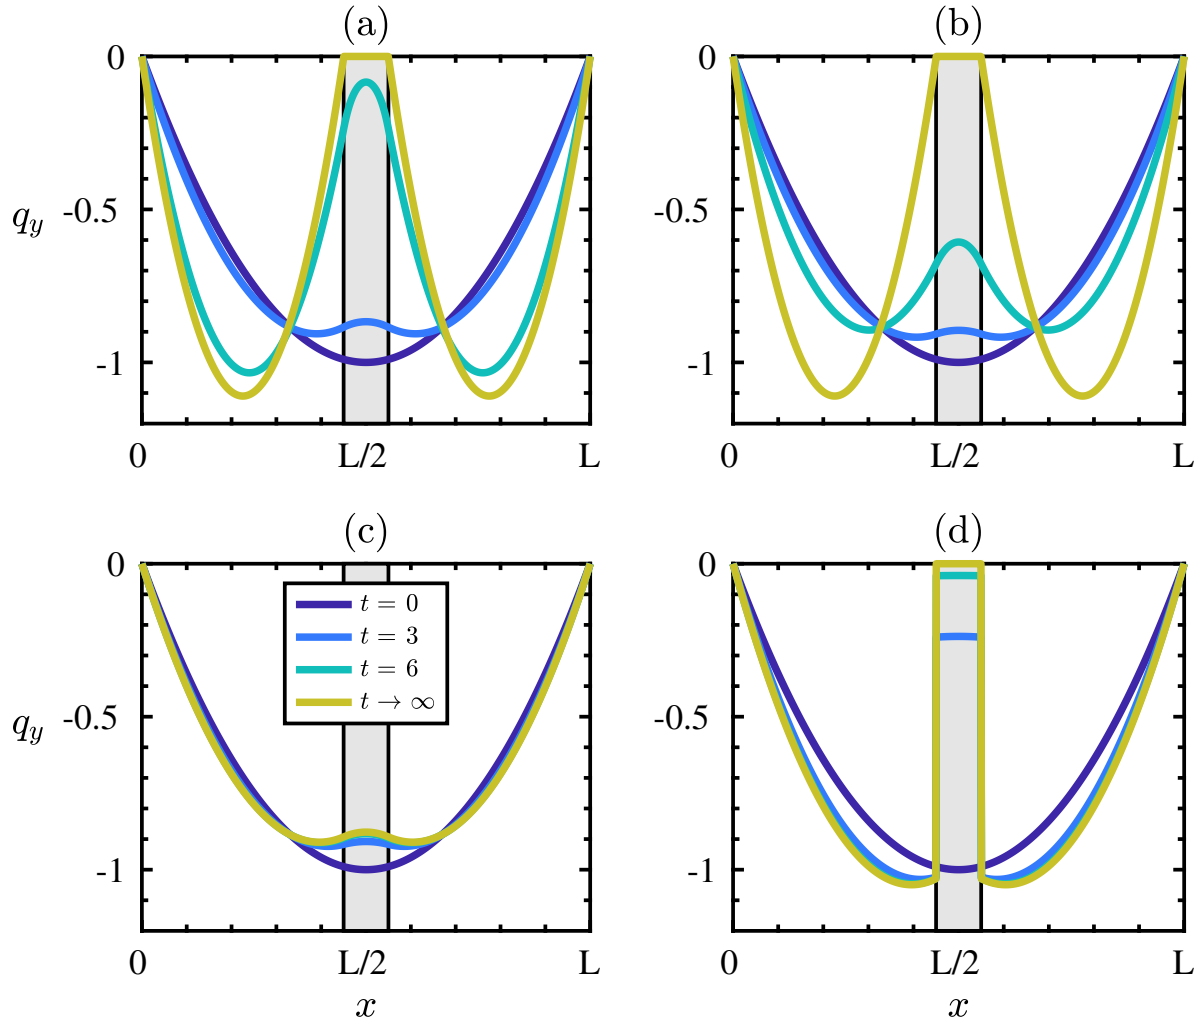


Figure 3.3: Comparison of different models. (a)–(c) use the stress in Eq. (2.47) with the friction term given by (a) Kozeny-Carman, (b) Hill function, and (c) the ‘biofilm’ term; (d) uses a conventional multiphase stress \mathbf{T}' in Eq. (3.13) with ξ_B . The coefficients h are scaled so as to make the three coefficients comparable in strength. The first two friction terms produce the desired no-slip on the membrane interface and the third, while affecting the fluid flow, does not generate the desired no-slip boundary condition. Additionally, the stress and friction combination used in (d), which is usually employed in multiphase models, does not yield the desired no-slip behavior on the fluid-structure interface. Reproduced from [36] with permission.




Figure 3.4: Effect of increasing precipitate thickness. As a percentage of the domain length, membrane width is (a) 5%, (b) 15%, and (c) 33%. The increasing maximum flow speed is due to the constant-flux constraint, and is analogous to that what would occur if the precipitate boundaries and interface conditions were prescribed *a priori* in a single-phase flow. Reproduced from [36] with permission.

presence of ψ_A and ψ_B . The membrane initially has zero growth rate due to the absence of ψ_C , and grows at a slower rate than ψ_C . This ordering on the growth rates matches our expectations from the eigenvalue analysis of Section 3.1.1. The q_y curve demonstrates the transition from one-channel to two-channel pipe flow by measuring the normalized value in the middle of the pipe as a function of time. By comparing q_y with θ_p , one can see the effect of the percolation threshold $\theta_s^* = 0.3$. After the solvent volume fraction declines past this value, the fluid velocity begins to respond strongly to precipitating membrane.

To be precise, we say a velocity displays “no-slip” behavior when, in the limit as a transition region between a fluid region ($\theta_s = 1$) and membrane region ($\theta_s = 0$) becomes discontinuous (i.e. a “sharp interface”), the velocity goes to zero as one approaches the interface from the fluid-region. In the case that the volume fraction is always discontinuous (as in this reduced model), this limit is achieved as the membrane fully develops, i.e. $\theta_s \rightarrow 0$, at the interface. Fig. 3.3 demonstrates the effect of using different stresses and friction coefficients. Both Figs. 3.3(a,b) demonstrate the desired no-slip behavior in the membrane limit by using ξ_{KC} and ξ_H , respectively. In Fig. 3.3(c) the effect

of the friction coefficient ξ_B is shown. While there is some effect on the flow profile, ξ_B does *not* demonstrate the no-slip condition on the membrane. While the ξ_B term was developed primarily for high permeability applications, our framework was developed to capture the transition from purely fluid behavior, to partially permeable, to a fully-developed impermeable solid. As demonstrated, this full transition requires either the ξ_{KC} or ξ_H friction coefficient. Fig. 3.3(d) displays the flow profile when using the multiphase stress tensor typically used in multiphase models:

$$\mathbf{T}' = \eta\theta_f(\nabla\mathbf{v}_s + \nabla\mathbf{v}_s^\top) \quad (3.13)$$

given in Drew [33], Cogan and Guy [23] as well as the biofilm friction coefficient. Even in long-time, the model using the above \mathbf{T}' does not produce the desired no-slip at the fluid-membrane interface. Using \mathbf{T}' with other friction terms produces inconclusive results. It was this behavior that motivated the authors' to modify the stress term to the form given in Eq. (2.47), which has the additional advantage that it allows for a reduction of the momentum equation to the Brinkman form.

Fig. 3.4 shows three flows with reaction regions of various sizes, and therefore different width of membranes. The initial flow profiles of all are equivalent, as the reaction has not yet occurred and no membrane is present. However, as membrane develops, the constant-flux condition requires that for regions with thicker membranes, the flow velocity must increase in the non-reacting regions to compensate for the loss of flux in the membrane region. These results demonstrate that, once the membrane is fully developed, the flow domain treats the membrane portion as a no-slip boundary and the prescribed constant-flux conditions lead to the expected results from single-phase fluids.

Alternative solution to chemical system. There exists an alternative solution to the chemostat chemical C concentration ODE. The notable difference of this solution method is that there is no need to solve Ricatti's differential equation. As a reminder, the system under consideration is

$$\dot{\psi}_C = a_1 - (\psi_C - a_2)\dot{\theta}_s/\theta_s \quad (3.14)$$

$$\dot{\theta}_s = a_3\psi_C\theta_s. \quad (3.15)$$

where a_1 , a_2 and a_3 are constant parameters. The previous solution method is to use Eq. (3.15) to transform Eq. (3.14) into a constant-coefficient Ricatti equation, which is exactly solvable. Then, because the antiderivative to this solution can be expressed using elementary functions, one can separate Eq. (3.15) to solve for θ_s exactly.

The alternative method of solution is that we can take the derivative of Eq. (3.15) to obtain

$$\ddot{\theta}_s = a_3(\dot{\psi}_C\theta_s + \psi_C\dot{\theta}_s). \quad (3.16)$$

Plugging in the expression for $\dot{\psi}_C$ from Eq. (3.14) allows us to cancel the dependence on ψ_C and thus we have a second-order, linear, constant-coefficient expression that can be solved for θ_s :

$$\frac{\ddot{\theta}_s}{a_3} = a_1\theta_s + a_2\dot{\theta}_s. \quad (3.17)$$

This is useful because now, given the parameters a_1 , a_2 and a_3 we can solve the membrane dynamics exactly without relying on the chemistry solution. Additionally, knowing θ_s exactly turns equation (3.14) into a variable-coefficient linear first-order ODE, which is also solvable. While the end result is the same, this alternative technique is reported for the interested reader.

3.2 Interpretation of multiphase drag

Until this point, the fluid-structure interactions properties of the model have only been shown to qualitatively match intuition. The ultimate truth, in any scientific field, is experiment. Experiments studying drag on porous bodies have been conducted [103, 118, 68, 53], but in this section we take a different approach. We give a quantitative analysis of the drag through porous bodies using two different models of porous media – one using our framework, and the other using an “exact” representation of a porous body as a collection of smaller, impenetrable circles. The goal is to compare how the drag produced by the current framework agrees with one version of the “true value”. The key quantity of interest will be drag, as that is something that is easily measurable from experiment and whose quantitative accuracy is of the utmost importance to those who employ fluid-structure interaction models.

3.2.1 Benchmark problem

We will analyze the drag on two dimensional porous objects. In order to avoid to Stokes Paradox [48, pp. 47-49], our domain will be bounded. Specifically, a domain like the one in Fig. 3.5 is considered: a circular object of radius r exists in a channel of length 6 and height 2. Parabolic inflow comes from the left and exits the domain on the right, while the top and bottom are no-slip walls. We first must validate the numerical methods by considering the full circle as an impenetrable solid, while later during the real test the circle will be a porous object. The inflow condition is parabolic:

$$u(y) = 1 - y^2, \quad (3.18)$$

chosen so that the maximum velocity at the inflow is 1; specifically we assume $\mu = 1$ and the pressure gradient have been chosen such that this condition occurs.




Figure 3.5: Benchmark domain for drag analysis. The circle (shaded) with radius r will either be represented by a porosity field (present framework) or by a collection of small “filling circles” (boundary integral equations).

Different numerical methods are used to simulate each model of porous media: a finite element method (FEM) using open-source software package eFEMpart [37] for the current framework, and a boundary integral method using a private package designed by Bryan Quaife and Nick Moore [84, 21] for the “exact” porous body composed of smaller circles. To validate the numerical methods, the radius of a single, impenetrable circle is fixed at 0.125 and the computed value of the drag on the circle, for a fixed inflow condition, is used as the quantity for comparison as the numerical domain is refined. Then, once the numerical packages are validated, a comparison for porous media will be conducted. First, a method for computing drag on multiphase objects is described.

3.2.2 Indirect method for computing drag

The drag \mathbf{D} on the surface is defined by the contour integral

$$\mathbf{D} = \int_C \boldsymbol{\sigma} \cdot \mathbf{n} \, ds \quad (3.19)$$

where C is the circle’s boundary. While that definition works for classical bodies with clearly defined interfaces, one runs into difficulties trying to apply it directly to “porous objects” defined within the framework. Specifically, there is no interface in our multiphase approach. Therefore, we give here an indirect approach to calculate the drag on multiphase bodies in Stokes flows. The drag on a generic body can be calculated by integrating the traction on the channel walls, inflow, and outflow, as opposed to the body itself; this is similar but not equivalent to the method of measuring

the pressure drop across a channel with the body and subtracting off the pressure drop that would occur without the body.

To begin the derivation, take Stokes equation $\nabla \cdot \boldsymbol{\sigma} = \mathbf{0}$, where $\boldsymbol{\sigma} = -p\mathbb{I} + 2\mu\mathbb{D}$, integrate over the domain Ω , and use the divergence theorem to obtain

$$\int_{\Gamma} \boldsymbol{\sigma} \cdot \mathbf{n} \, ds = \mathbf{0} \quad (3.20)$$

where \mathbf{n} is the unit normal vector – that points *outside* the domain and *into* our circle – and $\Gamma = \Gamma_{\text{in}} \cup \Gamma_{\text{out}} \cup \Gamma_{\text{wall}} \cup \Gamma_{\text{obj}}$ (see Fig. 3.5 for boundary indicators). The drag on the body is traditionally defined as

$$\mathbf{D} = - \int_{\Gamma_{\text{obj}}} \boldsymbol{\sigma} \cdot \mathbf{n} \, ds, \quad (3.21)$$

where the negative sign is due to our convention having \mathbf{n} point *into* the object. We can represent the drag on the body by a surface integral over the boundary

$$\mathbf{D} = \int_{\Gamma_{\text{wall}}} \boldsymbol{\sigma} \cdot \mathbf{n} \, ds + \int_{\Gamma_{\text{in}}} \boldsymbol{\sigma} \cdot \mathbf{n} \, ds + \int_{\Gamma_{\text{out}}} \boldsymbol{\sigma} \cdot \mathbf{n} \, ds. \quad (3.22)$$

The convenience of this equation comes from the fact that the geometry of the inner body has been completely removed from the calculation, and we are computing line integrals over straight lines along our axis coordinates. This technique is applicable using the current framework as long as the fluid volume fraction is unity on the boundary – in which case, our framework is equivalent to Stokes equations. As mentioned previously, this method will be even more convenient in the context of multiphase flows because the definition of the “body boundary” itself is unclear.

Example. As a sanity check, we compute the drag on a non-existent body; the drag should be zero. The velocity solution for an empty channel is given in in Eq. (3.18) with corresponding pressure solution $p = 2(3 - x)$. The symmetric rate-of-strain tensor for the channel flow can be written in cartesian form as

$$\mathbb{D} = \begin{bmatrix} 0 & -y \\ -y & 0 \end{bmatrix}.$$

Now all that is needed is to compute the drag on each of the four faces of the outer boundary. Remember, $\boldsymbol{\sigma} = -p\mathbb{I} + 2\mu\mathbb{D}$ with $\mu = 1$ for this benchmark. At the boundaries one obtains:

$$\int_{\Gamma_{\text{in}}} \boldsymbol{\sigma} \cdot \mathbf{n} \, ds = \begin{bmatrix} -24 \\ 0 \end{bmatrix}, \quad \int_{\Gamma_{\text{out}}} \boldsymbol{\sigma} \cdot \mathbf{n} \, ds = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \quad \int_{\Gamma_{\text{wall}}} \boldsymbol{\sigma} \cdot \mathbf{n} \, ds = \begin{bmatrix} 12 \\ 36 \end{bmatrix} + \begin{bmatrix} 12 \\ -36 \end{bmatrix} = \begin{bmatrix} 24 \\ 0 \end{bmatrix} \quad (3.23)$$

so that the total drag on our “imaginary” circle, via Eq. (3.22), is zero, just as we had hoped. We have learned that the drag on the walls cancels the drag on the inflow, while the outflow is

Table 3.1: Benchmark Comparison of the errors on a circle in a channel. The error refers to the relative error; the FEM error is computed by using the highest value from BIE as the “exact” value. For this benchmark, a no-slip condition on the circle is provided to the multiphase flow – this table does not reflect the capabilities of our model framework, only verifying that the code works as expected for a benchmark problem.

Boundary Integral			Finite Elements		
dof ^a	drag	error	dof	drag	error
8	10.4754984115	4.9e-7	1252	10.74769495	2.6e-2
16	10.4755035934	1.7e-9	2584	10.50925591	3.2e-3
32	10.4755035661	9.1e-10	4732	10.45909286	1.6e-3
64	10.4755035809	5.1e-10	10032	10.42472542	4.8e-3
128	10.4755035784	2.6e-10	30720	10.47527272	2.2e-5
256	10.4755035777	2.0e-10	113136	10.47627087	7.3e-5
512	10.4755035756	-	425856	10.47535576	1.4e-5

^aThe dof in the BIE refers to the discretization of the circle – it does not include the discretization of the outer boundary.

stress-free. Therefore, to compute the drag on the porous circle, one only needs to compute line integrals on the outside of the domain.

To validate the numerical methods we will use, a domain convergence study is conducted; the results are given in Table 3.1. One can see that the boundary integral formulation is much more accurate for all discretization levels; this is due to the spectral convergence properties of the the boundary integral formulation for closed, smooth boundaries. The finite element formulation, while being overall less accurate, shows the expected first-order convergence due to the pressure variable being represented by first-order finite element family; see Sec. 4.2 for a complete description of the FEM discretization. To summarize, this says that the numerical methods are validated in the measurement of drag. As a visual validation of our method, we see in Fig. 3.6 a visualization of the flow field for a flow around a “solid” circle, both using the multiphase framework and Stokes equations. In the multiphase model, the circle is represented by a discontinuous volume fraction scalar field, and the flow reacts accordingly, while in (b) a circle must be explicitly built into the mesh so that a no-slip boundary condition may be applied.

