

Fronts and Filaments: Methods for
Tracking and Predicting the Dynamical
Effects of Advection on Excitable
Reactions

by

Thomas D. Nevins

Submitted in Partial Fulfillment of the

Requirements for the Degree

Doctor of Philosophy

Supervised by

Professor Douglas H. Kelley

Department of Physics and Astronomy
Arts, Sciences and Engineering
School of Arts and Sciences

University of Rochester
Rochester, New York

2019

To my beloved wife, who braved this adventure with me,
and to my parents, who passed their love of knowledge, wisdom, and discovery
on to me.

Table of Contents

Biographical Sketch	vi
Acknowledgments	viii
Abstract	xii
Contributors and Funding Sources	xiv
List of Tables	xv
List of Figures	xvi
1 Introduction	1
1.1 Motivation	1
1.2 Context	6
1.3 Fluid Dynamics	14
1.4 Reaction Dynamics	17
1.5 Experimental Methods	19
1.6 Thesis Overview	22

2	Front Tracking Method for Quantifying Reaction-Diffusion	23
2.1	Introduction	23
2.2	Relating Thickness and Speed to Diffusivity and Reaction Rate	26
2.3	Front Tracking	31
2.4	Tracking Fronts in Validation Data	36
2.5	Tracking Fronts in Simulation Data	38
2.6	Tracking Fronts in Experimental Data	42
2.7	Summary and Future Work	48
2.8	Acknowledgments	50
3	Front Tracking Including Advection	52
3.1	Introduction	52
3.2	Separating Flow Velocity from Chemical Velocity	55
3.3	Front Tracking Algorithm	58
3.4	Tracking Fronts in Simulation Data	63
3.5	Experimental Advection-Reaction-Diffusion Devices	70
3.6	Tracking Fronts in Experimental Data	73
3.7	Summary and Future Work	82
3.8	Acknowledgments	85
4	Vertical shear alteration of chemical front speed in thin-layer flows	87
4.1	Introduction	87
4.2	Simulations of Single-Layer System	93
4.3	Simulations of Hele-Shaw System	104
4.4	Simulations of Two-Layer System	109

4.5	Experimental Apparatus	114
4.6	Experimental Results	117
4.7	Conclusions	124
4.8	Acknowledgements	127
5	Optimal Stretching and Blowout in Excitable Reactions	129
5.1	Introduction	129
5.2	Radial front growth enhancement	133
5.3	Observation of optimal stretching	137
5.4	Understanding optimal stretching	146
5.5	Summary and Implications	164
5.6	Acknowledgements	166
6	Conclusion	168
6.1	Future Directions	173
	Bibliography	178

Biographical Sketch

The author, Thomas Nevins, was born in Kalamazoo, Michigan before moving to Chippewa Falls, Wisconsin where he lived from a young age through high school. In 2010, the author began attending the University of Wisconsin - Eau Claire (UWEC). While at UWEC the author benefited from an abundance of research opportunities and worked with several different advisors at various times: Dr. Thomas Lockhart, Dr. Chris Ahrendt, and Dr. Paul Thomas of UWEC, Dr. Wei Jiang Yeh of the University of Idaho, and Dr. Robert Wagner of Argonne National Laboratory. In 2013, the author received the Goldwater Research Scholarship for these research experiences, and he was named the outstanding senior in Physics for UWEC. In 2014, the author graduated with Bachelor of Science degrees in Physics, and in Applied Mathematics as Summa Cum Laude. The author received departmental honors in both Physics and Math, and completed the University Honors Program.

Later in 2014, the author followed his passion for research to the University of Rochester for a Doctorate in Physics. During the first year he was supported by a teaching assistantship. After this the author joined Dr. Doug Kelley's lab and Dr. Pierre Gourdain served as his advisor within Physics. Following completion of the preliminary exam, the author received a Masters of Science in Physics in 2016. In fall 2016, the author began working on a Masters of Science in Technical Entrepreneurship and Management, which was completed in December of 2018.

In 2017 the author received both the Donald and Janet Bernard Fellowship, and the National Defense Science and Engineering Graduate Fellowship.

List of publications:

1. Nevins, T.D., and Kelley, D.H., “Optimal stretching in advection-reaction-diffusion,” *Physical Review Letters* **117**,164502 (2016).
2. Nevins, T.D., and Kelley, D.H., “Front tracking for quantifying advection-reaction-diffusion,” *Chaos* **27**, 043105 (2017).
3. Wang, J., Tithof, J., Nevins, T.D., Colon R.O., and Kelley, D.H., ”Optimal stretching in the reacting wake of a bluff body,” *Chaos* **27**, 123109 (2017).
4. Nevins, T.D., and Kelley, D.H., “Front tracking velocimetry in advection-reaction-diffusion systems,” *Chaos* **28**, 043122 (2018).
5. Plog, B.A., Mestre, H., Olveda, G.E., Sweeney, A.M., Kenney, H.M., Cove, A., Dholakia, K.Y., Tithof, J., Nevins, T.D., Lundgaard, I., Du, T., Kelley, D.H., Nedergaard, M., “Transcranial optical imaging reveals a pathway for optimizing the delivery of immunotherapeutics to the brain,” *JCI Insight* **3** (20),120922 (2018).
6. Nevins,T.D., and Kelley, D.H., “Vertical shear alters chemical front speed in thin-layer flows,” Under Review (2019).

Acknowledgments

I have been absolutely blessed with the people I work with, and who have supported me over the years. Many deserve thanks for helping me complete this thesis, and others deserve thanks for making me the person who could complete a thesis in the first place. I personally like thanking people so I'm going to thank as many as I can think of.

To start, I would like to thank my PhD advisor Dr. Doug Kelley and the Mixing Lab as a whole. Dr. Kelley is an incredible leader, organizer, scientist, and teacher. He has an infectious passion for learning, is eager to teach new students, listens to and encourages new ideas, and he helped turn a lab full of strangers into close friends. So thank you Dr. Kelley for introducing me to the weird worlds of fluids and nonlinear dynamics, for coaching me on research, presentations, and papers in the early days, for trusting me to do it myself in the late days, for being a mentor in life as well as a mentor in science, and for putting up with me calling you Dr. Kelley instead of Doug through it all.

In addition to always helping brainstorm across projects with me, the Mixing Lab team has been some of the best friends I've ever had. I want to thank Rakan Ashour from never shying away from big questions, and being one of the most admirable people I've ever known; Jeff Tithof for always being willing to share ideas, pass on know-how, and for his passion for creativity; Ibrahim Mohammad for bringing out the best in me, and his inspiring ambition; Bitong Wang for his

never ending kindness; Ruy Ibanez for helping me try new things, and for being a font of knowledge.

A number of undergraduate members of the Mixing Lab also helped make this thesis a reality, foremost among them being Jinge Wang whose problem solving skills and inquisitiveness was always an asset. Brian Knisely, Rony Colon, and Ben Martell also provided support to this thesis. I would also like to thank all the other undergraduates from the Mixing Lab for being a huge part of making the lab into a community: Logan Bashford, Rebeca Toro Garza, Mira Bodek, Andrew Fianu, Frederico Hama, Catherine Knox, Kimberly Llajaruna, Meghan Patrick, and Jack Quinlivan.

Also at the university, I would like to thank Jim Alkins who was always willing to provide advice for my build projects; Laura Blumkin for being a tireless advocate for Physics department graduate students, and handling all my administrative concerns; Scott Russell helped me with designs of experimental apparatus; Christine Pratt for help with Chematix and poster creation; Danielle Daniels for welcoming me to the Kearns Center, and guiding me through teaching Upward Bound workshops.

Just as I won't stop learning after this thesis, I did not start learning when I arrived at the University of Rochester. Any list of people who deserves thanks for this thesis would be incomplete without the teachers I had at the University of Wisconsin - Eau Claire (UWEC). I would like to thank Dr. Paul Thomas and Dr. Thomas Lockhart of the Physics department for cultivating my research skills and interests. I'd also like to thank most of the Physics department for always having their doors open, being willing to pass on their advice, and sharing some Laffy Taffy jokes, especially: Dr. Eric Hendrickson, Dr. Matt Evans, Dr. Nathan Miller, Dr. Kim Pierson, Dr. George Stecher, and Dr. Scott Whitfield.

But just as I didn't start learning when I reached the University of Rochester, I didn't start when I reached University of Wisconsin - Eau Claire. While profes-

sors at the undergraduate and graduate level often get applauded for their part in a thesis, they are the beneficiaries of the work of dozens of teachers and mentors before them. These teachers shape our values, work ethic, personality, and ambitions. So I would like to thank all the teachers who were crucial in getting me to this point: Mary Lange (Kindergarten), Sarah Daniels (Grade 1), Sandra Rassbach (Grade 2), Mohamed Elgindi (After-School Math Program), Mary Jeneman (Grade 3), Diane Decker (Grade 4), Cindy Lee (Grade 5), Mr. Borman (Grade 6, Science), David Hutzler (Grade 7, Science), Dawn Kuehn (Grade 8, Science), Jason and Michelle Dutton (Karate), Ronald Buckles (Grade 9-12, Choir), Mark Cloutier (Grade 9, Biology), Jeffrey Keding (Grade 9, Pre-Calculus), Daniel Loomis (Grade 9, Earth Science), Rolan Rohde (Grade 10, Chemistry), Timothy Fox-Meyer (Grade 8, Trigonometry and Grade 10, Calculus), Paul Nevins (Grade 11 and 12, Physics and also my dad), Jane Rubisch (Grade 12, Biology), Robert Kuchta (Grade 12, Chemistry), and all the other teachers I learned from growing up.

My brother, my parents-in-law, my aunts and uncles, and cousins were all very supportive of me and my wife as we attempted to get our PhD's at the same time. They also provided advice which was sometimes useful and sometimes hilarious. I especially want to thank my parents, Paul and Alison Nevins. Anyone who knows the two of them should be unsurprised that I made it this far. Before I was ten I was spending inservice days with mom in the clinic laboratory, where I was surrounded by microscopes and lab coats. Meanwhile, dad did not wait until I had him for AP Physics in high school to start teaching me anything I would want to learn about. The two of them supported me and challenged me every step of the way.

Finally, and most importantly, I want to thank my wife Mandy. It was a challenging road to go from getting married straight into graduate school for both of us, but we did it. She took the experience from stressful and difficult to exciting

and wonderful. She has even been my secret weapon when I'm particularly stuck on a research question. Above all I want to thank her for her bravery, support, and love.

Abstract

In advection-reaction-diffusion (ARD) systems, the spreading of a reactive scalar can be significantly influenced by the flow field in which it grows. These systems, which range from chemical reactions to ocean plankton populations, become very difficult to understand, predict, and control, because of the complex interaction between fluid flow and nonlinear chemical reactions. Greater understanding of these systems can be obtained through models that isolate the most important behaviors. I present results which model ARD systems as reaction fronts, or stretched filaments carried by the flow. I perform experiments and analysis on the excitable BelousovZhabotinsky (BZ) reaction in a nearly two-dimensional system with flows both steady and unsteady in time. I create and utilize a reaction front tracking tool, and use Lagrangian methods to understand the importance of the flow gradient to reaction growth.

To start, I consider the motion of the reaction fronts that lie at boundaries between reacted and unreacted which can quantify spreading. I present an algorithm to measure local speed and thickness of reaction fronts. These quantities can in turn measure diffusivity and reaction rate. The algorithm is then expanded to measure front speed in the presence of flow, and test reaction front models. After identifying that the chemical speed appears to depend on flow, I show that this is due to depth shear through the thin experiment layer, and I suggest improvements to make future experiments more two dimensional.

After that I investigate how BZ tends to gather at the same value of the Lagrangian stretching in both laminar and chaotic, time-varying flows. I hypothesize that the optimal stretching is due to moderate stretching promoting reaction, but large stretching causing blowout. This optimal stretching value is dependent on the reaction, but not fluid flow. I investigate the causes of this by considering three effects of the flow gradient: perimeter change, thickness change, and filament stretching.

My results offer new avenues to research ARD systems, and new insights to allow models to more closely match ARD systems, especially in time dependent flows.

Contributors and Funding Sources

This work was supported by a dissertation committee consisting of Professor Douglas H. Kelley (advisor) of the Department of Mechanical Engineering, Professor Pierre Gourdain, Professor Eric Blackman, and Professor Nicholas Bigelow of the Department of Physics and Astronomy, as well as Professor Jessica Shang of the Department of Mechanical Engineering. Professor Matthew Yates of the Chemical Engineering department was the chair of the defense committee.

Financial support for this thesis was received from the Department of Defense (DOD) through the National Defense Science and Engineering Graduate Fellowship (NDSEG) Program. Additional funding was provided through the University of Rochester's initial funding of Dr. Doug Kelley.

Dr. Kelley was the originator of the two camera experimental setup and created a few of the figures which are attributed to him accordingly. Jinge Wang contributed to Chapter 5, with his work in reactions in the wake of bluff bodies.

List of Tables

2.1	Parameters chosen by the algorithm's user.	33
4.1	Parameters measured from experiments and used for simulations.	98

List of Figures

1.1	Reaction-diffusion front in 1D.	9
1.2	(a) Phase diagram for excitable systems.	20
2.1	Example results of front tracking	32
2.2	Algorithm described for measuring front velocity.	34
2.3	Algorithm for measuring front thickness, shown using validation data.	35
2.4	Comparison of true and measured front speed and thickness.	37
2.5	Robustness of front tracking to noise.	39
2.6	Sensitivity of front tracking to input parameters.	40
2.7	Accuracy of diffusivity and reaction rate measurements made by tracking fronts in simulation results.	42
2.8	Histograms of front speeds (a) and thicknesses (b) gathered from an outward-propagating front.	44
2.9	Effects of ferroin concentration on (a) front speed, (b) front thickness, (c) diffusivity, and (d) reaction rate.	46
2.10	Measured concentration profile of chemical waves in the BZ reaction.	49
3.1	Example results of front tracking with flow, showing experimental product concentration snapshots (gray-scale).	57

3.2	Cartoon of a single step of front tracking.	61
3.3	Validation of front tracking by simulating the advection-reaction-diffusion equations in uniform flow.	66
3.4	Validation of front tracking by simulating an advancing front near a hyperbolic critical point.	68
3.5	Laboratory advection-reaction-diffusion devices and typical data.	72
3.6	Front velocity measurements are consistent with measured front locations in simulation and in three different experimental devices.	74
3.7	Variation of chemical speed's summary statistics.	76
3.8	Distributions of the change in chemical velocity components caused by changing the time step dt from 0.5 s to 1 s.	77
3.9	Concentration fields in vortex-dominated experimental flows with small (a) and large (b) characteristic forcing length scales.	79
4.1	Experimental measurements of chemical reaction fronts using front tracking.	91
4.2	Depth dependence of streamwise velocity in a uniform thin-layer flow.	96
4.3	Fronts propagating in supporting and opposing thin-layer flow, according to the Eikonal equation.	99
4.4	Front propagation in thin-layer flow, varying with flow speed, according to the Eikonal equation.	100
4.5	Apparent chemical front speed in a simulated thin-layer flow, determined using depth-averaged concentration.	102
4.6	Dimensionless front propagation in thin-layer flow.	104
4.7	Depth dependence of streamwise velocity in a Hele-Shaw style, pressure driven flow.	105

4.8	Front propagation in Hele-Shaw flow, varying with flow speed, according to the Eikonal equation.	106
4.9	Apparent chemical front speed in simulated Hele-Shaw flow, determined using depth-averaged concentration.	108
4.10	Dimensionless front propagation in the Hele-Shaw flow.	109
4.11	Depth dependence of streamwise velocity in a uniform thin-layer flow including a lubrication layer.	111
4.12	Front propagation in two-layer flow, varying with flow speed, according to the Eikonal equation.	112
4.13	Apparent chemical front speed in simulated two-layer flow, determined using depth-averaged concentration.	113
4.14	Dimensionless front propagation in the two-layer flow.	114
4.15	Diagram of the experimental apparatus.	116
4.16	Mean flow in a single-layer experiment.	118
4.17	Mean flow in a two-layer experiment.	119
4.18	Snapshots of the BZ reaction in a single-layer experiment (a - c) and a two-layer experiment (d - f).	120
4.19	Variation of apparent front speed with the front-normal flow speed $\mathbf{u} \cdot \hat{\mathbf{n}}$ in single-layer experiments and simulations.	122
4.20	Variation of signed front speed with the front-normal flow speed $\mathbf{u} \cdot \hat{\mathbf{n}}$ in two-layer experiments and simulations.	123
4.21	Variation of front thickness with the front-normal flow speed $\mathbf{u} \cdot \hat{\mathbf{n}}$ in single-layer experiments.	124
5.1	Schematic of the channel flow experiment, viewed from above. . .	132
5.2	Snapshots of ARD experiments with the BZ reaction.	134

5.3	Propagation of reacted regions toward vortex centers.	135
5.4	Speed toward vortex centers increases with Reynolds number for reaction fronts, but not for a passive scalar.	136
5.5	Observed reaction state and measured stretching S overlaid, for (a) $Re = 35$, (b) $Re = 74.9$, (c) $Re = 135$, and (d) $Re = 218$	140
5.6	Observed reaction state and measured stretching S overlaid, for the open flow experiment at $Re = 586$	141
5.7	Dependence of reaction state on stretching in the single layer vortex flow.	144
5.8	Dependence of the reaction state on stretching in the open flow experiment.	145
5.9	Results of simulations of the Oregonator growing in a strainless vortex.	151
5.10	Comparison of full ARD and Eikonal equation in the uniform strain situation.	154
5.11	Trends of front speed and front thickness in uniform strain simula- tions at Da_I from 1.5 to 3.	156
5.12	Results of two simulations of the Oregonator growing in a hyper- bolic strain in 1D.	160
5.13	Plots of simulation results near Da_{Ic}	162

1 Introduction

1.1 Motivation

A small amount of one chemical reactant is dropped into a pool of another. The first reactant diffuses into the bulk of the second, and they begin to react and create a product. If this was the end of the story, the reaction could take a very long time to complete: introductory chemistry classes would get tedious, cooking with more than two ingredients would be impossible, and the world's chemical industry would grind to a halt. Fortunately, mixing due to fluid flow can cause all of these processes to happen an order of magnitude faster than diffusion alone would allow. Intuitively it is easy to see why mixing should matter: chemical reaction cannot happen until the two reactant molecules can collide, and that cannot happen if the reactants are not mixed together, say by stirring them together. Mixing through fluid flow is so common that it can be easy to take for granted. Yet despite its intuitive and fundamental nature, mixing — especially mixing of substances that grow over time — has proven to be a challenging problem to analyze. Questions about predicting or controlling a reaction's growth are difficult to handle in even the simplest situations.

Mixing is such a difficult subject for a number of reasons. Reactions occur

when the reactants are within a molecular distance apart, but fluid flows are usually forced at a much higher length scale. Vigorous stirring is usually more effective for mixing reactants, so mixing situations of interest often involve all the complexities of modeling turbulence. For a reaction the turbulent cascade must reduce the reaction length scale all the way from a forcing length scale, down to the Batchelor scale, n , where diffusion can take over the mixing [1]. In aqueous reactions typically material diffusivity $D \sim 10^{-9}$ m²/s, while kinematic viscosity $\nu \sim 10^{-6}$ m²/s, meaning that this Batchelor scale is often even smaller than the Kolmogorov scale, where viscosity dominates flow eddies. This poses an incredible challenge to simulating these mixing problems. On top of this, chemical reactions are nonlinear most of the time [2]. This causes unusual behavior where the same stirring of a different reaction does not have the same effect. Changing the stirring can cause the reaction to transition between stable, periodic, and chaotic states or change the final reaction product's chirality from left to right handed [3]. It's even possible for increased stirring to result in a longer reaction time [4] or for fast stirring to reverse a reaction [5].

In spite of these complexities, the problems of understanding chemical reactions in fluid flows are common to many systems. Any system in which some scalar quantity is carried by a fluid flow, spreads out over time, and can grow has these same fundamental problems: scale separation, turbulence, nonlinear growth, and the interaction between reaction growth and flow. Systems that fall into this category are considered advection-reaction-diffusion (ARD) systems. They are made of scalar fields (typically material concentration fields) which change dynamically due to the combined effects of advection (the motion of the domain the field is in), reaction (the growth of the scalar), and diffusion (the spreading of the concentration field from high concentration to low). Mathematically the time evolution of these scalars is then written according to the ARD equation,

$$\frac{dc_i}{dt} = -\nabla \cdot (\mathbf{u}c_i) + D_i \nabla^2 c_i + k_i F_i(c_1, \dots, c_n) \quad (1.1)$$

where each c_i is a field that provides the concentration of a species at each x,y , and z position over time, t . In general, there can be multiple species which all interact with each other. Here each species may have a different diffusion coefficient D_i , reaction rate constant k_i , and reaction term F_i , which is a function of the local concentration of all the species. All species experience the same flow \mathbf{u} . For simplicity, I have assumed D_i is truly constant and therefore diffusion can be written this way. I have also assumed that the reaction term F_i depends only on the other concentration fields.

This equation is relevant to a variety of applications, each of which has important unanswered questions which require more general results and tools to solve. In industrial mixing [6], misunderstanding mixing results in inefficient mixing which wastes energy and can cause a failure to scale up new processes. In 1989 estimates of the cost of bad mixing to the chemical process industry were on the order of \$ 1 Billion [7]. For applications like the manufacturing of pharmaceuticals, complex chemistry can cause unusual behavior when stirred [8]. Microfluidic reactors also perform mixing of chemicals, so an understanding of interaction between reaction dynamics and flow dynamics is necessary to improve them [9]. In packed bed reactors, flows can be very complex and very important to reaction growth [10, 11]. Yet research in industrial reactive mixing is often restricted to detailed empirical studies which measure how well a given reactor design mixed the reaction [12, 13]. These studies cannot identify the cause of good or bad mixing. The field has developed a number of techniques to help measure mixing time after the reactor or simulation has been built [14], but general results about what flow mixing does to reaction growth are not rigorous. Related to these problems, Ascanio [14] says “A question rarely put forward by researchers and processors involves whether the method of mixing operation is optimum”. Working towards better understanding of mixing, and simplified modeling could help reduce waste, increase energy efficiency, and enable new industrial processes in the future.

Another area where the ARD equation applies is the study of combustion reactions. From gas flames [15] to the dynamics of wildfires [16, 17], the development of combustion can have strong dependence on the surrounding fluid flow, in this case the flow of air containing oxygen or fuel. This dependence can be very complex, as anyone who has ever tried to start a campfire knows. In starting a campfire, light fanning will help a spark catch and grow, but fast fanning will extinguish a small flame. To control wildfires or optimize combustion reactions in engines and other industrial processes, understanding the complex interaction between reaction and flow is necessary. Combustion reactions also have the added complexity that the ongoing reaction in combustion drives fluid flow, because combustion is strongly exothermic. I will be experimenting entirely with reactions that have little or no effect on the surrounding flow, but it may be possible for future researchers to generalize some or all the results of this thesis to combustion. There are an enormous variety of different possible flame configurations, but applications of many of them would benefit from a simplified way to model their growth, and ways to predict or control their dynamics over time.

Finally, a variety of research has gone into understanding plankton and algae species as ARD systems [18–25]. Plankton are carried by turbulent currents, they consume nutrients to grow and reproduce, and they spread out over time. It may seem odd that plankton can be thought of this way, because a single plankton is much larger than a single molecule. Approximating it as a continuum field may seem questionable. However, the separation in length scales between the flow forcing in lakes and oceans (> 10 km in the ocean) and the size of plankton (< 1 mm) is actually *the same* as the length scale separation between lab scale experiments (~ 10 cm) and the smallest molecules (0.1 nm), both separated by a factor of $\sim 10^7$. Due to the similarities between plankton growth and reactions, I will refer to the growth of plankton as a “reaction” even though the growth is biological in nature. Algae and phytoplankton are the biggest producers in the

ocean ecosystem, so the health of the oceans is dependent on how they grow. Phytoplankton in particular also make up the world's largest carbon sink [26]. Algae can sometimes have negative effects on marine environments through the growth of harmful algae blooms. Despite their importance, big open questions remain in this field such as the so called "Paradox of the plankton", which refers to the fact that many plankton species fill the same niche as other plankton species, and yet they coexist in the same ocean [27, 28]. Biological systems can be very complex, but flow structures do seem to have an important impact on growth [19–21], and further understanding of the interaction between flow and growth is needed to make more predictions possible.

While each of these systems have their own unique elements, the ARD equation will be similar for each. When the math is the same, the physics is also the same. In this thesis, I aim to test or develop simplified models of ARD systems and to thereby construct a more general understanding of ARD phenomena. I will do this by developing new experimental tools and analyses, characterizing the accuracy of the tools, and then comparing the results of my experiments to existing general models of ARD. To begin this introduction I will cover past results in ARD research, especially pertaining to the dynamical systems understanding of ARD systems, in section 1.2. Next in section 1.3 I will cover the basics of the flow field, \mathbf{u} , that are needed for the rest of the thesis. Similarly in section 1.4 I will cover the basics of chemical kinetics, and introduce excitable reaction kinetics which I will be using in experiments and focusing on. Excitable reactions are analogous to plankton growth [22] and combustion [2]. Then in section 1.5 I will cover basics of ARD experiments that I will use in all the chapters. Finally the introduction will conclude by outlining the remainder of the thesis, and stating the goals of the work.

1.2 Context

To understand why simplified modeling of ARD systems is important to these applications, consider the difficulty in directly simulating an ARD system. Without any simplifying assumptions, this means simulating at least two 3D fields over time (flow, \mathbf{u} , and reaction state, c). In real world situations, because the time it takes a reaction to proceed is often of interest, or because many ARD systems exist in turbulence which is dynamically changing, many questions about these 3D fields are about their transient dynamics. Chaotic advection and nonlinear reaction means these systems can be sensitive to initial condition, and this can make it hard to find general results. Numeric errors can also quickly lead to qualitatively different results, especially since we are interested in transient dynamics instead of characterizing final states. To obtain results which keep numeric errors small, a fine spatial grid is required, since diffusion is also very small (at least in aqueous environments). This in turn results in a very fine time-stepping for stability. The result of all these complications is that directly simulating eq. 1.1 in real environments is usually either computationally expensive, inaccurate in quantities such as reaction growth rate, or difficult to generalize. For this reason, many simulation [8, 29–32] and experiments [11, 33, 34] have to do with steady state behavior and critical transitions from one steady state to another in specially conceived situations that admit steady solutions. Some studies have investigated transient growth in simple flows in detail [35–37], and in more complex flows [38], but overall the study of transient dynamics of ARD has been restricted by the difficulty of finding the most important aspects.

The first step I will take to simplify the ARD system is to narrow my scope. I will focus my attention on just the behavior of the reaction given the fluid flow. This means ignoring any effect of the reaction which drives fluid flow, such as releasing heat to drive convection. I will also focus on two-dimensional (2D) sys-

tems, where chaotic advection is still possible, but measurements and simulations are easier to perform. This simplifying assumption is reasonable if I am modeling phytoplankton in the ocean, because flow on the ocean surface is roughly two-dimensional. Unlike plankton however, my reaction will not have any capability to move on its own, it only moves according to diffusion. While these are some significant assumptions, I still allow any laminar or turbulent flow, nonlinear reactions, and include the effects of diffusion. This niche still needs exploration, so it serves as a starting point upon which further complications can be added to reach the applications I mentioned above.

Within ARD systems with these assumptions, diffusion and advection both act through the gradient of c to change concentration. Meanwhile, reaction terms usually take a chemical state from an unreacted state towards a reacted state. The unreacted state can be linearly stable or unstable, and the reacted state is linearly stable or is a state that the system spends a long period of time near. All three terms tend to be near zero in areas that are completely reacted or unreacted. Most of the interesting behavior takes place at the interface between these two states, which is called the *reaction front*. Modeling the evolution of the field c as the motion of a reaction front surface drastically reduces the system complexity, and matches intuition. Reaction fronts specify that the reaction state inside them is approximately all reacted, and outside is approximately unreacted. To be valid, front modeling assumes that the reaction changes from unreacted to reacted in a distance much smaller than the curvature of the front, or the domain size. Complexity is reduced because the complex partial differential equation on an abstract field is reduced to a ordinary differential equation based on geometry. The combined effects of reaction and diffusion are modeled as a “front speed” which moves the front outwards. Intuition is improved, because while it is difficult to talk about how a flame’s temperature, reactants, and products vary over space, it is much easier to say “just the lab is on fire”. The statement “just the lab is on

fire” is implicitly a statement using reaction front modeling, because it identifies where the edge of a reacted region is located. Thinking of reactions this way allows me to identify a position for a reaction field, and also to observe how that position changes over time, which gives rise to a velocity. The reaction field near a front along a perpendicular profile can also be characterized to regain a little more of the reaction field’s details.

The modeling of reaction fields as fronts goes back to the studies of front propagation without flow. In the 1930s, two papers [39, 40] laid the groundwork by studying how reaction-diffusion fronts advanced in a simple one-dimensional (1D) situation. Fisher [39] was able to show that the solution for second order reactions was a traveling sigmoid wave of constant shape which moved at a constant speed, while Kolmogoroff *et al.* [40] showed that these fronts converged to a constant speed in many reaction types. From these results, the field of front modeling was born, and these fronts were named FKPP fronts. The solutions showed that only a few parameters were needed to give the overall behavior of a reaction-diffusion system: front position, front speed, and front profile width. Figure 1.1 shows an example front profile, which Fisher calculated. These results can then be expanded to higher dimensions, by noticing that in 3D ARD systems any point will be part of a level surface. Along that level surface, concentration is constant, so all the interesting front propagation occurs on a 1D line perpendicular to this surface. Along with a curvature correction [41, 42], this created the Eikonal equation for stagnant reaction front growth

$$\mathbf{w} = v_0 \hat{\mathbf{n}} + D\kappa. \tag{1.2}$$

Here I use the shorthand $\mathbf{w} = \partial \mathbf{x} / \partial t$ where x is the position of the front. The stagnant front speed derived by Fisher is v_0 , $\hat{\mathbf{n}}$ is the front normal, and the curvature is κ . In the experiments I will be doing, $D \approx 0.001 \text{ mm}^2/\text{s}$, so the curvature correction is negligible. However, it is worth mentioning that it tends to smooth out cusps in the reaction very quickly.

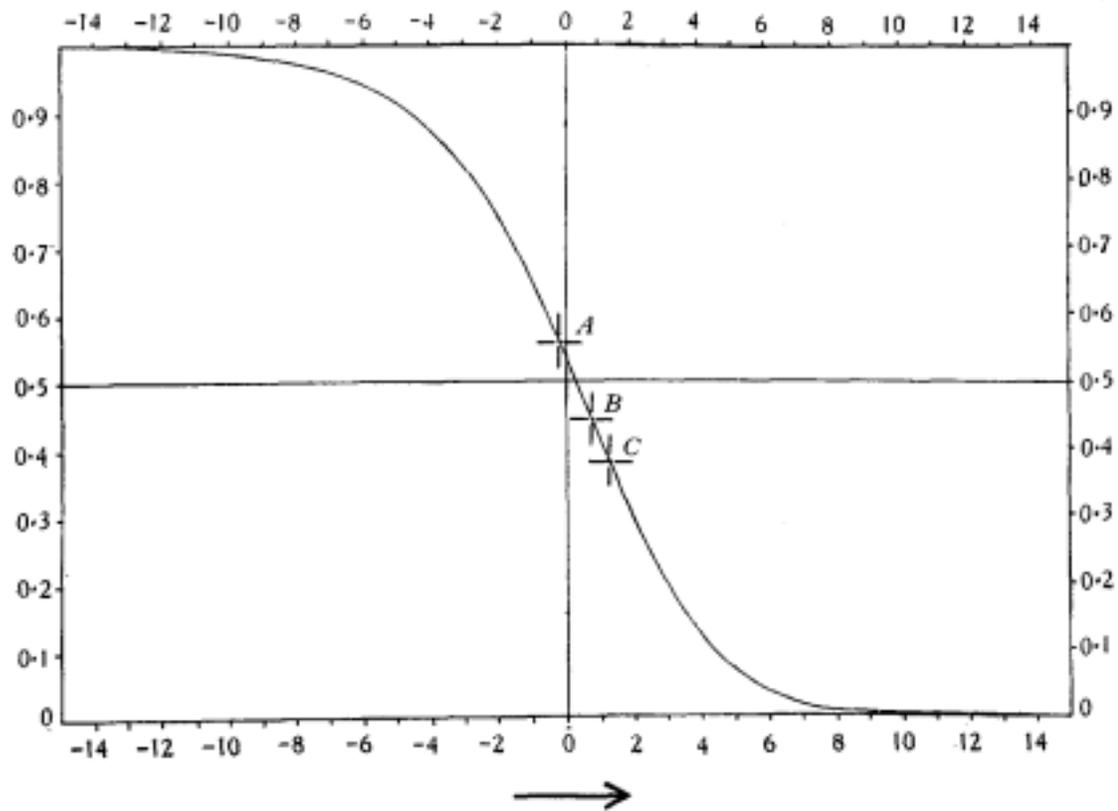


Figure 1.1: Reaction-diffusion front in 1D. The vertical axis is concentration, and the horizontal is position. The front advances over time in the direction indicated by the arrow. Work of Fisher [39].

When advection is added, other dynamics must be added to eq. 1.2, because the flow carries the reaction along. For example, if a reaction propagates in a spatially uniform flow it must be identical to a reaction in a stagnant medium observed from a moving reference frame, according to Galilean invariance. To account for this, the easiest adjustment is to simply add the advection velocity at the front position on top,

$$\mathbf{w} = v_0 \hat{\mathbf{n}} + D\kappa + \mathbf{u}. \quad (1.3)$$

This step is surprisingly subtle, and involves a large number of assumptions, such as that v_0 be independent of flow, that the spatial gradient of \mathbf{u} be small near the front, and that v_0 and κ don't change much across the front. For a derivation of this see Spangler and Edwards [43]. This is the Eikonal equation in flow, and represents an extremely simple way to model an ARD system.

To perform a simulation using the Eikonal equation it is useful to rewrite the equation in terms of differential equations for the position of a set of front elements, and their orientations. In a 2D ARD system this creates a 3D phase space (2 position, 1 angle), and in 3D the phase space is 5-dimensional (3 position, 2 angles). The methodology for doing this is outlined in Mitchell and Mahoney [44], Mahoney *et al.* [45], and Mahoney and Mitchell [46], in which they obtain the equations for a front in a 2D system as

$$\dot{x} = u_x + v_0 \sin \theta \quad (1.4)$$

$$\dot{y} = u_y - v_0 \cos \theta \quad (1.5)$$

$$\dot{\theta} = -\frac{\partial u_x}{\partial x} \sin \theta \cos \theta - \frac{\partial u_x}{\partial y} \sin^2 \theta + \frac{\partial u_y}{\partial x} \cos^2 \theta + \frac{\partial u_y}{\partial y} \sin \theta \cos \theta. \quad (1.6)$$

Here, θ gives the angle that the front tangent makes with the x axis. Specifically the tangent is such that it is a 90° counterclockwise rotation from the front normal, $\hat{\mathbf{n}}$, which points down the gradient. These equations can be extended to three-dimensions (3D) [47], but I will be using them only in two-dimensional systems.

These equations turn the ARD partial differential equation into a model system of ordinary differential equations.

Now for this model to be useful it should do three things, (1) be simpler than the original system, (2) provide insight into questions surrounding ARD dynamics, and (3) be an accurate representation of the real physics. By virtue of being an equation for dynamic motion, it is much simpler than the original ARD equation. Furthermore, since it provides easily digestible information about how a reaction “moves”, it meets the insight requirement. All that remains is to determine its validity. With regard to validity, there has already been a large body of work which uses the Eikonal equation to explain experimental results.

The first type of these experiments has to do with an implication of equations 1.4-1.6. With these Mitchell and collaborators were able to simply predict points in steady and time periodic flows where fronts tended to accumulate. From these burning fixed points, they observed the surfaces in the x, y, θ phase space where fronts would tend to move along the surface. This meant these surfaces were invariant manifolds of the system, and thus they got the name “Burning Invariant Manifolds” (BIMs). One can think of these manifolds as being where flow opposes a front in one direction such that the front cannot penetrate the manifold, but can slide along it. This barrier only works in one direction in the x, y plane, because in the other direction flow supports the front. BIMs are elegantly simple, and they have been observed to accurately predict barriers to fronts in disordered flows[34], flows with a headwind[48, 49], and periodic flows [50]. A series of experiments have also used the Eikonal equation to explain the freezing of reaction fronts in a packed granular bed with flow through it [10, 11], and past an obstacle [33].

All of these tests have three limitations in common. First, the flows are steady or time periodic, but never completely time dependent. Some work has been done to expand BIMs to time dependent systems [48, 51], but it has yet to face serious

experimental testing. Second, the tests all deal with where the front stops, and not its behavior during the transient time while it converges to a BIM. As such they do not truly test the accuracy of the Eikonal equation, but instead only deal in its ability to predict the BIMs. In order to use the Eikonal equation as a model for ARD systems, evidence that it can be used to predict transient behavior is also needed. BIMs can be useful for prediction of qualitative behavior, but lack the ability to make quantitative predictions, such as mixing time, or time to fully react. Third, flows in these experiments had low flow speeds, incapable of creating large flow gradients. As a result, the flow often does not significantly interact with the reaction, making constant v_0 a reasonable assumption. I will show in chapter 5 that this should not in general be the case. Interactions between flow gradient and reaction kinetics can supersede BIMs in importance to qualitative behavior in turbulent flows. Nonetheless, it may be that the Eikonal equation does match transient motion of a front, including fronts in turbulence, or that a small modification to the front speed could account for the biggest interactions with flow.

The Eikonal equation will be studied in detail in chapters 2 through 4. This will be done by creating a “front tracking” tool to measure reaction front speed and behavior in nearly 2D flows which can be either time dependent or independent. Front tracking makes these measurements at all points in space and time of a growing reaction, so it can test transient dynamics of the reaction. Front tracking can be directly compared to the Eikonal equation to test its applicability as flow speed increases, or flow becomes time dependent. Through this analysis we discover important new experimental considerations in chapter 4, and can analyze the cause of difference with the Eikonal equation when it appears.

Another important simplified model of front dynamics aims at understanding ARD specifically in more complex flows. The “filament model” also has created general results through a simplification of the ARD equation, but in this case the

simplifying assumption is that a reaction is pulled and stretched by a chaotic fluid flow, and therefore it quickly becomes long and skinny [52]. If one considers a reference frame moving with the filament, then the equation is symmetric along the length of the filament. The dynamics can be replicated by assuming that the results for a perpendicular cross section of the filament are the same all along it. The filament model can be traced back to the work of Kierstead and Slobodkin [24] and Skellam [25] which predicted the minimum viable habitat size for a reaction that diffuses into an inhospitable environment. No reaction can survive in a habitat smaller than the KiSS length $l_c = \pi\sqrt{D/\alpha}$, where α is a normalized reaction rate, because diffusion will cause the reaction to leave the safe zone faster than reaction growth. In my experiments there will be no locations that prohibit growth, but there will be fluid flow to carry diffusing reaction product away from the filament. It has been shown that diffusion could also cause a reaction to go extinct if there is sufficient strain rate to carry reaction away faster than reaction can replace it [18], and that this may impact phytoplankton populations [19].

This model has important predictive power, relevant to ARD systems in turbulence. Importantly, this allows the inclusion of chemical kinetics, which front dynamics lacks. Predictions from the model are easier to gather than in the full system, and are easier to understand. In particular, a number of papers using the filament model show that chaotic mixing induces a transition from a filament of constant width growing exponentially to a filament which decays in concentration until it goes extinct as the stirring strength is increased. This pattern appears for a variety of chemical reaction types including excitable reactions [53], oscillatory reactions [54], plankton population models [23] and autocatalytic or bistable reactions [30]. Its predictions are more independent of details than BIM modeling, as the pattern persists independent of initial conditions [30], and independent of whether the flow is open or closed [29]. Therefore, the filament model could potentially be very powerful in predicting optimal stirring parameters, because it is

the rare general result in ARD. However, this model suffers from not being utilized in explaining experiments. Part of the reason for this is that as a very simple 1D model it seems too approximate to be useful, but more importantly, prior to the work in this thesis there were no experiments providing the simultaneous measurement of reaction state and flow needed to study the filaments. There is also some question as to how best to parameterize the stirring rate used by Neufeld and collaborators, such that it is the most relevant for reaction dynamics.

In chapter 5, I present an experimental study relevant to the filament model. My experimental flow is physically generated with complex flow structure, and the results cannot be explained by front dynamics. Yet the results agree well with predictions from the filament model for excitable reactions [53]. By comparison to the front interpretation, I use the filament model to investigate the interactions between nonlinear reactions and flow gradient. These results point to adjustments to make to front modeling to expand its validity, and suggest the relevant stirring parameter for the filament model to be related to the Lagrangian stretching [55].

1.3 Fluid Dynamics

While I will take the fluid flow as given in most of my analysis, it is important to understand where the fluid flow comes from, and what properties it has. First of all, my experiments will be conducted in liquid phase, and the flow speeds involved are far less than the speed of sound in water. Therefore the fluid flow is incompressible,

$$\nabla \cdot \mathbf{u} = 0. \tag{1.7}$$

This fluid flow is driven by Lorentz forces, and the reacting fluid is a Newtonian fluid. Therefore the Navier-Stokes equation which governs the flow field along

with eq. 1.7 is

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\frac{1}{\rho} \nabla P + \nu \nabla^2 \mathbf{u} + \frac{1}{\rho} \mathbf{J} \times \mathbf{B} \quad (1.8)$$

where P is the pressure field, ρ is density, ν is the kinematic viscosity, \mathbf{J} is a current density passed through the fluid, and \mathbf{B} is the magnetic field. The experimental fluid is contained in a vessel with a free surface at the top, and no-slip boundaries on the sides and bottom.

The Navier-Stokes equation emerges from conservation of momentum, and the requirement that the physics be invariant to whether we fix our reference frame to be stationary or if we follow an infinitesimal fluid element as it is forced around by the flow. Watching the flow in the stationary reference frame is known as an Eulerian viewpoint, and following the flow itself is called a Lagrangian viewpoint. The left side of eq. 1.8 is known as a material derivative, and it ensures that the two perspectives are equivalent. In the case of chemical reactions in flow, the reaction is also carried along so the material derivative can be seen in eq. 1.1. Lagrangian perspectives will prove particularly useful for understanding my results, and indeed both the front dynamics and filament models are developed in a Lagrangian framework. The reason for this is that in attempting to find any general results about ARD, it is more relevant to consider how the same *patch of reaction*, not the *same spot in space* evolves in time.

The components of eq. 1.8 cause some complex dynamics, and tend to resist easy analysis. In particular this is due to the advective, or inertial, term $(\mathbf{u} \cdot \nabla) \mathbf{u}$, which provides a nonlinear forcing. This nonlinear forcing can lead to solutions that are neither steady nor periodic. In fact, as flow speed is increased, fluid flows usually change from a steady solution to an unsteady solution and eventually to fully turbulent flow. In his legendary paper, Reynolds [56] showed that these transitions occurred for a critical nondimensional number in pipe flow. This number would later be renamed the Reynolds number, $Re = UL/\nu$, where

U is a characteristic velocity, and L is a flow length scale. Other authors have characterized these critical transitions in flow configurations like the ones I will use, both theoretically [57, 58], and experimentally [59]. In these transitions, the steady flow gives way to increasingly complex periodic orbits, until it transitions to a chaotic system. As flow speed is further increased, the system can be called truly turbulent.

Turbulent systems are notoriously difficult to analyze. Within scalar mixing it is known that the length scale of variation of a scalar quantity is reduced by the turbulence, because the turbulent flow eddies are themselves cascading to smaller length scales. This cascade continues down to the Kolmogorov scale, η , after which the length scale may need to be further reduced before it reaches the Batchelor scale, n , where diffusion can mix the scalar [1]. This picture of turbulent mixing is very qualitative, and is not particularly useful for identifying where any particular reactive scalar is likely to be in the future, because the underlying fluid is itself difficult to predict. Yet for all the unpredictability of turbulent flows, many physical flows contain structures that persist for a long period of time, and are carried by the surrounding flows. For example Jupiter's Great Red Spot [60], or the Agulhas Rings in the ocean [61]. Recent work has identified flow structures like these as Lagrangian coherent structures (LCS) [62–64]. These structures can be measured in real flows [65], and possibly used to forecast sudden changes in scalar mixing situations like oil spills [66]. The LCS are a Lagrangian way to classify the parts of a turbulent flow, and are considered the most important barriers to non-reactive scalar mixing. These structures can be calculated through a variety of ways, but for practical considerations the finite time Lyapunov exponent (FTLE) is often used [67, 68].

Whereas the LCS provide a way to classify regions of the flow, FTLE's are scalar values for all points in space and time. Depending on the finite time interval, T chosen, these FTLE values quantify how much a given spot in the flow will be

stretched or deformed during that time interval. The math to calculate this is

$$\lambda = T^{-1} \log S \quad (1.9)$$

where S is the stretching field [55]. The stretching field is calculated from a Lagrangian representation of fluid flow known as the flow map, ϕ . ϕ maps each point \mathbf{x}_0 at time t to its location \mathbf{x} at time $t + T$. T can be positive or negative (meaning backwards in time). Distortion as experienced by the fluid element then can be described using the Cauchy-Green strain tensor [69],

$$\begin{aligned} (C_T(\mathbf{x}_0))_{ij} &= (\nabla\phi_T)_{ki} (\nabla\phi_T)_{kj}, \text{ where} \\ (\nabla\phi_T)_{ij} &= \frac{\partial x_i}{\partial x_{0,j}}. \end{aligned}$$

With this definition, stretching is then defined as the square root of the maximum eigenvalue of the Cauchy-Green strain tensor. The resulting value gives a way to quantify whether the point in the fluid will in the future (for positive T) remain close to its neighbors or diverge from them. T can also be negative to run the stretching time in reverse and examine whether the point recently diverged from its old neighbors. Since FTLEs can help make sense of nonreactive scalars in turbulent flows, their relevance along with LCS to reactive scalars is very interesting, and already they show signs of being important [8, 31, 32, 70].

1.4 Reaction Dynamics

The last term of eq. 1.1 that requires more explanation is the reaction term. F_i is the dimensional reaction function which can usually be written down based on the chemical equations. Reactions occur at a rate proportional to some dimensional reaction rate constant, k , and the concentration of all the constituent molecules that must interact. The more molecules are available, the faster the reaction. Multiple terms often appear in the reaction function due to multiple ways a single

chemical species may be made, and each chemical species involved has its own reaction function. For example consider the chemical reaction which happens at rate k ,



then a reaction function exists for each chemical species X, Y, Z . In this case, these functions would be

$$\begin{aligned}\frac{\partial X}{\partial t} &= -kaX^aY^b \\ \frac{\partial Y}{\partial t} &= -kbX^aY^b \\ \frac{\partial Z}{\partial t} &= kcX^aY^b.\end{aligned}$$

so the coefficients on each species serve to both multiply how much that species changes for each reaction, and also to raise the order of the reaction function. The only example of a linear system in chemical kinetics is if all the reactant species except one are in a large enough excess to be considered constant. Nonlinearity is the default behavior for chemical kinetics [2]. I will usually simplify the system by writing one chemical species in terms of the concentration of another, so I only consider one or two chemical species. To make this nondimensional, the reaction coefficient and the magnitude of the reaction will usually be factored out into a constant called α which has units of 1/s. This also means that concentrations range from 0 to 1. I call this nondimensionalized reaction function G_i .

There are numerous ways for chemical reactions to feed-back on themselves. The growth of the reaction may reduce the remaining concentration of reactants, the reaction could drive an increase in temperature which changes α , or a side reaction could produce a catalyst. In this thesis, my chemical reaction will be an excitable, autocatalytic reaction. Autocatalytic means that the reaction product is a catalyst for more reaction, and excitable means that the reaction does not

proceed until it experiences a perturbation above some finite threshold, at which point it goes on a long deviation from the unreacted state. The excitable reaction I am using is the Belousov-Zhabotinsky (BZ) reaction, for which a simple model of the kinetics [2, 71, 72] is

$$\frac{\partial c_1}{\partial t} = \alpha \left[c_1(1 - c_1) - \frac{c_1 - q}{q + c_1} f c_2 \right] \quad (1.10)$$

$$\frac{\partial c_2}{\partial t} = \alpha \epsilon [c_1 - c_2] \quad (1.11)$$

where f , ϵ , and q are nondimensional constants. Once a threshold is passed, c_1 grows rapidly, so it is the catalyst. As c_1 increases, c_2 rises, but since $\epsilon \approx 0.01$ this rise happens more slowly. Once c_2 is large enough it begins to inhibit the reaction, causing it to reset back to the original state. A phase space plot of this reaction is shown in fig. 1.2, which can be used to gain intuition about how this reaction behaves.

This reaction is fairly complicated, but the rewards for focusing on an excitable system are large. Many other systems experience the combination of a finite threshold for perturbations, a rapid growth, and then resetting to the original state. For instance, population modeling of phytoplankton in the ocean can be described as excitable with population of phytoplankton as the activator, and the predatory zooplankton as inhibitor [19]. Combustion reactions also share a lot of commonalities with excitable systems, because they also possess reaction thresholds (often called ignition points), and wildfires can be considered as refractory because the fuel grows back [17]. Other systems with excitable dynamics range from fibrillating hearts [73] to “The Wave” performed at sporting events [74].

1.5 Experimental Methods

Linking all the results in this thesis, are the unique experimental tools that I used to gather data and perform analysis. In every chapter there are two fields which

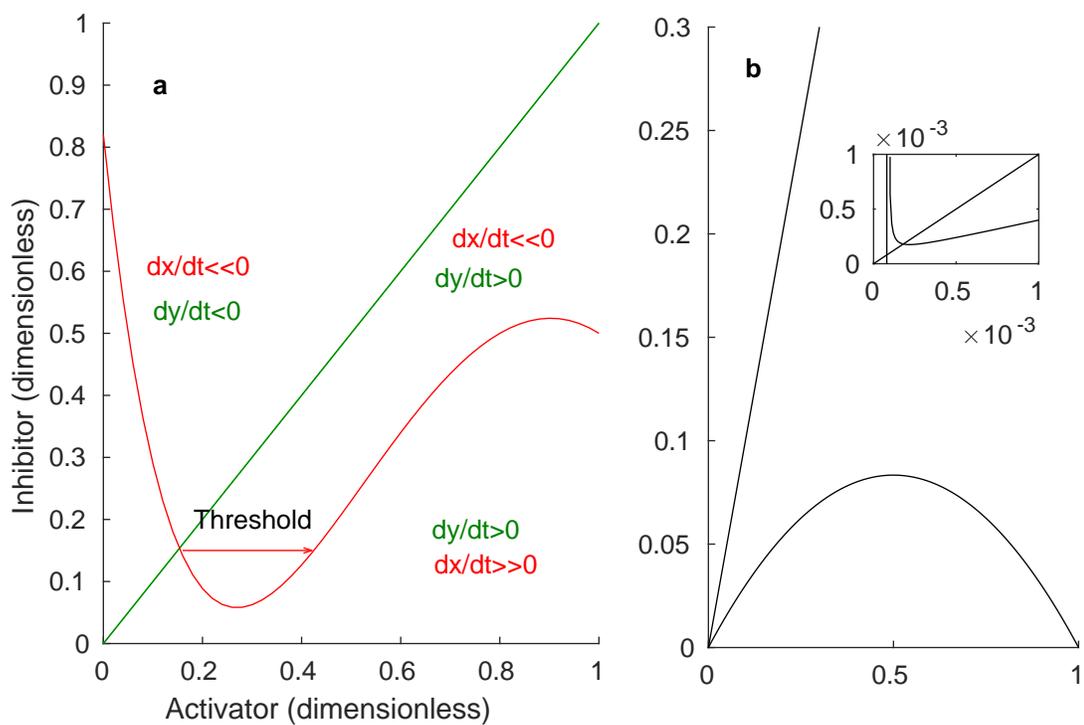


Figure 1.2: (a) Phase diagram for excitable systems. The lines represent nullclines along which the rate of change of one of the chemical species is zero. (b) The actual phase plot for the Oregonator shows that the threshold is very small, and the shape is closer to a quadratic than a cubic like in (a). Unpublished work of the author.

must be quantified in experiments: reaction and flow. To do this, I gather data from thin layer flow experiments using two hardware synchronized cameras. One camera gathers flow data by tracking fluorescent particles that float atop the layer, the other watches the reaction. Since I always have the fluid flow that the reaction actually experienced, I do not have to model the fluid flow from eq. 1.8, and I can analyze flows which are time dependent. The later of these benefits is a huge advantage, because beyond the transition from periodic flows to chaotic flows it is impossible to generate the same flow repeatably. If data is taken on a reaction growing in chaotic or turbulent flows, a second experiment in the same flow cannot be conducted to then measure the flow field. The synchronized cameras record high speed, high resolution data that is from the exact same flow. By applying a blue color filter to the reaction camera, and a color filter matching the fluorescence of floating tracer particles which track the flow to the other, I obtain two videos which contain either only particles or only reaction.

The flow tracking methodology is well established, and roughly the same in all my experiments. Flow is quantified by tracking fluorescent tracer particles in the images from the flow camera according to the Lagrangian Particle Tracking Velocimetry (PTV) method [75]. Particle tracking code originally comes from Dr. Nicholas Ouellette. Tracer particles fluoresce under the blue light we illuminate the apparatus with. They follow the flow closely, as can be shown by considering the ratio of the characteristic settling time of the particle to the characteristic time of the flow, which is the Stokes number St . For my experiments $St < 0.1$ which ensures that the particles always change to match the surrounding flow faster than the flow changes speed. In the thin layer system we also do not reach the flow speeds at which three-dimensional up-welling or down-welling occur [76, 77].

Due to all these aspects I am able to accurately reproduce the fluid flow and reaction at all times and places, even if the flow is unsteady. This setup was conceived by Dr. Doug Kelley, and I made continual improvements to the system over

several years, so each chapter in this thesis has some slightly different parameters.

1.6 Thesis Overview

The goal of this thesis is to perform some tests on modeling of ARD systems to better understand and predict their dynamics. Specifically, I will be focusing on experiments using the excitable BZ reaction, and thin layer flows (known as quasi-2D). To do this, in chapters 2-4, I develop a novel algorithm for tracking reaction fronts at all locations and all times. Tracking reaction fronts allows me and future researchers to obtain data on reaction front position and speed from ARD systems that is exactly analogous to the Eikonal approximation. In chapter 2, I start by introducing this algorithm without flow, and observe that front speed and front thickness can both be measured simultaneously, and they can be used to obtain accurate measurements of diffusion and reaction rate constants. In chapter 3, I enable the algorithm to handle fronts moved around by advection, and show it gathers accurate data in simulation. I also observe an apparent dependence of chemical speed on flow speed, but in chapter 4 I am able to explain much of this dependency on experimental considerations. I suggest the use of a lubrication layer for future ARD experiments. Finally in chapter 5 I discuss the observation of an optimal Lagrangian stretching for excitable chemical reactions. We observe that this optimal range makes a good predictor of reaction behavior in both laminar and unsteady flows, but cannot be explained by the Eikonal approximation. I show that the filament model helps to understand the optimal stretching behavior, and I use numerical and analytical methods to investigate other effects of strain-rate (which gives rise to Lagrangian stretching) on ARD systems.

2 Front Tracking Method for Quantifying Reaction-Diffusion

This chapter is based on Reference [78] and the author's qualifying exam submission. It has been reformatted and edited to fit within the thesis. The thesis author was the primary author on this work.

2.1 Introduction

The three physical processes advection (fluid flow), reaction, and molecular diffusion govern the dynamics of concentrations in many common systems. Because underlying fluid flow is often chaotic or turbulent, and because reaction often involves nonlinear chemical kinetics, complicated concentration dynamics are typical in advection-reaction-diffusion systems. When a reaction grows in the presence of a flowing fluid, the flow itself is different from one location to another at each instant. As a direct result of spatial complexity of flows, the reaction grows in different directions from one location to another, and can even grow at different rates from place to place (See Chapter 5). Studies on the interaction between reaction and advection have shown frozen reaction fronts in porous media [10, 11],

and reactions growing at different rates in Poiseuille flow based on the front shape of the reaction [35–37]. Significant global effects of fluid flow on enhanced reaction growth in the bulk are also well known [8, 53, 79]. Still other studies have found pattern formation among advection driven reactions [80].

However, in order to provide insight for the numerous applications in advection-reaction-diffusion (ARD), a simplified representation is needed. This model of ARD must then be sensitive to all the complexity listed above. One technique used for making a simplified representation is to choose a concentration threshold and identify the *reaction fronts* separating high-concentration reacted regions from low-concentration unreacted regions. The complex dynamics of ARD systems can then be thought of as the movement of these reaction fronts, which must be able to vary from point to point along the front. One application of fronts are the recently-developed theories of burning invariant manifolds (BIMs) and burning Lagrangian Coherent Structures (BLCS) [34, 44, 46, 48, 50, 51]. These identify barriers to propagation of fronts caused by opposing advection, enabling some prediction of reaction growth. BIMs and BLCS predict where fronts will advance and where they will not. Calculating those barriers requires knowing the speed at which fronts advance when only reaction and diffusion are present. The front tracking algorithm I will describe is a data analysis method to measure that front speed locally at many locations, providing a statistically robust measurement. Prior studies have measured front speed only along a single line [81, 82]. Instead, this algorithm works from time series of concentration fields, which can be gathered from either simulations or experiments. Also, since the algorithm gathers local data, I will be able to study how reaction growth changes with local flow behavior to cause the complex behaviors that have been observed.

The algorithm also determines the front thickness at many locations, that is, a specific distance which characterizes the distance over which product concentration varies from small (near zero) to large (near saturation). Through well-known

solutions to reaction-diffusion equation, front thickness along with front speed can be used to determine the diffusivity and reaction rate of the reacting species in reactions without flow. By determining front speed and thickness at many locations, the algorithm provides statistically robust measurements of those microscopic quantities as well. Prior studies have typically measured diffusivity and reaction rate at just one location [83, 84].

Though I am interested in modeling advection-reaction-diffusion systems, the discussion below centers on the theory of reaction-diffusion systems, in which advection is absent. Similarly, the validation and simulation data we use to test the algorithm involve no advection, and the experiments described below involve no advection. Instead in this chapter I focus on validating and verifying a front tracking method, and demonstrating that it has physical relevance. I have chosen reaction-diffusion systems because of their relative simplicity. Their physics is well understood, and that creates better intuition about what results to expect from the algorithm. Generalizing front tracking from reaction-diffusion to advection-reaction-diffusion requires only that front displacements due to advection be subtracted off before implementing the algorithm described below, then restored later. That this can be done will be shown in Chapter 3. In prior literature, the phrase “front tracking” sometimes refers to a simulation method, where the simulation works differently near a reaction front [85, 86]. Throughout this chapter, however, “front tracking” is specifically an algorithm for *measuring* a reaction front in data produced either in simulation or experiment.

This chapter proceeds as follows: In §2.2 we explain the theory of reaction-diffusion equations, and how diffusion and reaction rate determine front speed and thickness. In §2.3 we detail the algorithm and its parameters. Section 2.4 shows that the algorithm measures correct front speed and thickness, working from validation data in which those quantities are prescribed. Section 2.4 also shows the algorithm to be robust to random noise and insensitive to parameter choices of

the user. In §2.5, we apply the algorithm to simulation results, showing that the algorithm measures correct diffusivity and reaction rate, working from simulation results in which those quantities are prescribed. Section 2.6 shows the results of applying the algorithm to laboratory experiments with the Belousov-Zhabotinsky (BZ) reaction, in which front speed, front width, diffusivity, and reaction rate are all measured. Finally, §2.7 summarizes our results and considers implications for future work.

2.2 Relating Thickness and Speed to Diffusivity and Reaction Rate

The growth of a reacted region is governed by its initial condition and the material diffusivity and reaction rates. That is, the physical constants describing microscopic processes (diffusivity and reaction rate) must determine the macroscopic behavior (front speed and thickness). Conversely, measuring macroscopic behavior can reveal microscopic quantities, often through simple algebraic relationships. The differential equations allow us to understand chemistry through measuring the physics. Though this insight has been published before [2, 11], to my knowledge it has not been utilized before, because front thickness in the sense we describe has not been measured. However, local reaction profile length scales have been measured using maximum gradient methods in combustion [87, 88], and reaction profiles of the BZ reaction have been photographed [89]. This section will describe the mathematical connections between diffusivity, reaction rate, front speed, and front thickness.

Starting from the advection-reaction-diffusion equation,

$$\frac{dc}{dt} = -(\mathbf{u} \cdot \nabla)c + D\nabla^2c + kF(c), \quad (2.1)$$

where c is the concentration of the reaction product, \mathbf{u} is the underlying flow, D is the diffusion coefficient, k is a reaction rate constant, and $F(c)$ specifies the reaction kinetics. If the advection term is zero (or we consider a reference frame moving with the flow: a Lagrangian frame), eq. (2.1) simplifies to the reaction-diffusion equation,

$$\frac{dc}{dt} = D\nabla^2 c + kF(c). \quad (2.2)$$

Since diffusion spreads scalars down the gradient of their own concentration fields, it spreads reacted regions *perpendicular* to their concentration level curves. Since the reaction term $F(c)$ is local, it only serves to raise or lower concentration. Based on this intuition, a front is defined as a line (or in three-dimensional reactions, a surface), which separates reacted and unreacted regions. In particular, a front is a particular level curve or surface of the concentration field c . Therefore, if a reaction-diffusion front is defined to be a characteristic level curve of the reaction, it will appear to “move” down the concentration gradient, locally perpendicular to the front itself. Systems can be prepared such that fronts do not appear to move perpendicularly, but if we assume the reaction is fast enough that the front advancement is limited by diffusion ($kc_0 \gg D/l^2$ where c_0 is an equilibrium concentration, and l is the length scale of interest) then its level curves quickly become close together, and the front becomes “sharp”. Perfectly sharp fronts must evolve perpendicular to themselves, because the reaction term is zero except at the front and therefore it cannot create apparent front advancement by causing a far off region to increase to above the characteristic front level.

For suitably fast reactions, the changes from the unreacted to the reacted state occur in a very short spatial distance, so a change in the choice of which level curve is used results in only a small change in curve position. If the concentration profile is constant over time, this position shift does not lead to a shift in front velocity. Far from a front, whether it be thin or thick, very little change occurs in c , so the front is the dominant feature of a reaction. By understanding the

motion of these surfaces, I obtain a simple picture of what happens in an entire reaction volume, but retain many of the important features. Since the reaction causes growth by boosting an existing concentration above the threshold, front motion is an “apparent motion”, which emerges from the dynamics without being specified. As a result even simulations need a method to obtain front speed or thickness after a simulation is complete. The novelty of the algorithm lies not in locating fronts, which can be done by simple thresholding, but in quantifying their properties: velocity and thickness.

Moreover, because a reaction front is a level curve the spatial variation of concentration along it is zero, the diffusion term $D\nabla^2 c$ is dominated by gradients locally perpendicular to the front. Eq. (2.2) can then be written in terms of one spatial coordinate, x , locally perpendicular to the front. If the reaction term has second-order dynamics, the result is the Fisher Equation,

$$\frac{dc}{dt} = D \frac{\partial^2 c}{\partial x^2} + kc_0 c(1 - c), \quad (2.3)$$

where now c is dimensionless, having been normalized by c_0 , an equilibrium concentration. Considering the boundary conditions $c = 0$ at $x = x_0$ and $c = 1$ at $x = -x_0$, and choosing $x_0 \rightarrow \infty$, the solution to the Fisher Equation [39, 90, 91] is

$$c(x, y, t) = \frac{1}{(1 + e^{(x-vt)/L})^2}, \quad (2.4)$$

where the front thickness and speed are

$$L = \sqrt{\frac{6D}{\alpha}}, \quad (2.5)$$

$$v = \frac{5}{6} \sqrt{6D\alpha}, \quad (2.6)$$

respectively. Here D is the diffusion coefficient and $\alpha = kc_0$ is the reaction rate normalized by equilibrium concentration. This result is slightly different than the commonly used Kolmogoroff result [40], which says $v = -2\sqrt{D\alpha}$. This is because the Kolmogoroff result is based on numerous approximations and does not provide

us with the nice results with regard to diffusion and reaction rate that the Fisher solution does. Furthermore, the constant $\frac{5}{6}\sqrt{6} \approx 2.04$, making the Kolmogoroff result still very close. Equations (2.5) and (2.6) can be solved for D and α :

$$D = \frac{|v|L}{5}, \quad (2.7)$$

$$\alpha = \frac{6|v|}{5L}. \quad (2.8)$$

where I use the speed, $|v|$, to emphasize that this also works for eq. (2.3) under the opposite boundary conditions.

Specifying boundary conditions at infinity is an idealization, but the resulting solution closely approximates the true situation as long as $x_0 \gg L$. Also, though curvature of a front is known to alter its speed [42], solutions of Eqs. (2.3) and (2.1) match as long as the radius of curvature of the front is much larger than the length scale of spatial gradients. While under these assumptions the front is approximately planar, which means the solution to the 1-dimensional case is the solution.