3.2.3 Porous circle problem

Now that both numerical methods are validated, they can be used to investigate the real purpose of this section. We assume that we are given the geometry files for the “filling circles” representation




Figure 3.6: Comparison of flow fields using our model framework versus a classical no-slip boundary condition. Color represents the magnitude of the velocity field, and both images use the same color scaling. Comparison of the flow generated from (a) our multiphase Brinkman equation and (b) classical Stokes flow over a circle. In the multiphase model, the circle is represented by a volume fraction scalar field, while in (b) a circle must be explicitly built into the mesh so that a no-slip boundary condition may be applied.

of a porous body by a benevolent god.¹ We will examine only two parameters: volume fraction and drag on the body. For the “filling circle” representation of the porous object, due to the random distribution of these bodies, we will expect some variation in the drag result. See Fig. 3.7 for examples of porous media of various volume fractions. For the simulation that uses our framework, we assume the surface is discontinuous.

Drag comparisons. We first run the simulations on the “filling circles” geometries where the drag will be computed using the boundary integral method. Results are shown in Fig. 3.8(a). We immediately notice that there is wide variability, for a given volume fraction, in the drag on a “porous circle” composed of 40 filling circles. One notable observation is that the drag, even for extremely small volume fraction, is still about 90% of the drag on a solid circle of radius 0.125. While this does not agree with our intuition about how drag operates on objects, we must remember two things: (1) we are computing the drag in a two-dimensional setting, where the drag on a single object decays according to the inverse log of the shrinking radius [84, see Eq. (66)], which we can consider a proxy for volume fraction and (2) we have many different bodies, not a single body, so there is a compounding effect of many, different bodies affecting the drag on each other.

For the multiphase model, we look at Fig. 3.8(b). One can see that the drag is a smooth function of volume fraction, and approximately equals the drag on a solid body when the volume

¹CircGeometry.jl was developed for this purpose. See Appendix C for details on the full algorithm.

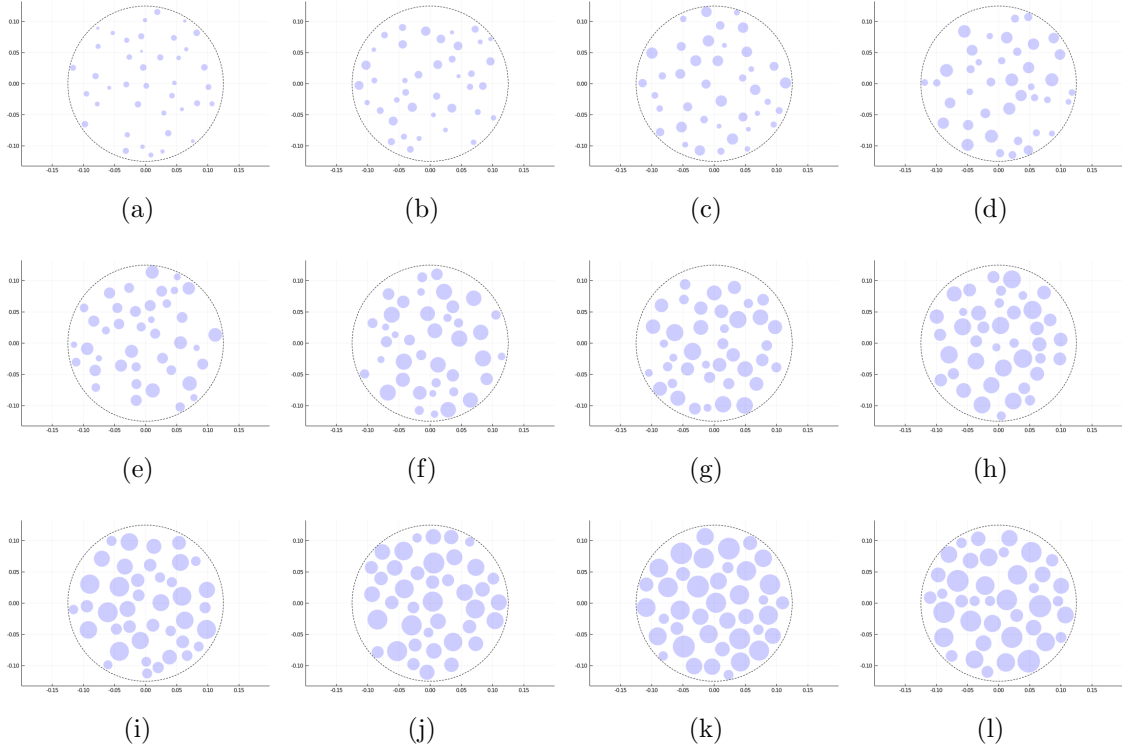


Figure 3.7: Examples from porous circles with volume fractions equal to (a) 5% (b) 10% (c) 15%, (d) 20% (e) 25% (f) 30% (g) 35% (h) 40% (i) 45% (j) 50% (k) 55% (l) 60%. Note that the actual volume fractions (cumulative area of filling circles divided by area of outline) are not exactly these values, but only approximate, due to the filling circle radii being taken from a random distribution. See Appendix C for more details.

fraction of the solid is 1 (i.e. the precipitate is fully developed). With this comparison, it is clear that, at least in two dimensions, the volume fraction interpretation of the precipitate membrane as being composed of small filling circles is incorrect. However, we do see that (1) the drag on the porous body varies smoothly from zero volume fraction (no circle present) to unity volume fraction (porous object is completely solid), and (2) the drag goes to the correct limits (approximately) in either case. While the quantitative accuracy of our multiphase model regarding porous objects is therefore in question, the same accuracy for fully-developed objects looks to be quantitatively accurate. This is excellent news, as it means that, once a precipitate is fully developed, the drag it exerts on a fluid behaves exactly like that of a solid object imposed *a priori*.

This section has validated the current model as a *quantitatively* accurate model for solids, but the quantitative accuracy for modeling drag on porous bodies is still in question. We also detailed an indirect method for computing the drag on a multiphase body, and demonstrated that we should

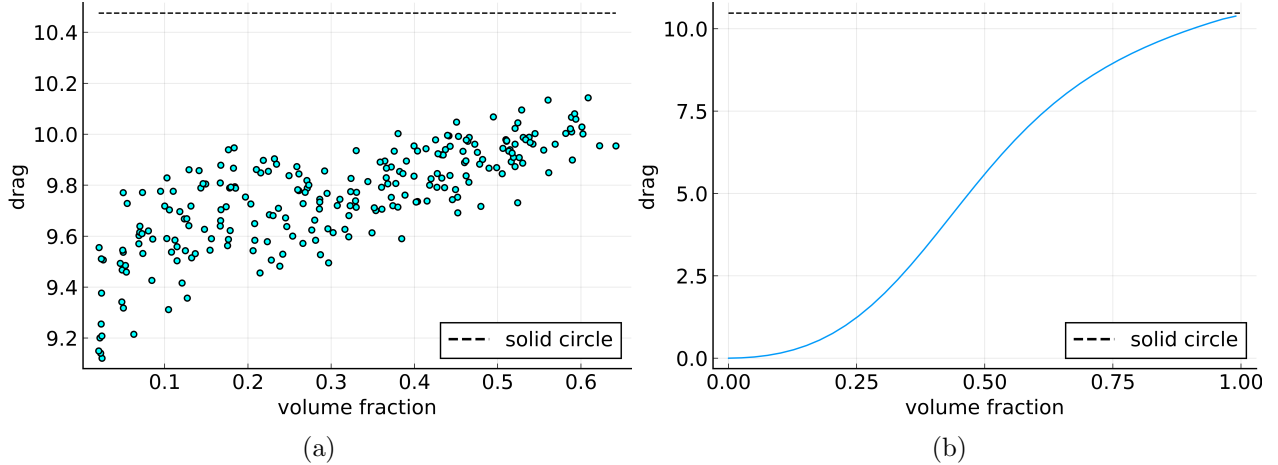


Figure 3.8: Drag on porous body versus volume fraction. The “porous body” is (a) composed of 40 “filling circles” and (b) represented by a volume fraction scalar field in our domain with $\xi^* = 100$, $\theta_{fluid}^* = 0.6$. The disparate scales of the vertical axis indicates that we should not interpret our porous objects as being composed of small circles, at least in two-dimensions.

not consider the multiphase solid material as being composed of small “filling circles”, at least in two dimensions. With the validation of the model for immobile solid bodies behind us, we now turn to how the framework can be used for applications where a slow-moving solid with complex rheology, such as occurs in geophysical and micro-scale biological phenomena.

3.3 Particles within a step channel

To see some simple effects of the discrete particle model, we demonstrate four separate simulations of a group of particles within a step channel. First, the effect of velocity only is shown – this means that the particles act as tracers, simply following the flow field. Then, the flow field is turned off, but the effects of cohesion, adhesion, and gravity are turned on. This demonstrates how the particles clump and interact with walls realistically. Then, all forces are turned on but only a one-way coupling is included; the fluid affects the particles, but there is no “solid volume fraction” to affect the velocity field in return. Finally, the two-way coupling is included to show how the velocity field within the step channel is affected by the presence of the porous media, and in turn this affects the dynamics of the particles. This final simulation demonstrates the full power of the discrete particle approach.

We will now describe the general algorithm that is used in the “full” case that couples fluid and structure. We begin in all cases by initializing random circles inside a domain, along with a

description of the walls. See Appendix C for a more detailed explanation of these steps, specifically the packages `CircGeometry.jl` and `StokesParticles.jl`.

Before the time-stepping routine of the simulation begins, we generate two cell lists. One cell list is to represent the particle positions for efficient computation of cohesion force, and the second cell list is for finite element mesh nodes for efficient computation of the seepage force. Once these cell lists are initialized, the particle cell list needs to be updated at each time step (because the particles are moving, and therefore might change cells), but the FEM cell list does not need to be updated. The fluid velocity field, being represented on a finite element mesh, must be interpolated for each particle position – see Appendix B for a description of a scheme to do this efficiently.

The walls in Figs. 3.9–3.11 are represented with finite width, which is apparent in the dynamics as well. Once the wall and particle initial conditions are given, we then solve for the initial fluid flow. Flow is coming into the left and exiting the right on the channel; the classical step flow is represented in Fig. 3.10 where we observe the particles acting as tracers for the fluid velocity. With only seepage force turned on, the particle velocity is simply proportional to the fluid velocity.

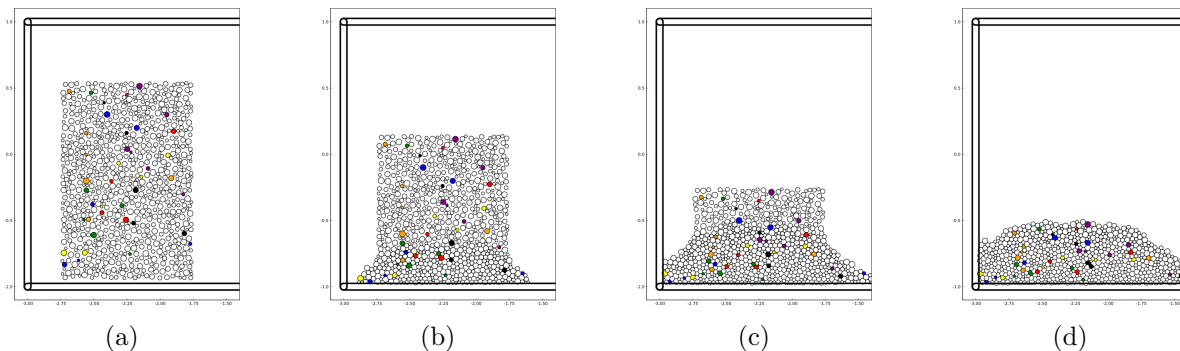


Figure 3.9: Particle dynamics for cohesion, adhesion and gravity, where fluid flow is not incorporated. The particles, while only being represented as points, exhibit concrete radii and allow themselves to pile up next to and on top of walls. Some particles are colored for convenience.

At each time step, we compute the gravity, cohesion, adhesion, and seepage forces on the particles, and use Eq. (2.60) to move the particles in time. In theory, we should be updating the fluid velocity field at each time step, as they are coupled, but in practice the particles move so slowly with each time step (a necessary condition of the stiff nature of the system due to the cohesion force), we only update the velocity field every 100 time steps. The reason for this is that re-building the brinkman system (which is necessary because the solid volume fraction field is constantly changing) and solving the resulting linear system would be too costly for the small

amount of accuracy that would be gained. It is hoped that, in the future, we may make this process more efficient by building an iterative solution scheme that uses past solutions as an initial condition starting point – which would necessarily be faster than solving the system directly every time, as well as build a faster way to update the system that is only changing by small perturbations to the variable coefficient.

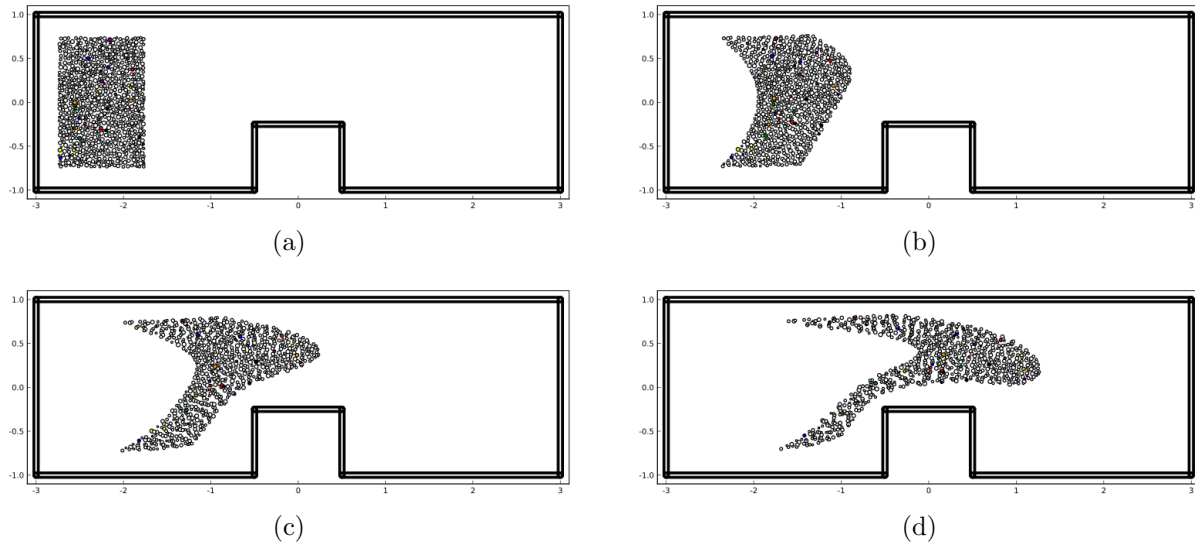


Figure 3.10: Particle dynamics for only fluid flow. In this simulation, particles behave simply as volumeless tracers without any steric interactions. They follow the fluid flow over the step.

The end result of this algorithm is a method for using discrete particle dynamics and, importantly, discrete description of forces coupled to a continuum description of flow. This algorithm will be particularly important for our scientific application of modeling sinkhole formation in Chapter 5. More generally, this algorithm can be applied to any scientific problem where the a solid with complex rheology is interacting with a surrounding fluid flow. We note the important missing piece so far has been that we have not yet coupled together our reaction, phase-change mechanism to the discrete particles. We have only made use of the discrete-to-continuum operator to map the (already existing) discrete particles onto a continuum field representing the volume fraction; we have not yet used the continuum-to-discrete operator that would allow one to fully take advantage of our model framework. In the next section, we give a simple validation of an algorithm to implement this continuum-to-discrete operator.




Figure 3.11: Particle dynamics for (a) one-way fluid coupling, including gravity, cohesion and adhesion and (b) all forces, including a two-way fluid-particle coupling. The cohesion of the particles allows clumps to move coherently as a single (or multiple) bodies. In (a) the top of the initial clump feels the full force of the fluid, and is pushed over the step; the bottom half of the initial clump is overtaken by gravity and falls before reaching the step. In (b) the entire initial clump is overtaken by gravity, unlike Fig. 3.11, because the fluid velocity inside the clump is significantly less due to the permeability of this clump affecting the fluid velocity

3.4 Benchmark for continuum-to-discrete operator

This section provides a validation test for the continuum-to-discrete operator \mathcal{G}^* defined in Sec. 2.3. Neither of the two-dimensional problems in chapters 4 or 5 will require this operator, so a one-dimensional test will suffice as a proof-of-concept. We begin by describing the gradient descent algorithm that is used to optimize the cost function in Eq. (2.65), as well as an analytic description of its gradient with respect to the parameters. This algorithm is tested on an example where the field F is itself generated from a single particle undergoing the forward transform \mathcal{G} . Then, \mathcal{G}^* is applied to a step function to see how it performs in a realistic context. Because the current

framework conserves mass, any error resulting from an application of \mathcal{G}^* violates this principle – in practice, one must consider how much mass loss due to numerical error is acceptable.

3.4.1 Gradient descent algorithm

The gradient descent algorithm is a foundational algorithm for optimization problems [90]. While it has many variations, only its simplest form will be considered here, which is sufficient for benchmark purposes. Needless to say, use of the operator \mathcal{G}^* in production-level code would require a more efficient optimization algorithm.

For a cost function J , parameters $\boldsymbol{\theta}^n$, and “learning rate” α , the optimization algorithm is to repeatedly apply the following formula:

$$\boldsymbol{\theta}^{n+1} \leftarrow \boldsymbol{\theta}^n - \alpha \nabla J(\boldsymbol{\theta}^n) \quad (3.24)$$

where the superscript n refers to the step index. For a sufficiently small α , and certain constraints on J , this algorithm is guaranteed to converge. Use of this algorithm requires the gradient of J with respect to $\boldsymbol{\theta}$. While it is possible to do this numerically using, for example, a finite difference method, that will introduce unnecessary errors into the calculation that could slow or even prevent convergence. Because J is analytically defined, we can find its gradient analytically.

Remember that $\boldsymbol{\theta}$ is essentially a container for (in one spatial dimension) N points x_i , N strength parameters s_i , and one “spreading” parameter h . To compute the gradient of J with respect to $\boldsymbol{\theta}$, we need to take the partial derivative of J with respect to each of these components. Each of these partial derivatives takes the same form:

$$\frac{\partial J}{\partial \square} = 2 \int_a^b [\mathcal{G}(\mathbf{x}, \mathbf{s}, h, y) - F(y)] \frac{\partial \mathcal{G}}{\partial \square} dy, \quad (3.25)$$

where $\square \in \{x_1, \dots, x_N, s_1, \dots, s_N, h\}$ and $\frac{\partial \mathcal{G}}{\partial \square}$ are defined to be

$$\frac{\partial \mathcal{G}}{\partial x_j} = -\frac{2}{h^2} (x_j - y) \exp\left(-\frac{(x_j - y)^2}{h^2}\right), \quad (3.26a)$$

$$\frac{\partial \mathcal{G}}{\partial s_j} = \exp\left(-\frac{(x_j - y)^2}{h^2}\right), \quad (3.26b)$$

$$\frac{\partial \mathcal{G}}{\partial h} = \frac{2}{h^3} \sum_{i=1}^N s_i (x_i - y)^2 \exp\left(-\frac{(x_i - y)^2}{h^2}\right). \quad (3.26c)$$

In practice, each of these partial derivatives must be evaluated at each point y in order to evaluate the integral. So, while this algorithm will give the desired result, it is computationally expensive to run in practice.

3.4.2 Single particle transform

For the most straightforward test, the inverse Gauss transform \mathcal{G}^* should undo the effect of the forward Gauss transform \mathcal{G} . To test this, we take a single particle defined by its center (x_0) and radius (combination of s_0 and h). We then run the above optimization problem from an initialization θ_0 to recover these parameter values. The initial conditions will be $x_0 = -0.3$, $s_0 = 0$, and $h = 0.1$. The “true values” to obtain through optimization are $x_0^* = 0.7$, $s_0^* = 1.2$, and $h^* = 0.5$. With a learning rate of $\alpha = 0.2$, the gradient descent algorithm takes 166 iterations until convergence, where we define convergence to be the iteration where J changes by less than 1×10^{-8} . This convergence is visually depicted in Fig. 3.12.

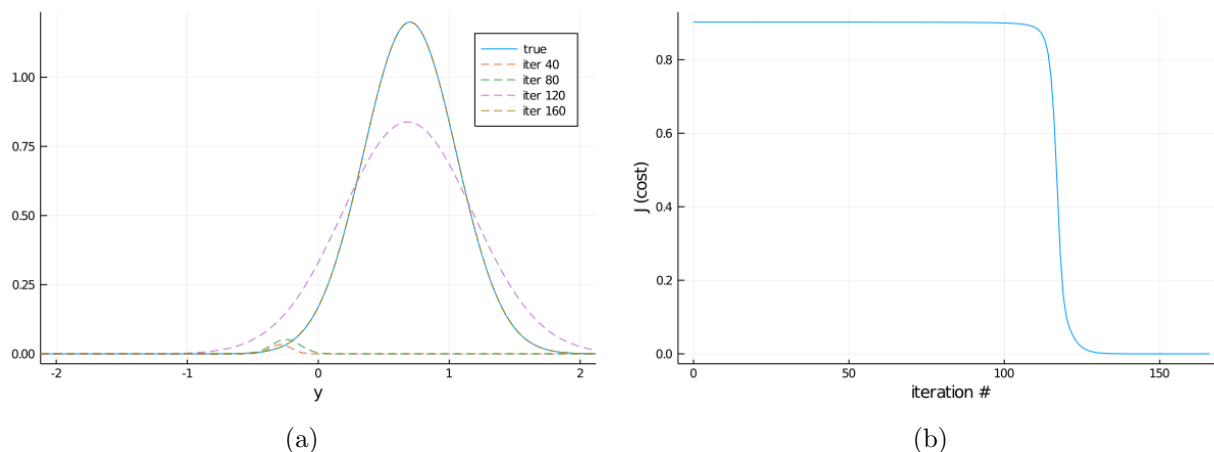


Figure 3.12: Convergence for single particle continuum-to-discrete transform. (a) Physical space convergence and (b) cost function for increasing iterations both demonstrate that the convergence experiences a “transition region” where, before and after, convergence is slow.

As the number of particles that are used to generate the field F increases, the gradient descent algorithm becomes more and more computationally expensive – the learning rate α must shrink more than linearly with the increasing number of particles. This reflects the fact that a simple gradient descent algorithm is used. To compute \mathcal{G}^* on a field generated by hundreds of particles, one should first choose a better optimization algorithm. Now that \mathcal{G}^* has been shown to work in principle, next it is applied to a field F that is *not* generated via \mathcal{G} .

3.4.3 Transform of step function

For the purpose of the eventual application of \mathcal{G}^* to problems using the model framework, we need to know how it performs on a field *not* generated from a Gauss transform of particles. This is

the type of situation that will arise in the context of the current framework; a precipitate field will form from a reaction and a transform to discrete particles will need to occur before the discrete forces description can be used as a model of a solid with complex rheology. For this simple test, we will examine how well \mathcal{G}^* performs on a step function:

$$S(x) = \begin{cases} 1 & -\frac{1}{2} < x < \frac{1}{2} \\ 0 & \text{otherwise} \end{cases}. \quad (3.27)$$

This function is one discrete representation of a Dirac delta distribution, which approximately represents the precipitation that will initially occur at a single point. The algorithm as describes does a *bad* job of approximating the step function – see Fig. 3.13 for one attempt with 20 particles. In addition to suffering from a manifestation of the well-known Gibb’s phenomena, it also converges with negative values of the strength parameter s_i which means the interpretation of the parameter θ loses meaning relating to particles of radius proportional to s_i and h . While the efficiency of \mathcal{G}^* applied to arbitrary fields F is beyond the scope of this dissertation, one way to improve this algorithm would be to define a different cost function that was regularized such that it become extremely large whenever s_i were negative.




Figure 3.13: N particles attempting to approximate a step function via the inverse transform \mathcal{G}^* .

To conclude, while the continuum-to-discrete operator \mathcal{G}^* works in principle for obtaining θ for fields F that are themselves generated with a Gauss transform \mathcal{G} , it performs poorly in practice. In addition to the simple gradient descent algorithm scaling poorly for field F generated with large N , the inverse does a bad job of approximating generic fields F which is the entire point of this operator in the context of this dissertation. Therefore, additional work will need to be done in order to employ this method in production-level simulations.

Part II

Applications

CHAPTER 4

PRECIPITATION WITHIN A MICROFLUIDIC CHAMBER

A specific precipitate reaction within a microfluid device was given in the introduction to this dissertation. This chapter is devoted to using the current model framework to numerically simulate this experiment, represented by Fig. 4.1. The finite element method (FEM) is used to discretize the PDE system. An operator splitting technique is used to efficiently handle multiple timescales. In particular, the reaction operation is approximated by a quasi-steady assumption; this allows one to resolve the reactions analytically with asymptotic accuracy, and the simulation step size can be orders of magnitude larger. The momentum of the fluid in the low Reynolds number regime is resolved via a time-independent, variable-coefficient Brinkman system which, as derived in Chapter 2 and validated in Chapter 3, provides a model for this experimental setup. Simulation software is provided open-source for interested readers; see Appendix C.



Figure 4.1: Microfluidic setup for precipitate reaction. Ionic species enter the tubes on the left, and exit the channel on the right. Reproduced from [109] with permission.

The rise of microfluidic devices offers an economical approach for chemical experiments at the micrometer scale. They have been used to explore precipitation reactions which characteristically produce a fluid-impeding structure. These experiments model, for example, the conditions on primordial earth necessary to create the first “life” molecules [31, 108, 110], as well as so-called chemical gardens

[8]. Although microfluidic experiments have shown great success in analyzing chemical and biological samples [30, 119], all possible permutations of domain geometry, chemistry, and flow conditions creates a parameter space that is prohibitively expensive to probe experimentally. Numerical simulation of a mathematical model offers an inexpensive alternative to experiments, and can guide researchers towards interesting parameter regimes while avoiding a costly trial-and-error approach.


The focus of this dissertation has been to derive a novel model framework applicable to precipitation reactions in microfluidic devices. This model conserves mass throughout the entire reaction, and preserves the fluid-structure interaction behavior resulting from the formation of precipitate

solids. In the present chapter, we develop an efficient numerical discretization scheme and implementation to simulate this PDE system. The algorithms are valid in either two or three spatial dimension, although only two-dimensional domains were simulated for the results. This simulation framework will be useful to experimenters who desire a model of their system that is quicker and cheaper to run than wet-lab experiments.

This chapter proceeds as follows: Sec. 4.1 explicitly states the model to be simulated using the framework from Chapter 2. Sec. 4.2 describes the numerical algorithms, including the algebraic reaction equations, and the FEM discretization of the PDEs; rigorous validation of all schemes is included. Sec. 4.3 shows simulations in various realistic microfluidic geometries and parameter regimes.

4.1 Model and timescale separation

The situation we consider is a precipitation reaction occurring within a microfluidic chamber where the Reynolds number is negligible. Specifically, aqueous species A and B react to form an aqueous product C , which precipitates to form a solid. The chemistry stoichiometry is given by



where the first, purely aqueous, reaction occurs within the solvent, and the second reaction describes the precipitation of the product from liquid to solid. The model is a multiphase PDE system that takes into account the spatial and temporal dynamics of all constituents. Note that the model framework can handle more complex reaction stoichiometry, as outlined in Sec. 2.1.2, but only the simple case is simulated here.

The mathematical system under consideration is a multiphase PDE system where advection-diffusion-reaction (ADR) equations model the aqueous chemistry and a variable-coefficient Brinkman system models the fluid-structure dynamics. We will first go through a nondimensionalization and timescale separation so that the fastest timescale does not place a constraint on the simulation time step size.

4.1.1 Nondimensionalization

Considering realistic parameter values for the above system, one discovers that there are three disparate timescales of the problem; from fastest to slowest, they are: reaction, advection, and diffusion. Because the reaction dynamics are the fastest timescale, we simplify our model considerably by making a quasi-steady state reaction assumption, described in detail in Sec. 4.1.2, that

allows mass transfer between components of the systems to be handled algebraically, as opposed to solving nonlinear ADR equations. This assumption still preserves a fundamental quality of the model: conservation of mass.

The first step is to nondimensionalize the governing equations from the model framework of Chapter 2. The main chemistry variable is $\psi_i^*(\mathbf{x}, t)$ for the molarity of chemical species i in position \mathbf{x} at time t . It is important to remember for physical interpretation that ψ_i^* is the number of i molecules per unit *fluid* volume, rather than unit *total* volume. This distinction is not always recognized in the literature, although it is important for phase-change

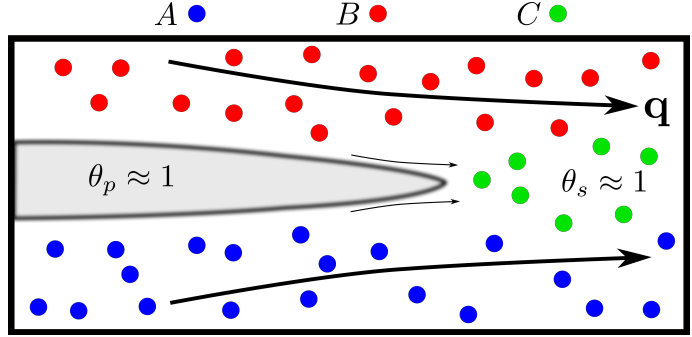


Figure 4.2: Schematic of variables and parameters, with a top-down view of the channel. Precipitate forms in the middle, which affects the surrounding fluid flow.

reactions; the “fluid” volume does not necessarily equal the “total” volume. Additionally, physical processes such as reaction rates are proportional to a species’ *molarity*, not concentration per unit total volume. In dimensional form, the governing system is given by [36]:

$$\theta_s \frac{\partial \psi_A^*}{\partial t^*} = \nabla^* \cdot (\kappa_A^* \nabla^* \psi_A^* - \psi_A^* \mathbf{q}^*) - ar \theta_s \psi_A^* \psi_B^* + \psi_A^* \frac{\partial \theta_s}{\partial t^*} \quad (4.2a)$$

$$\theta_s \frac{\partial \psi_B^*}{\partial t^*} = \nabla^* \cdot (\kappa_B^* \nabla^* \psi_B^* - \psi_B^* \mathbf{q}^*) - br \theta_s \psi_A^* \psi_B^* + \psi_B^* \frac{\partial \theta_s}{\partial t^*} \quad (4.2b)$$

$$\theta_s \frac{\partial \psi_C^*}{\partial t^*} = \nabla^* \cdot (\kappa_C^* \nabla^* \psi_C^* - \psi_C^* \mathbf{q}^*) + cr \theta_s \psi_A^* \psi_B^* + \psi_C^* \frac{\partial \theta_s}{\partial t^*} - \alpha \frac{\partial \theta_s}{\partial t^*} \quad (4.2c)$$

$$\theta_s + \theta_s = 1 \quad (4.2d)$$

$$\frac{\partial \theta_s}{\partial t^*} = \beta \psi_C^* \theta_s \mathcal{H}(\psi_C^* - \psi_C') / \rho_m \quad (4.2e)$$

$$\eta \nabla^{*2} \mathbf{q}^* - \beta^* (\theta_s) \mathbf{q}^* = \theta_s \nabla^* p^* \quad (4.2f)$$

$$\nabla^* \cdot \mathbf{q}^* = 0 \quad (4.2g)$$

where $\alpha = (\rho_p - \rho_s) / M_C$ and * superscript denotes the dimensional form of a variable. The κ_i^* can be written as $\kappa_i^* = \kappa_i^0 \kappa_i(\theta_s)$ to denote both a dimensional, characteristic “diffusivity” κ_i^0 as well as a nondimensional, functional dependence on volume fraction. Note that ψ' is the precipitation concentration threshold, and is dimensional. Using the standard *MLT* (mass, length,

time) formulation for dimensional analysis, the units of all of the variables and parameters are:

$$\begin{aligned} [\theta_i] &= 1, & [\psi_i] &= L^{-3}, & [\mathbf{q}] &= LT^{-1}, & [p] &= ML^{-1}T^{-2}, \\ [\kappa_i] &= L^2T^{-1}, & [a, b, c] &= 1, & [r] &= L^3T^{-1}, & [\beta] &= MT^{-1}, \\ [\alpha] &= L^{-3}, & [\rho_i] &= ML^{-3}, & [\eta] &= ML^{-1}T^{-1}, & [h] &= ML^{-3}T^{-1}. \end{aligned}$$

We introduce the nondimensional variables

$$\begin{aligned} \psi_i &= \frac{\psi_i^*}{\alpha_\psi}, & \mathbf{q} &= \frac{\mathbf{q}^*}{\alpha_q}, & p &= \frac{p^*}{\alpha_p}, \\ t &= \frac{t^*}{\alpha_t}, & \mathbf{x} &= \frac{\mathbf{x}^*}{\alpha_x}. \end{aligned}$$

Note that α_i are dimensional, characteristic constants, and are unrelated to the dimensional parameter α in our system. We also note the following derivative quantities:

$$\frac{\partial}{\partial t^*} = \frac{1}{\alpha_t} \frac{\partial}{\partial t}, \quad \nabla^* = \frac{1}{\alpha_x} \nabla, \quad \nabla^{*2} = \frac{1}{\alpha_x^2} \nabla^2 \quad (4.3)$$

Plugging these into the above system and multiplying through by α_t one obtains

$$\theta_s \frac{\partial \psi_A}{\partial t} = \frac{\alpha_t}{\alpha_x^2} \kappa_A^0 \nabla \cdot (\kappa_A(\theta_s) \nabla \psi_A) - \frac{\alpha_q \alpha_t}{\alpha_x} \nabla \cdot (\psi_A \mathbf{q}) - \alpha_\psi \alpha_t a r \theta_s \psi_A \psi_B + \psi_A \frac{\partial \theta_p}{\partial t}, \quad (4.4a)$$

$$\theta_s \frac{\partial \psi_B}{\partial t} = \frac{\alpha_t}{\alpha_x^2} \kappa_B^0 \nabla \cdot (\kappa_B(\theta_s) \nabla \psi_B) - \frac{\alpha_q \alpha_t}{\alpha_x} \nabla \cdot (\psi_B \mathbf{q}) - \alpha_\psi \alpha_t b r \theta_s \psi_A \psi_B + \psi_B \frac{\partial \theta_p}{\partial t}, \quad (4.4b)$$

$$\theta_s \frac{\partial \psi_C}{\partial t} = \frac{\alpha_t}{\alpha_x^2} \kappa_C^0 \nabla \cdot (\kappa_C(\theta_s) \nabla \psi_C) - \frac{\alpha_q \alpha_t}{\alpha_x} \nabla \cdot (\psi_C \mathbf{q}) + \alpha_\psi \alpha_t c r \theta_s \psi_A \psi_B + \psi_C \frac{\partial \theta_p}{\partial t} - \frac{\alpha}{\alpha_\psi} \frac{\partial \theta_p}{\partial t}, \quad (4.4c)$$

$$\theta_s + \theta_p = 1, \quad (4.4d)$$

$$\frac{\partial \theta_p}{\partial t} = \frac{\alpha_t \alpha_\psi \beta}{\rho_p} \psi_C \theta_s \mathcal{H}(\alpha_\psi \psi_C - \psi_C'), \quad (4.4e)$$

$$\nabla^2 \mathbf{q} - \frac{\alpha_x^2}{\eta} \beta^*(\theta_s) \mathbf{q} = \frac{\alpha_p \alpha_x}{\eta \alpha_q} \theta_s \nabla p, \quad (4.4f)$$

$$\nabla \cdot \mathbf{q} = 0. \quad (4.4g)$$

Note that there are roughly three components to the above ADR equations. Realistic parameter values suggest that both reaction processes occurs orders of magnitude faster than either advection or diffusion processes. This implies that there are two timescales. To make the dependence of the variables on these two timescales explicit, we first choose our timescale α_t to be on the order of the reaction, then create a new “slow” timescale, τ , which will capture the slower advection and diffusion processes. In the next section, we carry out this procedure in great detail.