Equations (2.7) and (2.8) make it possible to determine the microscopic physical constants D and α by measuring the macroscopic properties L and v — if the chemical kinetics are second order. Similarly, for third order kinetics,

$$\frac{dc}{dt} = D\nabla^2 c + kc_0^2 c^2(1 - c), \quad (2.9)$$

which is solved by [2, 11]

$$c(x, y, t) = \frac{1}{(1 + e^{-(x-vt)/L})}, \quad (2.10)$$

where the front thickness and speed are

$$L = \sqrt{\frac{2D}{\alpha}}, \quad (2.11)$$

$$|v| = \sqrt{\frac{D\alpha}{2}}, \quad (2.12)$$

$$(2.13)$$

respectively. Again it is straightforward to write D and α in terms of L and v :

$$D = L|v|, \quad (2.14)$$

$$\alpha = \frac{2|v|}{L}. \quad (2.15)$$

It seems likely that the diffusivity and reaction rate can be written in terms of macroscopic quantities for chemical kinetics of any order for which the appropriate one-dimensional reaction-diffusion equation can be solved analytically. In this chapter I consider experiments and simulations where the reactants are dissolved in water and the reaction is second-order. I suspect the method works even better for third-order reactions, because the form of the solution for third-order dynamics makes fits to measurements more straightforward. I have not attempted the method on any other reaction types, but I see no reason that it would not work on different chemistries, even solid- or gas-phase reactions.

In the experimental section, I consider the BZ reaction. For my purposes, the reaction kinetics of BZ can be simplified via the Oregonator model to two differential equations [2, 72]:

$$\epsilon \frac{da}{d\tau} = D_1 \nabla^2 a + a(1-a) - \frac{a-q}{q+a} fb \quad (2.16)$$

$$\frac{db}{d\tau} = D_2 \nabla^2 b + a - b. \quad (2.17)$$

Where a is a nondimensional value for concentration of Hydrogen-Bromate (HBrO_2), and b is a nondimensional concentration of oxidized ferroin (M_{ox} for short). The quantity q is a nondimensional ratio of reaction rates, and f is a constant from the reaction kinetics. In this reaction, HBrO_2 generates itself autocatalytically, and is eventually removed by growing M_{ox} . This creates an oscillating chemical reaction that can be tested many times without having to remix chemicals. Also, since ϵ is a small nondimensional number, changes in a

are much faster than b . Thus when a reaction first starts, b is nearly zero, and a grows quickly. This means the last term of eq. (2.16) is negligible, and a is a second-order reaction-diffusion system at the front when it first reacts.

Therefore L and $|v|$ are henceforth defined as in Eqs. (2.5 and 2.6). Throughout this paper I also consider the simplified case of reaction in a thin layer that can be approximately two-dimensional, in which reaction fronts are one-dimensional curves, not two-dimensional surfaces. This matches the type of data I will be collecting. However, the algorithm should readily generalize to three-dimensional systems if a suitable technique is found for gathering three-dimensional concentration fields.

2.3 Front Tracking

In this section I describe the algorithm to locally measure front thickness and speed in reaction-diffusion systems, using a series of concentration fields measured at regular intervals and on a regular spatial grid. The algorithm will be generalized to advection-reaction-diffusion systems as long as the velocity field \mathbf{u} is known in Chapter 3. Written in MATLAB, the open-source algorithm is freely available¹. An example of the end result is shown in Fig. 2.1.

The algorithm begins by locating reaction fronts, which are defined as the surfaces that separate regions of high product concentration $c > c_{\text{thresh}}$ from low product concentration $c < c_{\text{thresh}}$, where c_{thresh} is a user-defined concentration threshold. This corresponds to a concentration level curve at c_{thresh} . In our experiments and simulations these fronts are suitably thin so front position is weakly dependent on c_{thresh} .

After finding the edges of regions where $c > c_{\text{thresh}}$, the algorithm smooths those edges, because measuring concentration on a discrete grid makes them rough

¹<http://www.me.rochester.edu/projects/dhkelley-lab/>

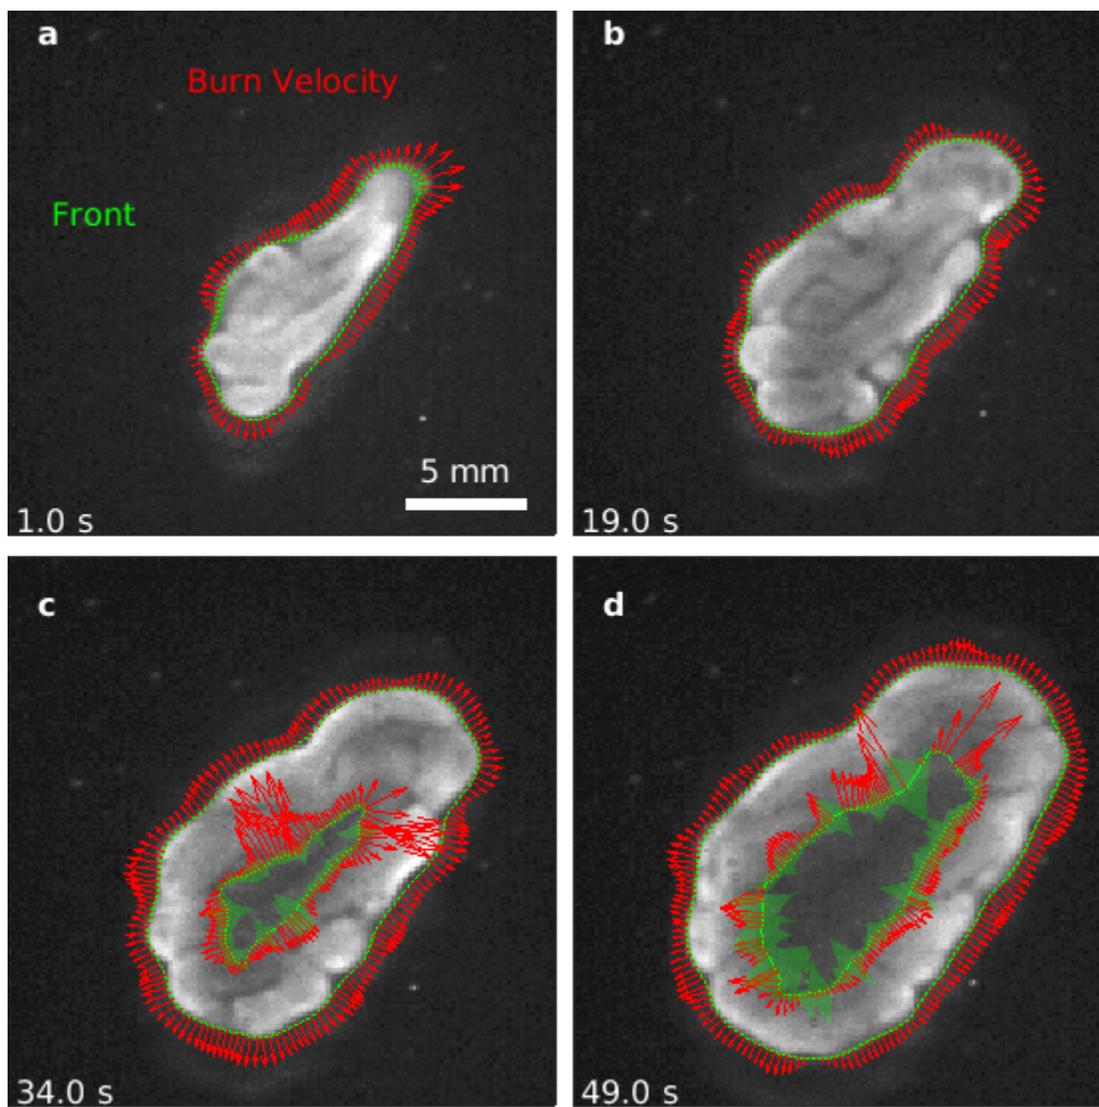


Figure 2.1: Example results of front tracking. Front location is indicated in green and thickness by the green area; front velocity is indicated by red arrows in reaction-diffusion experiments with the Belousov-Zhabotinsky reaction, described below. Work of the author published in [78].

Parameter	Description
c_{thresh}	Concentration threshold.
S	Span of sliding fit for smoothing fronts.
A_{min}	Minimum area of reacted regions.
M	Size of excluded margin.
d_{max}	Maximum front displacement.
n_{skip}	Number of concentration fields skipped when measuring speeds.
L_{fit}	Half profile fit length for determining front thickness.

Table 2.1: Parameters chosen by the algorithm’s user.

at grid scale and does not accurately represent the front shape. This smoothing is done using a sliding line fit spanning S points along the boundary. In this process S neighbors are fit to a line which is then used to adjust the central point’s location. Line smoothing is always necessary for accurate reconstruction of perpendiculars, since boundaries from a thresholded image always follow pixel boundaries. To further reduce noise effects and in anticipation of the presence of small tracer particles in future advection-reaction-diffusion experiments, the algorithm removes regions where $c > c_{\text{thresh}}$ whose area is smaller than A_{min} , another user-defined parameter. Finally the algorithm imposes margins by ignoring all front elements within a distance M of the edges of the field of view. This prevents field borders from being incorrectly identified as front boundaries. All algorithm parameters, including c_{thresh} , S , A_{min} , and M , are listed in Table 2.1.

Once the user has located fronts, they can measure their speed, and Fig. 2.2 illustrates the algorithm for doing so. First, the algorithm considers the fronts in two subsequent concentration fields. Then, it draws a perpendicular line extend-

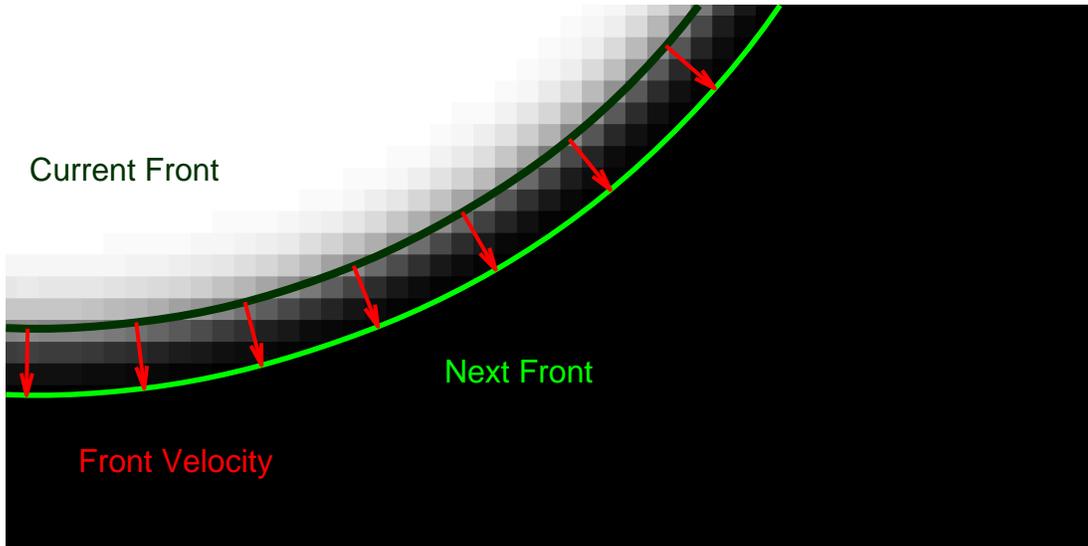


Figure 2.2: Algorithm described for measuring front velocity. First, the edges of regions where $c > c_{\text{thresh}}$ are, after smoothing, identified as fronts. The process is repeated for a subsequent concentration field. The perpendicular distance between the first and subsequent location of a front, along with the delay time between concentration fields, provides the local velocity. Work of the author published in [78].

ing a distance d_{max} in each direction from each point on the earlier front; d_{max} is specified by the user. The algorithm then finds the locations where the perpendicular intersects fronts in the subsequent field. The displacement between the original point and each intersection is calculated, and the smallest displacement is divided by the delay time between concentration fields, yielding the front speed at that one point. If the perpendicular intersects no front, a “not-a-number” value is recorded. The process is repeated for every point on the fronts in the earlier concentration field. To eliminate noise caused by high-frequency changes in the concentration field, the algorithm can implement a simple low-pass temporal filter by skipping n_{skip} subsequent concentration fields, as specified by the user. Assuming the grid spacing and interval between concentration fields are known, front speed can be expressed in physical units.

Algorithm users can also measure their thickness, and Fig. 2.3 depicts the

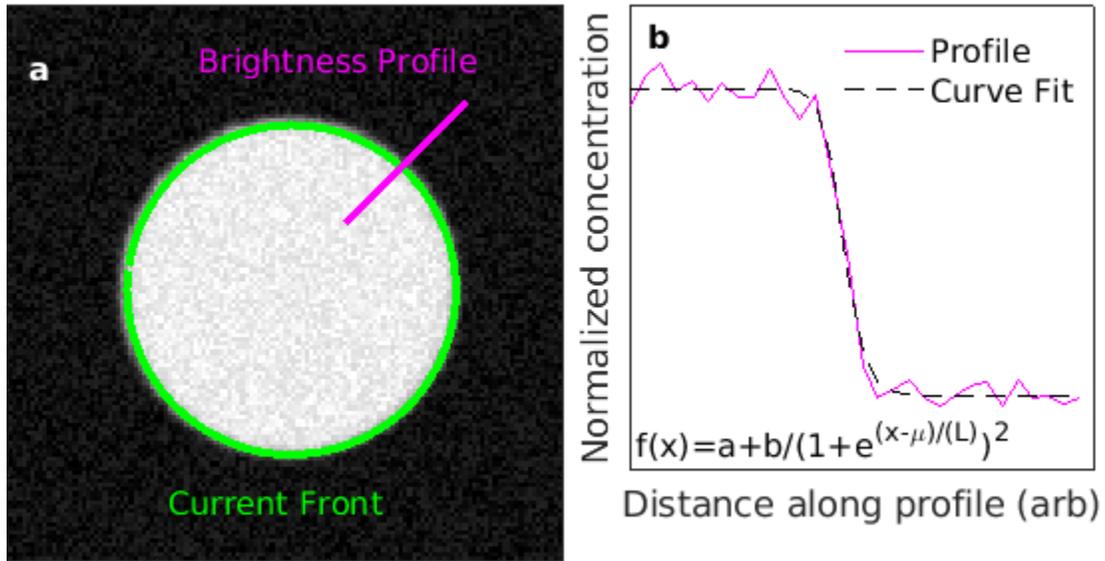


Figure 2.3: Algorithm for measuring front thickness, shown using validation data. First, the edges of regions where the local concentration c exceeds the chosen threshold c_{thresh} are, after smoothing, identified as fronts. Then the concentration profile along a line perpendicular to the front is fit to eq. (2.4), yielding the local thickness L . Work of the author published in [78].

algorithm for doing so. Only one concentration field need be considered. The algorithm again draws a perpendicular at each point along each front, as in velocity measurement, but with length L_{fit} , specified by the user. The algorithm interpolates the concentration on the regular grid to a set of points along the perpendicular to obtain a concentration profile as shown. That profile is fit to the corresponding one-dimensional solution to the reaction diffusion equation, eq. (2.4), yielding the local thickness L . Fitting allows the algorithm to determine L with sub-pixel accuracy, since more than one pixel is used in its determination. Assuming the grid spacing is known, front thickness can be expressed in real units. To ensure accurate fitting L_{fit} should be small enough not to overlap with a nearby front, and no smaller. To see the dependence on L_{fit} and other parameters see Fig. 2.6.

2.4 Tracking Fronts in Validation Data

To test the algorithm, I first verified that it could reproduce front speed and thickness, given validation data in which speed and thickness had been specified. I created a series of concentration fields in which were constructed fronts of varying shape, speed, and thickness. Each front had a reaction profile specified by eq. (2.4), with $0 \leq c \leq 255$. I also added in uniformly-distributed noise of magnitude 20 to each concentration field. All validation data are measured in units of pixels and frames. Figure 2.4 compares the true speeds and thicknesses to the speeds and thicknesses measured using the algorithm. The match is close. The greatest errors occur for a square front and are due to curve smoothing at the sharp corners. These corners move faster than the “true speed”. In any physical front, however, diffusion eliminates sharp corners. Measured thicknesses are accurate over a wide range of possible thicknesses. Certain extreme cases can disrupt this accuracy. In particular, as the thickness approaches the size of L_{fit} , errors increase. To improve these errors L_{fit} can be increased; here I chose a large fit, $L_{\text{fit}} = 20$ pixels, to reduce the errors. For annuli, as the fronts get thicker, the concentration profiles of the inner and outer boundaries overlap, so that $x_0 \rightarrow \infty$ is no longer a good approximation (see Sec. 2.2). In this case accuracy can be improved by decreasing L_{fit} . Increased resolution reduces error in both speed and thickness.

Next I quantified the robustness of the algorithm to noise in concentration fields. Starting with a circular front as shown in Fig. 2.4a, I added uniformly-distributed noise of varying magnitude, then used the algorithm to measure front speed and thickness. True values were chosen based on preliminary experimental values. Figure 2.5 shows the results. Speed measurements are essentially unaffected by noise, and errors in thickness measurements grow to just $\sim 3\%$ when the noise is almost 50% the size of the signal. One reason these results are robust to uncorrelated noise is that the average speed and thickness are calculated

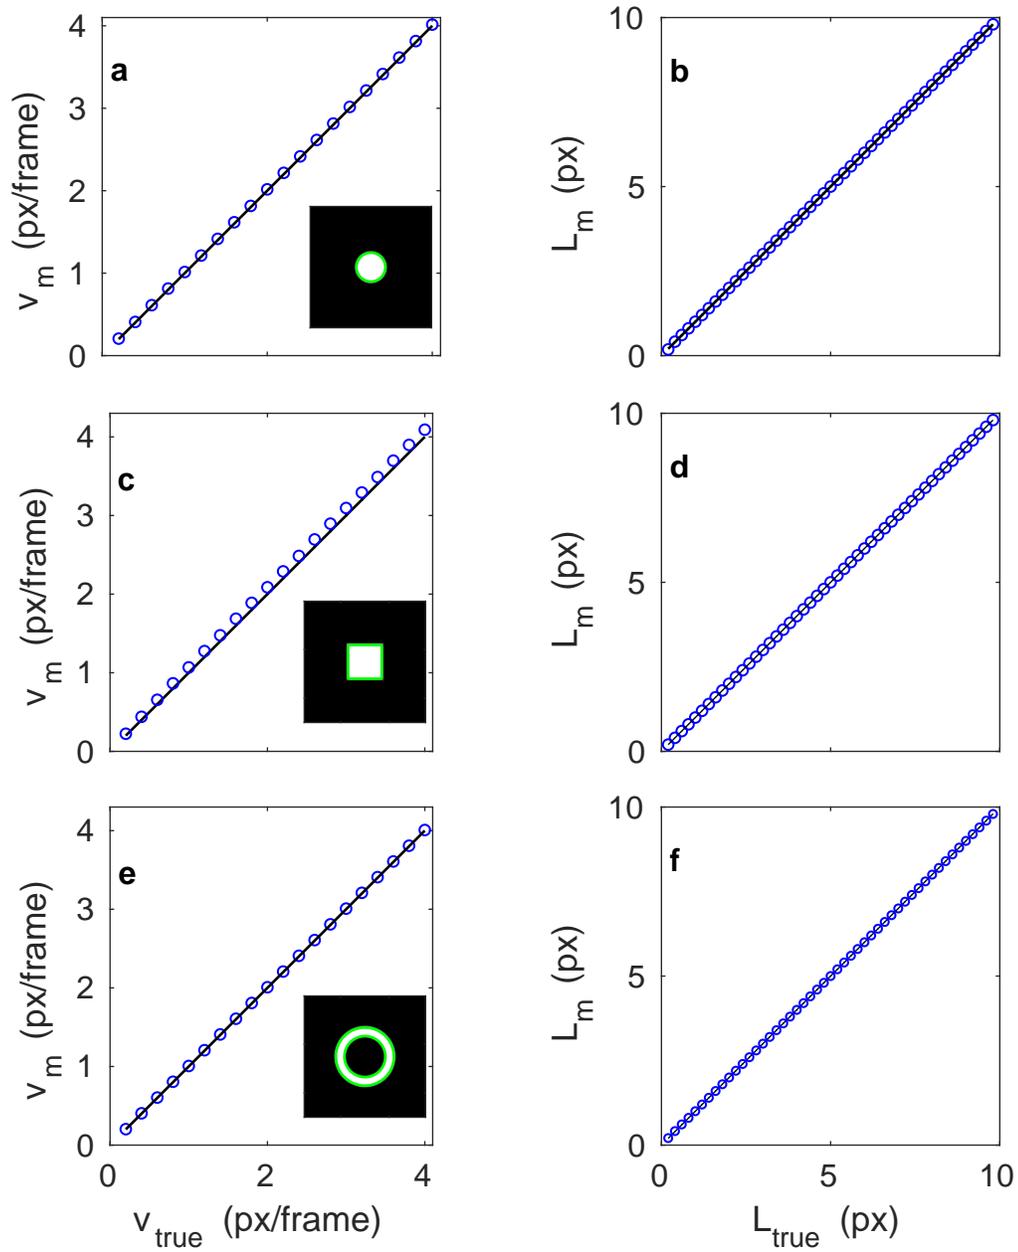


Figure 2.4: Comparison of true and measured front speed and thickness. (a–b) Measured median speed v_m (a) and median thickness L_m of a circular front propagating outward, compared to the input values. Analogous comparisons are shown in (c–f) for square and annular fronts, all propagating outward. Work of the author published in [78].

from many local measurements along the entire length of the front; local measurements provide a major advantage. Curve smoothing also helps to remove sudden, incorrect protrusions caused by noise.

Next I quantified the sensitivity of errors to user choices of the input parameters. Figure 2.6 shows how front speed and thickness vary with threshold c_{thresh} , profile fit length L_{fit} , and smoothing span S . Errors in speed were typically less than 1%. The only significant error occurs for extremely large smoothing span. Errors in thickness were less than 1% except for $L_{\text{fit}} < 7L$, in which case it seems there were too few point measurements for a good fit. Default parameter values were $c_{\text{thresh}} = 50\% = 128$, $L_{\text{fit}} = 40$ pixels, $S = 20$ pixels, $n_{\text{skip}} = 5$ frames, and $d_{\text{max}} = 20$ pixels, when each parameter wasn't the one being changed. The algorithm achieves high accuracy despite the fact that all data sets have a noise signal with a peak-to-peak amplitude of 8% of the dynamic range. For very small smoothing span there is no error in thickness, indicating that smoothing slightly biases thickness measurements. The algorithm has little sensitivity to the concentration threshold, staying accurate to about 1% for thresholds ranging from 10% to 90% of the maximum concentration.

I found that d_{max} , the maximum front displacement, matters only that it has to be long enough to intersect the next frame's front. Once it is long enough to cross, the accuracy is no longer effected by d_{max} . Likewise there is insignificant error sensitivity to user choices of A_{min} and n_{skip} .

2.5 Tracking Fronts in Simulation Data

Having found the algorithm successful in reproducing front speed and thickness from validation data, I set out to verify that the algorithm could reproduce reaction rate and diffusivity from the results of reaction-diffusion simulations.

I wrote a simple two-dimensional simulation that calculates concentration c

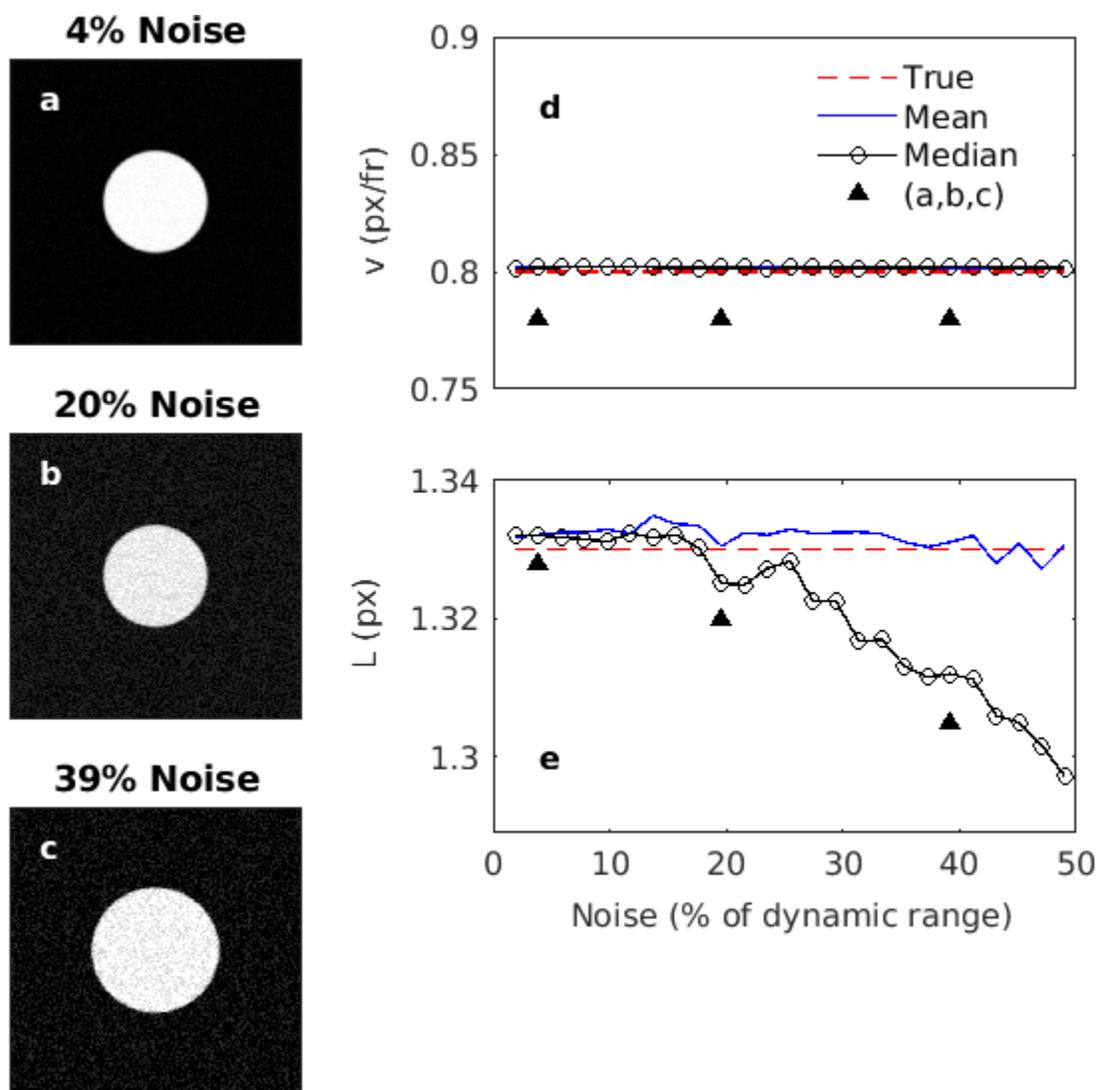


Figure 2.5: Robustness of front tracking to noise. (a–c) Example validation data with varying noise magnitude. (d) Noise changes front speed measurements negligibly. (e) Noise with magnitude up to half the dynamic range changes front thickness measurements by about 3%. Triangles indicate noise magnitudes corresponding to (a–c). Measurements were made using the same default values as in Fig. 2.6. Work of the author published in [78].

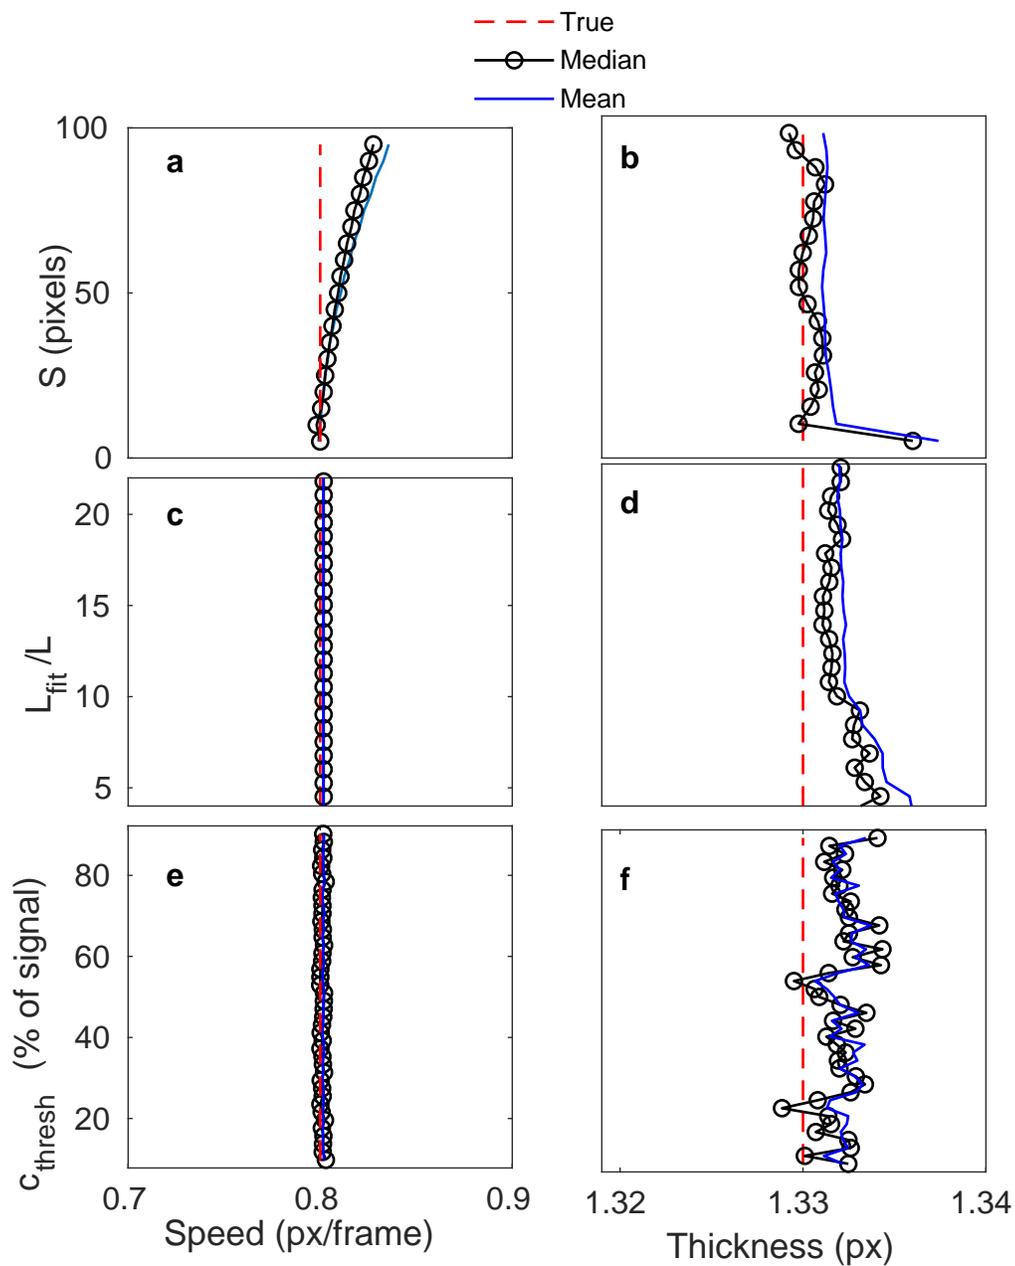


Figure 2.6: Sensitivity of front tracking to input parameters. (a–b), Variation of front speed (a) and thickness (b) with smoothing span S . Analogous variations are shown for profile fit length L_{fit} (c–d) and concentration threshold c_{thresh} (e–f). Work of the author published in [78].

as it evolves through time according to eq. (2.3), using central differences for spatial gradients, first-order time integration, and periodic boundary conditions. I simulated concentration in a two-dimensional domain $-3\pi \leq x < 3\pi$, $-\pi \leq y < \pi$, where x and y are now Cartesian coordinates measured in mm and discretized on a 2400×30 rectangular grid (7.85 mm grid spacing in the x -direction). The initial condition was a straight front of the form

$$c(x, y, 0) = \frac{1}{(1 + e^{-(x-\mu)/L})^2}, \quad (2.18)$$

invariant in y , with D and α specified. The y direction was only included to make the final result a field which could be processed cleanly by the algorithm, which is written for two-dimensional concentration fields. Invariance in y also eliminates curvature effects. To ensure convergence the time step was 9.25×10^{-6} s. Each simulation was run for 4,800,000 time steps, giving 44.4 s of effective data. That duration provided ample data and prevented the front from wrapping around the periodic domain. I then applied the front tracking algorithm to the concentration fields produced by the simulation, obtaining speeds and thicknesses. For each simulation, I recorded the average speed and thickness, then used them to calculate D and α according to equations (2.7,2.8).

Figure 2.7 compares input reaction rates and diffusivities to the values measured using the algorithm. Only one combination D and α gives greater than 5% error, and it is for a case with a low D . Errors in diffusivity are larger than in reaction rate, and diffusivity is consistently over-estimated. I attribute this systematic error not to a problem with the tracking algorithm but to the presence of numerical diffusion in the simulation, caused by estimating continuous derivatives as discrete differences. Percentages are then higher for low D , because the same numeric error becomes a bigger percentage of the measurement. In fact, when we increased the number of grid points to make numerical gradients more accurate, the error in diffusivity dropped. Larger reaction rate also tends to exacerbate er-

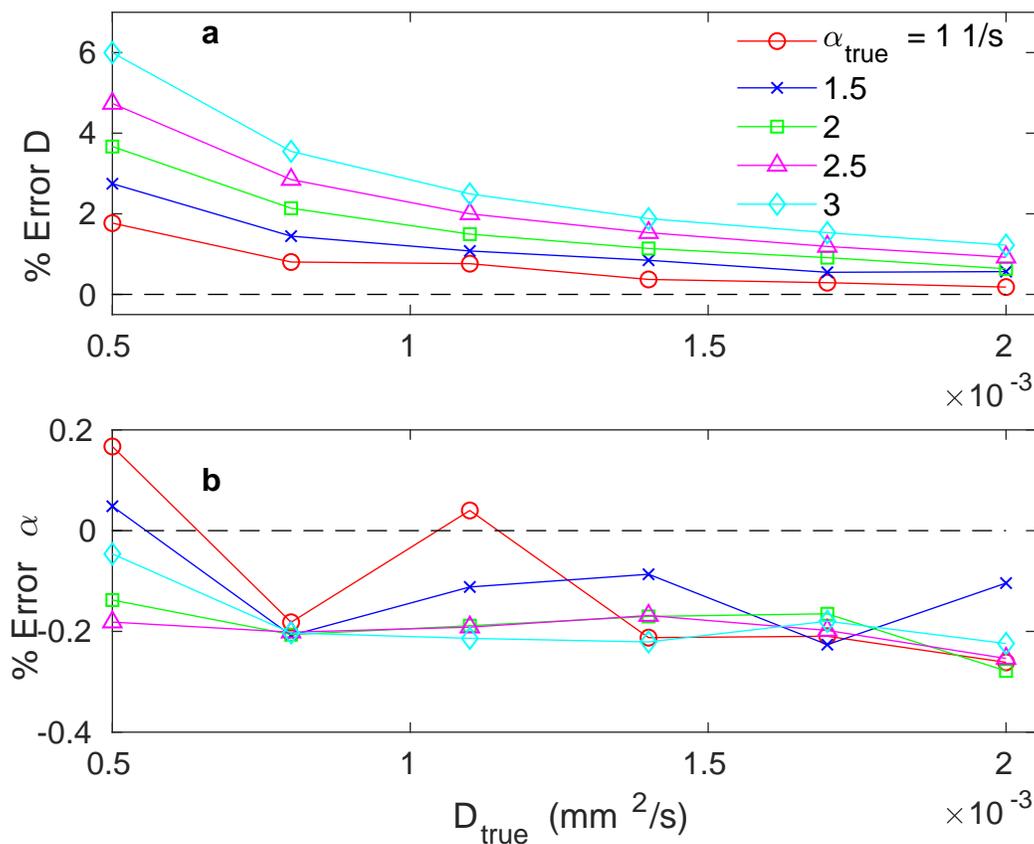


Figure 2.7: Accuracy of diffusivity and reaction rate measurements made by tracking fronts in simulation results. The true diffusivity is D_{true} , and the true reaction rate is α_{true} . Errors are small and primarily due to numerical diffusion in the simulation, not errors in front tracking. Work of the author published in [78].

rors produced by numerical diffusion. Errors in reaction rate are small, and may also be due to numeric diffusion.

2.6 Tracking Fronts in Experimental Data

Having found the algorithm successful with both validation data and simulation results, it was applied to laboratory experiments using the Belousov-Zhabotinsky (BZ) reaction [2, 91–94]. BZ is a complicated oscillating chemical reaction, often

parameterized with the Oregonator model [72]. BZ oscillates between states by passing electrons back and forth with a catalyst. In our case the catalyst is ferroin indicator, so the BZ oscillates between blue and red. The reactions that turn BZ blue are well-modeled as second-order because higher-order reaction terms are negligible, as discussed in §. 2.2 and the book by Scott [2]. However, the reactions that turn BZ back from blue to red are not second-order. Thus we would expect “fronts” and “backs” to show measurably different macroscopic behavior when tracked.

I prepared a 2-mm deep layer of BZ solution using the same chemical concentrations as Gowen & Solomon[48], then immersed a silver wire for ~ 20 s to initiate reaction. I measured the temperature to be 20°C with a thermocouple, and found the temperature to be constant over the course of each experiment. I illuminated experiments with blue light-emitting diodes and filmed BZ using a gray-scale, high-speed camera (Emergent HS-4000M) behind a blue band-pass filter. The higher the concentration of blue ferroin, the more light is passed through the BZ and reflected by a white background. The filter blocks red light from unreacted regions, making reacted regions appear bright and unreacted regions appear dark. The transmittance is related to concentration [92], so the brightness measured by the camera provides an idea of the local concentration of the reacted state. I recorded image sequences for 1 to 5 minutes, until spurious fronts from the boundaries of our vessel invaded the field of view. The value of n_{skip} is set to take measurements every 1 s; frame rates were 3 - 5 Hz. The spatial resolution was 0.1218 mm/px for experiments with circular fronts, and 0.0726mm/px for experiments with planar fronts. The front-tracking algorithm is finally used on the image sequences to measure front speed, front thickness, diffusivity, and reaction rate from these experiments.

Figure 2.8 shows distributions of speed and thickness obtained from an experiment with a clear front and back, both growing out from a point trigger. The

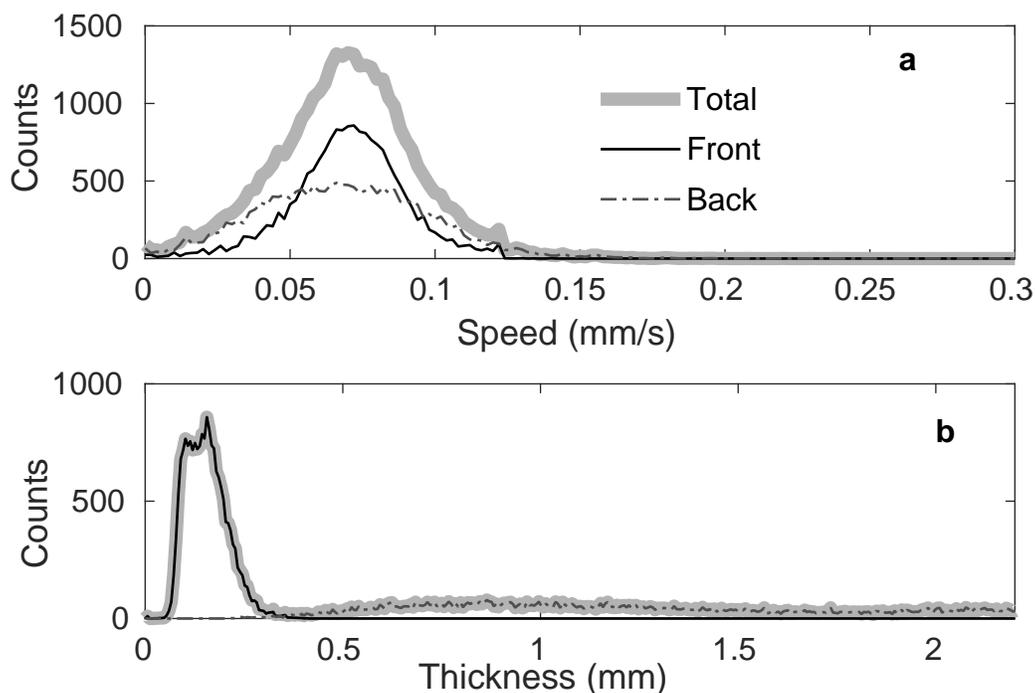


Figure 2.8: Histograms of front speeds (a) and thicknesses (b) gathered from an outward-propagating front. Backs have a much wider distribution of thicknesses. Work of the author published in [78].

front and back have nearly the same mean speed, though the back has a broader speed distribution. Measuring identical speeds for fronts and backs is consistent with physical expectations. Because BZ oscillates with a fixed duration, each back must have a speed matching the front that came before it. However, the front and the back differ starkly in their thickness. Nearly every measurement of front thickness is close to $168 \mu\text{m}$, whereas measurements of back thickness vary from 0.5 mm to to 2.5 mm. As expected, the different chemical kinetics of fronts and backs yield different macroscopic behavior.

Considering just the front, which is well-modeled by second-order kinetics, I can extract well-defined values of v and L from histograms like those shown in Fig. 2.8, then calculate D and α from v and L . Because BZ is actually a complex chain reaction, the values measured correspond to rate-limiting steps. I

used the modal value for v and L (at the peak of each distribution), because the distributions are not symmetric; their large skew cause mean and median to be poor indicators of typical front dynamics.

In practice, ferroin concentration is often varied during laboratory experiments with BZ, because its colors fade slowly over time, and adding more ferroin restores the contrast. Prior studies found that BZ front speed varies little with ferroin concentration [71, 92], but to my knowledge thickness, diffusivity, and reaction rate have not been considered. I performed two sets of experiments with varying ferroin concentrations, one with circular fronts and one with planar fronts, then extracted v and L , and calculated D and α . For all sets, at least 1,500 data points are taken, however for circular most data sets contain over 8000 points, and for planar all data sets have at least 20,000 points, and most have over 70,000. The results are shown in Fig. 2.9. In both cases, and especially with linear fronts, we see essentially no variation of v with ferroin concentration. In six of seven experiments, $71 \mu\text{m/s} \leq v \leq 73 \mu\text{m/s}$. For comparison, prior studies found $v = 95 \mu\text{m/s}$ at 25°C [71], $65 \mu\text{m/s}$ at 18°C [92], and $70 \mu\text{m/s}$ at 20°C [34], all using slightly different recipes for BZ. Our recipe matches Bargteil & Solomon [34] and is similar to Field & Noyes [71]. Considering the reaction's strong temperature sensitivity [92] the agreement is good, and variations are probably related to temperature not recipe. And though front speed is known to depend on curvature [42], we measure no obvious variation with curvature in these experiments. However, we measure the mode of the speed and thickness, which is insensitive to large curvature areas which occupy little of any data set. Therefore it is clear that the algorithm has accurately reproduced prior results.

To my knowledge, neither front thickness nor back thickness has been measured before in BZ (though profiles have been imaged before [89]). We find that front thickness increases with ferroin concentration, as shown in Fig. 2.9. However, since ferroin concentration can be increased between experiments, but never decreased

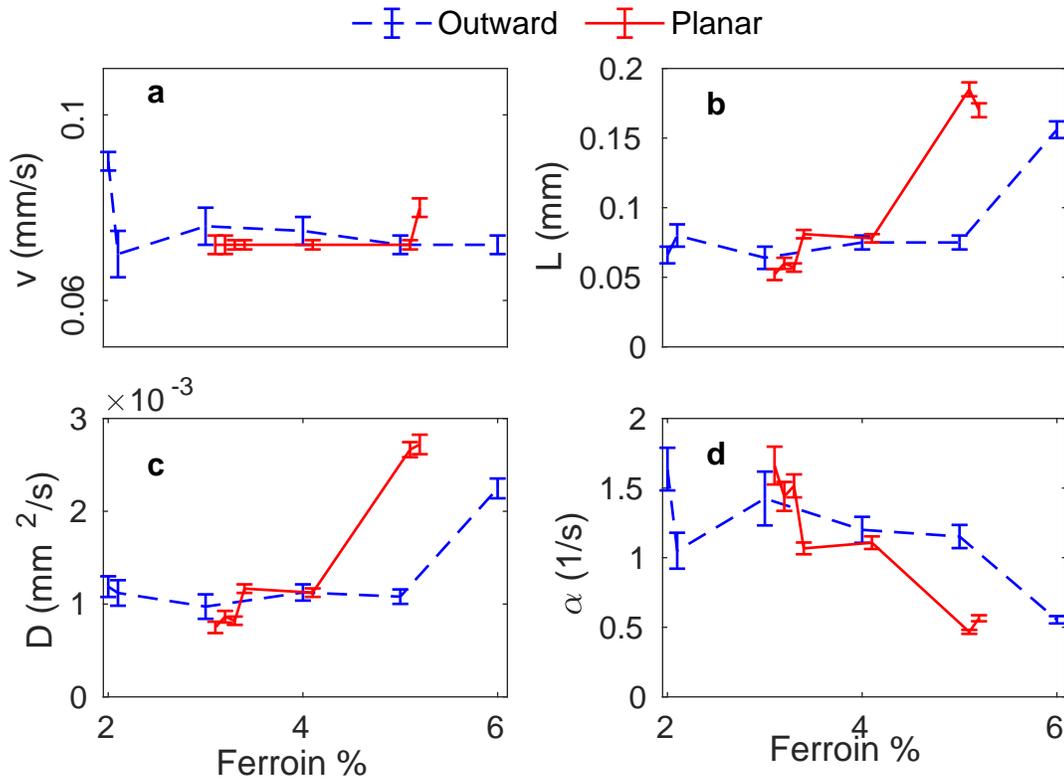


Figure 2.9: Effects of ferroin concentration on (a) front speed, (b) front thickness, (c) diffusivity, and (d) reaction rate. In (a) and (b), errorbars indicate the bin size of the histograms whose peaks are the measurements; other error sources are smaller. In (c) and (d), indicated errors are propagated from the errors in (a) and (b). In all cases, ferroin concentrations were integer percentages, but plot symbols have been offset slightly for readability. Work of the author published in [78].

except by mixing a new batch of BZ solution, ferroin concentration necessarily correlates with the age of the reaction. When the data points shown in Fig. 2.9 are arranged in chronological order instead, the trend is maintained even if several data points were taken with the same ferroin concentration. I found thickness to depend on ferroin concentration (or age) for both planar and circular fronts. In both cases thickness is steady around $70 \mu\text{m}$, until the BZ has aged or gotten enough ferroin. At this point its thickness jumps to around $170 \mu\text{m}$. Further work would be needed to understand what causes the jump from the “low-thickness” data to the “high-thickness” data.

I measured diffusivity and reaction rate to be of the same order of magnitude found in prior studies, which measured $D = 2 \times 10^{-3} \text{ mm}^2/\text{s}$ [2], $D = 1.8 \times 10^{-3} \text{ mm}^2/\text{s}$ [71], and $D = 1.3 \times 10^{-3} \text{ mm}^2/\text{s}$ [95]. The average “low-thickness” diffusivity we measured is $D = 1.0 \times 10^{-3} \text{ mm}^2/\text{s}$, matching Kuhnert et al[95] most closely. For the “high-thickness” data sets this average is $D = 2.5 \times 10^{-3} \text{ mm}^2/\text{s}$. The measured $\alpha = 1.0\text{s}^{-1}$ and for BZ, $\alpha = k[\text{BrO}_3^-][\text{H}^+]$. Using our measurements of α and taking $[\text{BrO}_3^-] = 0.06\text{M}$ and $[\text{H}^+] = 0.8$ from the literature [2], we estimate the reaction rate constant in our experiments to be $k = 27 \text{ M}^{-2}/\text{s}$ for low thickness, and $k = 11 \text{ M}^{-2}/\text{s}$ for large thickness. The first value resolves a contradiction in the literature. Field and Noyes[71] developed a theoretical relationship for front speed which involves the coefficient $(4Dk)^{1/2}$. They measured it to be $24.75 \text{ M}^{-1} \text{ mm min}^{-1}$, but predicted $509 \text{ M}^{-1} \text{ mm min}^{-1}$ based on available values of D and k . We repeat the calculation using our measurements and arrive at the value $22 \text{ M}^{-1} \text{ mm min}^{-1}$, closely matching their experimental value and showing that their estimates of D and k may have been mistaken. The “high-thickness” results come out to a value of $20 \text{ M}^{-1} \text{ mm min}^{-1}$. Further confirmation for our reaction rate comes from a measurement by [95], who found $k = 30 \pm 4 \text{ M}^{-2}\text{s}$ at 30°C .

Beyond measuring coefficients, local front tracking also allows quantitative

comparison of the concentration profiles of BZ fronts and BZ backs. While backs are not second-order fronts, the shape of the solution for second order roughly approximates their shape. Distributions of our thickness measurements, like the one shown in Fig. 2.8, always show two peaks, exactly corresponding to fronts and backs. At fixed time, the thickness is the only parameter for a reaction profile. Figure 2.10 shows profile curves of BZ fronts and backs, using the form given by eq. (2.4), with thicknesses corresponding to the most likely values for fronts and backs, taken from Fig. 2.10. Considering the middle 50% of thickness values for either of the peaks, we can predict that the true profile shape (with fixed μ) has a 50% chance of lying in the range, which is shaded in Fig. 2.10. The range is much wider for backs, probably because second-order reaction kinetics are a poorer approximation than for fronts. Regardless, whereas similar curves were hypothesized and qualitatively sketched in prior publications [2], with local front tracking we have been able to provide a quantitative, statistical characterization of the concentration profile of chemical waves in BZ. To our knowledge, ours is the first such quantitative characterization.

2.7 Summary and Future Work

This chapter presented a new algorithm for tracking the fronts that separate reacted regions from unreacted regions in reaction-diffusion systems, and for measuring the speed and thickness of those fronts. By repeating the process at many points along fronts, the algorithm can measure spatial variations in front speed and thickness, as well as providing large numbers of measurements appropriate for statistical analysis. Further, the algorithm also provides measurements of diffusivity and reaction rate. We demonstrated that the algorithm made accurate measurements when applied to validation data (where macroscopic quantities of front speed and thickness were specified) and simulation results (where micro-

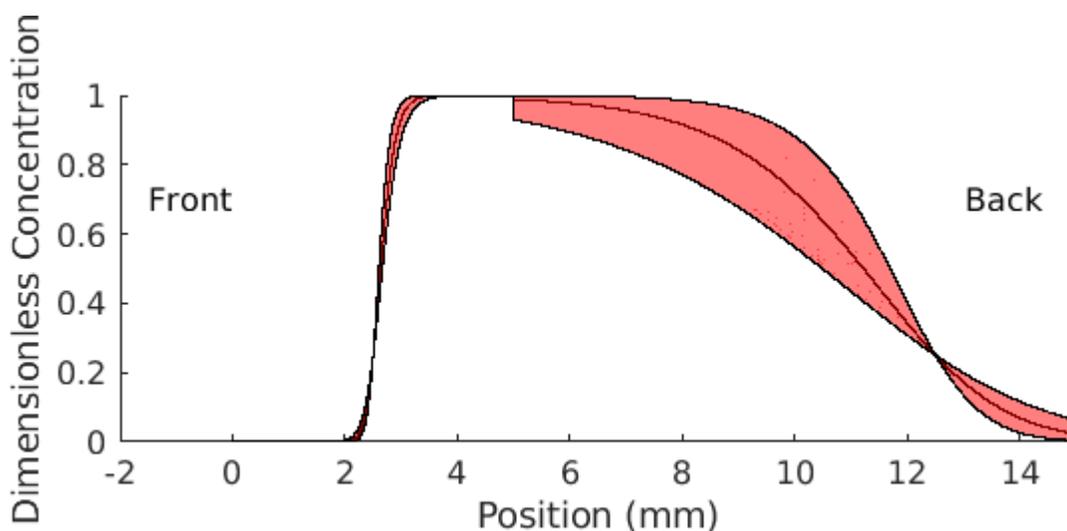


Figure 2.10: Measured concentration profile of chemical waves in the BZ reaction. The curve shows the most likely profile, and the shaded region represents the interquartile (middle 50%) of possible profiles. Front and back are shown farther apart than their actual separation, which depends on the reaction’s oscillation time and is not indicated here. Work of the author published in [78].

scopic quantities of diffusivity and reaction rate were specified). Results varied less than 1% over a wide range of the user-defined parameters the algorithm requires. Finally I applied the algorithm to laboratory experiments with BZ and obtained statistics with measurements similar to published values.

With this algorithm I am able to simplify reaction-diffusion in terms of the dynamics governing the motion of reaction fronts. Since it can detect local variation, it also shows promise for studying the complexities inherent in ARD systems, while being a great simplification. In addition, this new algorithm offers practical advantages over existing methods for microscopic properties. Reaction rate and diffusion are regularly measured with *in situ* chemical tests, but specialized equipment like spectrometers and magnetic resonance imaging machines are typically required. Front tracking measures diffusivity and reaction rate with just a camera, making it less expensive and less invasive. However, it is also compatible

with concentration fields produced with more complicated instruments and simulations. The method is independent of length scale and timescale; it requires only that camera speed and magnification be sufficient to capture the front dynamics. Front tracking also enables future investigations that were not possible before. We have measured the front speed and thickness in one chemical system, the BZ reaction; speed and thickness can immediately be measured in many other chemical systems as well. Diffusivity and reaction rate may also be measurable in other chemical systems, as can the concentration profiles of fronts in those reactions (as in Fig. 2.10).

Having built confidence that the algorithm accurately tracks fronts in reaction-diffusion systems, I next apply it to advection-reaction-diffusion by simultaneously measuring the flow field and accounting for its affect on fronts in Chapter 3. Local measurements of front speed and flow velocity can then be used to develop or test simplified models of advection-reaction-diffusion in terms of the fronts. In fact, separating reacted from unreacted with a front allowed me to make the observations which comprise the remainder of this thesis. Without the definition of a front, its speed, its thickness and the machinery to measure it, the determination that anything abnormal is occurring with a reaction's growth in a flow is difficult to notice and impossible to define.

2.8 Acknowledgments

Reproduced from Nevins, T.D. and Kelley, D.H., "Front tracking for quantifying advection-reaction-diffusion", *Chaos* **27**, 043105 (2017), with the permission of AIP Publishing, which has been modified to fit this thesis. The thesis author is the primary author and was responsible for writing and troubleshooting the algorithm described, the validation section, the simulation section, the experiment section, writing the paper, and preparing the figures. Co-author Douglas Kelley wrote

the initial front tracking algorithm, provided guidance during the development, provided guidance on figure formation, and edited the paper and figures. Insightful conversations with N. T. Ouellette and J. G. Puckett also helped to inspire this Chapter.

3 Front Tracking Including Advection

This chapter is based on Reference [96]. It has been reformatted and edited to fit within the thesis. The thesis author was the primary author on this work.

3.1 Introduction

In Chapter 2, I showed an algorithm for locating concentration fronts in a reaction-diffusion system and measuring their dynamics. Chapter 2 also established the physical relevance of reaction fronts as a way to understand a changing concentration field in a simplified way. This chapter will discuss an extension of the front tracking algorithm to full advection-reaction-diffusion (ARD) systems in two dimensions. The algorithm enables local measurement of front growth dynamics in the presence of flow. It is our hope that by making this analysis technique possible, future researchers will be able to make observations necessary to develop models of front growth that accurately replicate real ARD behavior. For instance, front tracking allows investigation of: What flow characteristics alter the reaction front speed? How could we optimize or control the growth of a reactive scalar? When can ARD simulations be avoided in favor of front simulations, which are much less computationally intense?

Adding a flow to the dynamics of reaction-diffusion significantly complicates the motion of reacted regions because variation of the flow in space and time makes the usual solutions [39, 43] to the reaction-diffusion equation incorrect. With flow, reacted regions grow in different directions and at different speeds from place to place. For instance, frozen fronts are known to persist despite headwinds exceeding the front speed, in contexts including porous media [10, 11], flow past a bluff body [33, 97], and vortex flows [51, 98]. Poiseuille flow has been shown to alter front speed through a change in the front shape [35–37]. In some cases, gentle flow enhances reaction and violent flow inhibits reaction [5, 29, 99]. Being able to identify the spatial dependence of front speed, and the cause of that front speed, is essential to determining the root cause of each of these phenomena.

The dynamics of ARD, which depend on a partial differential equation (PDE), can be approximated with a simpler dynamics that depends only on ordinary differential equations (ODEs), by considering instead the motion of reaction fronts. These fronts are still defined as concentration level curves at a characteristic level separating reacted from unreacted. The ODE which gives these simplified dynamics is known as the Eikonal equation [43]. The Eikonal equation relies on the local front velocity being a superposition of the flow velocity, \mathbf{u} , and the velocity of a front in a stagnant fluid, \mathbf{v} , which is the “chemical velocity”. In Chapter 2, the chemical velocity was what was measured, and I showed that front dynamics provided a great deal of information about reaction diffusion systems. Burning invariant manifolds (BIMs) and burning Lagrangian Coherent Structures (bLCS), are developed from such an approximation, and predict when and where a front will advance [34, 44–46, 48, 50, 51].

The Eikonal approximation and BIM/bLCS theory traditionally assume that the chemical velocity of a front is unaffected by flow. That is to say $|v|$ is constant. This assumption serves as a starting point for analyzing ARD systems using front dynamics. By measuring chemical velocity in real experiments, we wanted to test

this assumption and determine its limits and possible models for its improvement. Since the ODEs of the Eikonal equation are much simpler than the PDE, determining when ODEs are accurate would indicate when full ARD simulations can be avoided in favor of simpler front advection simulations. If ODEs are found to be inaccurate in some flow regimes, front tracking observations suggest corrections. For example, the Eikonal assumption can include a front curvature correction [42] that has previously been measured. While front behavior changing as a function of local flow properties is expected, the front dynamics required to account for it have not been tested.

Tracking chemical velocity separately should also inform existing experimental observations of reaction front growth. In recent studies of a reaction pinned to a flow obstacle [33], it was found that the Eikonal approximation could explain the selection of frozen fronts but failed to predict front shape. A deeper understanding of chemical front velocity is needed to understand what caused this difference.

In this chapter, I generalize our earlier front tracking algorithm in order to quantify the growth of reacted regions in *flowing* fluid. I put special emphasis on separating the advection of the front from motion of the front relative to the fluid. The algorithm can measure front dynamics in data from experiments or simulations. Throughout this chapter, “front tracking” is specifically an algorithm for *measuring* a reaction front in data produced either in simulation or experiment. This is contrary to prior literature where the phrase “front tracking” has sometimes referred to a simulation method [85, 86].

I begin the explanation of front tracking with flow in §3.2, where I will explain how front velocities can be separated into flow velocities and chemical velocities, and the physical arguments that justify that separation. In §3.3 I outline the front tracking algorithm. In §3.4 I validate that front tracking provides accurate measurements of total front velocity and chemical front velocity in two simulations: uniform flow and hyperbolic flow. In §3.5 I describe the experimental apparatus for

measuring reaction and flow simultaneously, thereby enabling front tracking with flow in time-dependent ARD experiments. I also describe experiments using the excitable Belousov-Zhabotinsky reaction in time-independent and time-dependent flows. In §3.6 I show results of front tracking in experiments, demonstrating that front tracking self-consistently maps fronts to the following time step. I also demonstrate the algorithm’s resilience to user parameter choice. My results suggest an unexpected correlation between flow and chemical speed, which we will investigate in Chapter 4. Finally, in §3.7 I describe strengths and weaknesses of the algorithm, opportunities for improvement, and future research questions that may be answered by this algorithm.

3.2 Separating Flow Velocity from Chemical Velocity

The goal in developing this algorithm is to extend the study of reaction front dynamics in ARD systems. Specifically, when a reacted region is subject to simultaneous local reaction, diffusion, and advection by a flow, its edge — the reaction front — moves. Mathematically, the front is a characteristic concentration level curve, for which the motion is determined by the changes of the underlying concentration field of the chemical species, c , given by the ARD equation,

$$\frac{\partial c}{\partial t} + (\mathbf{u} \cdot \nabla)c = D\nabla^2 c + kF(c). \quad (3.1)$$

Here \mathbf{u} is the advection velocity, D is the diffusion coefficient of the chemical species, k is the reaction rate, and $F(c)$ specifies the reaction kinetics. Note that this is identical to Eq. 2.1). The $F(c)$ term ensures the global integral of any specific c is not conserved over time. In reacting systems this is generally because the reactant is converted into a different species. In Chapter 2 I covered the algorithm for capturing the front dynamics in a system with $\mathbf{u} = 0$. By measuring

normal displacement of the front at a large number of points I replicated known front speed, and via known solutions to the reaction-diffusion equation [39, 90, 91], we replicated the diffusion coefficient and reaction rate as well.

Here we present a method for obtaining a front velocity for ARD systems, in which $\mathbf{u} \neq 0$. This change represents a significant complication. Fluid flows relevant to reactive mixing could be anything from laminar to turbulent, and reaction terms are typically nonlinear. What this means is, even if the reaction has no feedback to the flow (as I will assume in all experiments for this thesis) the flow effect on concentration is not separable from the reaction term. While this means there are numerous possible interactions to consider, I will start with the most obvious. I will assume that flow only serves to displace a reaction front.

Even focusing on only flow displacement is a suitable challenge, because it can move material in an arbitrary direction. The velocity of a front element, \mathbf{w} , can then be approximated as

$$\mathbf{w} = \mathbf{u} + \mathbf{v} \tag{3.2}$$

where \mathbf{v} is chemical velocity. The direction of \mathbf{v} is always locally normal to the front because fronts are concentration level curves, and diffusion moves material down the concentration gradient, normal to the level curves. An example of the result of front tracking, in which \mathbf{u} , \mathbf{v} , and \mathbf{w} have been measured, is shown in Fig. 3.1. Equation 3.2 is the Eikonal approximation minus the curvature correction, which underlies BIM theory [50, 51] and bLCS theory [45]. (For a derivation, see Spangler and Edwards, 2003 [43].) BIM theory with constant chemical speed has previously been used successfully to understand front propagation in time independent and periodic flows [34, 48, 49], and the Eikonal approximation has proven useful for understanding frozen fronts in adverse flows [10, 33]. Measuring chemical velocity locally could allow quantification of the range of applicability of the Eikonal approximation.

This approach will allow me to measure the chemical front speed as it is

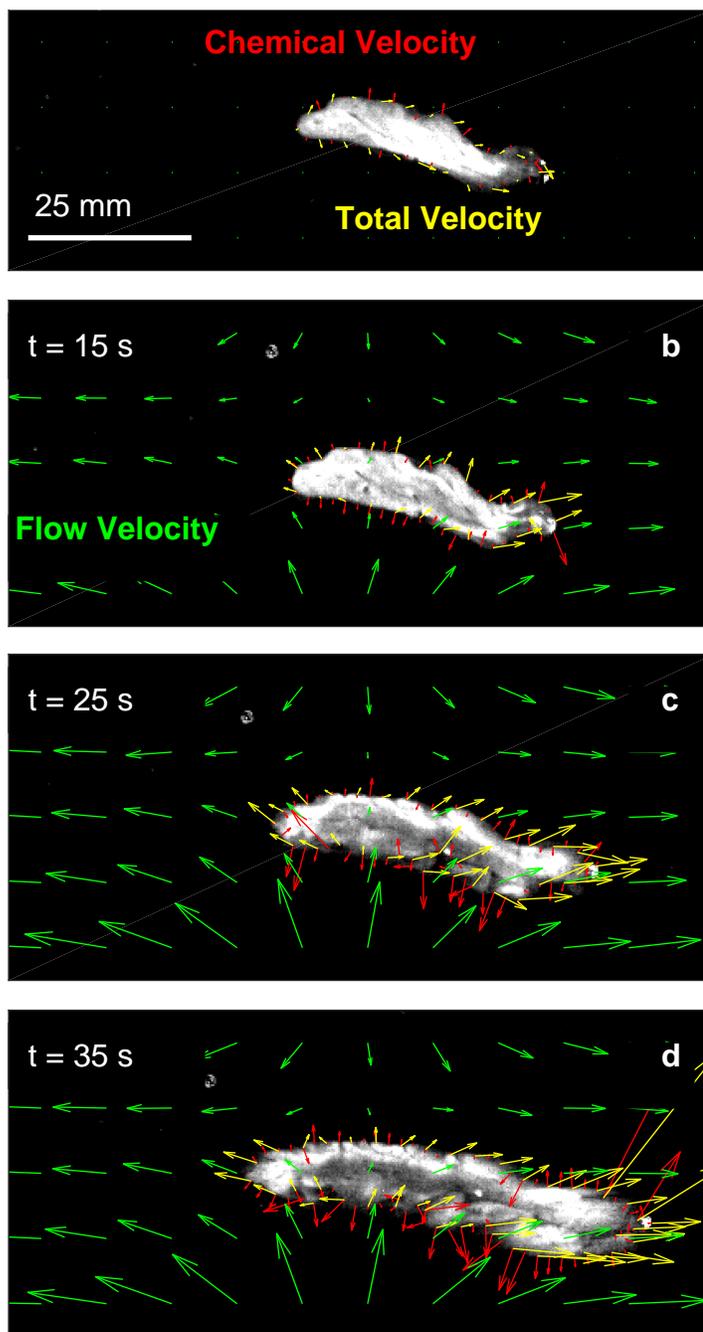


Figure 3.1: Example results of front tracking with flow, showing experimental product concentration snapshots (gray-scale). The flow (green), measured total velocity (yellow), and measured chemical velocity (red) are overlaid (and down-sampled). Flow was initiated after $t = 5$ s and was fully developed by $t = 25$ s. Work of the author published in [96].

typically used in the Eikonal approximation. There are additional benefits to this approach though. It effectively removes front displacement from the overall motion of the front, leaving only the ordinary reaction growth and any interactions between the reaction and the flow. For instance, if turbulence in some area is leading to apparent increased front speed, a front tracking user will get a higher chemical front speed in that area. Furthermore, this separates all the front motion that does not lead to overall reaction growth. If a front is adequately sharp, then the total reacted quantity is determined by the volume within the reaction front. This volume changes according to the flux equation

$$\frac{dV}{dt} = \oint \mathbf{w} \cdot \hat{\mathbf{n}} dS \quad (3.3)$$

which can be written,

$$\frac{dV}{dt} = \oint (\mathbf{u} + v\hat{\mathbf{n}}) \cdot \hat{\mathbf{n}} dS \quad (3.4)$$

$$= \iiint \nabla \cdot \mathbf{u} dV + \oint v\hat{\mathbf{n}} \cdot \hat{\mathbf{n}} dS \quad (3.5)$$

$$= \oint v dS \quad (3.6)$$

where I have used Eq. 1.7 to eliminate \mathbf{u} . Thus the v defined and the front surface are the only thing that needs to be measured in order to understand the time evolution of the total reaction volume.