4.1.2 Timescale separation

We first note that the momentum equations (4.4f) and (4.4g) are independent of time, so we do not need to consider them in the following analysis. To simplify things, we can substitute (4.4e) into (4.4a),(4.4b),(4.4c) to simplify the number of equations we need to consider in our timescale analysis. Additionally, without loss of generality, we do not consider there to be any precipitation threshold, i.e. $\psi'_C = 0$, so that we drop the Heaviside function. To choose our reaction timescale, let $\alpha_t = (\alpha_\psi r)^{-1}$ while assuming that $\beta \approx r\rho_m$ to obtain

$$\theta_s \frac{\partial \psi_A}{\partial t} = \frac{\alpha_t}{\alpha_x^2} \kappa_A^0 \nabla \cdot (\kappa_A \nabla \psi_A) - \frac{\alpha_q \alpha_t}{\alpha_x} \nabla \cdot (\psi_A \mathbf{q}) - a \theta_s \psi_A \psi_B + \psi_A \psi_C \theta_s, \quad (4.5a)$$

$$\theta_s \frac{\partial \psi_B}{\partial t} = \frac{\alpha_t}{\alpha_x^2} \kappa_B^0 \nabla \cdot (\kappa_B \nabla \psi_B) - \frac{\alpha_q \alpha_t}{\alpha_x} \nabla \cdot (\psi_B \mathbf{q}) - b \theta_s \psi_A \psi_B + \psi_B \psi_C \theta_s, \quad (4.5b)$$

$$\theta_s \frac{\partial \psi_C}{\partial t} = \frac{\alpha_t}{\alpha_x^2} \kappa_C^0 \nabla \cdot (\kappa_C \nabla \psi_C) - \frac{\alpha_q \alpha_t}{\alpha_x} \nabla \cdot (\psi_C \mathbf{q}) + c \theta_s \psi_A \psi_B + \psi_C^2 \theta_s - \frac{\alpha}{\alpha_\psi} \psi_C \theta_s. \quad (4.5c)$$

Now let $\epsilon = \alpha_q \alpha_t / \alpha_x$ be a small parameter ($\epsilon \ll 1$) to reduce the nondimensional coefficient for advection. Then, we can also represent the nondimensional diffusion coefficient in terms of ϵ by introducing a new nondimensional parameter called the Péclet number

$$\text{Pe}^i = \frac{\alpha_q \alpha_x}{\kappa_i^0} \quad (4.6)$$

which represents the relative effects of advection and diffusion. The system is now

$$\theta_s \frac{\partial \psi_A}{\partial t} = \frac{\epsilon}{\text{Pe}^A} \nabla \cdot (\kappa_A(\theta_s) \nabla \psi_A) - \epsilon \nabla \cdot (\psi_A \mathbf{q}) - a \theta_s \psi_A \psi_B + \psi_A \psi_C \theta_s, \quad (4.7a)$$

$$\theta_s \frac{\partial \psi_B}{\partial t} = \frac{\epsilon}{\text{Pe}^B} \nabla \cdot (\kappa_B(\theta_s) \nabla \psi_B) - \epsilon \nabla \cdot (\psi_B \mathbf{q}) - b \theta_s \psi_A \psi_B + \psi_B \psi_C \theta_s, \quad (4.7b)$$

$$\theta_s \frac{\partial \psi_C}{\partial t} = \frac{\epsilon}{\text{Pe}^C} \nabla \cdot (\kappa_C(\theta_s) \nabla \psi_C) - \epsilon \nabla \cdot (\psi_C \mathbf{q}) + c \theta_s \psi_A \psi_B + \psi_C^2 \theta_s - \frac{\alpha}{\alpha_\psi} \psi_C \theta_s, \quad (4.7c)$$

so that as long as Pe^i is not large, there are only two timescales in the problem. To make this formal, let $\tau = \epsilon t$ and expand the time derivatives using the multivariable chain rule, and group by

orders of ϵ to obtain

$$\mathcal{O}(1) \begin{cases} \theta_s \frac{\partial \psi_A}{\partial t} = -a\theta_s \psi_A \psi_B + \psi_A \psi_C \theta_s \\ \theta_s \frac{\partial \psi_B}{\partial t} = -b\theta_s \psi_A \psi_B + \psi_B \psi_C \theta_s \\ \theta_s \frac{\partial \psi_C}{\partial t} = c\theta_s \psi_A \psi_B + \psi_C^2 \theta_s + \frac{\alpha}{\alpha_\psi} \psi_C \theta_s \end{cases} \quad (4.8)$$

$$\mathcal{O}(\epsilon) \begin{cases} \frac{\partial}{\partial \tau} (\theta_s \psi_A) = \frac{1}{\text{Pe}^A} \nabla \cdot (\kappa_A(\theta_s) \nabla \psi_A) - \nabla \cdot (\psi_A \mathbf{q}) \\ \frac{\partial}{\partial \tau} (\theta_s \psi_B) = \frac{1}{\text{Pe}^B} \nabla \cdot (\kappa_B(\theta_s) \nabla \psi_B) - \nabla \cdot (\psi_B \mathbf{q}) \\ \frac{\partial}{\partial \tau} (\theta_s \psi_C) = \frac{1}{\text{Pe}^C} \nabla \cdot (\kappa_C(\theta_s) \nabla \psi_C) - \nabla \cdot (\psi_C \mathbf{q}) \end{cases} \quad (4.9)$$

This is form is where one can apply the quasi-steady assumption, such that the $\mathcal{O}(1)$ equations can be sufficiently approximated by steady states after undergoing the reaction to completion. The $\mathcal{O}(\epsilon)$ equations on the τ timescale will be numerically simulated. The practical effect is that only one timescale needs to be included in the simulation, and three nonlinear, coupled equations have become uncoupled and linear which has considerable numerical benefits.

The nondimensional variable-coefficient Brinkman equation is

$$\nabla^2 \mathbf{q} - \beta(\theta_s) \mathbf{q} = \theta_s \nabla p \quad (4.10a)$$

$$\nabla \cdot \mathbf{q} = 0 \quad (4.10b)$$

where β is the reciprocal of the effective viscosity (which includes the friction relation), and p is the spatially-averaged pressure which has already been rescaled by the (negligible) Reynolds number.




Figure 4.3: Schematic representing idea that aqueous chemical dynamics reach steady state on every reaction-diffusion time step $\Delta\tau$.

4.2 Numerical discretization

The numerical framework discretizes a domain into finite elements. We first describe the algorithm that allows us to handle the reaction operations algebraically, as described theoretically in the previous section. Next the solution of the advection-diffusion differential equations is described. Finally, the method for simulating a variable-coefficient Brinkman equation is detailed. For the FEM discretizations, we provide validation using the method of manufactured solutions.

4.2.1 Quasi-steady reaction

As stated previously, and formally derived in 4.1.2, the reaction occurs orders of magnitude faster than any other process. This allows the Quasi-Steady State Reaction (QSSR) assumption to be made. Under this assumption we can make the approximation that, within each finite time interval $\Delta\tau$, the reaction at each point in space goes to completion (see Fig. 4.3). In other words, all of the mass from aqueous chemicals A and B react into chemical C mass, and then (if the precipitation threshold is exceeded) precipitate mass. The obvious advantages of this approximation is that only one timescale needs to be simulated, and that the nonlinear reaction terms can be handled algebraically, instead of being solved as a part of the differential equation system. The final algorithm is given in pseudo-code in Alg. 3.

First, we must determine whether the reaction occurs at all. To do this, determine whether there is any solid volume fraction; pre-existing precipitate acts as a “nucleation site” whose presence overrides the threshold requirement. If there is no precipitate, then one must further check whether the ψ_C exceeds its precipitation threshold. If either of the above conditions are satisfied, proceed with the reaction calculation.

To perform the calculation, one must determine the limiting reagent. Compare the values of ψ_A/a and ψ_B/b ; the smaller quantity is the limiting reagent. To simplify the following derivation, we will assume that A is the limiting reagent, i.e. $\psi_A/a \leq \psi_B/b$. The algorithm for the case that B is the limiting reagent follows by symmetry. We now summarize the derivation. Because A is the limiting reagent, and we have made the QSSR assumption, there will be no A molecules after a time step. This initial number of molecules of A will determine the number of B molecules that undergo reaction, adjusting for stoichiometry. Finally, the combined mass from A and B that undergoes reaction will eventually become precipitate mass. In summary, because we make the QSSR assumption, our problem is accounting for “mass before” and “mass after” the time step without worrying about the intermediate dynamics. See Alg. 1 for these steps listed in pseudo-code.

Algorithm 1: aqueous_reaction

```

input :  $\{\psi_i^n\}, \{M_i^n\}, a, b$ 
 $\mathcal{M} \leftarrow M_A\psi_A^n + M_B\psi_B^n + M_C\psi_C^n;$ 
if  $\psi_A^n/a \leq \psi_B^n/b$  then
     $\psi_A^{n+1/2} \leftarrow 0;$ 
     $\psi_B^{n+1/2} \leftarrow \psi_B^n - \frac{b}{a}\psi_A^n;$ 
     $\psi_C^{n+1/2} \leftarrow \frac{1}{M_C} \left( \mathcal{M} - M_B\psi_B^{n+1/2} \right);$ 
else
     $\psi_A^{n+1/2} \leftarrow \psi_A^n - \frac{a}{b}\psi_B^n;$ 
     $\psi_B^{n+1/2} \leftarrow 0;$ 
     $\psi_C^{n+1/2} \leftarrow \frac{1}{M_C} \left( \mathcal{M} - M_A\psi_A^{n+1/2} \right);$ 
end
output:  $\{\psi_i^{n+1/2}\}$ 

```

Algorithm 2: precipitate_reaction

input : $\{\psi_i^n\}, \{\theta_i^n\}, \{M_i\}, \{\rho_i\}, a, b$
 $\Delta\rho \leftarrow \rho_p - \rho_s$ $\mathcal{M} \leftarrow$
 $\Delta\rho\theta_p^n + (M_A\psi_A^n + M_B\psi_B^n + M_C\psi_C^n)\theta_s^n;$
if $\psi_A^n/a \leq \psi_B^n/b$ **then**
 $\psi_A^{n+1} \leftarrow 0;$
 $\phi_B^{n+1} \leftarrow (\psi_B^n - \frac{b}{a}\psi_A^n)\theta_s^n;$
 $\psi_C^{n+1} \leftarrow 0;$
 $\theta_p^{n+1} \leftarrow (\mathcal{M} - M_B\phi_B^{n+1})/\Delta\rho;$
 $\psi_B^{n+1} \leftarrow \phi_B^{n+1}/(1 - \theta_p^{n+1})$
else
 $\phi_A^{n+1} \leftarrow (\psi_A^n - \frac{a}{b}\psi_B^n)\theta_s^n;$
 $\psi_B^{n+1} \leftarrow 0;$
 $\psi_C^{n+1} \leftarrow 0;$
 $\theta_p^{n+1} \leftarrow (\mathcal{M} - M_A\phi_A^{n+1})/\Delta\rho;$
 $\psi_A^{n+1} \leftarrow \phi_A^{n+1}/(1 - \theta_p^{n+1})$
end
 $\theta_s^{n+1} \leftarrow 1 - \theta_p^{n+1};$
output: $\{\psi_i^{n+1}\}, \{\theta_i^{n+1}\}$

Now we describe the conversion of aqueous mass to precipitate mass; psuedo-code is given in Alg. 2. First, we introduce the notation $\phi_i = \psi_i\theta_s$ for the number of molecules of i per unit *total* volume. This will allow us to think about the number of molecules undergoing reaction, without worrying about the phase dynamics (i.e. volume fraction changes). In words, we know the variables at time τ_n , and wish to know the variables $\psi_A^{n+1}, \psi_B^{n+1}, \psi_C^{n+1}$, and θ_s^{n+1} at time $\tau_{n+1} = \tau_n + \Delta\tau$. Immediately, we know $\psi_A^{n+1} = 0$ because A is the limiting reagent, and $\psi_C^{n+1} = 0$ because all reactions proceed quickly enough to wind up as membrane without any intermediate step, within the discrete time step taken in a simulation.

Therefore, we really only need to solve for ψ_B^{n+1} and θ_s^{n+1} . Equivalently, of course, we could find θ_p^{n+1} . In words, we know that

$$\frac{\text{moles of } A \text{ lost}}{a} = \frac{\text{moles of } B \text{ lost}}{b},$$

and

$$\text{mass of } A \text{ lost} + \text{mass of } B \text{ lost} = \text{mass of precipitate gained}.$$

We know that, because A is the limiting reagent and the reaction proceeds until A is gone, the number of moles of A lost is simply the initial amount, i.e. ϕ_A^n . The moles of B lost is the difference in moles between time steps, or $\phi_B^n - \phi_B^{n+1}$. Therefore we can write Eq. (4.2.1) as

$$\frac{\phi_A^n}{a} = \frac{\phi_B^n - \phi_B^{n+1}}{b} \quad (4.11)$$

which we can solve for our unknown:

$$\phi_B^{n+1} = \phi_B^n - \frac{b}{a}\phi_A^n. \quad (4.12)$$

So now our job is to find the amount of mass being transferred, and then we will be able to work backwards to find θ_p^{n+1} . To do this, we have a simple accounting where mass before equals mass

after. Or, written out,

$$\rho_s \theta_s^n + \rho_p \theta_p^n + \sum_i M_i \psi_i^n \theta_s^n = \rho_s \theta_s^{n+1} + \rho_p \theta_p^{n+1} + \sum_i M_i \psi_i^{n+1} \theta_s^{n+1}.$$

where ρ_s and ρ_p are the mass densities of the pure fluid and solid, respectively. Knowing that $\psi_A^{n+1} = \psi_C^{n+1} = 0$, we can use Eq. (4.12) to obtain

$$\theta_p^{n+1} = \frac{J - M_B \phi_B^{n+1}}{\rho_m - \rho_p}, \quad J = (\rho_m - \rho_p) \theta_p^n + (M_A \psi_A^n + M_B \psi_B^n + M_C \psi_C^n) \theta_s^n \quad (4.13)$$

Note that this reaction means we cannot have $\rho_m = \rho_p$. Does this make physical sense? Yes, because our model assumption is that the membrane material is composed of sequestered solvent and precipitated chemicals. Therefore, the mass density of the solvent will always be more than the mass density of the membrane. If they are equal, then there are no chemicals to react, and none of this procedure applies anyway.

One thing that we have not considered so far, is the possibility that after the reaction $\theta_p > 1$. This corresponds to the unphysical situation where there was so much chemical A and B in the solvent that, if the reaction were to run to completion, then we do the same procedure but assume that $\psi_A^{n+1} \neq 0$ and $\theta_p^{n+1} = 1$. This gives us

$$\phi_A^{n+1} = \frac{Q - M_B (\phi_B^n - \frac{b}{a} \phi_A^n)}{M_A \left(1 - \frac{b}{a} \frac{M_B}{M_A}\right)} \quad (4.14)$$

$$\phi_B^{n+1} = \phi_B^n - \frac{b}{a} \quad (4.15)$$

where we now have

$$Q = \rho_m \theta_p^n + \rho_p \theta_s^n + M_A \phi_A^n + M_B \phi_B^n$$

Note that, because $\psi_i = \phi_i / \theta_s$, and $\theta_s = 0$, this does not make any mathematical sense and the variable ψ_i has lost meaning. Still, this result holds. In practice, the advective and diffusive fluxes would account for the impermeable (or semi-permeable) solid, and would prohibit this situation from arising.

4.2.2 Weak formulation of PDEs

Advection-diffusion equation. The advection-diffusion equation, written generically for an unknown scalar function c is

$$\mathbf{w} \cdot \nabla c - \nabla \cdot (\kappa \nabla c) = f \quad (4.16a)$$

$$c = g_D \quad \text{on } \Gamma_D \quad (4.16b)$$

$$(\mathbf{w}c - \kappa \nabla c) \cdot \mathbf{n} = g_R \quad \text{on } \Gamma_R \quad (4.16c)$$

Algorithm 3: QSSR complete reaction

```

input :  $\{\psi_i^n\}, \{\theta_i^n\}, \{M_i^n\}, a, b$ 
 $\theta_s^n = 1 - \theta_p^n;$ 
 $J = (\rho_m - \rho_p)\theta_p^n + (M_A\psi_A^n + M_B\psi_B^n)\theta_s^n;$ 
if  $\theta_p = 0$  then
   $\{\psi_i^{n+1/2}\} \leftarrow \text{aqueous\_reaction}(\{\psi_i^n\}, \{M_i^n\}, a, b);$ 
  if  $\psi_C^{n+1/2} \geq \psi'_C$  then
     $\{\psi_i^{n+1}\}, \{\theta_i^{n+1}\} \leftarrow \text{precipitate\_reaction}(\{\psi_i^n\}, \{\theta_i^n\}, \{M_i\}, \{\rho_i\}, a, b)$ 
  end
else
   $\{\psi_i^{n+1}\}, \{\theta_i^{n+1}\} \leftarrow \text{precipitate\_reaction}(\{\psi_i^n\}, \{\theta_i^n\}, \{M_i\}, \{\rho_i\}, a, b)$ 
end
output:  $\{\psi_i^{n+1}\}, \{\theta_i^{n+1}\}$ 

```

where κ is the variable diffusivity and the flow field \mathbf{w} is incompressible ($\nabla \cdot \mathbf{w} = 0$). The boundary Γ_D corresponds to Dirichlet conditions and Γ_R has the Robin boundary condition. The above PDE can be represented in weak form with scalar trial function v by

$$\int_{\Omega} v \mathbf{w} \cdot \nabla c \, d\mathbf{x} + \kappa \int_{\Omega} \nabla c \cdot \nabla v \, d\mathbf{x} = \int_{\Omega} f v \, d\mathbf{x} + \int_{\Gamma_R} v g_R \, ds. \quad (4.17)$$

A detailed derivation can be found in any introductory textbook on the finite element method; for example see [54, 39]. Note that we do *not* employ the now-standard streamline-diffusion method that provides stability for advection-dominated problems. Simulation of advection-dominated flows should use some stabilization scheme.

Time-stepping. Upon discretization of the time-dependent advection-diffusion equation above, one obtains a linear system of the form

$$B\dot{\mathbf{u}} + A\mathbf{u} = \mathbf{f} \quad (4.18)$$

where A, B are matrices and \mathbf{f} is a vector. This can be solved in time by the implicit Crank-Nicholson (CN) method [54, p. 154]:

$$\left(B + \frac{\Delta\tau}{2}A\right) \mathbf{u}^n = \left(B - \frac{\Delta\tau}{2}A\right) \mathbf{u}^{n-1} + \Delta\tau (\mathbf{f}^n - \mathbf{f}^{n-1}) \quad (4.19)$$

which is $\mathcal{O}(\Delta\tau^2)$ in time and unconditionally stable; this absence of a stability criteria makes (CN) preferred over explicit methods such as Forward Euler or even Runge-Kutta algorithms.

Variable-coefficient brinkman FEM. Our model for the fluid-structure dynamics is a multiphase form of the zero-Reynolds number momentum equation. Due to the solid phase being

assumed immobile and an unconventional choice of stress tensor, one can end up with a variable-coefficient version of the well-known Brinkman equation; see Chapter 2 or [36] for a discussion on the validity of this choice. To simulate this equation, we use the weak form of Eq. 4.10:

$$-\nabla^2 \mathbf{q} + \beta(\theta_s) \mathbf{q} + \theta_s \nabla p = \mathbf{f} \quad (4.20a)$$

$$-\nabla \cdot \mathbf{u} = 0 \quad (4.20b)$$

which is given by

$$\begin{aligned} \int_{\Omega} \nabla \mathbf{q} : \nabla \mathbf{v} \, d\mathbf{x} + \int_{\Omega} \beta(\theta_s) \mathbf{q} \cdot \mathbf{v} \, d\mathbf{x} - \int_{\Omega} p(\nabla \theta_s \cdot \mathbf{v}) \, d\mathbf{x} - \int_{\Omega} \theta_s p(\nabla \cdot \mathbf{v}) \, d\mathbf{x} \\ = \int_{\Omega} \mathbf{f} \cdot \mathbf{v} \, d\mathbf{x} + \int_{\partial\Omega} (\nabla \mathbf{q} \cdot \mathbf{v}) \cdot \mathbf{n} \, dS - \int_{\partial\Omega} (\theta_s p \mathbf{v}) \cdot \mathbf{n} \, dS, \end{aligned} \quad (4.21a)$$

$$-\int_{\Omega} (\nabla \cdot \mathbf{u}) \phi \, d\mathbf{x} = 0 \quad (4.21b)$$

which can be solved for \mathbf{q} and p , with test functions \mathbf{v} and ϕ from appropriate function spaces. The forcing \mathbf{f} is preserved in order to use the method of manufacturing solutions to validate our numerical code in the next section.

4.2.3 Validation

This section contains validations for the finite element methods described above. There is no numerical validation for the QSSR algorithms. Because the reaction equations are algebraic, one simply needs to verify that mass before and mass after a reaction is conserved up to machine precision.

Steady advection-diffusion equation. For validation purposes, the domain is a rectangle $(x, y) \in [-2, 2] \times [-1, 1]$, and all boundaries are Dirichlet, i.e. $\Gamma_R = \emptyset$. Using the method of manufactured solutions, the solution, velocity parameter, and forcing function are:

$$c(x, y) = x^4 y^4 \quad (4.22a)$$

$$w_x(x, y) = 2y(4 - x^2) \quad (4.22b)$$

$$w_y(x, y) = -2x(1 - y^2) \quad (4.22c)$$

$$\kappa = 2.3 \quad (4.22d)$$

$$f(x, y) = 8x^3 y^5 (4 - x^2) - 8x^5 y^3 (1 - y^2) - 12\kappa x^2 y^2 (x^2 + y^2) \quad (4.22e)$$

where we note our flow field $\mathbf{w} = (w_x, w_y)$ is incompressible. The convergence results are shown in Fig. 4.4(a), where we see the expected convergence rates. These convergence rates are for the

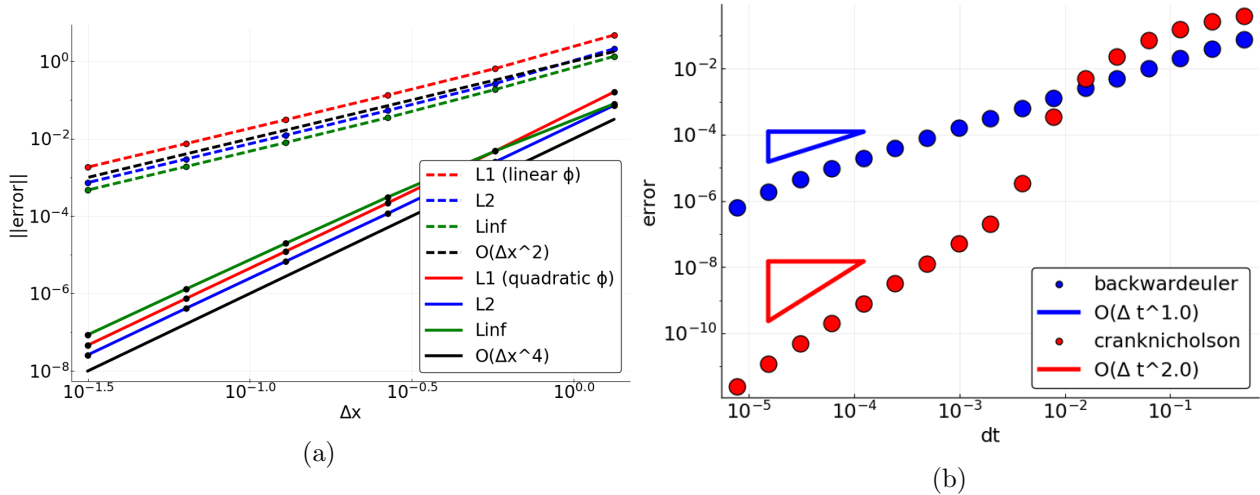


Figure 4.4: Validation of time-dependent advection-diffusion equation. (a) Convergence study for steady advection-diffusion equation using method of manufactured solution on rectangular domain. (b) Convergence study for time stepping on diffusion equation (no advection). This convergence study was conducted in the “Y-shaped” domain, and all errors were taken after one time unit of simulation. Both errors are relative.

steady state advection-diffusion equation. Because we do not use a stabilization term, increasing Pe will increase the error for a given mesh.

Time stepping. Two schemes are tested to approximate the time derivative in the ADR equations. We test Backward Euler and Crank-Nicholson schemes, which are first and second order accurate, respectively. Using the relative error in a Y-shaped domain (see Fig. 4.6) demonstrates these expected convergence rates – see Fig. 4.4(b). The CN scheme is used in the simulations of Sec. 4.3.

Brinkman equation. The Brinkman code is tested with the following manufac-




Figure 4.5: Convergence study for the variable-coefficient multiphase Brinkman equation, using the method of manufactured solutions. Results show second-order convergence in space for the velocity variable, with first-order convergence for the pressure variable, as expected.

tured solution:

$$\begin{aligned}
 q_x(x, y) &= -x^4y^2, & q_y(x, y) &= 4x^3y^3/3, \\
 p(x, y) &= x^3y^3, & \beta(\theta_s)(x, y) &= 0.5 + 0.4 \cos(xy), \\
 \theta_s(x, y) &= 0.5 + 0.4 \sin(7xy), & f_i(x, y) &= -\nabla^2 q_i + \beta(\theta_s)q_i + \theta_s \partial_i p,
 \end{aligned}$$

where $\mathbf{q} = \langle q_x, q_y \rangle$ and $\mathbf{f} = \langle f_x, f_y \rangle$. The convergence results are shown in Fig. 4.5. Because the velocity variables use quadratic elements, and pressure uses linear elements, their respective convergence rates are second and first order, respectively, as the mesh is refined. The pressure being chosen one degree less than the velocity is chosen for stability criterion of the Taylor-Hood finite element family. This family is commonly employed for Stokes equations, and has here been adapted to the Brinkman system.

4.3 Results

This section will detail simulation results in a realistic, Y-shaped domain with wavy outflow. This domains are visualized in Fig. 4.6. All of the simulations of this chapter were implemented and run using the open-source packages; see Appendix C for more details.

4.3.1 Wavy domain geometry

To showcase the ability of our framework and numerical method to handle complex, realistic geometries, we will run a full simulation in a “wavy” Y-shaped domain. Simulation results are shown in Fig. 4.7. Flow is pumped from left to right; aqueous chemicals A and B are input on the bottom and top inflow channels, respectively. The Pe for this simulation is 30, reflecting an advection-dominated scheme that is nonetheless within our accuracy requirements for our discretization method. As the aqueous chemicals meet in the middle of the channel, they react and, because of the lack of a precipitation threshold, instantly precipitate to form solid. As the solid volume fraction increases, the flow velocity field responds by going “around” the developed membrane. the precipitate follows the curves of the domain, as is expected. For this first simulation, we do not include a volume-fraction-dependent diffusivity, which leads to bidirectional membrane growth.

This would be a very difficult problem to simulate using existing mathematical methods detailed in Sec. 1.2 that required explicit, *a priori* knowledge of where the precipitate would appear. Because the emphasis of this domain is on the method itself, and not necessarily on scientific investigation, we do not comment further on the solution. In the next section we will give a brief description of

a way simulations *can* be used for scientific inquiry to aid experimentalists in answering questions that are either too technically involved, or impossible, to conduct with current experiments.

4.3.2 Hypothesis for unidirectional growth

One important experimental observation given in Sec. 1.1 was that, for certain precipitate reactions, lateral membrane growth is asymmetric. This is attributed to the precipitate membrane being selectively permeable to negatively charged ions. Because the permeability of the membrane cannot be measured directly, this is an excellent case study for using the current model framework to aid scientific investigation.

A straight Y-shaped domain is used. This is similar similar to the previous section but without the wavy outflow. In this simulation the diffusion term κ_i will be variable. This diffusion term should be linked to the volume fraction, but the exact relationship is unclear and for us to decide. One method to obtain unidirectional growth is to allow the aqueous chemical to diffuse across solid. The most straightforward expression for Péclet number is then:

$$\text{Pe}' = \frac{\text{Pe}}{\theta_p + z\theta_p}, \quad (4.23)$$

where Pe' is the “normalized” Péclet number and $z \ll 1$. This definition is useful because, for a pure fluid we have $\text{Pe}' = \text{Pe}$, but for pure precipitate we have $\text{Pe}' = \text{Pe}/z$. Therefore, if $z = 0$, the precipitate is impermeable (infinite Pe implies zero diffusivity) and for $0 < z < 1$ diffusion is hindered but non-zero.

An alternative hypothesis for how one obtains unidirectional growth is the following, which we will call hypothesis 2 (the currently accepted hypothesis we call hypothesis 1). Suppose the membrane is permeable to all ions, but in unequal amounts. This is different than being impermeable, as in isolation either ion could pass through unhindered. The reason this, potentially, could result in asymmetric growth is that a reaction is always occurring at the interface, so even though in isolation the slower moving ions could pass through, they react to form the product before they are able to cross.




Figure 4.6: Schematic of wavy domain. The generality of the FEM scheme allows us to run simulations on domains with realistic geometries.

Our exploration of these two hypothesis are using $z = 0.5$ for species A , and $z = 0$ (hypothesis 1) and $z = 0.25$ (hypothesis 2) for species B . In the simulations of hypothesis 1, we see unidirectional growth in the direction of the B -side of the channel; this agrees with both experiment and our expectation of what should happen. In hypothesis 2, which is the real object of our small study, we see uneven *bidirectional* growth – growth occurs faster in the direction of B , but there is still some lateral membrane growth in the direction of the A side of the channel. These results lend further evidence to the hypothesis that the membrane really is impermeable to one charge. This study, while simple, would be impossible to conduct with experiment as one cannot modify diffusivity experimentally (using the same reactants). Therefore, our framework and simulations have lended aid to experimentalists and domain-specific scientific inquiry.

4.4 Conclusion

These simulations validate the framework derived in Chapter 2 for the specific application of precipitation reactions in microfluidic channels. While time-splitting of advection-diffusion-reaction equations is not a new idea [113, 29], to our knowledge, this is the first use of QSSR applied to ADR equations while also handling the reaction terms algebraically. This significantly decreased the computational cost of simulation for two reasons: (1) simulation of the model equations is only done on the advection-diffusion timescale, skipping over the much shorter reaction timescale while achieving asymptotically similar results and conserving mass exactly, and (2) the nonlinear, coupled PDEs were made linear and independent. The benefit of using this framework and numerical discretization is demonstrated in this section. In particular, its ability to simulate both chemistry and fluid-membrane interaction in realistic microfluidic geometries. This chapter also provided further evidence to the hypothesis for how unidirectional growth is achieved in a manner that would be impossible to conduct experimentally. This chapter supports the use of our model in microfluidic experiments by experimentalists to aid their intuition and probe questions that are impossible, difficult, or simply expensive to conduct experimentally.




Figure 4.7: Flow at 4 different time points for wavy-shaped domain. Left side represents solid volume fraction (red = 1, blue = 0) along with velocity vector field, right side shows velocity magnitude; rows are at equivalent times. One can see the velocity field going to no-slip as the precipitate develops. Additionally, the precipitate follows the curves of the domain, as is expected. This would be a very difficult problem to simulate using a method that required explicit knowledge of where the precipitate would appear.

CHAPTER 5

SINKHOLE FORMATION

In this chapter, the model framework is adapted beyond the motivating experiment. This will demonstrate the framework’s ability to model porous media with complex rheology. Several famous sinkholes demonstrate the enormous effect these geological phenomena can have on urban developments. Significantly, sinkhole formation can be aggravated by human factors. Understanding these factors is a major motivation for scientific inquiry [59, 121].

In 2010 a sinkhole opened up in Guatemala City that swallowed a three-story factory [105]. In 1980 a drill poked a small hole into a salt mine, which unfortunately was lying below a pond. The salt was dissolved, forming a massive sinkhole in Lake Peigneur [6]. Sinkholes have been the focus of PBS documentaries¹ and have even formed on the White House lawn [60].

Sinkholes form when surface water (either from rain, rivers, or underground aquifers) seeps through cracks in the bedrock. This flow of water gently erodes the bedrock over time. When the bedrock is weakened to the point that it cannot support the earth above it, a collapse occurs, which we see as a small – or large! – divit on the earth’s surface.

The study of geophysical processes has been conducted numerically [3, 2, 4] and experimentally [83, 115, 102]. Fieldwork on active sinkholes is possible using ground-penetrating waves (laser, radar) to detect the sub-terranean growth of known sinkholes in real time [47, 52]. Labwork has also been conducted to examine how sinkholes might form in model setups.²

In this chapter the model framework of Chapter 2 is adapted to study geophysical processes, specifically sinkhole formation. Notably, this simulation will not include reaction terms, but will include the discrete particle approach to modeling complex media. The chapter proceeds as follows: Sec. 5.1 describes how to adapt the model framework of Chapter 2 to study geophysical phenomena. In Sec. 5.2 the specific algorithms to make simulations feasible are described. Finally in Sec. 5.3 simulation results are shown.

¹Sinkhole documentary: <https://www.pbs.org/wgbh/nova/video/sinkholesburied-alive>

²Video of one experimental setup: https://www.youtube.com/watch?v=1d_5yLnlpA0

5.1 Geophysical model

The model will consist of two phases: fluid and porous material. The porous material is a combination of soil, clay, sand, and rock; these materials are all combined into a single phase for model simplicity. Each point in space is occupied by some fraction of these components such that

$$\theta_s + \theta_p = 1, \quad (5.1)$$

where θ_s and θ_p represent the volume fraction of fluid (or “solvent”) and porous material, respectively. The fluid will be represented by a continuum, and the porous material will be represented by discrete particles. We assume that all material motion is slow enough that a quasi-immobile assumption on the porous material is valid, and therefore the fluid Darcy velocity will be incompressible and obey the multiphase Brinkman equation described in Chapter 2 with the KC friction term:

$$\nabla^2 \mathbf{q} - h \frac{\theta_s^2}{(1 - \theta_s)^2} \mathbf{q} = \theta_s \nabla p, \quad (5.2a)$$

$$\nabla \cdot \mathbf{q} = 0, \quad (5.2b)$$

where \mathbf{q} is the fluid Darcy velocity, h is the friction coefficient, and p is the common pressure experienced by all multiphase components. The porous material, modeled with discrete particles, is subject to the force of gravity, cohesion, adhesion, and seepage.