3.3 Front Tracking Algorithm

To track fronts, we must first locate them. As in Chapter 2, first I find the edges of binary field $c > c_{\text{thresh}}$ at a time t , choosing a threshold c_{thresh} that separates reacted regions from unreacted regions. For ARD systems in which characteristic reaction times are much smaller than characteristic diffusion times (that is, the second Damköhler number $Da_{\text{II}} = kl_D/D \gg 1$, where $l_D = 1$ mm is a length scale and k is reaction rate), fronts are sharp and results are insensitive to the choice

of c_{thresh} . For the experiments and simulations discussed below, $Da_{\text{II}} = 1300$. Typically the concentration field c is measured at discrete locations on a grid (pixels), so the edges of the binary field $c > c_{\text{thresh}}$ are piecewise-linear curves which follow pixel boundaries and have sharp corners. The true concentration field has no corners because diffusion smooths singularities, so I smooth the piecewise-linear curves using a sliding line fit. Smoothing also uses data from adjacent pixels to achieve sub-pixel accuracy in locating the fronts. To correct small-area fluctuating noise in the concentration field, we eliminate $c > c_{\text{thresh}}$ regions that are smaller than a minimum area, which is set by the user. I advise choosing a minimum area that is larger than the data resolution δx but smaller than the smallest structures of interest. The smoothed boundaries of the remaining regions are taken to be the original fronts. The process repeats at a time dt later to locate the subsequent set of fronts, keeping the same value for c_{thresh} . Thickness of both sets of fronts can be obtained using the same algorithm described previously [5].

The fronts at time $t + dt$ differ from the original fronts at time t by a displacement $d\mathbf{x}_w$ that varies in space and time. Eq. 3.2 indicates that the instantaneous displacement is due to a combination of flow velocity and chemical velocity:

$$d\mathbf{x}_w = d\mathbf{x}_u + d\mathbf{x}_v, \quad (3.7)$$

where $d\mathbf{x}_u$ is the displacement due to flow and $d\mathbf{x}_v$ is the chemical velocity. Assuming \mathbf{u} is known, either from simulation or from measurements, then for a fluid element at location \mathbf{x} at time t ,

$$d\mathbf{x}_u = \int_t^{t+dt} \mathbf{u}(\mathbf{x}) dt', \quad (3.8)$$

where t' is an integration variable and the integration is performed in a Lagrangian sense, tracking the fluid element as it moves. By integrating all points on the original fronts forward in time according to Eq. 3.8, I calculate the shapes and positions those fronts would take at time $t + dt$ if they marked high-concentration regions of

a passive scalar with zero diffusivity. I refer to the result as the “advected original fronts”.

The displacement between the advected original fronts and the subsequent fronts is $d\mathbf{x}_v$ and can result only from the chemical velocity, which acts in a direction normal to the original fronts, as discussed above. To calculate $d\mathbf{x}_v$ locally, front tracking draws a normal line at a point on the advected original fronts, and locates the nearest location where it intersects one of the subsequent fronts; the distance to the intersection is $d\mathbf{x}_v$. I then estimate the local chemical velocity as

$$\mathbf{v} = \frac{d\mathbf{x}_v}{dt}; \quad (3.9)$$

I emphasize that velocity is approximated by the quotient of the displacement and the time step, both of which are discrete. Finally, I calculate the total displacement $d\mathbf{x}_w$ at each point using Eq. 3.7 and estimate the total velocity as $\mathbf{w} = d\mathbf{x}_w/dt$. The intersection is unlikely to occur at one of the discrete points identified as being on the subsequent front; intersections typically occur between points. Sometimes, fronts fade or change shape such that the nearest intersection lies on a different front. To eliminate those non-physical measurements, front tracking discards all displacements larger than a user-defined maximum. Usually I choose a maximum displacement many times larger than the chemical velocity in stagnant fluid. Figure 3.2 shows an example in which the original front, advected original front, subsequent front, flow velocity, chemical velocity, and total velocity are all sketched.

I estimate the integral in Eq. 3.8 using an integration scheme that is first-order in time, and I choose delay time dt which need not match the shortest timescale dt_{\min} at which an experiment is imaged or simulation results are calculated. In the results described below, I often choose $dt \neq dt_{\min}$ because $|\mathbf{u}| > |\mathbf{v}|$. Accurately measuring flow velocity requires a higher frame rate (smaller time step) than measuring chemical velocity. In fact, chemical velocity usually causes fronts to

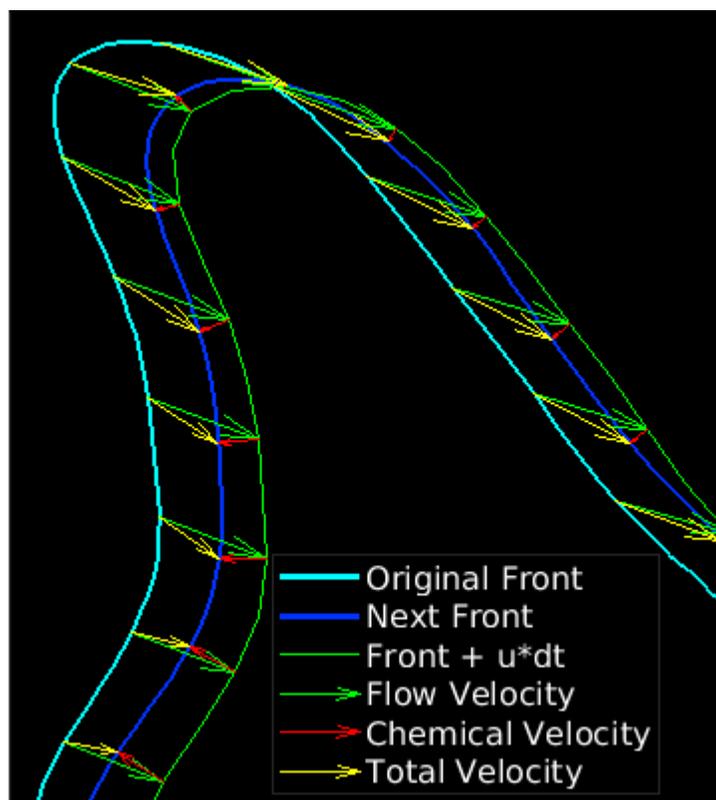


Figure 3.2: Cartoon of a single step of front tracking. Front tracking advects points on the original front by Lagrangian integration of flow data to determine flow displacement. Where the advected original front differs from the next front (measured one time step later), the difference is attributed to a chemical velocity locally normal to the advected original front. Thus the total front velocity is the vector sum of flow velocity and chemical velocity. Work of the author published in [96].

advance only a fraction of a pixel in the time dt_{\min} : $|v| \ll \delta x/dt_{\min}$, where δx is the size of a pixel (the grid size). Minor noise fluctuations in pixel brightness can therefore appear to cause rapid front motion that is not physical. Choosing $dt > dt_{\min}$ reduces the noise. There is a tradeoff however, dt must be kept small enough that the flow does not alter front perimeters too drastically within a time step. Choosing dt_{\min} small makes the integration in Eq. 3.8 more accurate, as I will discuss further below. In my experiments, typically $dt \approx 0.5$ s while 0.1 s $> dt_{\min} > 0.025$ s.

The algorithm then is equivalent to assuming that during each delay time dt , the front advances first by the local flow velocity, then by the local chemical velocity. In reality, both processes act simultaneously. Similar approximations are common in simulations, which often calculate different dynamical terms in sequence. In the limit $dt \rightarrow 0$, the results of the algorithm converge to match the results of both processes acting simultaneously, since diffusion is instantaneously perpendicular to the front. Still, when dt is finite, advancing the front first by flow velocity and then by chemical velocity introduces small errors to front tracking, due in part to the continuous reorientation of local front elements. The algorithm might also be written to assess the effects of advection and reaction in the opposite order, by advecting the current front *backwards* in time, then attributing the offset between that advected front and the *previous* front to chemical velocity. Evaluating the effects in the opposite order would change the sign of some inaccuracies of numerical integration and differentiation, but the precision would otherwise be similar, and measurements would still converge as $dt \rightarrow 0$.

Written in MATLAB, the open-source algorithm is freely available¹.

¹<http://www.me.rochester.edu/projects/dhkelley-lab/>

3.4 Tracking Fronts in Simulation Data

I used two series of simulations to test the front tracking algorithm in the presence of flow. First, I performed direct numerical simulations of an ARD system on a second-order chemical reaction in the presence of a uniform flow. This system is simple enough that I have clear, *a priori* expectations for the front velocity. The reaction term is given by

$$F(c) = c(1 - c), \quad (3.10)$$

where c is made dimensionless such that it varies between 0 and 1, so its scaled reaction rate constant becomes $\alpha = kc_0$. In the special case of reaction-diffusion ($\mathbf{u} = 0$) we have the equation covered in Chapter 2.

$$\frac{\partial c}{\partial t} = D \frac{\partial^2 c}{\partial x^2} + \alpha c(1 - c). \quad (3.11)$$

Equation 3.11 has an analytic solution [39, 90, 91] in which an initial band of high concentration grows in both directions at a constant speed given by

$$v_f = \frac{5}{6} \sqrt{6D\alpha}. \quad (3.12)$$

In Chapter 2 I showed that front tracking accurately reproduced this front speed in simulation and experiment in the $\mathbf{u} = 0$ case. Including a uniform flow $\mathbf{u} = U\hat{\mathbf{x}}$ (where U is a constant speed and $\hat{\mathbf{x}}$ is a unit vector in the direction normal to the front) and keeping all other conditions the same, I would expect to observe a propagating band whose edges have fronts with total speeds $w = U + v_f$ and $w = U - v_f$. This new situation is a Galilean transformation, and it is mathematically identical to filming with a camera moving at $-U\hat{\mathbf{x}}$.

I simulated Eq. 2.1 in a two-dimensional periodic domain, using Adams-Bashforth time stepping and central differences for spatial gradients. Initially, the concentration was zero everywhere except in a band where c varied with x according to

the analytic solution [39] of Eq. 3.11,

$$c(x, 0) = \begin{cases} \frac{1}{\left(1+e^{-(x+1)/L}\right)^2} & x \leq 0 \\ \frac{1}{\left(1+e^{(x-1)/L}\right)^2} & x > 0, \end{cases} \quad (3.13)$$

where

$$L = \sqrt{6D/\alpha} \quad (3.14)$$

is the front thickness. The initial concentration did not vary with y , making the problem one-dimensional, consistent with Eq. 3.12. I chose values of D and k near those of the Belousov-Zhabotinsky reaction used in the experiments described in §3.6. The grid size was $5 \mu\text{m}$, and for stability, I used a $4.11 \mu\text{s}$ time step, saving one concentration field every 0.16 s. I used those concentration fields, along with flow fields $\mathbf{u} = U\hat{\mathbf{x}}$, as inputs for the front tracking algorithm.

Figure 3.3 shows the results of simulation and subsequent front tracking. In snapshots from one simulation, the reacted region widens and moves to the right as expected. I repeated the simulation, first varying the flow speed over the range $0 < U < 2 \text{ mm/s}$, then varying D to adjust the chemical speed $0.05 \leq v_f < 0.2 \text{ mm/s}$ according to Eq. 3.12. In all cases, the flow speeds and chemical speeds measured via front tracking closely match the true values. Measurement errors in chemical speed are usually less than $\delta x/dt = 0.03 \text{ mm/s}$, and caused by uncertainty when tracking fronts at sub-pixel accuracy; errors do not grow with true flow speed or with true chemical speed. To confirm, I repeated the tracking, setting dt to be three times as large, and found that the error in chemical speed dropped by a factor of three as expected. Reducing the grid spacing δx would also reduce measurement errors. Though the chemical speed is measured at many locations, only a few distinct values appear in Fig. 3.3. As mentioned above, the algorithm typically locates fronts with sub-pixel accuracy via spatial smoothing, but in this unusual case of strictly straight fronts, smoothing has no effect, and

fronts remain on (discrete) pixel boundaries, so the calculated chemical speeds are discretized.

Considering a one-dimensional front in a uniform flow proved helpful for validating the algorithm, because I could form *a priori* expectations of front speed in true advection-reaction-diffusion dynamics governed by Eq. 2.1. Such a simple case, however, captures no effects caused by complexities in the flow field, such as flow velocity gradients or front turning. To further validate the algorithm, I considered another example for which we have *a priori* expectations, but in which there are nonzero velocity gradients and curved fronts, though the fronts are no longer governed by Eq. 2.1. The example is described in §IIIB1 of the paper by Mitchell and Mahoney [44]: a reaction front propagates with constant chemical speed v_0 in flow around a hyperbolic critical point at the origin, such that the front dynamics is reduced to a set of ordinary differential equations,

$$\begin{aligned} w_x = \frac{\partial x}{\partial t} &= -Ax + v_0 \sin \theta, \\ w_y = \frac{\partial y}{\partial t} &= Ay - v_0 \cos \theta, \\ \frac{\partial \theta}{\partial t} &= 2A \cos \theta \sin \theta. \end{aligned} \tag{3.15}$$

Here x and y specify the location of a front element in two dimensions, θ specifies front tangent orientation with respect to the x axis, and A is a positive constant with units of inverse time that sets the speed of the hyperbolic flow.

The front dynamics can be calculated from Eq. 3.15 and are described in the paper [44]: a front propagating from the left is blocked by a BIM at $x = v_0/A$, while the reacted region continues to spread vertically. Furthermore, the dynamics of a front element with orientation $\theta = \pi/2$ can be solved analytically via separation of variables, yielding

$$x = \frac{v_0}{A} + \left(x_0 - \frac{v_0}{A}\right) e^{-At},$$

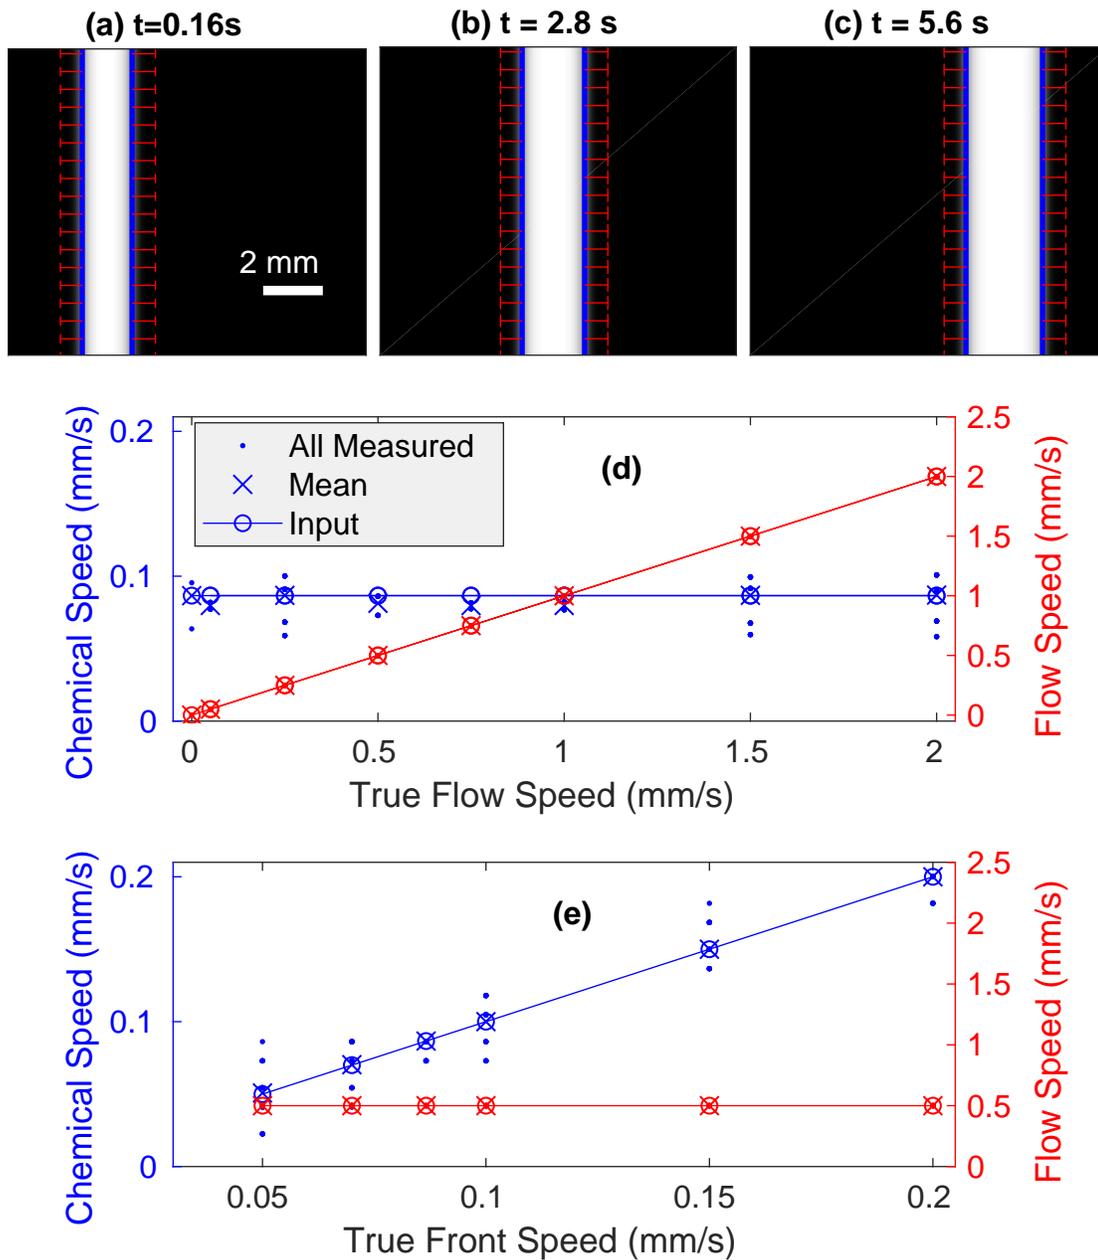


Figure 3.3: Validation of front tracking by simulating the advection-reaction-diffusion equations in uniform flow. (a-c) Snapshots of a reacted region, showing fronts and chemical velocities determined by the front tracking. (d) Flow speed and chemical front speed measurements, for a range of true flow speeds, as simulated. (e) Measurements for a range of true chemical speeds, as simulated. Agreement between measurements and specified values are within 1 pixel/frame for front speed. Work of the author published in [96].

where x_0 gives one component of the position of the front element at $t = 0$. The front element approaches $x = v_0/A$ at long times, as expected. Front elements with other orientations approach $\theta = \pi/2$ or $\theta = 3\pi/2$ over time because those are the stable values of θ in Eq. 3.15, such that their long term behavior is the same. Differentiating, we find that in the $\theta = \pi/2$ case, the total front speed decays exponentially over time:

$$w_x = \frac{\partial x}{\partial t} = -A \left(x_0 - \frac{v_0}{A} \right) e^{-At}. \quad (3.16)$$

The results of simulating Eq. 3.15 with $v_0 = 1$ mm/s and $A = 1$ s⁻¹ are shown in Fig. 3.4. A front that is initially circular and positioned left of the hyperbolic point approaches the BIM and stretches laterally, as expected. After simulating with Eq. 3.15, concentration fields are assembled which have high intensity inside front boundaries and low intensity outside. Applying the front tracking algorithm, I measure $\langle w_x \rangle$ as it varies in time, finding close agreement with Eq. 3.16. Measurements of chemical speed v are more complicated, and illustrate the balance required in choosing input parameters, because of the finite resolution of the data. As Fig. 3.4 shows, when I choose $dt = dt_{\min}$, the distribution of chemical speed measurements has a sharp peak at 1 mm/s, matching the true value. However, the distribution also has long tails: at some locations, the measured chemical speed differs substantially from the true value. As discussed above, choosing $dt > dt_{\min}$ can reduce noise in the chemical speed measurements. Choosing $dt = 10dt_{\min}$ shortens the tails as expected, but also moves the peak of the distribution away from the true chemical speed by a few percent.

Choosing $dt > dt_{\min}$ also moves the mean value $\langle v \rangle$ away from the true value by a few percent, as shown. That discrepancy is caused by the nonzero velocity gradient of this flow, and by using discrete approximations for derivatives. In this example, most measurements occur near the BIM, where the total front speed is zero, so the true chemical velocity is equal and opposite to the true flow velocity.

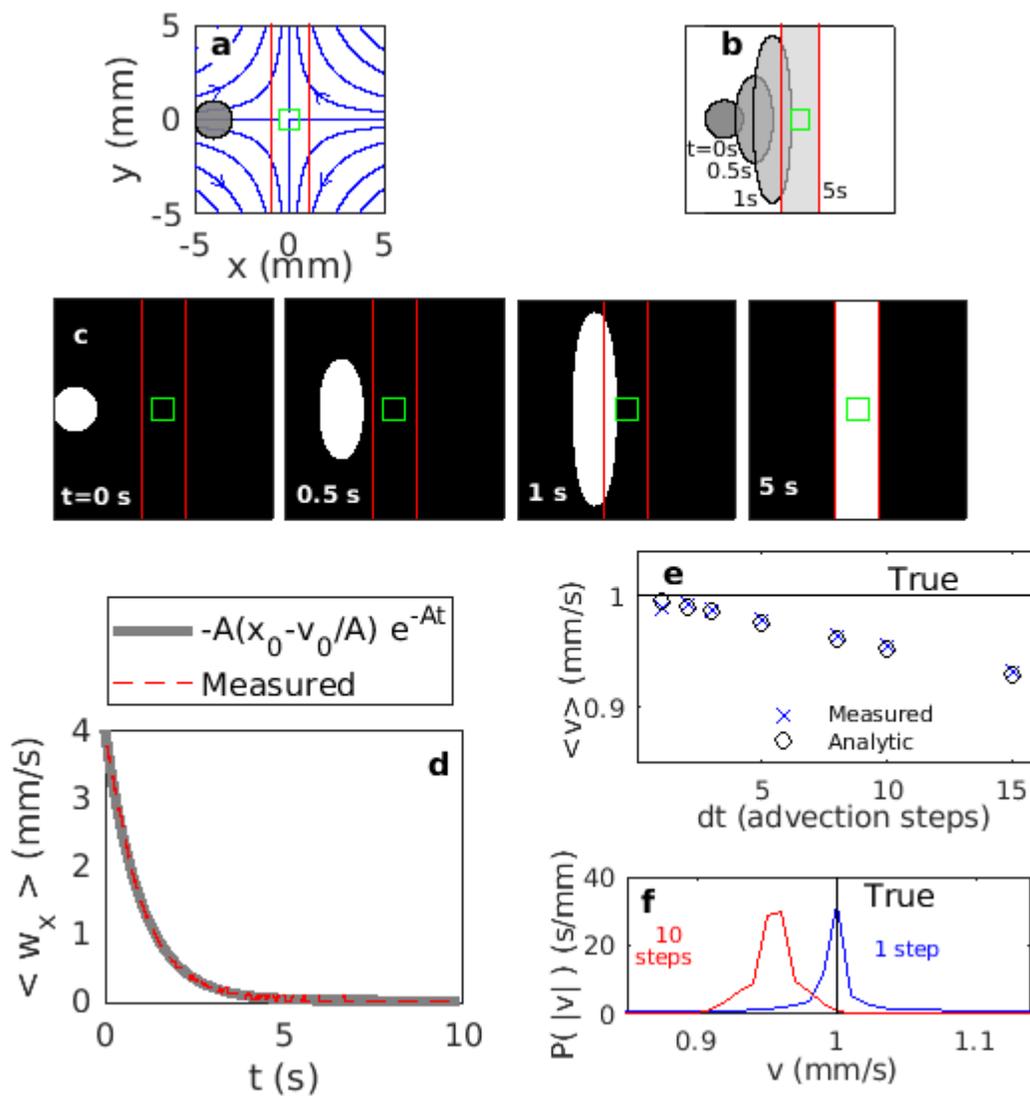


Figure 3.4: Validation of front tracking by simulating an advancing front near a hyperbolic critical point. (a) Flow field (blue streamlines), hyperbolic point (green square), BIMs (red lines), and initially-reacted region (gray). (b) Simulated front at four subsequent times. (c) Corresponding concentration fields at the same times. (d) Measured mean total front speed $\langle w_x \rangle$ closely matches Eq. 3.16, as expected. (e) When dt is small, the distribution of local chemical speeds is very near the true value, but (f) the distribution has long tails. Increasing dt shifts the peak but shortens the tails. Work of the author published in [96].

To measure chemical velocity, the algorithm first uses Eq. 3.8 to calculate $d\mathbf{x}_u$, the displacement of the front over time dt by flow alone. Since $d\mathbf{x}_w = 0$ at the BIM, $d\mathbf{x}_v = -d\mathbf{x}_u$, according to Eq. 3.7. Finally, the algorithm calculates the chemical speed using Eq. 3.9, which estimates a derivative by division. If $dt = dt_{\min}$ then Eq. 3.8 reduces to $d\mathbf{x}_u = \mathbf{u}(\mathbf{x}) dt$ (an integral estimated by simple multiplication) and the algorithm measures the true chemical speed exactly. Similarly, if the velocity gradient is zero, the velocity remains constant over the integral in Eq. 3.8, which again yields $d\mathbf{x}_u = \mathbf{u}(\mathbf{x}) dt$, so the true chemical speed is again measured exactly. However, if $dt > dt_{\min}$ and the velocity gradient is nonzero, a small systematic error occurs. In this case, the front is advected in the $-\hat{\mathbf{x}}$ direction, where flow speed is slightly slower. Then the front is advected at this slower speed, and the process is repeated for dt seconds. Thus it works out that $\mathbf{v} = -\langle \mathbf{u} \rangle$, where the average is over the integration time steps. This is systematically lower magnitude than the \mathbf{u} located at the BIM. Therefore, estimating the integral in Eq. 3.8 by summing over multiple steps of size dt_{\min} produces a higher-precision result than Eq. 3.9. However front speed measurements can deviate from expectations by a few percent in this example as a tradeoff, as shown in Fig. 3.4.

Setting $dt = dt_{\min}$ eliminates the discrepancy illustrated in Fig. 3.4 but increases measurement noise caused when the chemical speed is less than $\delta x/dt$, as discussed above. Thus I leave dt as a user-specified parameter instead of setting $dt = dt_{\min}$ everywhere, and allow users to make their own decisions about the trade-off. The best solution is to minimize δx by increasing the spatial resolution of the simulations or experiments to which the front tracking algorithm is applied. Moreover, the simulation shown in Fig. 3.4 represents an extreme case, since the front aligns with the BIM, where this error is maximized. In experiments, I observe much smaller errors of this type, as described in §3.6.

3.5 Experimental Advection-Reaction-Diffusion Devices

Having verified that the results of tracking fronts in a simple, simulated ARD system closely match expectations, I set out to track fronts in ARD experiments. Below I show results from three different devices that drive three different flow patterns, but use the same chemistry and imaging apparatus, similar to what I will use in Chapter 5. We track fronts produced by the Belousov-Zhabotinsky (BZ) reaction [2, 91–94]. BZ is an excitable redox reaction catalyzed by ferroin indicator. We prepare BZ using the recipe of Bargteil and Solomon [34], adding ferroin throughout the experiment to maintain contrast. Ferroin is known to have a negligible effect on front speed [92]. I initiate reaction at a chosen location by immersing a silver wire for ~ 20 s. A large number of inert $98 \mu\text{m}$ fluorescent green tracer particles (Cospheric UVPMS-BG-1.025) float on the surface of the fluid layer and make its motion visible. These particles are small enough to follow the flow accurately, as indicated by their small Stokes number $St \leq 0.5$.

Experiments are illuminated with blue LEDs, which excite the tracer particles, causing them to fluoresce green. The BZ solution is blue in its reacted state and red in its unreacted state. I used a magenta background beneath the fluid layer which reflects blue LED light only in reacted regions so they appear bright. I image BZ experiments with two hardware-synchronized cameras (Emergent HS-4000M), as sketched in Fig. 3.5. One camera images the concentration field and is equipped with an optical filter that passes blue light but blocks green, to accentuate the blue reacted regions while minimizing visibility of the green tracer particles. The brightness of its images measure BZ product concentration. The other camera images particle positions and is equipped with an optical filter that passes green light but blocks blue, to accentuate particles while minimizing visibility of the reaction state. The magenta background helps dim the brightness of the blue

reaction in the green camera. Particle tracking [75] provides the flow field. Though the cameras are mounted close together, their fields of view are not identical. Before taking data I photograph a calibration grid with both cameras, which allows me to determine their relative shift and rotation, as well as the resolutions of both. The shifting algorithm was written by Doug Kelley. In the experiments described below, resolutions are typically about $150 \mu\text{m}/\text{pixel}$. Each camera records images that are 2048×2048 pixels. I vary the frame rate depending on the flow speed, sometimes imaging as rapidly as 40 Hz, but more typically imaging around 10 Hz. Together, the two cameras measure the concentration field c and flow field \mathbf{u} throughout space and time, providing the necessary inputs for front tracking.

In all three experimental devices, we drive flow electromagnetically by passing electrical current through the fluid layer above permanent magnets. Two of the devices are built with magnets whose orientation alternates like a checkerboard and are sketched in Fig. 3.5. In small-length-scale vortex flow, cylindrical NdFeB magnets with diameter 12.7 mm are spaced 25.4 mm on center, and the vessel has lateral dimensions $254 \text{ mm} \times 254 \text{ mm}$, producing a 10×10 square array of alternating vortices. In the large-length scale vortex flow, rectangular ceramic magnets with lateral dimensions $152.4 \text{ mm} \times 101.6 \text{ mm}$ are tiled tightly beneath the same vessel, producing a 2×2 rectangular array with larger length scale. Both devices produce flows dominated by vorticity (not shear), but the characteristic flow length scales differ substantially. The third device is a water channel that produces a wake behind a bluff body, which has been previously studied [99]. The water channel experiment was built and run by laboratory undergraduate Jinge Wang with help from Rony Colon. The bluff body is an equilateral triangle with side length 45 mm, positioned symmetrically about the channel center with one vertex pointing upstream. The channel flow is dominated by shear (not vorticity) except in the wake behind the bluff body. In all three devices, the depth h is sufficiently small that the flow is quasi-two-dimensional, though the front tracking

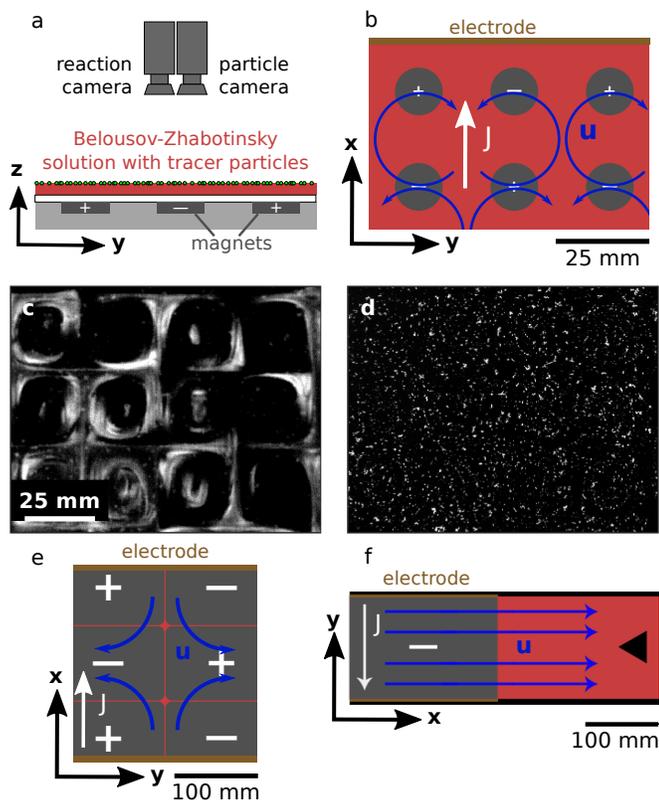


Figure 3.5: Laboratory advection-reaction-diffusion devices and typical data. (a) Two hardware-synchronized cameras image a thin layer of BZ solution containing green fluorescent tracer particles. One camera images product concentration, while the other camera images particles. (b) Flow is driven by an electric current J in the x direction, which interacts with a magnetic field in the z direction. In the small-length-scale vortex flow device, the field is produced by a checkerboard array of small magnets below the fluid that drive vortices. (c) Example snapshot of the BZ reaction, used to measure product concentration. (d) Simultaneous snapshot of tracer particles, used to measure flow. (e) The large-length-scale vortex flow device uses larger magnets to produce vortices with a larger characteristic length scale. (f) The water channel device uses magnets to drive flow past a bluff body. Work of Douglas Kelley published in [96].

algorithm could be adapted to three-dimensional systems as well.