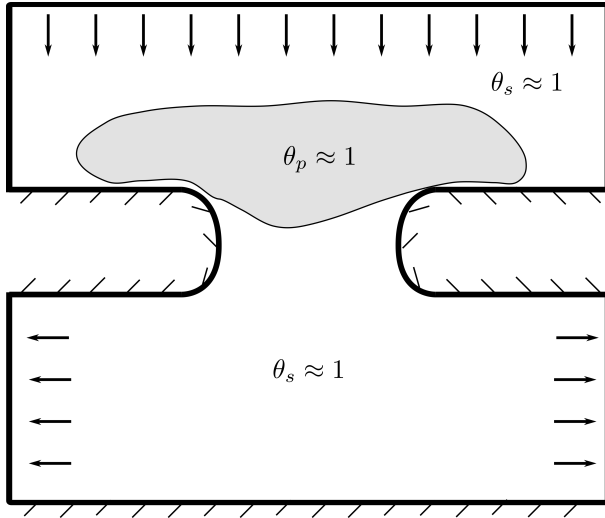


Figure 5.1: Sinkhole schematic with boundary and initial conditions. Not to scale.

The domain itself will be two-dimensional and idealized to simplify analysis. A sketch of the domains and boundaries is given in Fig. 5.1 and the finite element discretization of the domain can be seen in Fig. 5.2. The middle and bottom boundaries will be considered walls from which particles will experience an adhesion force. For the fluid, the middle and bottom boundaries will be no-slip, no-penetration walls. The left and right boundaries will be no-flux and the top will have uniform inflow. The strength of this inflow will be one parameter that is varied to influence sinkhole formation.

The variable unknowns, for all space \mathbf{x} and time t , are θ_s , θ_p , p , and \mathbf{q} . The associated parameters needed are the mass densities of fluid and porous material, ρ_s , and ρ_p , respectively, the friction coefficient h , the discrete cohesion strength coefficient s_p , the discrete particle interaction distance ϵ , the discrete adhesion strength coefficient s_w , and the coefficient of gravity g . Next a solution algorithm to run the full simulation is provided.

5.2 Numerical discretization

This section details the discretization techniques that will allow us to solve the full sinkhole system in time and space. The fluid system is solved on a finite element mesh (see Fig. 5.2) using the previously-described and validated weak formulation of Brinkman’s equation (see Sec. 4.2). The point particles, including their forces and dynamics, were previously described in Sec. 2.2.2. The porous material volume fraction field θ_p will be interpolated from the position of their point particles using the discrete-to-continuum transform \mathcal{G} , as described in Sec. 2.3. The simulation algorithm proceeds as follows. Assume all variables are known at time t_n . Consider N particles \mathbf{p}_j , $j = 1, 2, \dots, N$ and the finite element mesh composed of M nodes \mathbf{x}_i , $i = 1, 2, \dots, M$. Repeat the following steps:




Figure 5.2: Sinkhole mesh after finite element discretization.

1. Use the discrete-to-continuum transform \mathcal{G} of Sec. 2.3 on the discrete particle field $\mathbf{p}_j(t_n)$ to compute the volume fraction fields $\theta_p(\mathbf{x}_i, t_{n+1})$.
2. Compute $\theta_s(\mathbf{x}_i, t_{n+1})$ from $\theta_p(\mathbf{x}_i, t_{n+1})$ using Eq. (5.1).
3. Solve Eq. (5.2) for $\mathbf{q}(\mathbf{x}_i, t_{n+1})$.
4. Use the algebraic equation for particle velocity in Eq. (2.60). Use this particle velocity to obtain $\mathbf{p}_j(t_{n+1})$ by applying a forward difference scheme for position.

5.3 Results

A single initial condition configuration is used for the particles, where they are stacked above the boundary “hole” in the middle of the domain. This initial condition was created by running a

simulation with particles, generated with CircGeometry.jl (see Appendix C), being set above the hole and zero inflow velocity and running until there is no more movement. In other words, the initial condition was generated by letting particles settle into place. This single configuration which will allow the isolation of the effect of varying fluid input strength, v_{\max} , and “structural integrity”, as modeled through removing particles in a radial arc from the center of the domain, r_{cut} . See Fig. 5.3 for an initial condition where the “arc” of removed particles, r_{cut} , is small.




Figure 5.3: Zoomed-in view of initial condition for particle distribution. Successively larger chunks, radially from the center, will be removed. This will simulate sinkholes that have varying degrees of “structural integrity”.

The only results presented in this section will be to search a two-dimensional parameter space where the first parameter is the strength of the input fluid velocity v_{\max} , and the second parameter is the radial distance, from the center, of particles that are removed before running the simulation, r_{cut} . Additionally, gravity is set to zero so as to isolate the effect of seepage force. Before running the simulation, one might imagine that the effect of the seepage force will be to “whisk away” the particles closest to the opening in the boundary, particle after particle, until the amount of particles being removed weakens the “structural support” of the opening and a bifurcation occurs where particles fall until there are no more particles above the hole. We will call this type of collapse “slow-to-sudden”. Under this collapse regime, a bifurcation in parameter space is expected – if the seepage force is too small compared to the cohesion force of the particle, or the radial removal of particles is too small so that the structural support is greater than the seepage force, we will expect no particles to fall, and therefore no sinkhole to form. However, if this combination of effects is large enough, a sinkhole will develop. Therefore, the current simulations seek to (1) verify that this bifurcation exists, and (2) identify the relationship between parameters that might suggest whether one effect is more significant than another. This second point is what is important

for fieldwork where scientists would want to predict which underwater caverns are most likely to develop damaging sinkholes.

One representative simulation is given in Fig. 5.5. This type of collapse was representative of all simulations run in this chapter. All particles above the hole fall simultaneously, contrary to the hypothesis that collapse will start with individual particles at the hole falling before a collapse occurs. This type of behavior is contrary to experimental evidence, and therefore suggests that more fine tuning to the model is necessary before a simulation analogous to experimental setups can be achieved.

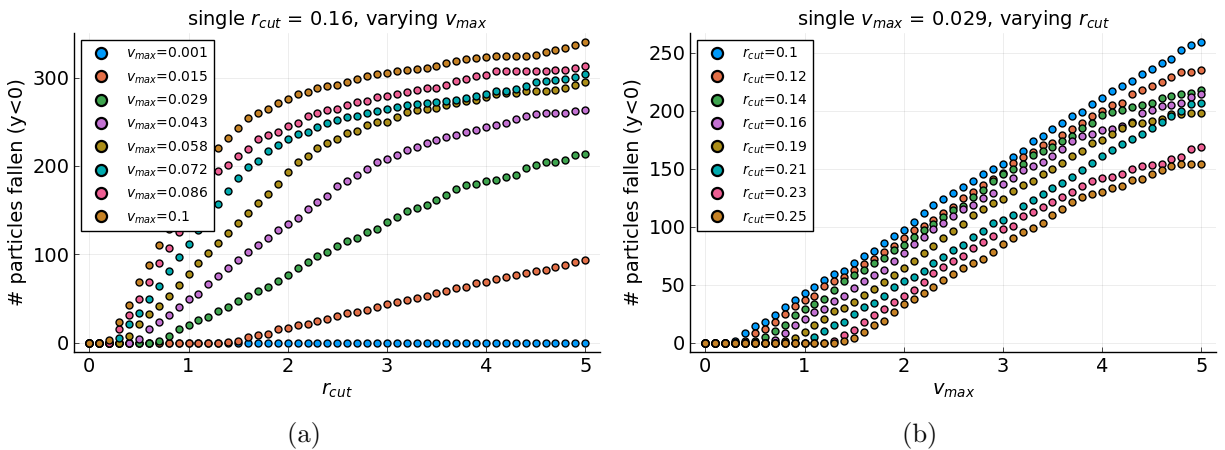


Figure 5.4: Varying two sinkhole parameters. Vertical axis signifies number of particles who have passed through the hole at $y = 0$, and horizontal axis is time. (a) Varying v_{max} leads to interesting differences in dynamics while (b) varying r_{cut} does not.

Varying the two parameters results help explain why the “slow-to-sudden” collapse does *not* appear in the simulations. In Fig. 5.4(a) the radial amount of cut away particles, r_{cut} , is held fixed while the seepage velocity is allowed to vary, while in Fig. 5.4(b) the opposite occurs. It is clear that r_{cut} , a proxy for structural integrity, does not impact whether or not a sinkhole forms or, furthermore, the rate of collapse when it does occur. Varying v_{max} , on the other hand, both impacts whether a collapse occurs (the lowest v_{max} does not lead to sinkhole formation) and, if it does occur, impacts the rate of formation. This further demonstrates that the current model must be amended to better represent structural integrity of sinkhole formation, and eventually obtain the “slow-to-sudden” collapse type that has been demonstrated in laboratory experiments [102].




Figure 5.5: Sinkhole formation for $r_{\text{cut}} = 0.16$ and $v_{\text{max}} = 0.3$. Particle and flow profiles are shown for (a) $t = 0$, (b) $t = 1.7$, (c) $t = 3.3$ and (d) $t = 5$. All particles fall simultaneously, contrary to the hypothesis that collapse will start with individual particles at the hole falling before a collapse occurs.

5.4 Conclusion

The model framework presented in this dissertation was adapted to represent geophysical flows. Sinkhole dynamics were examined by varying two parameters, r_{cut} and v_{max} , representing a proxy for structural integrity and magnitude of seepage force, respectively. It was found that varying seepage force created a bifurcation in sinkhole dynamics, while varying r_{cut} did not. This suggests

that the model should be amended to better represent structural integrity in order to accurately reflect what is known about sinkhole dynamics from lab experiments. In summary, although further refinements are needed, this chapter showcased the model framework's ability to handle geophysical flows, specifically it used the discrete particle representation of complex porous media as well as the discrete-to-continuum transform at production scale.

CHAPTER 6

DISCUSSION AND CONCLUSION

This dissertation was devoted to developing a multiphase framework for handling complex chemistry and the generation of precipitate structures which impede fluid flow. Numerical discretizations and simulations were conducted to demonstrate the applicability of this model framework. Open-source code is available online for the interested reader to make use of this framework for their own problems. We hope that the reader has learned something and will make use of this framework in their own scientific area. This final chapter reviews the achievements of this dissertation and detail two more scientific problems that we believe could use this framework successfully.

This dissertation was split into two parts; the first was for theory and validation, while the second was used for detailing two specific applications. In Chapter 1 a novel model framework was motivated by discussing a microfluidic experiment that could not be modeled by existing methods. We detailed these methods, and provided a bird's-eye view of what would be accomplished over the course of the dissertation. In Chapter 2, the most important chapter of the dissertation, the model framework was derived using rational principles. In Chapter 3, the model framework was validated, specifically by using a reduced model that admitted analytic solutions; we examined the interpretation of multiphase drag while detailing an indirect method that is useful for computing drag on multiphase bodies; we provided initial simulations using the discrete particle method for modeling complex rheology; and we presented a one-dimensional validation of the continuum-to-discrete operator.

The second part of the dissertation were devoted to full simulations of two application areas: microfluidic precipitation experiments and sinkhole dynamics. The microfluidic experiment was the motivating experiment for this entire dissertation, and so its successful simulation justified the development of the model framework. As an example of the usefulness of this model framework to experimentalists, simulations produced evidence for a hypothesis for why lateral membrane growth is asymmetric. In the sinkhole problem, we were able to showcase how our framework could be employed for scientific domains outside of the microfluidics community, in particular by using the discrete particle approach for modeling flows containing solids with complex rheology.

6.1 Future applications

This section details two specific scientific areas that we believe would benefit from this model framework. For researchers outside of these areas, we hope that reading this section will give insight into how our framework could be adapted to your scientific domain. Possible extensions to the model to extend it to other scientific areas are also discussed.

6.1.1 Ultrasonic drug delivery

Therapeutic ultrasound techniques are non-invasive and potentially beneficial in a number of clinical applications, including treating prostate cancer, breaking up blood clots and drug delivery through the blood-brain barrier [28, 63, 27]. All of these techniques are affected in one way or another by the physical phenomenon of inertial cavitation. When a fluid is subject to acoustic forcing, local pressure fluctuations cause the formation of gas bubbles. These bubbles oscillate under the effect of the acoustic field and, under certain conditions, collapse violently in a process called inertial cavitation to produce a pressure and temperature shock wave. This shock wave is viewed by medical practitioners as either beneficial or harmful depending on the application.

While our framework, in its present form, does not include the ability to add gas phase dynamics, there does exist the possibility of including a lipid vesicle to deliver drugs under the effect of ultrasound. These vesicles have complex rheology [114] and so it would be possible to implement them with our discrete particle approach. The effect of a passing ultrasound wave to temporarily “dismember” the vesicle, in order to release drugs, would be possible by temporarily affecting the cohesion strength s and buffer distance ϵ of the discrete particles representing vesicles. While mathematical models specific to these lipid vesicles exist [82, 64], our model would include the ability to couple fluid-structure interaction and drug reactions in a single, unified framework.

6.1.2 Biofilms in water filtration

Water filtration is incredibly important for the modern world. The UN has estimated that 1.6 billion people live in communities where water is “economically scarce”, referring to countries that “lack the necessary infrastructure to take water from rivers and aquifers” [76]. Additionally, a changing climate will exacerbate this issue in the future.

Before filtering, surface water contains impurities that are dangerous and unhealthy for consumption, including pathogenic microorganisms and potentially carcinogenic chemical waste. A common method for water filtration is to pass the unfiltered water through a physical filter, which

stops particles above a certain size. While this process works in practice, it is also unsustainable. Biofilms grow and “foul” the membrane over time, such that it needs to be replaced. One method for combating this biofilm fouling is through the technique known as electroflocculation [18], where a sacrificial reactive surface (typically made of aluminum or iron) is placed in a pipe before the filter such that, upon electrification, impurities react and coagulate into “flocs”, which experience a buoyant force and drift upward into a collection duct separate from the filter itself. The efficient application of this technique can reduce the replacement times required by filters by orders of magnitude.

This problem is an ideal situation for our model framework, employing the more advanced use of complex rheology. In a single problem, it includes reaction (electrified aluminum combined with impurities), transport of materials of complex rheology (flocs), and fluid-structure interaction (fouling of the filter via accumulation of impurities).

6.1.3 Framework extensions

We end this dissertation by summarizing the possible extensions to our model to make it applicable to scientific domains not considered in detail. First, we have never mentioned the possibility of electromagnetic forces. Because many scientific and industrial applications include, or even rely on, these forces, their exclusion from this dissertation certainly leaves an open area of research whose inclusion would extend the number of applications this framework could be applied towards.

We have also neglected any gas dynamics, which could be important for applications that wish to include, for example, air pokes in geophysics or gas bubbles in ultrasonic therapies. Additionally, while some of the numerical methods were efficient and used modern techniques, many did not, as that was

not the primary concern of this dissertation. This leaves room for considerable improvement in implementation. Finally, implementing these methods in three dimensions is conceptually sound but has not been implemented yet in any of the open-source code provided in appendix C.




Figure 6.1: Images demonstrating bi-directional, complex membrane growth from microfluidic experiments. Reproduced from [109] with permission.

For a final possible extension, see Fig. 6.1 from a microfluidic experiment very similar to the one described in Sec. 1.1. The difference with this experiment is that the membrane growth is bi-directional and not “simple”. We see two very different membrane growths on either side of the membrane; these complex membrane growths are not possible to achieve using the current method, as we have “homogenized” over growth space. However, if one was determined to achieve this level of resolution in the membrane crystal growth, here is one possible approach: make the diffusion coefficient in the ADR equations a tensor quantity that depends not only on volume fraction, but also gradients in volume fraction. One can imagine that this will allow preferential diffusive crystal growth. That being said, this makes no mention of the incredible computational cost associated with resolving this spatial scale, so resolving both the full channel advection with this crystal growth is likely not advisable.

To summarize, the results of this dissertation are (1) a rigorous derivation of a model framework that conserves mass and incorporates fluid-structure interaction, and (2) numerical methods for solving the full PDE system, which have been implemented in open-source software. The micrometer-scale resolution of our numerical simulations for microfluidic problems is a huge success. The model framework developed here allows simulations that would not have been possible, or at least computationally intractable, using existing methods. The framework has applications towards precipitate reactions where the precipitate greatly affects the surrounding flow, a situation appearing in many laboratory and geophysical contexts including the hydrothermal vent theory for the origin of life. More generally, this framework can be used to address fluid–structure interaction problems that feature the dynamic generation of structures.

APPENDIX A

SOLVING RICATTI'S DIFFERENTIAL EQUATION

In Sec. 3.1, the reduced model formulation led to a Ricatti differential equation with constant coefficients, which is exactly solvable. The solution technique to Ricatti's differential equation [see 104, p. 97] is not very well known, so the derivation is stated here for the interested reader. A homogeneous 1st order ODE is called a Ricatti equation if it is quadratic in the unknown, i.e.

$$y'(x) = q_0(x) + q_1(x)y(x) + q_2(x)y^2(x). \quad (\text{A.1})$$

If $q_0(x) \equiv 0$, this reduces to Bernoulli's equation [see 104, pp. 95-96]. In general, one can transform Ricatti's equation to an equivalent 2nd order linear differential equation. In this appendix, we detail this transformation and explicitly solve for the case of constant coefficients q_0 , q_1 and q_2 .

First, define a new variable v by

$$v = q_2 y \quad (\text{A.2})$$

so that Eq. (A.1) becomes

$$v' = v^2 + Rv + S \quad (\text{A.3})$$

where $R = q_1 + q_2'/q_2$ and $S = q_0 q_2$. Then, introduce another variable, u , related to v via a Cole-Hopf transform:

$$v = -\frac{u'}{u} \quad (\text{A.4})$$

and now the original equation, in terms of u , becomes

$$u'' - Ru' + Su = 0. \quad (\text{A.5})$$

For constant coefficients R and S , Eq. (A.5) can be solved exactly. Using notation from Sec. 3.1, the constant coefficient Ricatti equation is

$$\dot{\psi}_C = cr\psi_A\psi_B - \frac{\alpha\beta}{\rho_p}\psi_C + \frac{\beta}{\rho_p}\psi_C^2 \quad (\text{A.6})$$

where $\psi_C = \psi_C(t)$ is the unknown variable and r , c , ρ_p , α , β , ψ_A , and ψ_B are fixed parameters. Relating this to Eq. (A.1), let $y = \psi_C$, $q_0 = cr\psi_A\psi_B$, $q_1 = -\alpha\beta/\rho_p$, and $q_2 = \beta/\rho_p$. Therefore, in the final equation, $R = q_1 + q_2'/q_2 = -\alpha\beta/\rho_p$ and $S = q_0 q_2 = cr\beta\psi_A\psi_B/\rho_p$.

The solution to Eq. (A.5) when R and S are constant depends on the eigenvalues of its corresponding characteristic equation. More specifically, it depends on the sign of the determinant of the root of the characteristic equation $R^2 - 4S = \chi/\rho_p^2$, where χ , defined in Eq. (3.4), was determined to be positive for physically realistic parameter values.

The solution to this case – where the characteristic equation to Eq. (A.5) has two real, distinct roots – can be found in any introductory book on ODEs [see 13, pp. 137 - 143], but for completeness the solution is detailed here with homogeneous initial conditions $y(0) = 0$:

$$y(t) = \frac{\gamma_1 \gamma_2}{q_2} \left[\frac{e^{\gamma_2 t} - e^{\gamma_1 t}}{\gamma_2 e^{\gamma_1 t} - \gamma_1 e^{\gamma_2 t}} \right]. \quad (\text{A.7})$$

where

$$\gamma_1 = \frac{R + \sqrt{R^2 - 4S}}{2}, \quad \gamma_2 = \frac{R - \sqrt{R^2 - 4S}}{2} \quad (\text{A.8})$$

$$R = q_1, \quad S = q_0 q_2. \quad (\text{A.9})$$

and the antiderivative of this solution is given by

$$\int y(t) dt = -\frac{1}{q_2} \log(\gamma_1 e^{\gamma_2 t} - \gamma_2 e^{\gamma_1 t}) + C \quad (\text{A.10})$$

where C is an arbitrary constant of integration.

APPENDIX B

TREE-BASED CELL LISTS FOR FAST INTERPOLATION

Due to the velocity flux “seepage” force experienced by the discrete particles in a fluid flow, we require a fast method to interpolate our velocity field at a point not on the mesh. In theory, this is not a difficult algorithm to implement, as the basis functions for our finite elements are known [93]. However, the process of actually finding which element a point lies within can be costly, and boils down to an unstructured search algorithm. To speed up this search, one can use a cell-list approach supported by a tree-structure to find an enclosing element. The process has two components; one component needs to be run once to initialize the data structure (the cell list), and the second component needs to be run every search (the tree structure). Another benefit of this data structure is that it extends easily to three dimensions.




Figure B.1: Example of finite elements (black solid) whose indices will be associated with the cell (red dash). An element is associated with a cell when at least one of its nodes lies within the cell. Elements can be associated with more than one cell.

The Cell List: Cell lists are data structures by which a domain (two-dimensional in this case) is divided into rectangular “cells” (see Fig. B.1). Each of these cells has a list associated with it. In the case of the present situation, each list will be an array of indices for the elements that are, at least in part, inside that cell. The benefit of this data structure is that there are only a few elements within each cell. For example, it is much faster to scan over a list of 12 elements (all components of the same cell), rather than search over the entire domain, which might include tens of thousands of elements.

The Tree Structure: The above procedure still leaves open the question of how the point-containing cell will be found. One could, of course, simply iterate over the cells. This is easier than iterating over elements because (1) cells are rectangular, so it is computationally cheaper to check for a point within a rectangle oriented parallel to the coordinate axes rather than an arbitrary quadrilateral, and (2) there are less cells than

elements. However, there is an easier approach as long as care is taken initially to align the cells perfectly as iterations of a “tree structure”. This tree has different levels, each level subdividing the one above into four equivalent rectangles. For example, enclose the (possibly complex) domain into a bounding rectangle. This rectangle is level 0. At level 1, divide level 0 into four equivalent rectangles. Repeat the process until you have the refinement in cells that you require; this will be problem-dependent. At the end of the day, this search algorithm will be logarithmic, as opposed to linear, in the number of cells. One disadvantage is that the number of cells at the lowest level must be a power of four, but that is a small price to pay for the efficiency gained at run-time.

APPENDIX C

SOFTWARE PACKAGES

This appendix details four software packages that I have developed over the course of my graduate career. They are all open-source, actively developed, and I encourage the reader to use them to explore this model framework. If I may prosthelytize for two sentences, open-source software is a backbone of modern computation. I encourage everyone to release open-source packages, especially pertaining to novel algorithms and methods, so that students and researchers might learn without requiring the substantial time and energy investment of coding an unfamiliar algorithm oneself.

eFEM.jl. This software package is finite element code in the Julia language focused on fluid-dynamics applications. It emphasizes easy problem setup and solutions, without worrying about the “backend” about how the problem is actually discretized and solved. This makes it ideal for modeling applications, where small changes in parameters or geometries occur frequently, and computational efficiency is de-emphasized. Currently, it can solve the following time-independent PDE systems:

- Poisson: $-\nabla \cdot (\alpha \nabla u) = f$,
- Advection-Diffusion: $-\mathbf{u} \cdot \nabla c = \nabla \cdot (\kappa \nabla c) + f$,
- Stokes: $-\nabla \cdot (\mu \mathbb{D}(\mathbf{u})) = \nabla p + \mathbf{f}$, $\nabla \cdot \mathbf{u} = 0$,
- Brinkman (multiphase): $-\alpha_1 \nabla^2 \mathbf{u} + \alpha_2 \mathbf{u} + \alpha_3 \nabla p = \mathbf{f}$, $\nabla \cdot \mathbf{u} = 0$.

All of the above systems are formed via operators, so unsteady problems (such as the ones solved in this dissertation) are possible as well. Additionally, all equations except Brinkman are available in an axisymmetric form. The package has the capability to generate its own rectangular meshes, but can also take in meshes generated in the open-source finite element meshing software Gmsh [41]. It is available in a Github repository¹ and works on all operating systems.

StokesParticles.jl. This software package allows one to simulate “Stokes particles”, meaning inertia-less particles, subject to discrete forces. It also contains functionality for using the fast Gauss transform [71], in case one wants to generate a continuum field from the discrete particles. It uses


¹<https://github.com/pseastham/eFEM.jl>

Lennard-Jones-like potential forces for particle-particles and particle-wall interactions, and includes cell list algorithms for the efficient evaluation of these forces. It is available in a Github repository² and is only available on Linux operating systems.

eFEMpart.jl. This software package is finite element code in the Julia language focused on complex fluid-dynamic and porous-media applications, with the possibility of including a particle simulator in the framework of the discrete element method. The 'eFEM' component allows the use of Finite Elements discretizations to solve common problems in fluid dynamics, and the 'part' refers to mesh-free particle methods (discrete element method) primarily aimed at granular-media simulations where continuum constitutive laws are unavailable. It is built on eFEM.jl and StokesParticles.jl, described above, while including code to couple these two packages together. It is available in a Github repository³ and is only available on Linux operating systems.

CircGeometry.jl. CircGeometry.jl takes in some basic information related to the object one is trying to approximate with "filling circles", and exports a file in a custom "circ" format. The following is the algorithm used to generate these geometries in an efficient manner: the first step is to take in both the outline geometry, the volume fraction, the number of "filling" circles, and an inner "buffer" percentage that prevents circles from being too close together, even though they won't overlap. The outline geometry can be a circle, a rectangle, or an arbitrary polygon (see Fig. C.2). The volume fraction can be any number between 0 and around 0.7; 0.7 is close to the theoretical maximum for circle packing in a plane. From the given data, one can compute the "ideal radius" for the filling circles to be packed into the outline. Then we create a random distribution of radii about this "ideal radius", and generate the random assortment of radii, which are then ordered largest to smallest.

Now, the placing begins in an iterative fashion. First, the circle is placed and checked for whether the filling circle is both inside the outline, and then checked whether it is overlapping any of the previously placed circles (while accounting for buffer). If both checks pass, then that circle placement is permanent, and the iteration proceeds. If the filling circle is outside the outline, then



CircGeometry.jl

Figure C.1: Logo for CircGeometry.jl.

²<https://github.com/pseastham/StokesParticles.jl>

³<https://github.com/pseastham/eFEMpart.jl>

a new random initial position is chosen. If the filling circle is inside the outline, but overlapping a previously placed circle, then we allow the current filling circle to “wobble” away from any adjacent circles or, after a fixed number of attempts, a new random initial starting location is chosen and the process starts over. This whole process goes on until either every filling circle is successfully placed, or a max number of iterations is reached where the process fails and an error message is printed. A successful placement is then either printed to a circ file to be used by a later application, or an figure in the svg format is saved.

The time it takes to generate a circ file will depend on these parameters, but will generally only be a couple minutes for production-level files (e.g. volume fraction of 0.4 with 800 circles). This appendix ignores some details for the sake of expositional clarity. As the source code is open-source, the interested reader is encouraged to look directly at the program files themselves for any further questions about implementation. It is available in a Github repository⁴.




Figure C.2: Examples of porous geometries that can be generated using CircGeometry.jl. The software works by filling a (a) circle, (b) rectangle, or (c) arbitrary polygon with circles of random radii so that a given volume fraction is achieved.

⁴<https://github.com/pseastham/CircGeometry.jl>

REFERENCES

- [1] Vishal Agarwal and Baron Peters. Solute precipitate nucleation: A review of theory and simulation advances. *Advances in Chemical Physics: Volume 155*, pages 97–160, 2014.
- [2] Djamil Al-Halbouni, Eoghan Holohan, Abbas Taheri, et al. Geomechanical modelling of sinkhole development using distinct elements: model verification for a single void space and application to the dead sea area. *Solid Earth*, 9:1341–1373, 2018.
- [3] Djamil Al-Halbouni, Eoghan P Holohan, Abbas Taheri, Robert A Watson, Ulrich Polom, Martin PJ Schöpfer, Sacha Emam, and Torsten Dahm. Distinct element geomechanical modelling of the formation of sinkhole cluster within large-scale karstic depressions. *Solid Earth*, 10:1219–1241, 2019.
- [4] Fernando Alonso-Marroquín and Yucang Wang. An efficient algorithm for granular dynamics simulations with complex-shaped objects. *Granular Matter*, 11(5):317–329, 2009.
- [5] Philippe Angot. Analysis of singular perturbations on the brinkman problem for fictitious domain models of viscous flows. *Mathematical methods in the applied sciences*, 22(16):1395–1412, 1999.
- [6] AP. Cave-in suit: \$219 million. *New York Times*, 1981. URL <https://www.nytimes.com/1981/11/24/business/cave-in-suit-219-million.html>.
- [7] Laura M Barge and Lauren M White. Experimentally testing hydrothermal vent origin of life on enceladus and other icy/ocean worlds. *Astrobiology*, 17(9):820–833, 2017.
- [8] Laura M Barge, Silvana SS Cardoso, Julyan HE Cartwright, Geoffrey JT Cooper, Leroy Cronin, Anne De Wit, Ivria J Doloboff, Bruno Escribano, Raymond E Goldstein, Florence Haudin, et al. From chemical gardens to chemobionics. *Chemical reviews*, 115(16):8652–8703, 2015.
- [9] Bruno C Batista and Oliver Steinbock. Growing inorganic membranes in microfluidic devices: Chemical gardens reduced to linear walls. *The Journal of Physical Chemistry C*, 119(48):27045–27052, 2015.
- [10] Gordon S Beavers and Daniel D Joseph. Boundary conditions at a naturally permeable wall. *Journal of fluid mechanics*, 30(1):197–207, 1967.
- [11] Steven Berline and Clark Bricker. The law of mass action. *Journal of Chemical Education*, 46(8):499, 1969.
- [12] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017.
- [13] Boyce and Richard C DiPrima. *Elementary Differential Equations*. Wiley, 2008.

- [14] Christopher JW Breward, Helen M Byrne, and Claire E Lewis. A multiphase model describing vascular tumour growth. *Bulletin of mathematical biology*, 65(4):609–640, 2003.
- [15] HC Brinkman. A calculation of the viscous force exerted by a flowing fluid on a dense swarm of particles. *Flow, Turbulence and Combustion*, 1(1):27, 1949.
- [16] Helen Byrne and Luigi Preziosi. Modelling solid tumour growth using the theory of mixtures. *Mathematical medicine and biology: a journal of the IMA*, 20(4):341–366, 2003.
- [17] Raymond Chang and Ken Goldsby. *Chemistry*. McGraw-Hill, 11th edition, 2013.
- [18] Shankararaman Chellam and Mutiara Ayu Sari. Aluminum electrocoagulation as pretreatment during microfiltration of surface water containing nom: A review of fouling, nom, dbp, and virus control. *Journal of hazardous materials*, 304:490–501, 2016.
- [19] Stephen Childress. Viscous flow past a random array of spheres. *The Journal of Chemical Physics*, 56(6):2527–2539, 1972.
- [20] Stephen Childress, Michael Shelley, and Jun Zhang. Fluid-structure interactions: Research in the courant institute’s applied mathematics laboratory. *Communications on Pure and Applied Mathematics*, 65(12):1697–1721, 2012.
- [21] Shang-Huan Chiu, M. N. J. Moore, and Bryan Quaipe. Viscous transport in eroding porous media. *Journal of Fluid Mechanics*, 893:A3, 2020. doi: 10.1017/jfm.2020.228.
- [22] NG Cogan and Shankar Chellam. Incorporating pore blocking, cake filtration, and eps production in a model for constant pressure bacterial fouling during dead-end microfiltration. *Journal of Membrane Science*, 345(1-2):81–89, 2009.
- [23] NG Cogan and Robert D Guy. Multiphase flow models of biogels from crawling cells to bacterial biofilms. *HFSP journal*, 4(1):11–25, 2010.
- [24] NG Cogan and James P Keener. The role of the biofilm matrix in structural development. *Mathematical medicine and biology: a journal of the IMA*, 21(2):147–166, 2004.
- [25] NG Cogan and James P Keener. Channel formation in gels. *SIAM Journal on Applied Mathematics*, 65(6):1839–1854, 2005.
- [26] John Horton Conway and Neil James Alexander Sloane. *Sphere packings, lattices and groups*, volume 290. Springer Science & Business Media, 2013.
- [27] CC Coussios, CH Farny, G Ter Haar, and RA Roy. Role of acoustic cavitation in the delivery and monitoring of cancer treatment by high-intensity focused ultrasound (HIFU). *International journal of hyperthermia*, 23(2):105–120, 2007.
- [28] Constantin C Coussios and Ronald A Roy. Applications of acoustics and cavitation to non-invasive therapy and drug delivery. *Annu. Rev. Fluid Mech.*, 40:395–420, 2008.

- [29] Clint N Dawson and Mary F Wheeler. Time-splitting methods for advection-diffusion-reaction equations arising in contaminant transport. Technical report, 1992.
- [30] Andrew J Demello. Control and detection of chemical reactions in microfluidic systems. *Nature*, 442(7101):394–402, 2006.
- [31] Yang Ding, Bruno Batista, Oliver Steinbock, Julyan HE Cartwright, and Silvana SS Cardoso. Wavy membranes and the growth rate of a planar chemical garden: Enhanced diffusion and bioenergetics. *Proceedings of the National Academy of Sciences*, 113(33):9182–9186, 2016.
- [32] Earl H Dowell and Kenneth C Hall. Modeling of fluid-structure interaction. *Annual review of fluid mechanics*, 33(1):445–490, 2001.
- [33] Donald Allen Drew. Mathematical modeling of two-phase flow. *Annual review of fluid mechanics*, 15(1):261–291, 1983.
- [34] Francis AL Dullien. *Porous media: fluid transport and pore structure*. Academic press, 2012.
- [35] L Durlafsky and JF Brady. Analysis of the Brinkman equation as a model for flow in porous media. *The Physics of Fluids*, 30(11):3329–3341, 1987.
- [36] P. S. Eastham, M. N. J. Moore, N. G. Cogan, Q. Wang, and O. Steinbock. Multiphase modelling of precipitation-induced membrane formation. *Journal of Fluid Mechanics*, 888: A20, 2020. doi: 10.1017/jfm.2020.52.
- [37] Patrick S Eastham. eFEMpart. <https://github.com/pseastham/eFEMpart>, 2019.
- [38] Patrick S Eastham and Kouros Shole. Axisymmetric squirmers in stokes fluid with nonuniform viscosity. *Physical Review Fluids*, 5(6):063102, 2020.
- [39] Howard C Elman, David J Silvester, and Andrew J Wathen. *Finite elements and fast iterative solvers: with applications in incompressible fluid dynamics*. Numerical Mathematics and Scientific Computation, 2014.
- [40] José Furtado. *Al mezclar una solución de sulfato cúprico (CuSO_4) con una de hidróxido sódico (NaOH), se forma sulfato de sodio (Na_2SO_4) e hidróxido de cobre (II) ($\text{Cu}(\text{OH})_2$), que al ser insoluble se precipita*, 2020 (accessed June 15, 2020). URL https://commons.wikimedia.org/wiki/Category:Solid_precipitation#/media/File:Precipitado.jpg.
- [41] Christophe Geuzaine and Jean-François Remacle. Gmsh: A 3-d finite element mesh generator with built-in pre-and post-processing facilities. *International journal for numerical methods in engineering*, 79(11):1309–1331, 2009.
- [42] William Gilpin, Matthew Storm Bull, and Manu Prakash. The multiscale physics of cilia and flagella. *Nature Reviews Physics*, pages 1–15, 2020.