3.6 Tracking Fronts in Experimental Data

The experimental flows produced by our three ARD devices are more complicated than the flows considered in the simulations described in §3.4, so I have no *a priori* expectations for front velocities. Nonetheless, I can use experimental measurements to show that the front tracking algorithm is self-consistent. I can also use local measurements of front velocity to explore ARD phenomena beyond the Eikonal approximation, such as variation of chemical speed with flow speed.

Figure 3.6 verifies the self-consistency of the algorithm in four cases, using data from the uniform-flow simulation and from all three experimental devices. In each snapshot, a concentration field is overlaid with the fronts located in that concentration field, the fronts located in the next concentration field, and the result of advecting the original front. Finally, the figure also shows original fronts displaced by $d\mathbf{x}_u + d\mathbf{x}_v$, accounting for both the local flow velocity and the local chemical velocity. According to Eq. 3.7 and the definition of $d\mathbf{x}_w$, those displaced fronts should exactly match the fronts in the next concentration field. The close agreement shown in Fig. 3.6 shows that the algorithm is implemented correctly. In practice, the match is not perfect because continuous, curved fronts are approximated by straight lines between discrete points, and the number of front points does not remain constant over time. Nonetheless, the close match shows that with sufficient spatial resolution the error in total velocity is small in four different flows that vary in complexity, length scale, speed, and whether their domain is open (as in the simulation and behind the bluff body) or closed (as in the vortex flows).

In addition to velocity and concentration fields, the front tracking algorithm requires user-specified parameters. It is important to know whether small variations

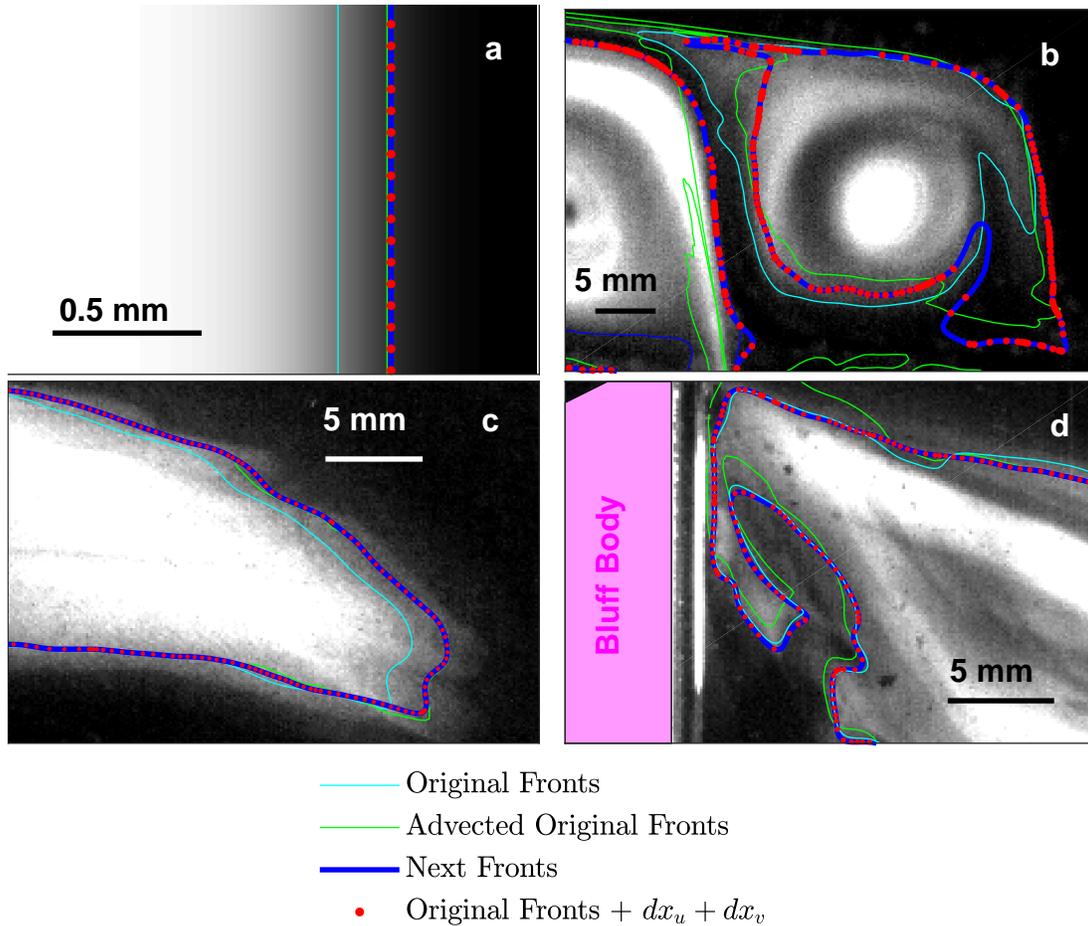


Figure 3.6: Front velocity measurements are consistent with measured front locations in simulation and in three different experimental devices. (a) Simulation of uniform flow. (b) Vortex-dominated flow with small length scale. (c) Vortex-dominated flow with large length scale. (d) Wake of a bluff body. The original front, displaced by $dx_u + dx_v$, closely matches the next front, indicating that the algorithm is self-consistent and that discretization errors are negligible at the given spatial resolution. Work of the author published in [96]. Concentration field in (d) provided by Jinge Wang.

in those parameters, especially the brightness threshold c_{thresh} and the time step dt , significantly affect the results of front tracking. I repeatedly tracked fronts in the same experiment, varying those two parameters over a wide range, and Fig. 3.7 shows the results. The change in median and root-mean-square chemical speed is much smaller than the interquartile range (the range between the 25th and 75th percentile), indicating that a wide range of c_{thresh} and dt choices provide similar results. One exception is excessively low thresholds cause measured front locations to change suddenly with minor brightness variations, artificially inflating the chemical speed. Excessively high thresholds cause some fronts to go undetected. Time steps significantly smaller than the typical time required for a front to advance by one pixel also introduce noise that artificially inflates the chemical speed. Increasing the spatial resolution allows reducing dt further. Large time steps may reduce accuracy via the mechanism described above and shown in Fig. 3.4.

Because front tracking provides local measurements of flow and chemical velocity, it can gather statistics on a single data set. For example, the chemical velocity can be measured at many locations along a front in order to determine its mean and moments with great accuracy when the only likely source of variation is random noise, as in Chapter 2. Perhaps more interestingly, front tracking can measure the variation of chemical velocity and front thickness in space and time when physical mechanisms may cause that variation. So it is interesting to also consider how changing parameters changes local measurements of these velocities. In Fig. 3.8 I plot probability density functions for chemical velocity in the x and y directions. To do this, I find front points at the same location and time in a front tracking with dt equal to that in Fig. 3.9, and dt twice as large. I find the most likely local front velocity difference is zero, indicating good match between the two locally. Errors caused by choosing $dt > dt_{\text{min}}$, which shifted the peak of the chemical velocity distribution by a few percent in the front propagating through

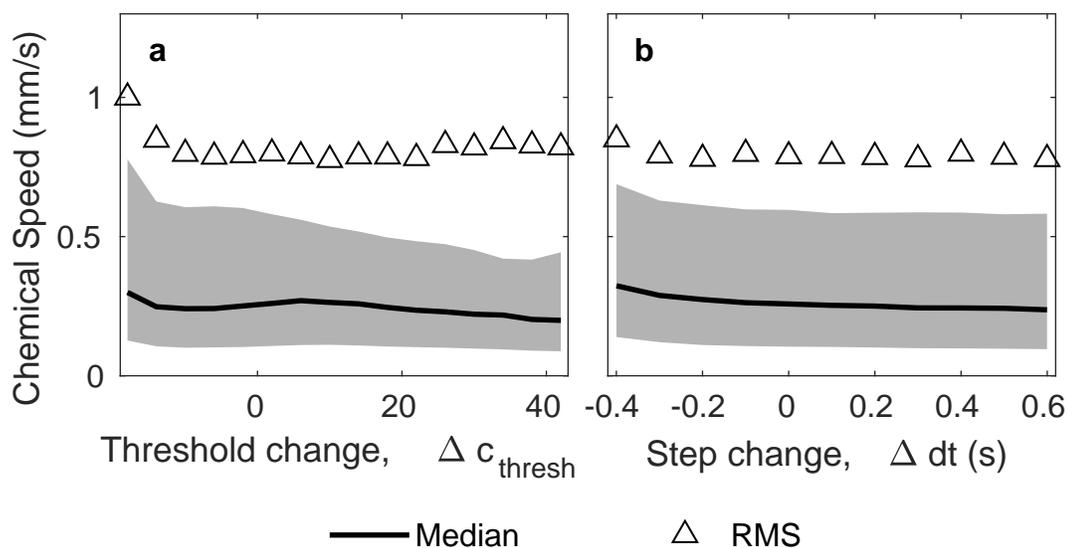


Figure 3.7: Variation of chemical speed's summary statistics. (a) Δc_{thresh} is the difference between a test threshold and the reference threshold used to produce Fig. 3.9. The gray region indicates interquartile range, and triangles indicate root-mean-square (RMS) chemical speed. (b) Variation of statistics of chemical speed with time step dt . Work of the author published in [96].

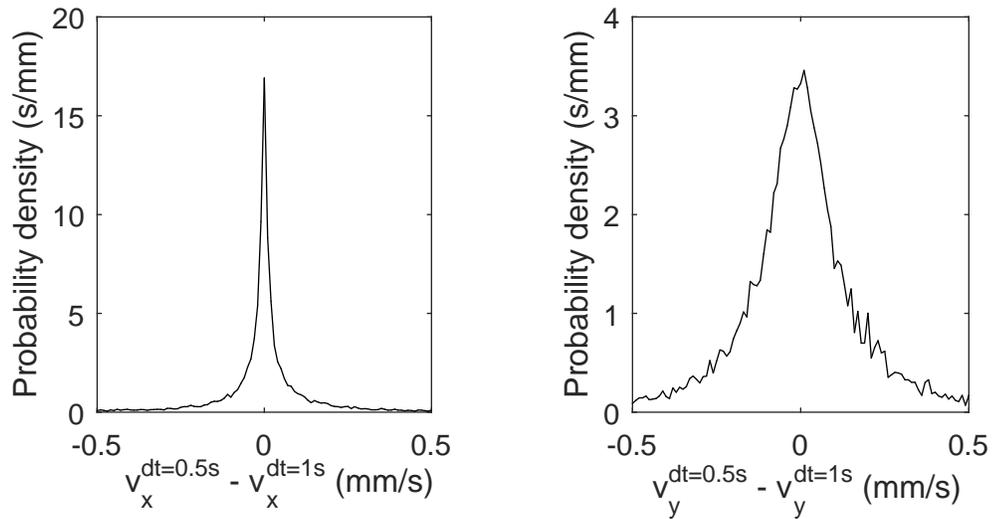


Figure 3.8: Distributions of the change in chemical velocity components caused by changing the time step dt from 0.5 s to 1 s. Both distributions are symmetric and peak at zero, indicating a match. Work of the author published in [96].

a hyperbolic flow example (§3.4) are negligible in this experimental example. Differences in v_y are somewhat larger than those in v_x because in this experiment, the front advances primarily in the y direction, resulting in higher available speeds in the y direction.

Having validated the front tracking algorithm in simulation and experiment, and having shown that the algorithm is robust to variation of input parameters, I measured distributions of chemical speed in both the small-scale and large-scale vortex flows. Figure 3.9 shows a typical concentration field from each set of experiments. I measured the chemical speed at between 8×10^4 to 2×10^6 locations over durations ranging from 36 s to 800 s. In the small-scale device, I varied the flow speed (characterized by the root-mean-square, RMS, speed u_{rms}) over two orders of magnitude, by changing the magnitude of the constant current used for each experiment. I found that the distribution of *chemical* speed changes with *flow* speed. The probability distributions vary nearly monotonically, growing a longer tail as u_{rms} increases, which signifies a greater prevalence of high front

speeds. In the large-scale device, because the magnets are weaker, I am unable to vary u_{rms} as widely. Still, over a range of flow speeds, I observe the same trend: regions where fronts have high chemical speed become more common as the flow becomes faster. These results contradict the basic Eikonal assumption, in which chemical speed is constant [43], because the constant front speed does not capture the full dynamics of these experiments. This is entirely unexpected, especially considering the success the Eikonal approximation has had in explaining at least long term behavior of a variety of reactive mixing systems [33, 48].

Seeing the same trend in experiments using two different devices suggested that a comparison might provide insight. I reduced each distribution to a single measurement, the RMS chemical speed v_{rms} of all fronts at all locations and times in the experiment. Surprisingly, plotting v_{rms} versus u_{rms} on the same axis for experiments in both the small-scale and large-scale device, as in Fig. 3.9e, shows that the two curves collapse to show a single trend. Neither scaling nor fit parameters are required for this collapse. It is also noteworthy that the RMS chemical speeds are far above the expected chemical speed, which for BZ is measured at $72 \mu\text{m/s}$ (See Chapter 2).

Many mechanisms might explain the observed deviation from the Eikonal approximation. The Eikonal approximation requires that the second Damköhler number, the Thiele number, and the first Damköhler number are all large. In these experiments, the second Damköhler number is $Da_{\text{II}} = 1300$, so I do not expect Da_{II} to explain the variation of chemical speed with flow speed. The Thiele number $\Phi^2 = Da_{\text{I}} Pe$ relates advection, diffusion, and reaction times, where $Pe = UL/D$ and $Da_{\text{I}} = kL/U$ (L is flow length scale). In these experiments, $\Phi^2 > 7.7 \times 10^5$, so I do not expect Φ^2 to be the reason, either. The first Damköhler number relates reaction rate to advection rate. These experiments span the range $2 \leq Da_{\text{I}} \leq 183$, in which $Da_{\text{I}} \gg 1$ is not strictly true, and I expect the small values of Da_{I} account for some deviation from the Eikonal approximation. In

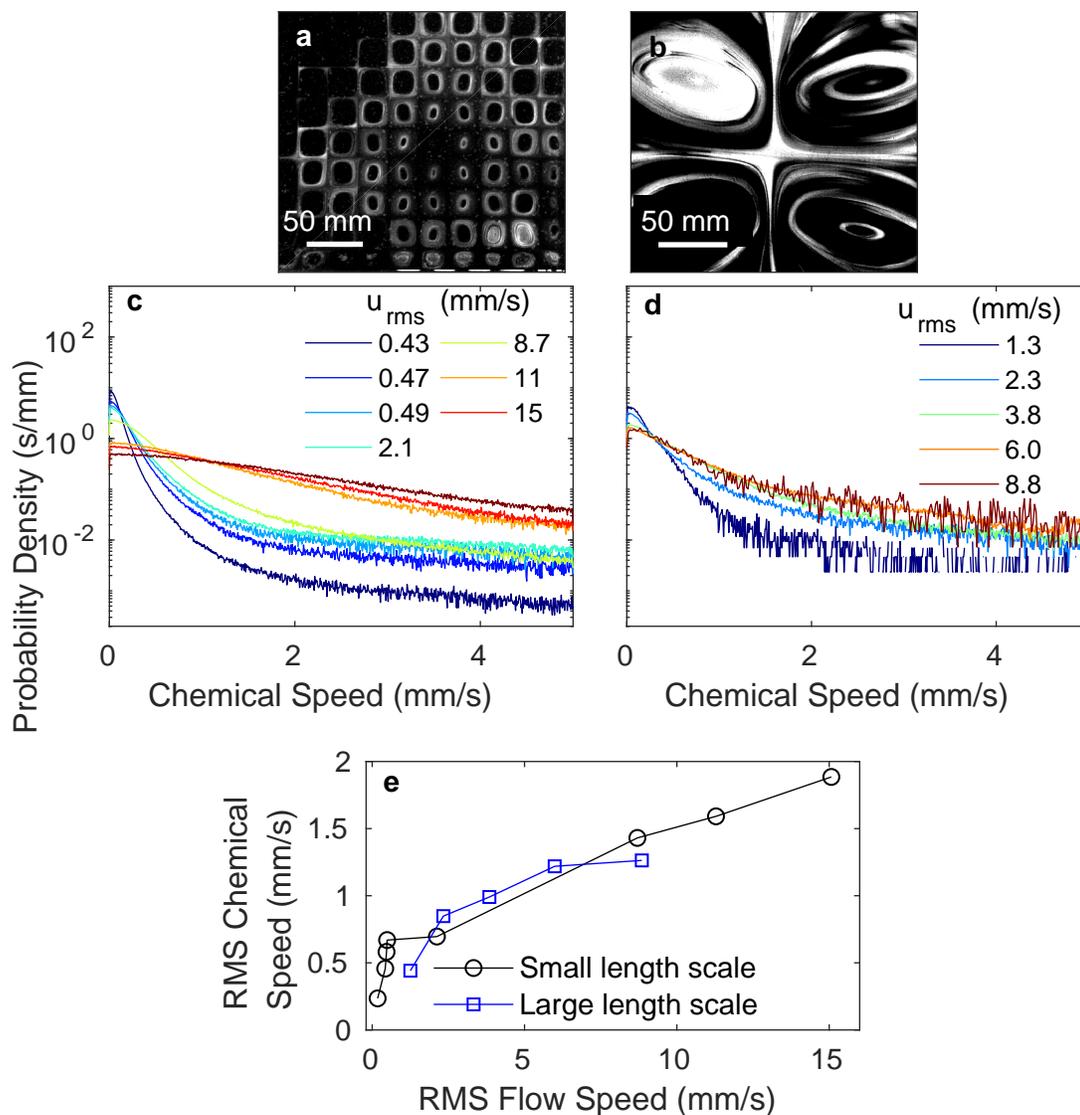


Figure 3.9: Concentration fields in vortex-dominated experimental flows with small (a) and large (b) characteristic forcing length scales. (c) Distributions of chemical velocity in small characteristic forcing scale vortex flows, at various RMS flow speeds. (d) Distributions of chemical velocity in large characteristic forcing scale vortex flows, at various RMS flow speeds. (e) RMS reaction speed increases with RMS flow speed, and measurements from experiments with two different length scales nearly collapse. Work of the author published in [96].

fact, I believe testing the $Da_I \sim 1$ regime to be particularly interesting, because neither advection nor reaction can be neglected, and because it is likely relevant for applications like modeling marine phytoplankton, whose growth rate nearly matches the tidal timescale. Plus this is the regime where improved modeling of fronts with a changing front speed could yield the most benefit. That said, many prior studies have considered the $Da_I \gg 1$ case, and they expect that chemical speed is nearly the stagnant speed in that case, which I do not find here. Scaling arguments cannot provide a satisfying explanation of why chemical front speed depends on flow speed.

One might also suspect front curvature of causing these front speeds. Front curvature has been observed to affect front speed [42]. However, the curvature correction term to the Eikonal approximation is proportional not only to curvature but to diffusivity. Considering the effective diffusivity of BZ products and the typical speed of BZ fronts, the curvature effect doubles front speed only when curvature reaches 72mm^{-1} , or approximately, when fronts have a radius of curvature less than $15\ \mu\text{m}$. Curvature effects cannot explain the observations partly because the apparatus described in §3.5 cannot resolve such small scales. More importantly, at the scales observed, the vast majority of chemical velocity measurements come from locations where fronts are nearly zero curvature. High curvature regions are necessarily small, and short lived, so we would not expect them to be powerful enough to sway the statistics of chemical speed.

Another possible excuse for the Eikonal approximation is that this front tracking measured fronts speeds of both fronts and “backs”. In Chapter 2 I discussed that the front and back of a BZ reacted region have different properties because the chemical kinetics during the transition from unreacted to reacted state (at reaction fronts) differ from those during the return to the unreacted state (at reaction backs). In stagnant fluid, both fronts and backs advance with the same chemical velocity, as required by the fixed refractory time that separates them.

But fronts and backs have different thicknesses, which may in turn be affected differently by an underlying flow. However, examination of a single front with no back in Fig. 3.7 still showed a wide chemical speed distribution, and summary statistics much faster than the speed of BZ fronts in stagnant fluid. Tracking fronts and backs separately in the future is likely to be interesting, but it is not the source of our chemical speed variation.

While trying to uncover the cause of the chemical speed variation, I also speculated about the possibility of unmeasured flow driving reaction front propagation. Specifically Ekman pumping arises from the presence of our no-slip boundary condition at the bottom of our vessel. Ekman pumping is a secondary flow that results from the fact that rotation generates a radial pressure gradient. That pressure gradient persists in the boundary layer, but flow there is too slow to balance the pressure gradient. The result is upwelling at vortex centers and downwelling at vortex edges, both of which cause mixing [100] and could affect measured front speed. In these experiments, I sometimes observe high front speed where fronts propagate into vortices, as would be caused by Ekman pumping. However, though I would expect the effects of Ekman pumping to differ in vortices having different sizes (and therefore different angular momentum), I have not observed such differences. Essentially, if Ekman pumping were causing this effect, the two different length scale data sets should not have collapsed onto a single curve.

The most interesting possible cause for the chemical front speed variation is new physics. The stretching and collision of fronts might cause interactions which increase front speeds. While a recent study found behavior consistent with the chemical speed of BZ fronts being unaffected by flow [34], they rarely measured flow speeds exceeding 1 mm/s, whereas the experiments described above often involved faster flows which still maintain $Da_I \gg 1$. Higher flow speeds increase the effects of local Lagrangian stretching, which may either enhance or inhibit reaction the nonlinear reaction term, depending on the rate of stretching, causing

reacted regions to grow or shrink. This and other interactions between stretching and nonlinear reaction will be discussed in Chapter 5. The Eikonal approximation of constant chemical speed assumes a lack of strain across the front, which is not true in the presence of strong stretching [43], so stretching may cause front speed to vary.

However, there is another possible physical explanation, which I claim to be the cause of these measurements. I will cover this explanation, and all of its consequences in Chapter 4. For now, note that these experiments are meant to approximate two-dimensional flow, but the actual flow varies vertically and drops to zero at the vessel floor, which is a no-slip boundary. Reaction fronts are therefore sheared vertically, which causes three-dimensional concentration variations and corresponding brightness variations that differ from those expected for purely two-dimensional fronts. An immiscible lubrication layer between the reacting fluid and the vessel floor to greatly reduce shear, and therefore three dimensional concentration gradients. The results of Chapter 4 will show that this is the likely cause of the front speed variation, which indicates that the front tracking algorithm is accurately tracking bright regions in experiments, and furthermore that analysis of this type opens up possibilities of observing new physics, as it has already done so for my research.

3.7 Summary and Future Work

This chapter presented a new algorithm for measuring the chemical velocity and total velocity of reaction fronts in ARD systems, expanding the prior algorithm measuring front velocity and thickness in stagnant reaction-diffusion systems. This algorithm requires knowledge of the underlying fluid flow. I validated the algorithm by tracking fronts in ARD simulations with uniform flow and simple enough initial conditions that we had clear *a priori* expectations for the front ve-

locity. I further validated the algorithm by tracking simulated fronts approaching a burning invariant manifold (BIM) in a hyperbolic flow region. I also tracked experimental fronts, using three different laboratory devices to show that the algorithm is self-consistent and that errors due to spatial discretization are negligible. I found experimental front tracking to be insensitive to user parameter choices. Finally, I showed one example of the sort of investigation enabled by front tracking, which makes a large number of local measurements of front velocity, throughout space and time. In a series of experiments in small-scale vortex flow and large-scale vortex flow, I observed that the distribution of chemical speeds varied with flow speed: faster flow caused more measurements of high chemical speed. Comparing experiments from the two different devices, I found that plots of RMS chemical speed varying with RMS flow speed roughly collapsed onto the same curve.

Front tracking is applicable to either experimental measurements or simulation results, in steady or unsteady flows. The algorithm is independent of the size and speed of the ARD system, since the calculations are done in camera units before being converted to physical units. This algorithm can accommodate situations where flow cannot be removed, such as ocean flows or combustion. This method also makes local measurements, providing far more information than prior methods, which measured the asymptotic, global front speed in steady flows [10, 37, 98].

The initial experiments yielded an immediately unexpected result, so it is likely that using front tracking to ask questions about reactive ARD systems will be a productive area of future research. Indeed, the algorithm allows testing new hypotheses about ARD dynamics. First, local measurement of front speed could show whether fronts in experiments slow as they approach BIMs or bLCS, as they do in the simulation I used for validating the algorithm. Second, other barriers to front propagation may also be possible, and front tracking could locate them. For example, experiments detailed in Chapter 5 will discuss how excitable reactions are promoted most in regions where the Lagrangian stretching falls in an optimal

range. Outside this range, fronts propagate slowly, or are driven to extinction by a process analogous to blowing out a flame. Third, front tracking opens up opportunities to test the dynamic causes of frozen front phase transitions observed in open flows [10, 11, 33]. Finally, other authors have recently speculated that front speed may depend on front thickness [33], a hypothesis that can be tested directly with our front tracking algorithm.

The algorithm can also have important practical applications. I have considered only the BZ reaction, but future measurements of chemical speed and its variation with flow speed in other chemical systems could reveal interesting and useful chemical behavior. Industrial reactors are intended to maximize overall reaction rate while minimizing the energy required by the mixers that drive flow. By locating regions of maximum front speed, front tracking could lead to improved reactor design. Ocean phytoplankton blooms occur at length scales inaccessible in the laboratory and depend on many complex factors, including ocean currents, temperature, nutrient availability, and plankton species. Determining phytoplankton front speed and its dependence on those factors might allow improved forecasting and even control of phytoplankton blooms. Also, variation of front speed might account for the genetic diversity of phytoplankton [27]. Applying front tracking to combustion might reveal how different flames respond to wind, whether a flame reaches all the fuel, what chemical mixes make the fastest flames, and which parts of flames grow fastest.

The algorithm leaves opportunity for future improvement as well. For example, we have implemented it only for two-dimensional domains, but the algorithm could generalize directly to three dimensions. Higher order time-stepping may also be possible. It would be especially helpful to account for the continuous change in the orientation of front elements. These further generalizations of front tracking could increase its usefulness and allow front tracking to provide insight into the growth dynamics of a multitude of systems.

While there are several strong pieces of evidence that front tracking in flow is replicating the dynamics of the ARD system, further studies were needed to ensure that front tracking was accurately measuring chemical front speed. Chapter 4 will detail our explanation of the depth shear phenomenon in experiments. This following chapter involves an analytic prediction derived from the Eikonal approximation, which then accurately predicts front tracking results in new experiments, and indicates that the anomalous front speed measured in this chapter is physical. While Chapter 4 is required to confirm the validity of the results in this chapter, it is worth noting that all the experiments of the following chapter flow from the analysis enabled by front tracking: measuring chemical front speeds and observing that they are not the value expected. The following chapter also demonstrates the importance of shearing to reaction growth, and the usage of a lubrication layer to improve two-dimensionality of an ARD experiment. I believe the combination of a front tracking algorithm and an understanding of depth shear will enable numerous future studies by making easy to operate experimental analogues for two-dimensional reactive systems, which can provide a wealth of data.

3.8 Acknowledgments

Reproduced from Nevins, T.D. and Kelley, D.H., “Front tracking velocimetry in advection-reaction-diffusion systems”, *Chaos* **28**, 043122 (2018), with the permission of AIP Publishing, which has been modified to fit this thesis. The thesis author is the primary author and was responsible for writing, troubleshooting, and improving the algorithm described, the simulation section, the experiment section, writing the paper (with the exception of a few paragraphs), and preparing the figures (except Fig. 3.5). Co-author Douglas Kelley conceived and wrote the initial front tracking algorithm, provided guidance during the development,

provided guidance on figure formation, wrote paragraphs 2 and 3 of the algorithm section, created Fig. 3.5, and edited the paper and figures. Data in Fig. 3.6d was used with permission of Jinge Wang, but the tracking results were assembled by the thesis author. Insightful conversations with N. T. Ouellette and J. G. Puckett inspired this work. Experimental design help from R. S. Russell and J. Tithof was used to gather the data. The thesis author was supported by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program throughout the work for this chapter.

4 Vertical shear alteration of chemical front speed in thin-layer flows

This chapter is based on Reference [101], which is currently under review for publication. It has been reformatted and edited to fit within the thesis. The thesis author was the primary author on this work.

4.1 Introduction

To study the growth of chemical reactions in fluid flows, I developed the algorithms of chapter 2 and 3. The front tracking algorithm tracks edges separating reacted and unreacted – fronts – and calculates local quantities to describe their behavior. Local perspective on the reaction front growth of a chemical reaction in a fluid flow can uncover features that prior studies on global average reaction front growth cannot [102–105]. Local understanding is the first step toward prediction and control, with goals like maximizing or minimizing reaction spreading. However, in Chapter 3 the output of the algorithm was puzzling, and demanded explanation to prove the measurements were real.

In order to explain the measurements of chapter 3 I took a closer look at the

growth of a reactive scalar advected in quasi-two-dimensional (quasi-2D) flow, common in reactive mixing experiments. Though three-dimensional (3D) reactive mixing experiments have recently been undertaken [47], quasi-2D systems simplify experiment, instrumentation, and analysis while capturing rich dynamics, including mixing barriers [48], sensitivity to boundary conditions [106], and the optimal flow properties I will discuss in chapter 5. Two-dimensional data is sometimes gathered by averaging over one spatial dimension [107], and more often gathered by driving reactive mixing in a thin layer [5, 33, 48, 99, 106]. However, even thin-layer fluid experiments have at least some three-dimensionality; their flows and reaction fields depend on depth, and that dependence goes unmeasured. As I will show, unmeasured 3D processes, especially shear, can cause surprising results. I will also suggest methods for minimizing the effects of shear and three-dimensionality.

To see why the results of chapter 3 were surprising, first recall that the growth of a scalar concentration field (such as a chemical reaction product) $c(\mathbf{x}, t)$ depends on advection (flow), reaction, and diffusion. (Here $\mathbf{x} = x\hat{\mathbf{x}} + y\hat{\mathbf{y}} + z\hat{\mathbf{z}}$ is position and t is time.) c is normalized to range from 0 (unreacted) to 1 (reacted), and in experiments, is approximated by measuring the local color of the reacting solution [92]. Consider the case where the local reaction rate depends only on existing concentration, and diffusion is independent of concentration, then the field c obeys the advection-reaction-diffusion (ARD) equation

$$\frac{\partial c}{\partial t} + (\mathbf{u} \cdot \nabla)c = D\nabla^2 c + \alpha G(c), \quad (4.1)$$

where t is time, \mathbf{u} is the velocity of the fluid in which c is growing, D is the diffusion coefficient, α is a reaction rate coefficient, and $G(c)$ is a dimensionless reaction term. Chemical reactions that involve multiple species are modeled by systems of ARD equations, one for each chemical species c_i with its own reaction term $F_i(c_1, c_2, \dots, c_n)$. I will focus on the simple case of a single reaction product. To complete the system of equations, the flow \mathbf{u} obeys the Navier-Stokes equation,

which I presume is independent of c , so that it need not be solved simultaneously with equation 4.1.

Though equation 4.1 rigorously governs the behavior of a reactive scalar field, simplifying the equation can lead to new intuition about reaction dynamics. In particular, reaction fronts, have proved conceptually useful. For many reactions $G(c = 0) = G(c = 1) = 0$, and $G(c) > 0$ for $0 < c < 1$, such that regions of moderate concentration quickly react and become nearly saturated ($c \approx 1$). Typically, $c \approx 0$ or $c \approx 1$ over much of the domain. In that case, equation 4.1 shows that $\partial c / \partial t \approx 0$ except in the small regions where ∇c is large, since the advection and diffusion terms in equation 4.1 both involve spatial variation. Thus the fronts that separate $c \approx 0$ regions from $c \approx 1$ regions locate the essential dynamics of the system, and front propagation is a natural descriptor of system evolution. The same reasoning holds for excitable chemical systems in which $G(c) > 0$ only in the smaller range $c_0 < c < 1$ (where c_0 is a dimensionless excitation threshold).