- [43] Kenneth M. Golden. *Percolation Models for Porous Media*, pages 27–43. Springer New York, New York, NY, 1997. ISBN 978-1-4612-1920-0. doi: 10.1007/978-1-4612-1920-0_2. URL https://doi.org/10.1007/978-1-4612-1920-0_2.
- [44] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [45] Leslie Greengard and John Strain. The fast gauss transform. *SIAM Journal on Scientific and Statistical Computing*, 12(1):79–94, 1991.
- [46] Boyce E Griffith and Neelesh A Patankar. Immersed methods for fluid–structure interaction. *Annual Review of Fluid Mechanics*, 52, 2020.
- [47] Francisco Gutiérrez, Alfonso Benito-Calvo, Domingo Carbonel, Gloria Desir, Jorge Sevil, Jesús Guerrero, Adrián Martínez-Fernández, Theodoros Karamplaglidis, Ángel García-Arnay, and Ivan Fabregat. Review on sinkhole monitoring and performance of remediation measures by high-precision leveling and terrestrial laser scanner in the salt karst of the ebro valley, spain. *Engineering geology*, 248:283–308, 2019.
- [48] John Happel and Howard Brenner. *Low Reynolds number hydrodynamics: with special applications to particulate media*, volume 1. Springer Science & Business Media, 2012.
- [49] Reghan J Hill, Donald L Koch, and Anthony JC Ladd. The first effects of fluid inertia on flows in ordered and random arrays of spheres. *Journal of Fluid Mechanics*, 448:213–241, 2001.
- [50] M. F. Horstemeyer. *Multiscale Modeling: A Review*, pages 87–135. Springer Netherlands, Dordrecht, 2010. ISBN 978-90-481-2687-3. doi: 10.1007/978-90-481-2687-3_4. URL https://doi.org/10.1007/978-90-481-2687-3_4.
- [51] Gene Hou, Jin Wang, and Anita Layton. Numerical methods for fluid-structure interaction—a review. *Communications in Computational Physics*, 12(2):337–377, 2012.
- [52] Emanuele Intrieri, Giovanni Gigli, Massimiliano Nocentini, Luca Lombardi, Francesco Mugnai, Francesco Fidolini, and Nicola Casagli. Sinkhole monitoring and early warning: An experimental and successful gb-insar application. *Geomorphology*, 241:304–314, 2015.
- [53] Graham W Jackson and David F James. The permeability of fibrous porous media. *The Canadian Journal of Chemical Engineering*, 64(3):364–374, 1986.
- [54] Claes Johnson. *Numerical solution of partial differential equations by the finite element method*. Courier Corporation, 2012.
- [55] Rimantas Kačianauskas, Robertas Balevičius, Darius Markauskas, and Algirdas Maknickas. Discrete element method in simulation of granular materials. In Peter Eberhard, editor, *IU-TAM Symposium on Multiscale Problems in Multibody System Contacts*, pages 65–74, Dordrecht, 2007. Springer Netherlands. ISBN 978-1-4020-5981-0.

- [56] James P Keener, Sarthok Sircar, and Aaron L Fogelson. Kinetics of swelling gels. *SIAM Journal on Applied Mathematics*, 71(3):854–875, 2011.
- [57] Deborah S Kelley, Jeffrey A Karson, Donna K Blackman, Gretchen L FruÈh-Green, David A Butterfield, Marvin D Lilley, Eric J Olson, Matthew O Schrenk, Kevin K Roe, Geoff T Lebon, et al. An off-axis hydrothermal vent field near the mid-atlantic ridge at 30 n. *Nature*, 412(6843):145–149, 2001.
- [58] Ebrahim M Kolahdouz, Amneet Pal Singh Bhalla, Brent A Craven, and Boyce E Griffith. An immersed interface method for discrete surfaces. *Journal of Computational Physics*, 400:108854, 2020.
- [59] Eve L Kuniandy, David J Weary, and James E Kaufmann. The current status of mapping karst areas and availability of public sinkhole-risk resources in karst terrains of the united states. *Hydrogeology Journal*, 24(3):613–624, 2016.
- [60] Marissa J Lang. There’s a sinkhole on the white house lawn. and it’s growing. *Washington Post*, 2018. URL <https://www.washingtonpost.com/news/local/wp/2018/05/22/theres-a-sinkhole-on-the-white-house-lawn-and-its-growing/>.
- [61] Karin Leiderman and Aaron L Fogelson. Grow with the flow: a spatial–temporal model of platelet deposition and blood coagulation under flow. *Mathematical medicine and biology: a journal of the IMA*, 28(1):47–84, 2011.
- [62] Karin Leiderman and Aaron L Fogelson. The influence of hindered transport on the development of platelet thrombi under flow. *Bulletin of mathematical biology*, 75(8):1255–1283, 2013.
- [63] Timothy Leighton. *The acoustic bubble*. Academic press, 2012.
- [64] John S Lowengrub, Andreas Rätz, and Axel Voigt. Phase-field modeling of the dynamics of multicomponent vesicles: Spinodal decomposition, coarsening, budding, and fission. *Physical Review E*, 79(3):031926, 2009.
- [65] Rabih Makki, Mohammed Al-Humiari, Sumana Dutta, and Oliver Steinbock. Hollow microtubes and shells from reactant-loaded polymer beads. *Angewandte Chemie International Edition*, 48(46):8752–8756, 2009.
- [66] Daniel Mark, Stefan Haeberle, Günter Roth, Felix Von Stetten, and Roland Zengerle. Microfluidic lab-on-a-chip platforms: requirements, characteristics and applications. In *Microfluidics Based Microsystems*, pages 305–376. Springer, 2010.
- [67] William Martin, John Baross, Deborah Kelley, and Michael J Russell. Hydrothermal vents and the origin of life. *Nature Reviews Microbiology*, 6(11):805, 2008.
- [68] Jacob H Masliyah and Marcel Polikar. Terminal velocity of porous spheres. *The Canadian Journal of Chemical Engineering*, 58(3):299–302, 1980.

- [69] Masayo Matsue, Masaki Itatani, Qing Fang, Yushiro Shimizu, Kei Unoura, and Hideki Nabika. Role of electrolyte in Liesegang pattern formation. *Langmuir*, 34(37):11188–11194, 2018.
- [70] MNJ Moore. Riemann-Hilbert problems for the shapes formed by bodies dissolving, melting, and eroding in fluid flows. *Communications on Pure and Applied Mathematics*, 70(9):1810–1831, 2017.
- [71] Vlad I. Morariu, Balaji Vasan Srinivasan, Vikas C. Raykar, Ramani Duraiswami, and Larry S. Davis. Automatic online tuning for fast gaussian summation. In *Advances in Neural Information Processing Systems (NIPS)*, 2008.
- [72] Vlad I Morariu, Balaji V Srinivasan, Vikas C Raykar, Ramani Duraiswami, and Larry S Davis. Automatic online tuning for fast gaussian summation. In *Advances in neural information processing systems*, pages 1113–1120, 2009.
- [73] Ole G Mouritsen. *Life-as a matter of fat*. Springer, 2005.
- [74] Lin Mu, Junping Wang, and Xiu Ye. A stable numerical algorithm for the Brinkman equations by weak Galerkin finite element methods. *Journal of Computational Physics*, 273:327–342, 2014.
- [75] Antonio A Munjiza. *The combined finite-discrete element method*. John Wiley & Sons, 2004.
- [76] United Nations. Water Scarcity. <https://www.un.org/waterforlifedecade/scarcity.shtml>, 2014. [Online; accessed 29-July-2020].
- [77] Graham Neale and Walter Nader. Practical significance of Brinkman’s extension of Darcy’s law: coupled parallel flows within a channel and a bounding porous medium. *The Canadian Journal of Chemical Engineering*, 52(4):475–478, 1974.
- [78] Pamela N Nge, Chad I Rogers, and Adam T Woolley. Advances in microfluidic materials, functions, integration, and applications. *Chemical Reviews*, 113(4):2550–2583, 2013.
- [79] Jace W. Nunziato and Edward K. Walsh. On ideal multiphase mixtures with chemical reactions and diffusion. *Archive for Rational Mechanics and Analysis*, 73(4):285–311, Dec 1980. ISSN 1432-0673.
- [80] J Alberto Ochoa-Tapia and Stephen Whitaker. Momentum transfer at the boundary between a porous medium and a homogeneous fluid—ii. comparison with experiment. *International Journal of Heat and Mass Transfer*, 38(14):2647–2655, 1995.
- [81] Borys I Ostapienko, Domenico Lopez, and Svetlana V Komarova. Mathematical modeling of calcium phosphate precipitation in biologically relevant systems: scoping review. *Biomechanics and modeling in mechanobiology*, pages 1–13, 2018.

- [82] William G Pitt, Ram N Singh, Krystian X Perez, Ghaleb A Husseini, and Daniel R Jack. Phase transitions of perfluorocarbon nanoemulsion induced with ultrasound: A mathematical model. *Ultrasonics sonochemistry*, 21(2):879–891, 2014.
- [83] Lionel Pournin, Marco Ramaioli, Patrick Folly, and Th M Liebling. About the influence of friction and polydispersity on the jamming behavior of bead assemblies. *The European Physical Journal E*, 23(2):229, 2007.
- [84] Bryan D Quaife and MNJ Moore. A boundary-integral framework to simulate viscous erosion of a porous medium. *Journal of Computational Physics*, 375:1–21, 2018.
- [85] D. C. Rapaport. *The Art of Molecular Dynamics Simulation*. Cambridge University Press, 2 edition, 2004. doi: 10.1017/CBO9780511816581.
- [86] Jacopo Riccati. *Animadversiones in aequationes differentiales secundi gradus*, volume 8. apud Joh. Grossii haeredes & J.F. Gleditschium Akkerman, 1724.
- [87] Leif Ristroph, Matthew NJ Moore, Stephen Childress, Michael J Shelley, and Jun Zhang. Sculpting of an erodible body by flowing water. *Proceedings of the National Academy of Sciences*, 109(48):19606–19609, 2012.
- [88] Laszlo Roszol and Oliver Steinbock. Controlling the wall thickness and composition of hollow precipitation tubes. *Physical Chemistry Chemical Physics*, 13(45):20100–20103, 2011.
- [89] S.I. Rubinow. *Introduction to Mathematical Biology*. John Wiley & Sons, Inc, 1975.
- [90] Sebastian Ruder. An overview of gradient descent optimization algorithms, 2016.
- [91] Philip Geoffrey Saffman. On the boundary condition at the surface of a porous medium. *Studies in applied mathematics*, 50(2):93–101, 1971.
- [92] Kouros Shoele and Patrick S Eastham. Effects of nonuniform viscosity on ciliary locomotion. *Physical Review Fluids*, 3(4):043101, 2018.
- [93] Gustavo HC Silva, Rodolphe Le Riche, Jérôme Molimard, and Alain Vautrin. Exact and efficient interpolation using finite elements shape functions. *European Journal of Computational Mechanics/Revue Européenne de Mécanique Numérique*, 18(3-4):307–331, 2009.
- [94] Sarthok Sircar, James P Keener, and Aaron L Fogelson. The effect of divalent vs. monovalent ions on the swelling of mucin-like polyelectrolyte gels: Governing equations and equilibrium analysis. *The Journal of chemical physics*, 138(1):014901, 2013.
- [95] Hrvoje Smoljanović, Nikolina Živaljić, Željana Nikolić, and Ante Munjiza. Numerical analysis of 3d dry-stone masonry structures by combined finite-discrete element method. *International Journal of Solids and Structures*, 136:150–167, 2018.
- [96] IC Sorribes, MNJ Moore, HM Byrne, and Harsh Vardhan Jain. A biomechanical model of tumor-induced intracranial pressure and edema in brain tissue. *Biophysical Journal*, 2019.

- [97] Marina Spivak, Shravan K Veerapaneni, and Leslie Greengard. The fast generalized gauss transform. *SIAM Journal on Scientific Computing*, 32(5):3092–3107, 2010.
- [98] John C Strikwerda. *Finite difference schemes and partial differential equations*, volume 88. Siam, 2004.
- [99] Wanda Strychalski and Robert D Guy. Intracellular pressure dynamics in blebbing cells. *Biophysical journal*, 110(5):1168–1179, 2016.
- [100] Wanda Strychalski, Calina A Copos, Owen L Lewis, and Robert D Guy. A poroelastic immersed boundary method with applications to cell biology. *Journal of Computational Physics*, 282:77–97, 2015.
- [101] Mark Sussman, Peter Smereka, Stanley Osher, et al. A level set approach for computing solutions to incompressible two-phase flow. *PhD Thesis*, 1994.
- [102] Xiaohu Tao, Ming Ye, Xiaoming Wang, Dangliang Wang, Roger Pacheco Castro, and Jian Zhao. Experimental and numerical investigation of sinkhole development and collapse in central florida. *National Cave and Karst Research Institute Conference Paper*, 2015.
- [103] Geoffrey Ingram Taylor. Fluid flow in regions bounded by porous surfaces. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 234(1199):456–475, 1956.
- [104] Morris Tenenbaum and Harry Pollard. *Ordinary Differential Equations*. Dover, 1985.
- [105] Ker Than. Guatemala sinkhole created by humans, not nature. *National Geographic*, 2010. URL <https://www.nationalgeographic.com/news/2010/6/100603-science-guatemala-sinkhole-2010-humans-caused/>.
- [106] Anna-Karin Tornberg. *Interface tracking methods with application to multiphase flows*. PhD thesis, Numerisk analys och datalogi, 2000.
- [107] Steven Vogel. *Life in moving fluids: the physical biology of flow*. Princeton University Press, 1996.
- [108] Qingpu Wang and Oliver Steinbock. Flow-assisted self-organization of hybrid membranes. *Chemistry–A European Journal*, 25(44):10427–10432, 2019.
- [109] Qingpu Wang and Oliver Steinbock. Flow-assisted self-organization of hybrid membranes. *Chemistry–A European Journal*, 25(44):10427–10432, 2019.
- [110] Qingpu Wang and Oliver Steinbock. Materials synthesis and catalysis in microfluidic devices: Prebiotic chemistry in mineral membranes. *ChemCatChem*, 12(1):63–74, 2020.
- [111] Qingpu Wang, Megan R Bentley, and Oliver Steinbock. Self-organization of layered inorganic membranes in microfluidic devices. *The Journal of Physical Chemistry C*, 121(26):14120–14127, 2017.

- [112] E Weinan. *Principles of multiscale modeling*. Cambridge University Press, 2011.
- [113] MF Wheeler and CN Dawson. An operator-splitting method for advection-diffusion-reaction problems. Technical report, 1987.
- [114] Steven P Wrenn, Stephen M Dicker, Eleanor F Small, Nily R Dan, Michał Mleczko, Georg Schmitz, and Peter A Lewin. Bursting bubbles and bilayers. *Theranostics*, 2(12):1140, 2012.
- [115] Bin Xu, Ming Ye, Shuning Dong, Zhenxue Dai, and Yongzhen Pei. A new model for simulating spring discharge recession and estimating effective porosity of karst aquifers. *Journal of Hydrology*, 562:609–622, 2018.
- [116] Changjiang Yang, Ramani Duraiswami, Nail A Gumerov, and Larry Davis. Improved fast gauss transform and efficient kernel density estimation. In *null*, page 464. IEEE, 2003.
- [117] Xiaogang Yang, Yuezheng Gong, Jun Li, Robert S Eisenberg, and Qi Wang. Quasi-incompressible multi-species ionic fluid models. *Journal of Molecular Liquids*, 273:677–691, 2019.
- [118] AP Yarlagadda and AP Yoganathan. Experimental studies of model porous media fluid dynamics. *Experiments in Fluids*, 8(1-2):59–71, 1989.
- [119] Jun Zhang, Sheng Yan, Dan Yuan, Gursel Alici, Nam-Trung Nguyen, Majid Ebrahimi Warkiani, and Weihua Li. Fundamentals and applications of inertial microfluidics: a review. *Lab on a Chip*, 16(1):10–34, 2016.
- [120] Tianyu Zhang and Isaac Klapper. Mathematical model of biofilm induced calcite precipitation. *Water Science and Technology*, 61(11):2957–2964, 2010.
- [121] Edward D Zisman and Daniel J Clarey. Problems associated with the use of compaction grout for sinkhole remediation in west-central florida. *National Cave and Karst Research Institute Conference Paper*, 2013.

BIOGRAPHICAL SKETCH

Patrick was born in Atlanta, GA in 1992. He attended Chamblee Charter High School and went to college at Florida State University intending to study political science. After spending his first year living abroad at the FSU campus in London, UK, he decided to switch majors to Psychology. In his third year, he added a dual major with Applied and Computational Mathematics. In his senior year, Patrick participated in an NSF REU with Arash Fahim working on financial mathematics. He graduated in 2015 after completing an honors thesis in Zuoxin Wang's neuroscience lab studying histone acetylation in prairie voles.

Patrick decided to pursue mathematics at the graduate level, entering the Biomathematics program at Florida State. In the fall of his second year, Patrick began a fruitful collaboration as a Research Assistant with Kourosh Shoele at FSU's College of Engineering to study the effect of variable viscosity on ciliary locomotion [92, 38]. In the spring of his second year, Patrick was awarded a NSF GRFP Fellowship. After obtaining his Masters degree, Patrick chose co-advisors Nick Moore and Nick Cogan due to their similar interests in biologically-motivated fluids problems. Patrick obtains his PhD in 2020 with the current dissertation. Patrick will continue in academia at a postdoctoral position in the FSU Birdsong Research Group with Frank Johnson.

Besides mathematics, Patrick enjoys reading, pursuing miscellaneous hobbies, and spending time with his wife and daughters.