To understand the velocity of any differential element of a reaction front, first consider what the motion must be in stagnant and uniform flows. The total velocity, \mathbf{w} , must move outward at the chemical speed v in stagnant. Then in uniform flow the exact solution is obtained by simply adding the flow velocity. The Eikonal equation for ARD systems is then created by assuming this simple summation works for non-uniform flows as well: [43]

$$\mathbf{w} = \mathbf{u} + v\hat{\mathbf{n}}. \quad (4.2)$$

Here $\hat{\mathbf{n}}$ is the local unit normal to the front. By definition, $\hat{\mathbf{n}}$ points in the direction of decreasing c (opposite ∇c). This equation is much simpler than Eq. 4.1 to understand and implement numerically, and therefore it is a powerful tool for predicting ARD systems. It also provides a point of comparison for further studies on front dynamics, where any deviation from the Eikonal equation indicates an interaction between reaction and advection.

The motion of reaction fronts, and the Eikonal approximation for describing that motion, were used to great effect in the development of the theory of burning invariant manifolds (BIMs) [45, 50]. BIMs are dynamical barriers to front propagation, located where $\mathbf{u} \cdot \hat{\mathbf{n}} = -v$ with $\hat{\mathbf{n}}$ as the unit normal of the BIM. The $\hat{\mathbf{n}}$ direction of the BIM causes these to be one-sided barriers. There, the flow velocity is fast enough in the direction opposite to the front propagation, such that \mathbf{w} is strictly tangent to the BIM according to equation 4.2. BIMs attract fronts and predict where fronts can propagate and where they cannot. Seeing the power of fronts for capturing the dynamics of reactive mixing, Dr. Kelley and I developed an automated algorithm for measuring the instantaneous, local chemical velocity $v\hat{\mathbf{n}}$ and the instantaneous, local total front speed \mathbf{w} which were presented in chapters 2 and 3. Fig. 4.1 demonstrates the front tracking technique. The algorithm is written for 2D simulations or quasi-2D experiments and works by identifying reaction fronts, advecting them forward according to the known flow \mathbf{u} , and measuring the perpendicular displacement between the advected front and the front observed at a later time. In this way it uses knowledge of \mathbf{u} and the nature of the Eikonal equation to measure $v\hat{\mathbf{n}}$ and \mathbf{w} .

Like BIM theory, front tracking is built on equation 4.2, but without the simplifying assumption that the chemical speed is a constant $v = v_0$. Instead I will measure a v_a speed, which can be positive (growth) or negative (shrinkage). However, the constant front speed assumption has been broadly applicable, despite neglecting higher-order effects like curvature, Ekman pumping, and shear. Curvature is known to affect chemical speed according to $v = v_0 + D\kappa$, where κ is the curvature [41]. Curvature effects have been observed in experiments having length scales on the order of microns [42] but are generally negligible in aqueous experiments at larger length scales because $D \sim 10^{-3}\text{mm}^2/\text{s}$. Inclusion of curvature can improve predictions in the pinning of reaction fronts [33]. BIM theory (using the $v = v_0$ assumption) agrees closely with experimental measurements in a variety of

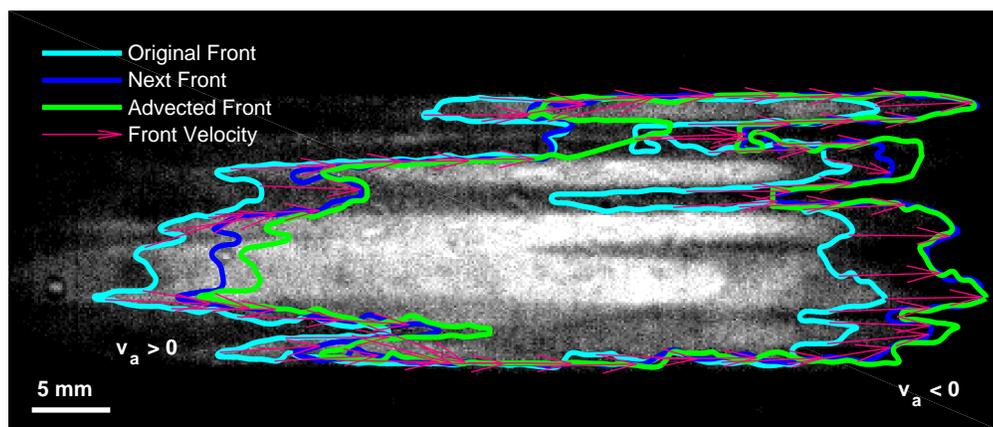


Figure 4.1: Experimental measurements of chemical reaction fronts using front tracking. A thin layer of reacting fluid is imaged from above, appearing bright where product concentration is high and dark where product concentration is low. Fronts are identified as borders on regions that are brighter than a user-defined threshold. To measure the local front speed, a front (Original Front) is advected using the known flow velocity field and compared to a front at a later time (Next Front). The perpendicular distance that separates them is attributed to the measured chemical velocity $v_a \hat{n}$, and the total velocity \mathbf{w} is calculated using equation 4.2. Work of the author in [101].

flows with a variety of parameters [47–49]. Other studies have predicted [102–105] or observed [38, 106] variation in the total velocity \mathbf{w} or the space-averaged chemical speed, while using that the local chemical speed v is essentially constant. This implies that stirring could accelerate the growth of chemical reactions without enhancing the true chemical speed.

However, in chapter 3, I observed in fig. 3.9 chemical speeds far in excess of v_0 , apparently contradicting the $v = v_0$ assumption. The effect was much larger than could be explained by curvature [42] or Ekman pumping [100]. I performed the experiments in a thin, horizontal layer of reacting solution and measured flow velocity \mathbf{u} , but I was able to measure only horizontal components of flow, and only at the top of the layer. Therefore, vertical flow was a potential source of complications, though I was careful not to drive flows fast enough to cause the onset of substantial vertical motions [76, 77]. Moreover, the apparently high chemical speeds persisted even in directions misaligned with electric field [108, 109]. These puzzling observations could be due to either a physical mechanism not included in equation 4.2, or a complication associated with doing experiments in reactive mixing.

In this chapter I focus on the second possibility: an experimental complication. I will show that vertical shear can change both the magnitude and sign of *apparent* the chemical speed v_a in front tracking experiments, depending on the orientation of the front relative to the flow. I will compare experimental measurements of v_a to simulations that assume the true chemical speed is v_0 but account for front deformation by vertical shear. Close agreement between simulation and experiment will show that vertical shear explains nearly all of the apparent deviation from v_0 . Vertical shear is caused by the no-slip boundary condition at the base of the thin layer and smears reaction fronts in a way that makes measuring the local concentration c more complicated. In most experiments, concentration is accessed optically, by measuring the depth-averaged concentration, but shear changes the

vertical concentration profile. Many prior and future quasi-2D reactive mixing experiments are subject to similar shear effects, so shear must be considered when interpreting results. The effect persists even if $\mathbf{u} \cdot \hat{\mathbf{z}} = 0$ everywhere (where $\hat{\mathbf{z}}$ points upward). Shear originating from a no-slip boundary condition has been shown to have a major effect before, especially in the case of Poiseuille flow [35, 37]. I will further show that vertical shear can be substantially reduced, allowing reactive mixing experiments that are more nearly 2D, if the reacting layer is separated from the no-slip boundary below it by a lubrication layer.

In section 4.2 I describe the results of simulating front propagation throughout the depth of a single thin layer, and extract the chemical speeds that would be measured in an experiment using optical methods. In section 4.4 and section 4.3 I repeat the discussion for different systems. First a two-layer system in which the thin, reacting layer is bounded below by an immiscible, dielectric, lubrication layer to reduce shear. Second the Hele-Shaw style of experiment in which the reacting layer is sandwiched between two solid plates and flow is driven by a pressure difference. In section 4.5 I detail the experimental apparatus and methods I developed to test the predictions of sections 4.2 —4.3. In section 4.6 I present the results of one- and two-layer experiments, showing good agreement with the simulations, including a significant reduction in apparent deviation from $v = v_0$ in the two-layer system. The chapter closes with a summary and outlook in section 4.7.

4.2 Simulations of Single-Layer System

I begin my investigation by simulating reactive mixing in a single, thin layer of fluid, with flow driven horizontally, a common configuration for quasi-2D experiments [5, 33, 48, 78, 96, 99, 106]. The layer is subject to a no-slip boundary condition at the solid floor that supports it, and may be subject either to a

no-slip condition or a no-penetration condition at its top surface. I will consider the no-penetration condition first before briefly visiting the other case when I consider Hele-Shaw experiments. In the laboratory, it is convenient to study chemical reactions whose products have a different color than the reactants, so that concentration c can be measured optically. Common candidates include the Belousov-Zhabotinsky reaction [2, 92–94], the iodate arsenous acid reaction [10, 33], and acid-base reactions [79]. In thin-layer experiments, the local color depends on the average concentration across the reacting layer, and reaction fronts can be defined as the locations where the brightness crosses some user-defined threshold. For simple forms of the reaction term $G(c)$, eq. 4.1 can be solved analytically in the one-dimensional case and predicts concentration profiles that maintain their shape while propagating, so that front speed measurements are insensitive to the choice of threshold. If the concentration is uniform across the layer depth, fronts located this way are unambiguous, and three-dimensionality need not be considered. However, if the concentration varies with depth, more care is required. The sensitivity of concentration measurements to variation with depth contrasts with measurements of other quantities. For example, tracer particles used for measuring the flow \mathbf{u} can be selected with a density that causes them to float atop the reacting layer. Their motion samples only the top of the layer, without depth-averaging, so 3D effects are negligible until upwellings and downwellings become substantial [76, 77]. When locating reaction fronts, however, three-dimensionality must always be considered carefully. I begin with a simple simulation.

4.2.1 Velocity Profile

In this section I simulate the implications of vertical shear for reactive mixing experiments. To simplify the problem, I consider a steady, planar flow that is uniform at the free surface of the layer: $\mathbf{u}(x, y, z = h) = U\hat{\mathbf{x}}$. Here h is the layer thickness, and I choose $z = 0$ at the bottom of the layer. The flow must

vary with depth, but we will assume that the forcing is such that we can separate the planar motion from its z dependence [110–112]. I expect this construction to provide a reasonable approximation for the non-uniform flows as well, because a similar approach by Suri *et al.* [110] closely matched experiments with non-uniform flows. Considering uniform flow also gives clear expectations for front velocity, since uniform flow differs from the well-studied $\mathbf{u} = 0$ case only by a Galilean transformation.

Such a flow can be generated by applying a vertical magnetic field $\mathbf{B} = B(z)\hat{\mathbf{z}}$ and passing a uniform, horizontal electrical current with density $\mathbf{J} = J\hat{\mathbf{y}}$ through the reacting layer. Including the resulting Lorentz force, the Navier-Stokes equation that governs the flow is

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla)\mathbf{u} = -\frac{1}{\rho}\nabla P + \frac{\mu}{\rho}\nabla^2\mathbf{u} + \frac{JB}{\rho}\hat{\mathbf{x}} - g\hat{\mathbf{z}}, \quad (4.3)$$

where ρ is the density of the fluid, P is the pressure, μ is the dynamic viscosity, and g is the gravitational acceleration. I consider a magnetic field that varies vertically, as it does in experiments [111] with magnets arranged below the reacting layer:

$$B(z) = B_0 e^{-\lambda z}. \quad (4.4)$$

Here B_0 and λ are empirical constants measured for my apparatus and listed in table 4.1. To satisfy $\nabla \cdot \mathbf{B} = 0$, the magnetic field must also have a horizontal component that varies with z , but the vertical force it produces is negligible compared to gravity, so I will not discuss it further. The flow occurs in a layer of infinite extent in x and y . At $z = 0$, I impose a no-slip boundary condition $\mathbf{u} = 0$. At $z = h$, I impose a no-penetration boundary condition $\mathbf{u} \cdot \hat{\mathbf{z}} = 0$ and require that the shear be zero: $\partial u_x / \partial z = 0$. Solving eq. 4.3, I found

$$\mathbf{u} = u_x \hat{\mathbf{x}} = \frac{JB_0}{\mu\lambda} \left(\frac{1}{\lambda} - \frac{e^{-\lambda z}}{\lambda} - e^{-h\lambda} z \right) \hat{\mathbf{x}}. \quad (4.5)$$

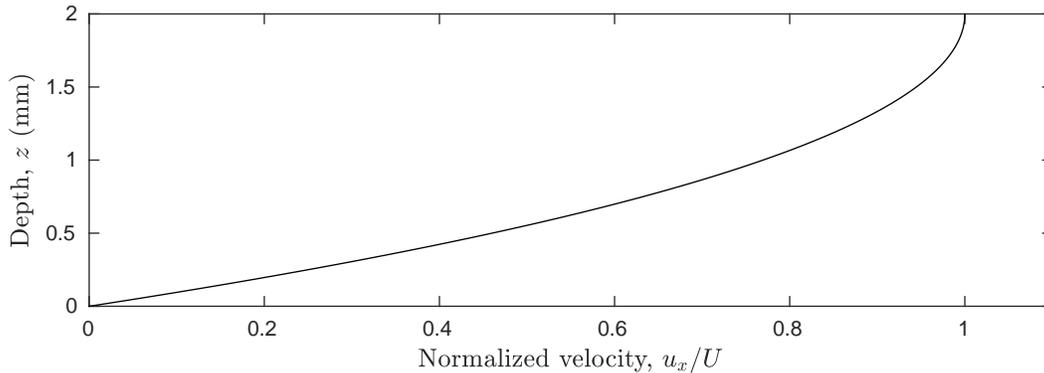


Figure 4.2: Depth dependence of streamwise velocity in a uniform thin-layer flow. Fluid properties and layer thicknesses are given in table 4.1, with J set to normalize the maximum speed. Work of the author in [101].

The velocity profile is shown in fig. 4.2. At $z = h$, where flow is typically measured in experiments, the velocity is

$$\mathbf{u}(x, y, h) = U \hat{\mathbf{x}} = \frac{JB_0}{\mu\lambda} \left(\frac{1}{\lambda} - \frac{e^{-\lambda h}}{\lambda} - e^{-h\lambda} h \right) \hat{\mathbf{x}}.$$

Fig. 4.2 and eq. 4.5 make it clear that flow in these experiments will vary with depth. Flow variation with depth ensures that concentration will also vary with depth and take on aspects of three-dimensionality.

4.2.2 Simulation

Knowing the velocity profile $u_x(z)$, I can simulate the propagation of reaction fronts having constant chemical speed $v = v_0$, then determine if the results are consistent with the Eikonal approximation (eq. 4.2), where \mathbf{u} is independent of depth. The Eikonal approximation leads to a set of ordinary differential equations governing the position and angle θ of a reaction front in the vertical $x - z$

plane Mitchell and Mahoney [44]:

$$\frac{\partial x}{\partial t} = u_x + v_0 \sin \theta, \quad (4.6)$$

$$\frac{\partial z}{\partial t} = -v_0 \cos \theta, \quad (4.7)$$

$$\frac{\partial \theta}{\partial t} = -\frac{\partial u_x}{\partial z} \sin^2 \theta. \quad (4.8)$$

Consistent with Mitchell and Mahoney [44], I define the angle, θ , as the angle a front element makes with the x axis, oriented such that the front advances in the $\hat{\mathbf{n}}$ direction when $\theta = \pi/2$. Accordingly, a front element with $\theta = \pi$ is horizontal and propagates upward with $\hat{\mathbf{n}} = \hat{\mathbf{z}}$.

There are a number of useful features of these equations. First, they mandate that front elements move with a speed that is the sum of the local flow speed and v_0 , a constant I choose. Second, the vertical shear $\partial u_x / \partial z$ appears explicitly and has the effect of changing the front angle θ . Eq. 4.7 shows that when $\theta \neq \pi/2$ and $\theta \neq -\pi/2$, the front element has a vertical component to its propagation. Even if $\theta = \pi/2$ initially for all front elements, because $\partial u_x / \partial z \neq 0$, the positions and angles of front elements evolve differently over time than if they were acted on by a uniform flow of magnitude equal to the flow speed at their starting height. Third, front curvature and Ekman pumping are absent and cannot be confused with the effect of shear on any apparent changes in v . Finally, it should also be noted that equations 4.6-4.8 are in an easy to numerically solve form, making them a preferable tool if they accurately model experiments.

I used eq. 4.6 - 4.8 to simulate reaction fronts with constant chemical speed v_0 in two flows: supporting flow $\mathbf{u} = u_x \hat{\mathbf{x}}$ as given by eq. 4.5, and opposing flow $\mathbf{u} = -u_x \hat{\mathbf{x}}$. Parameter values were chosen to match laboratory experiments and are listed in table 4.1. I initiated fronts at time $t = 0$ with $x = 0$ and $\theta = \pi/2$ for 200 front elements spaced evenly over $0 \leq z \leq h$. The front was advanced through time using a fourth-order Runge-Kutta method with a time step corresponding to 0.001 s. After each Runge-Kutta step, I interpolated along the front to relocate its

Parameter	Value	Description
h	2 mm	Single layer depth
h_e	3 mm	Reacting layer thickness in two-layer system
h_d	5 mm	Lubrication layer thickness in two-layer system
v_0	72 $\mu\text{m/s}$	Reaction front speed
B_0	0.02 T	Magnetic field surface strength
λ	230 1/m	Decay coefficient of magnetic field
ρ	1000 kg/m^3	Density of reacting layer
μ	1.60×10^{-3} Pa s	Dynamic viscosity of reacting layer
μ_d	1.42×10^{-3} Pa s	Dynamic viscosity of lubrication layer
γ	0.5	Concentration threshold locating apparent fronts

Table 4.1: Parameters measured from experiments and used for simulations.

elements at the original depths, preventing loss of resolution through the bottom and top of the domain, then recalculated θ from x and z for self-consistency.

The boundary conditions for θ are subtle. The concentration outside the domain is always $c = 0$, and the front propagation direction $\hat{\mathbf{n}}$ points in the direction of decreasing c , by definition. Since $c < 0$ is non-physical, a front element cannot emerge from outside the domain, but can vanish into the edge of the domain. Accordingly, I impose the boundary condition $-\pi/2 \leq \theta \leq \pi/2$ at $z = 0$. In practice, the condition must be imposed only for opposing flow; supporting flow rotates front elements in the allowable direction. At $z = h$, the initial angle $\theta = \pi/2$ remains unchanged because $\partial u_x / \partial z = 0$ there, so fronts never emerge from outside the domain. At $z = h$, and in supporting flow at $z = 0$, θ is determined using a front element interpolated outside the domain.

Fig. 4.3 shows simulated front evolution over time. Eikonal fronts in a thin layer do not remain straight and vertical, but have positions that vary with depth, even in an entirely horizontal flow (eq. 4.5), because of shear. Also, I observe a

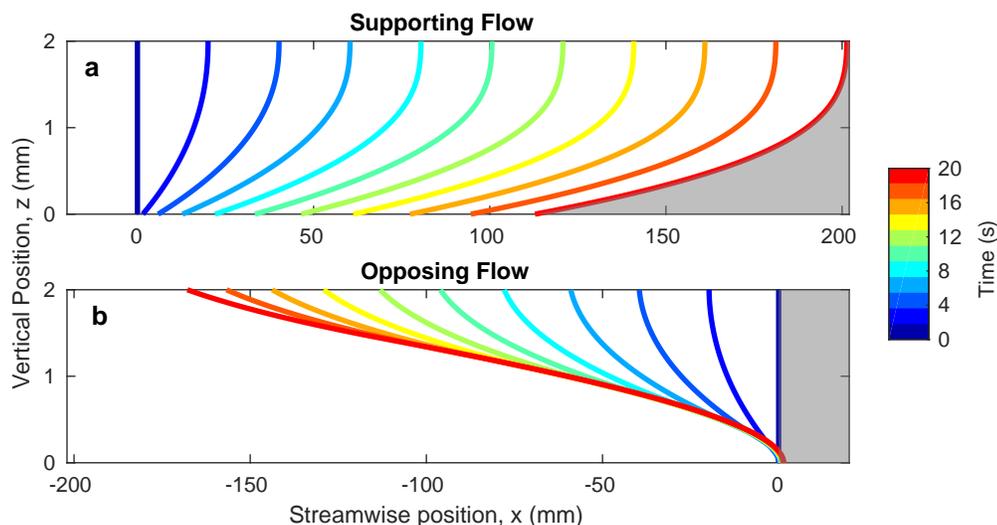


Figure 4.3: Fronts propagating in supporting and opposing thin-layer flow, according to the Eikonal equation. Each curve represents the location of a front moving at $v_0 = 72 \mu\text{m/s}$, throughout the depth of a thin layer. Different colors indicate different times. Flow speed is 10 mm/s , and fronts are initialized with $\theta = \pi/2$. The flow is directed to be either (a) supporting the front (\hat{x}) or (b) opposing the front ($-\hat{x}$). Work of the author in [101].

symmetry difference between fronts in supporting flow (fig. 4.3(a)) and opposing flow (fig. 4.3(b)), which causes a change in front shape. Fronts in opposing flow are pinned at the solid boundary, maintaining a point that resists flow, whereas no such pinning occurs in supporting flows. Because they are pinned, fronts in opposing flow are sheared more strongly. This difference results in higher apparent chemical speeds in opposing flow than supporting flow.

I also simulated Eikonal fronts for 20 s of front growth in flows with $-20 \text{ mm/s} \leq U \leq 20 \text{ mm/s}$. Fig. 4.4 shows all the fronts at two times, $t = 5 \text{ s}$ and $t = 10 \text{ s}$. As one might guess, larger flow magnitudes cause greater displacements and slopes that deviate more from vertical. Given the condition of fixed chemical speed v_0 in a single-layer system of fixed depth, I can focus on two important variables for fronts growing in a single layer: time and flow intensity.

In order to compare to experiments, I use the depth-averaged concentration

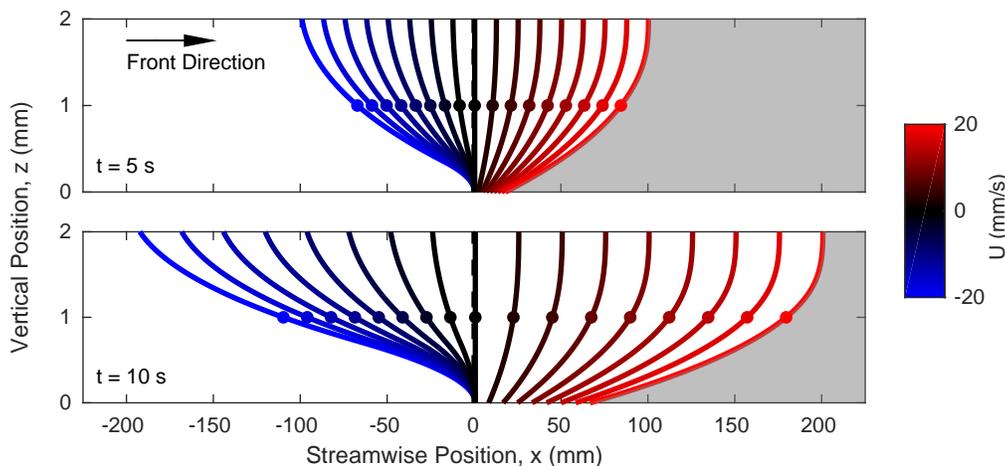


Figure 4.4: Front propagation in thin-layer flow, varying with flow speed, according to the Eikonal equation. Color indicates the flow speed U at the top surface ($z = h$). The white side indicates the reacted side, and the growth direction is to the right. The upper plot shows fronts after propagating 5 s, and the lower plot shows the same fronts after 10 s. All fronts were initialized as vertical lines at $x = 0$. Work of the author in [101].

of each simulation to assign an apparent front location, as would be done in the laboratory. When fronts deviate from being strictly vertical, the depth-averaged concentration varies more gradually in space; fronts appear to be smeared by shear. Apparent fronts can be assigned where the depth-averaged concentration crosses a user-defined threshold. For sharp, vertical fronts, the choice of threshold is irrelevant for front location. For fronts smeared by vertical shear, the apparent front location depends on both the choice of threshold and the actual front profile. If vertical shear changes the actual front profile over time, apparent front speed is also affected. Using the results described above, I chose threshold $\gamma = 50\%$ of the maximum depth-averaged concentration to locate an apparent front at each time t in each simulation, naming that location $x_f(t)$. I calculated the total velocity of the apparent front

$$\mathbf{w} = \frac{x_f(t + dt) - x_f(t)}{dt} \hat{\mathbf{x}},$$

where dt is the time step. Then I calculated the apparent chemical speed v_a using

eq. 4.2 and the known surface flow. Fig. 4.5 shows the results. We find that shear can cause the apparent front speed to be orders of magnitude larger than $v_0 = 72 \mu\text{m/s}$. Similar results can be found in Leconte *et al.* [37] for Poiseuille advection. However, those results pertain to the apparent total speed, w , missing the apparent chemical speed, and the focus was long-term behavior neglects interesting time dynamics.

These simple simulations confirm that apparent front speeds in quasi-2D experiments can vary and be anomalously large, even when the underlying dynamics have a constant front speed. The simulations also show that the apparent front speed converges to an asymptotic limit at long times, though the limit depends on the flow direction. For supporting flow, the apparent speed v_a converges to the true chemical speed v_0 . For opposing flow, the apparent speed converges to $v_0 + U$, the sum of the actual chemical speed and the maximum *flow* speed. In either case, the apparent speed approaches the asymptotic limit from below.

The lower panel of fig. 4.5 suggests an explanation. At $t = 0$, $v_a \propto U$ with the same slope for all values of U , but at later times, v_a and U are related by a piece-wise function with two linear parts, each having a slope that increases over time toward an asymptotic value. The initial slope can be explained by observing that at $t = 0$, the 50% position has $\theta = \pi/2$, so it moves perfectly in the \hat{x} direction. Thus for any U , the measured front speed at $t = 0$ will be $v_a = v_0 + u_x(z = h/2) - U$. Factoring out the surface speed and direction I am plotting against, the slope can be identified as $\tilde{u}_x(z = h/2) - 1$, where $\tilde{u}_x(h/2) = u_x(z = h/2)/U$. Convergence of the apparent chemical speed depends on the concentration profile converging to its long-term shape (fig. 4.3), which occurs as information propagates at the true chemical speed across the layer depth, starting at the leading edge of the front and finally reaching the depth γh . The process is analogous to the downstream widening of a boundary layer or narrowing of an entrance region in pipe flow. Opposing fronts have leading edges at the

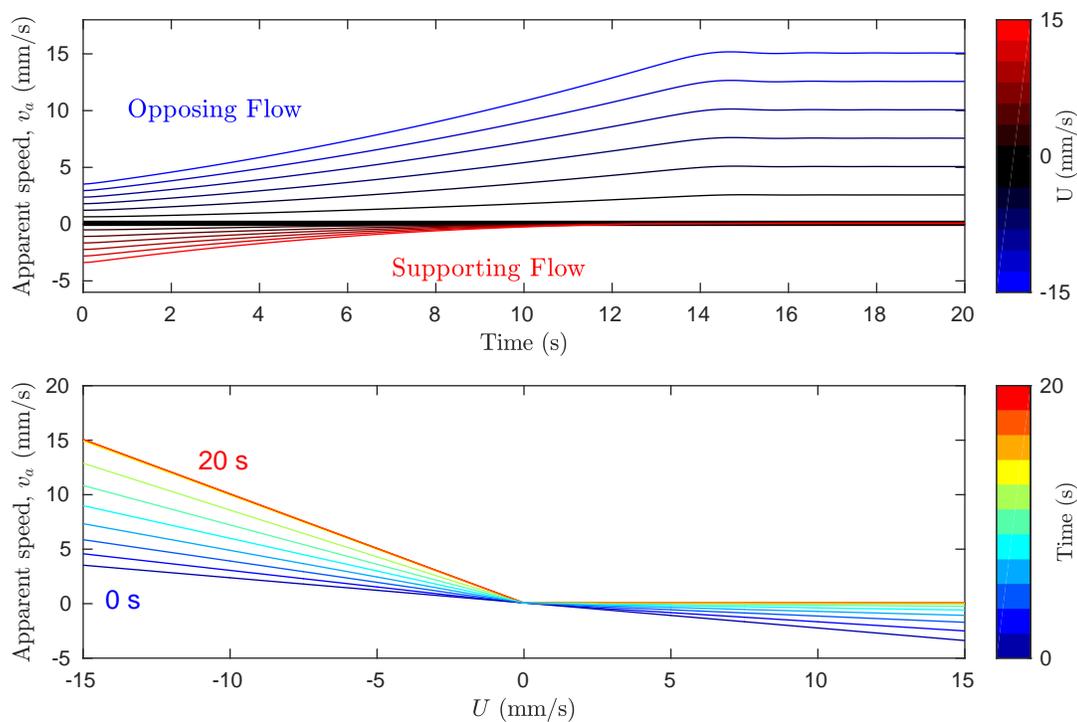


Figure 4.5: Apparent chemical front speed in a simulated thin-layer flow, determined using depth-averaged concentration. The upper panel shows apparent speed over time, for different flow speeds. The lower panel shows the variation of apparent speed with flow speed, at different times, in 2 s intervals. All chemical velocities were obtained by tracking the position where 50% of the layer is reacted, and subtracting surface flow speed, mimicking the procedures for tracking fronts in experiments. Over time, apparent front speed converges to v_0 for supporting flow and $v_0 + U$ for opposing flow. Work of the author in [101].

bottom of the layer, so they converge to the speed of the front at the bottom, which is $v_a = v_0$. Their characteristic time for convergence is $\tau_o = \gamma h/v_0$. Supporting fronts have leading edges at the top of the layer, so they converge to the speed at the surface, which is $v_a = v_0 + U$. Their characteristic time for convergence is $\tau_s = (1 - \gamma)h/v_0$. In the case of $\gamma = 0.5$, both convergence times are the same, and for the parameters listed in table 4.1, $\tau_o = \tau_s = 13.9$ s. Choosing $\gamma < 0.5$ results in faster convergence for opposing fronts than for supporting fronts; choosing $\gamma > 0.5$ results in faster convergence for supporting fronts. The convergence time does not depend on flow speed or structure. The value of γ does not change the asymptotic apparent chemical speed.

Further insight can be gained if I consider front propagation in dimensionless form. Normalizing simulation results like those shown in fig. 4.4 with velocity scale U , length scale h , and timescale h/U produces the dimensionless front profiles shown in fig. 4.6. The simulations plotted there differ only in the ratio $\tilde{v}_0 = v_0/U$. When $|\tilde{v}_0|$ is large, the front propagates further in the \hat{x} direction and extends across a shorter region (in dimensionless units): fronts are less distorted by vertical shear. Note that since all distances are non-dimensionalized with respect to h , that increasing depth will not reduce distortion. Depth's effect is to change how long in physical time a given dimensionless time is, but it cannot change the profile at that time. On the other hand, when $|\tilde{v}_0|$ is small, fronts converge to the shapes one would expect in passive scalar mixing. The curves for a non-reactive scalar are also mirror images of each other, because the front has no directionality in the $\tilde{v}_0 = 0$ case. Front speed is what breaks the symmetry between supporting and opposing flows. Though flow speeds in the experiments (described in section 4.5) are often much faster than reaction front speeds, I nonetheless observe clear deviation from the $\tilde{v}_0 = 0$ case.

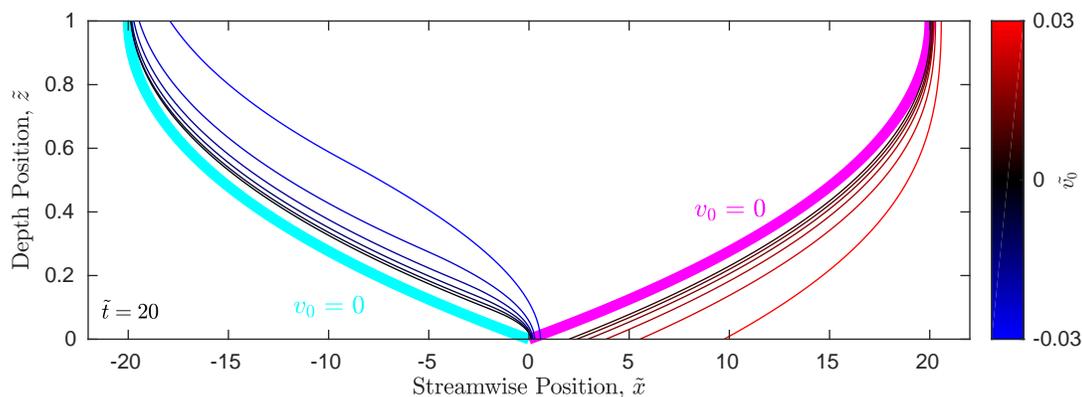


Figure 4.6: Dimensionless front propagation in thin-layer flow. Front speed, v_0/U , varies between different lines, and is normalized with the flow speed U and the layer depth h , at dimensionless time $\tilde{t} = 20$. As v_0/U decreases, the fronts approach the passive case for both opposing and supporting flows. Work of the author in [101].

4.3 Simulations of Hele-Shaw System

Another common type of experiment for thin layer flows is to use a single reacting layer, sandwiched between two solid boundaries and driven by a pressure difference [10, 33, 113]. One major advantage of the Hele-Shaw system is that it can be made much thinner than the free surface systems. However, the resulting no-slip conditions on both the top and bottom of the experiment cause roughly double the shear for the same layer thickness. Regardless, in many systems where a free surface is impractical the Hele-Shaw method is the only option. As with thin layer experiments, the depth shear has an effect, and here I quantify what that effect would be.

4.3.1 Velocity Profile

Flow in a Hele-Shaw system is driven by pressure difference ∇P with no-slip boundary conditions at $z = \pm h/2$. Solving eq. 4.3 results in Poiseuille flow. If the

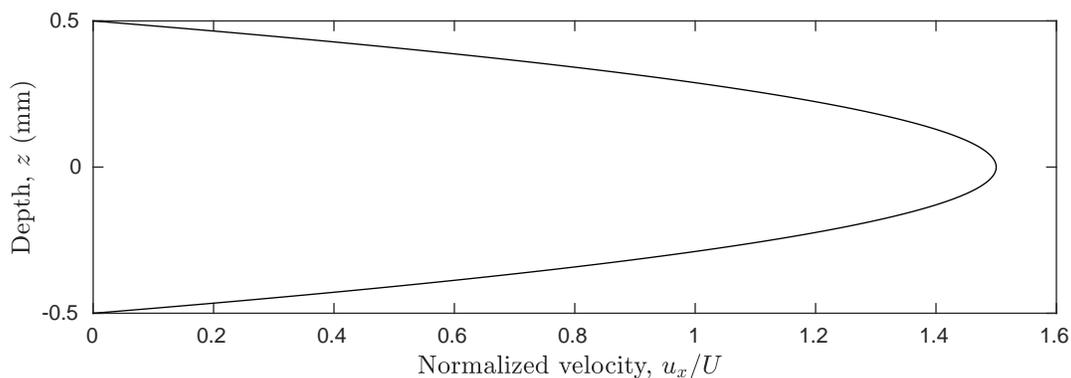


Figure 4.7: Depth dependence of streamwise velocity in a Hele-Shaw style, pressure driven flow. Fluid properties and layer thicknesses are given in table 4.1, with ∇P set to normalize the maximum speed. Work of the author in [101].

flow is in the $\hat{\mathbf{x}}$ direction, then the velocity profile is parabolic:

$$u_x(z) = (\nabla P) \frac{z^2 - (h/2)^2}{2\mu}, \quad (4.9)$$

as shown in fig. 4.7. The profile is symmetric across about $z = 0$, and the halves have the same form as the single-layer profile shown in fig. 4.2. Flow measurement in Hele-Shaw is also measured differently than free surface systems. Particles cannot float on the surface so, the average flow speed through the depth is frequently used instead. Since this is the flow speed of interest, I will define U equal to this average. Integration of eq. 4.9 yields,

$$U = -\nabla P \frac{h^2}{12\mu}.$$

Should an experiment be done using maximum flow speed ($U = u_x(z = 0)$) instead, the results presented here are essentially indistinguishable from those for a single layer system with layer thickness of $h/2$.

4.3.2 Simulation

I repeated the simulations described in section 4.2.2 with u_x given by eq. 4.9, increasing the values of ∇P so that values of U fall in the same range as before.

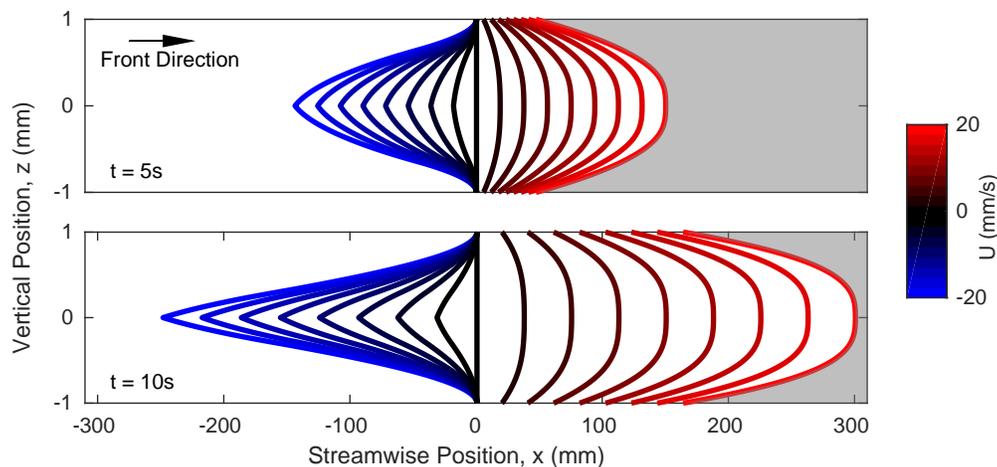


Figure 4.8: Front propagation in Hele-Shaw flow, varying with flow speed, according to the Eikonal equation. Color indicates the flow speed U equal to the average flow speed throughout the depth. The white side indicates the reacted side, and the growth direction is to the right. The upper plot shows fronts after propagating 5 s, and the lower plot shows the same fronts after 10 s. All fronts were initialized as vertical lines at $x = 0$. Work of the author in [101].

Fronts from simulations at different U values are presented in fig. 4.8. The two no-slip boundaries enhance the differences between fronts in supporting and opposing flow. Fronts in opposing flows now pin to both the top and the bottom of the domain, while the middle is now the trailing edge. A sharp cusp forms there on fronts in opposing flow, whereas fronts in supporting flow are not only smooth at the center of the layer, but flattened there.

As with the single layer, the apparent chemical front speed v_a shows strong deviation from v_0 . Fig. 4.9 shows apparent front speeds at different flow speeds and times. The shapes of the plotted curves differ from single-layer systems with a free boundary. This difference arises entirely from defining U as the average, which

is lower than the maximum flow speed in the layer. If maximum is used instead, these results differ from the single layer results only in the time to converge. In supporting flow, v_a converges to $v_0 - U + u_x(z = 0)$, because the leading edge in supporting flows is at $u_x(z = 0) \neq U$. In the previous section the flow velocity at this leading edge corresponded to U , and thus supporting flow converged to v_0 . The convergence time changes in Hele-Shaw because information propagates both up and down: $\tau_s = (1 - \gamma)h/2v_0$, which gives $\tau_s = 6.94$ s for these parameters. In the case of opposing flow, v_a converges to $v_0 + U$ as it did in single layer. The convergence time for opposing flow is analogous to the single-layer system: $\tau_o = \gamma h/2v_0$, which is also 6.94 s for these parameters. The initial slope of v_a versus U in the Hele-Shaw system has the opposite sign and is less steep than in the single-layer system: $\tilde{u}_x(z = h/2) - 1 = 0.13$ instead of -0.23 for a single layer with a free surface. However the lines that Hele-Shaw converges to are at least as far from $v_a = v_0$ as they were in single layer, and the convergence times are strictly smaller. Therefore, Hele-Shaw transitions to the convergent state faster than the single layer. Taken together, this implies that the effects of depth shear are at least as important in Hele-Shaw systems as they are in single layer systems, and using an average flow speed does not cause the measurement of front speed to match the expected value.

Fig. 4.10 presents front propagation in dimensionless form for the Hele-Shaw system. I still see the cusp formation in the dimensionless system. Fronts in the Hele-Shaw system can reach around 60 times wider than they are tall, whereas fronts in the single-layer system reached 20 times as wide as tall. This difference is due to the fact that a Hele-Shaw layer of the same thickness as a single layer must reach its maximum speed in half the distance due to the reflection across the x axis.

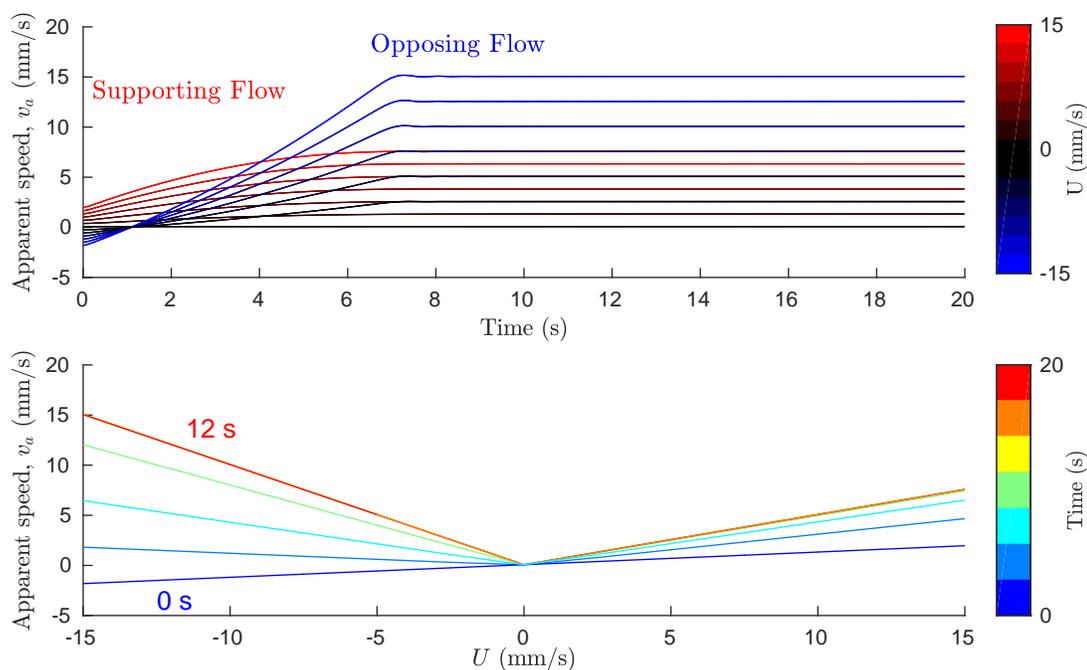


Figure 4.9: Apparent chemical front speed in simulated Hele-Shaw flow, determined using depth-averaged concentration. The upper panel shows apparent speed as it evolves over time, for different flow speeds. The lower panel shows the variation of apparent speed with flow speed, at different times, at 2 s intervals. All chemical velocities were obtained by tracking the position where 50% of the layer is reacted, and subtracting the average flow speed, mimicking the procedures for front tracking experiments. Over time, the apparent front speed converges to $v_0 - U + u_x(z = 0)$ for supporting flow and $v_0 + U$ for opposing flow. Work of the author in [101].

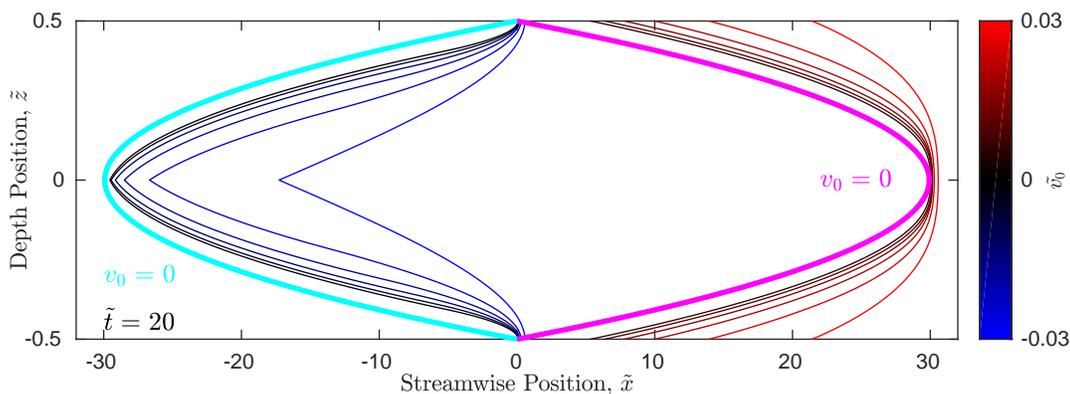


Figure 4.10: Dimensionless front propagation in the Hele-Shaw flow. Front speed, v_0/U , varies between different lines, and is normalized with the flow speed U and the layer depth h , at dimensionless time $\tilde{t} = 20$. As v_0/U decreases, the fronts approach the passive case for both opposing and supporting flows. Work of the author in [101].

4.4 Simulations of Two-Layer System

As I showed in section 4.2, shear in thin-layer flows changes the apparent chemical front speed, if it is measured using depth-averaged concentration, even in simulations that exclude higher-order mechanisms like curvature and Ekman pumping. Section 4.3 showed that shear can also change the results of a Hele-Shaw style experiment. Studies of front propagation with thin-layer experiments must therefore be undertaken with care. To ensure that the apparent front speed gives a good approximation for the true front speed in 2D experiments, shear must be minimized. One strategy for minimizing shear is to add an immiscible, dielectric lubrication layer below the reacting layer. In this section I will repeat the analysis and simulation of section 4.2 for such a two-layer system. I will show that the two-layer system produces apparent front speeds much closer to the true front speed, though shear still has some effect.

4.4.1 Velocity Profile

To determine the velocity profile of the two-layer system, first consider each layer separately. In the lubrication layer that occupies the region $0 \leq z \leq h_d$, current cannot flow: $J = 0$. In the reacting layer that occupies the region $h_d \leq z \leq h_d + h_e = h$, the current density is unchanged from the single-layer case. Flow at the interface between the layers must be continuous and stress free, so the boundary conditions at h_d require that u_x be continuous and

$$\mu_d \frac{\partial u_x}{\partial z} \Big|_{h_d^-} = \mu \frac{\partial u_x}{\partial z} \Big|_{h_d^+},$$

where μ_d and μ are the viscosities of the lubrication layer and reacting layer, respectively. Other parameters and boundary conditions remain unchanged. Solving eq. 4.3, I found

$$u_x(z) = \begin{cases} \frac{JB_0 z}{\mu_d \lambda} e^{-\lambda h_d} (1 - e^{-\lambda h_e}), & 0 \leq z \leq h_d \\ \frac{JB_0}{\mu \lambda} \left(h_d e^{-\lambda h_d} \left[\frac{\mu}{\mu_d} - \frac{\mu}{\mu_d} e^{-\lambda h_e} + \frac{1}{\lambda h_d} + e^{-\lambda h_e} \right] - \frac{e^{-\lambda z}}{\lambda} - z e^{-\lambda(h_d+h_e)} \right), & h_d \leq z \leq h_d + h_e \end{cases} \quad (4.10)$$

This velocity profile agrees closely with the results of Suri *et al.* [110], with differences arising from my use of an exponentially decaying magnetic field (eq. 4.4); they used a linear decay. Fig. 4.11 shows the velocity profile, for the parameters shown in table 4.1. Because $\mu \sim \mu_d$, the slope discontinuity (kink) at $z = h_d$ is weak. I define U in a two-layer system to be the velocity at $z = h_d + h_e$. As expected, only a small fraction of the vertical variation in velocity, about 20%, occurs in the reacting layer.

4.4.2 Simulation

I repeated the simulations described in section 4.2.2 with u_x given by eq. 4.10, increasing the values of J so that values of U fall in the same range as before. Fronts at different U values are presented in fig. 4.12. Fronts in the two-layer system are

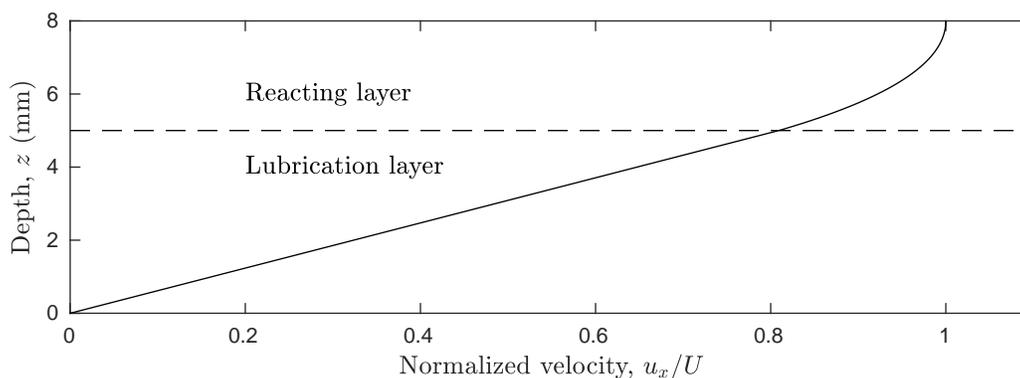


Figure 4.11: Depth dependence of streamwise velocity in a uniform thin-layer flow including a lubrication layer. Fluid properties and layer thicknesses are given in table 4.1, with J set to normalize the maximum speed. Work of the author in [101].

distorted by vertical shear much less than fronts in the single-layer system. The asymmetry between fronts in supporting and opposing flow is smaller in two-layer systems because the bottom of the reacting layer is no longer subject to a no-slip boundary condition, so pinning is eliminated. Some asymmetry remains, however, because front elements can vanish into the edge of the domain but not emerge from it (see section 4.2).

This reduction of distortion has a major effect on the apparent chemical front speed v_a . Fig. 4.13 shows apparent front speeds at different flow speeds and times. The shapes of the plotted curves closely resemble those of the single-layer system (fig. 4.4). However, the magnitude of v_a for any given flow speed U is only about 20% as large as in the single-layer system. For example, the initial slope is once again equal to $\tilde{u}_x(z = h/2) - 1$, but where this value was about -0.23 for single layer, it is -0.04 for two layers. Therefore there is less dependence of apparent chemical speed on flow speed in two layers. In the case of supporting flow, v_a converges to v_0 , just as in the single-layer system. The convergence time is analogous: $\tau_s = (1 - \gamma)h_e/v_0$, which gives $\tau_s = 20.8$ s for these parameters. In the case of opposing flow, v_a does not converge to $v_0 + U$, as in the single-layer

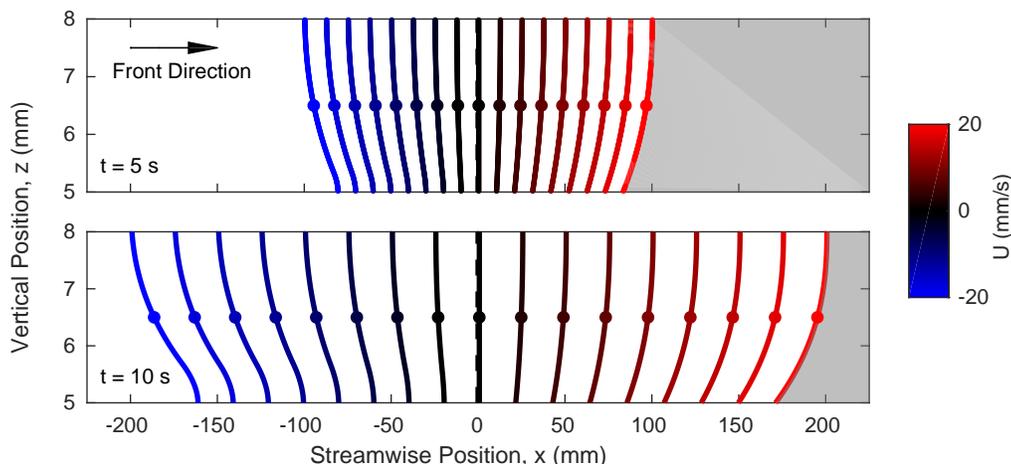


Figure 4.12: Front propagation in the two-layer flow, varying with flow speed, according to the Eikonal equation. Color indicates the flow speed U at the top surface ($z = h_d + h_e$). The white side indicates the reacted side, and the growth direction is to the right. The upper plot shows fronts after propagating 5 s, and the lower plot shows the same fronts after 10 s. All fronts were initialized as vertical lines at $x = 0$. Work of the author in [101].

system, but to $v_0 + U - u_x(z = h_d)$. That is, the apparent chemical speed exceeds the true chemical speed by the difference in flow speed at top and bottom. The same is true for the single-layer system, since the flow speed at the bottom of a single layer is zero. The convergence time for opposing flow in the two-layer system is analogous to the single-layer system: $\tau_o = \gamma h_e / v_0$, which is also 20.8 s for these parameters.

I can also consider front propagation in dimensionless form in the two-layer system, and results are shown in fig. 4.14. As in the single-layer system, fronts are distorted least when \tilde{v}_0 is large, and have shapes like fronts bounding a passive scalar when \tilde{v}_0 is small. The effect of changing \tilde{v}_0 is much weaker than in the single-layer system, however. Fronts in the two-layer system are never more than 4 times as wide as they are tall, whereas fronts in the single-layer system can be 20 times as wide as tall.

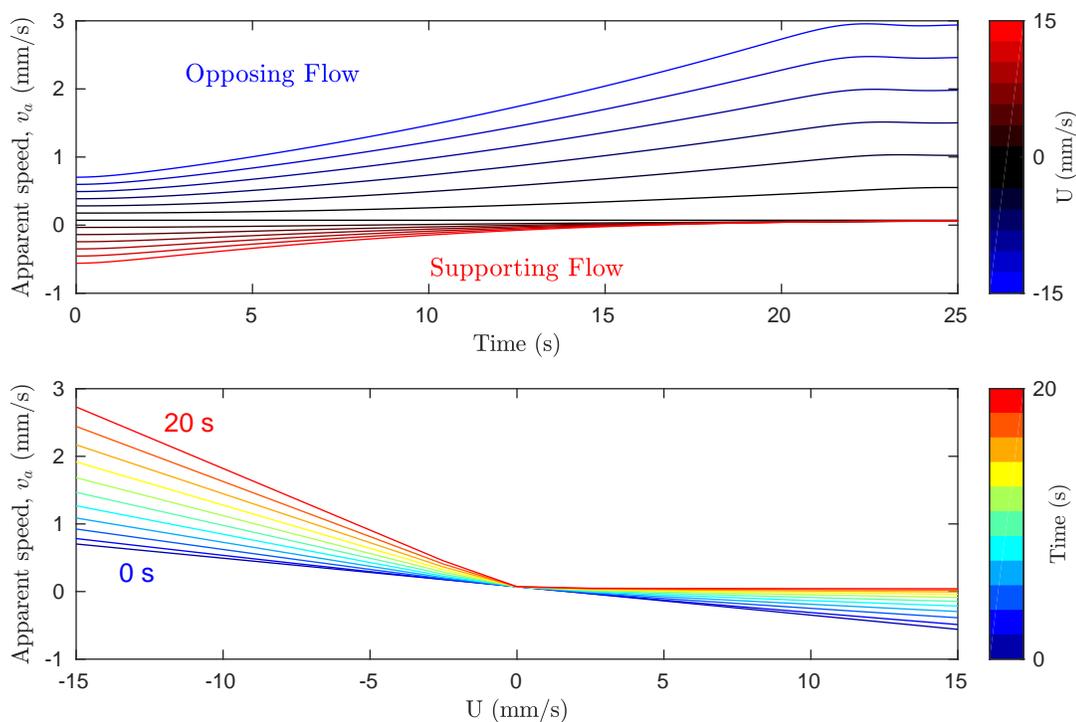


Figure 4.13: Apparent chemical front speed in simulated two-layer flow, determined using depth-averaged concentration. The upper panel shows apparent speed as it evolves over time, for different flow speeds. The lower panel shows the variation of apparent speed with flow speed, at different times, in 2 s intervals. All chemical velocities were obtained by tracking the position where 50% of the layer is reacted, and subtracting the flow speed at the surface, mimicking the procedures for front tracking experiments. Over time, the apparent front speed converges to v_0 for supporting flow and $v_0 + U - u_x(z = h_d)$ for opposing flow. Work of the author in [101].

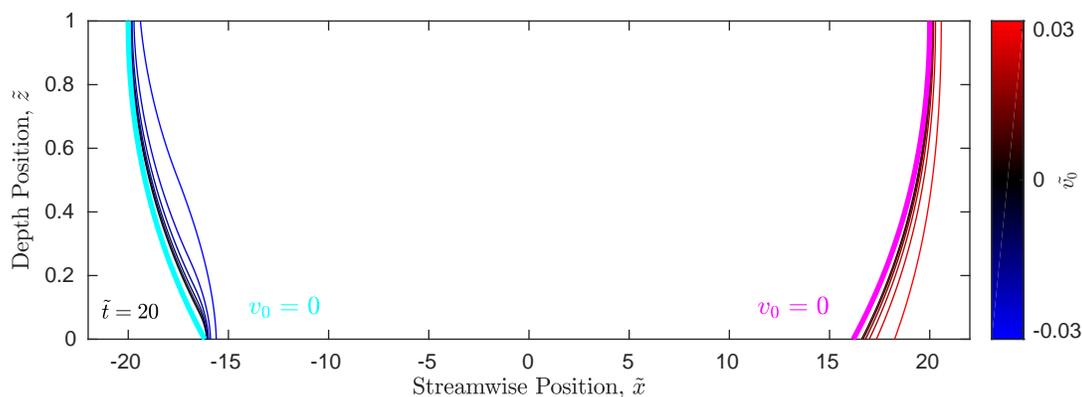


Figure 4.14: Dimensionless front propagation in the two-layer flow. Front speed, v_0/U , varies between different lines, and is normalized with the flow speed U and the layer depth h , at dimensionless time $\tilde{t} = 20$. As v_0/U decreases, the fronts approach the passive case for both opposing and supporting flows. Changing v_0/U has a much smaller effect than in the single-layer system (fig. 4.6), because fronts are not pinned by a no-slip boundary condition at the bottom of the layer. Work of the author in [101].

4.5 Experimental Apparatus

Simulations show that vertical shear causes the apparent front speed v_a to deviate from the true front speed v_0 in thin-layer experiments, assuming that the Eikonal approximation (eq. 4.2) holds and that reaction fronts propagate according to eq. 4.6, 4.7, and 4.8. Those assumptions generally hold in the limit of thin fronts, when the characteristic reaction rate is much faster than the characteristic rates of advection and diffusion. That is, I expect those assumptions to hold when the first Damköhler number $Da_I = \alpha h/U$ and second Damköhler number $Da_{II} = \alpha h^2/D$ are both large. I find $Da_{II} \geq 2600 \gg 1$ for all simulations described above, satisfying the necessary condition. However, while Da_I goes to infinity at $U = 0$, it drops below 10 at a flow speed of only $U = 0.25$ mm/s, which means a very small portion of the simulated domain falls in the high Da_I range. The validity of the Eikonal approximation is not obvious. To test the validity, and determine whether the simulations accurately predict the effect of shear on apparent front

speed, I performed reactive mixing experiments in thin-layer flows.

Experiments were configured to match the simulations. I drove flows that were nearly uniform and nearly unidirectional by passing electrical current through a thin layer of reacting fluid in the presence of a magnetic field that varied little in the horizontal directions and decayed exponentially in the vertical direction. I varied u_x by changing the current amplitude J . All flows were steady after an initial transient. I performed experiments both with and without a lubrication layer below the reacting fluid.

To accurately measure the dynamics of reactive mixing, I use an experimental apparatus which can measure both depth-averaged concentration and flow velocity simultaneously, as shown in fig. 4.15. A pair of cameras (Emergent HS-4000M) images a thin layer of the reacting Belousov-Zhabotinsky (BZ) solution [2, 92–94]. It changes color from blue to red as it reacts, producing fronts with chemical speed $v_0 = 72 \text{ } \mu\text{m/s}$ as shown in chapter 2. Single-layer experiments were recorded with spatial resolution 0.163 mm/pixel; two-layer experiments, 0.140 mm/pixel. On top of the reacting layer, red fluorescent tracer particles (Cospheric UVPMS-BR-0.995, 75-90 μm) follow the flow closely, with Stokes number $St < 0.1$. The rectangular flow channel is 263 mm long, and its width can be varied from 59 to 68 mm wide. It is placed above a large ceramic magnet (152 mm \times 102 mm) which provides unidirectional flow within the channel when current is induced from two bounding electrodes. Fluid recirculates outside the viewing area (fig. 4.15 (b)). In the two-layer system, immiscible Fluorinert FC-3283 is placed below the reacting layer. We illuminate the apparatus with blue LED light, which passes through the BZ solution to reflect off a white background and which matches the absorption wavelength of the particles. One camera has a blue-pass optical filter to see the light reflected from the blue reacted regions, and the other has a red-pass filter which sees the fluorescent particles; fig. 4.15(c)-(d) shows examples.

The BZ reaction is an excitable redox reaction catalyzed by ferroin indicator.

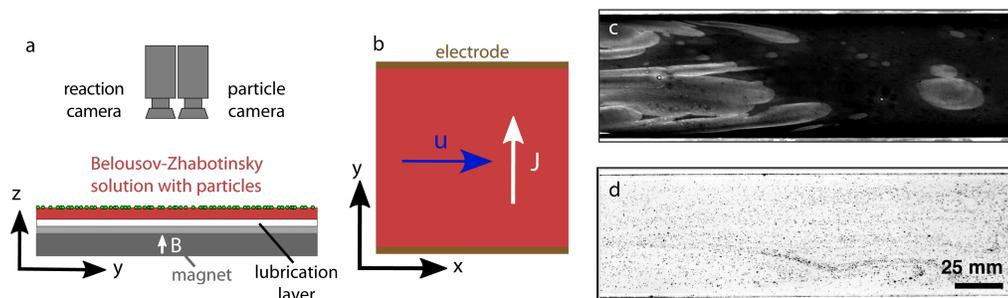


Figure 4.15: Diagram of the experimental apparatus. (a) Two cameras image a thin layer of reacting Belousov-Zhabotinsky solution with red tracer particles. For single-layer experiments there is no lubricating layer. (b) Flow is driven by a current in the $+\hat{y}$ direction, and a magnetic field in the $+\hat{z}$ direction, resulting in a channel flow in the $+\hat{x}$ direction. (c) Typical image from the reaction camera, and (d) from the particles camera (inverted to enhance visibility). I use particle tracking to measure flow velocity and front tracking to measure front speeds. Work of Douglas Kelley.

The reaction oscillates, allowing the same batch to be used for several different experiments. While the BZ reaction is chemically complex, I will measure only the early stages of the reaction within each experiment so it is well-modelled as a second-order autocatalytic reaction [2, 78]: $F(c) = c(1 - c)$. I produce the BZ solution using a recipe similar to one described in Bargteil and Solomon [34]. In a fume hood, I mix a beaker of liquid to concentrations: 0.22 M sulfuric acid (H_2SO_4), 0.12 M malonic acid ($\text{C}_3\text{H}_4\text{O}_4$), 0.12 M sodium bromide (NaBr), and 0.34 M sodium bromate (NaBrO_3). After mixing and the addition of ferroin indicator, a silver wire is immersed in the BZ layer for 20 s to catalyze the reaction, which is allowed to grow for 30 s before flow is started. Ferroin indicator is added repeatedly throughout experiments to maintain strong contrast, which would otherwise fade over the 2-3 hours in which a batch of solution continues reacting. Color does not provide a calibrated measurement of depth-averaged concentration, nor is it accurate to make quantitative comparisons between color at one time and color at a much later time. However, color variation in an image does accurately locate reaction fronts.

I use particle tracking velocimetry to produce velocity fields from particle motions [75]. Then I use front tracking velocimetry to produce chemical front velocities from the concentration fields and velocity fields as described in chapters 2 and 3. I varied the frame rate among experiments, from 15 Hz for slow flow to 45 Hz for fast flow, to ensure good particle tracking. Fig. 4.16 shows an example velocity field from a single-layer experiment, and fig. 4.17 shows an example velocity field from a two-layer experiment. Both experiments produced flows nearly uniform and unidirectional (in the \hat{x} direction) over most of the field of view. The two-layer experiment produced a slightly more uniform flow, partially because the magnet was better aligned during that experiment. Uniformity is not essential since we measure both the local flow velocity and the local apparent front velocity directly. Once the current J is switched on, two-layer experiments take longer to reach full speed than single-layer experiments because of the added inertia of the lubrication layer.

In the simulations the front always had horizontal orientation such that \hat{n} pointed in the \hat{x} direction. In our experiments the front is not always oriented in the streamwise direction. Initially, reacted regions are roughly circular, so the leading edge experiences supporting flow, while the trailing edge experiences opposing flow. When a flow is not strictly one-dimensional, the speeds of front elements are affected by the component of the flow locally normal to the front, $\mathbf{u} \cdot \hat{n}$. This dot product is positive for supporting flow and negative for opposing flow. Tangential front displacement has no effect, since concentration is constant along the tangential direction.

4.6 Experimental Results

Using the apparatus described in section 4.5, I measured the growth of reacted regions under the influence of uniform flows with different speeds. Experimental

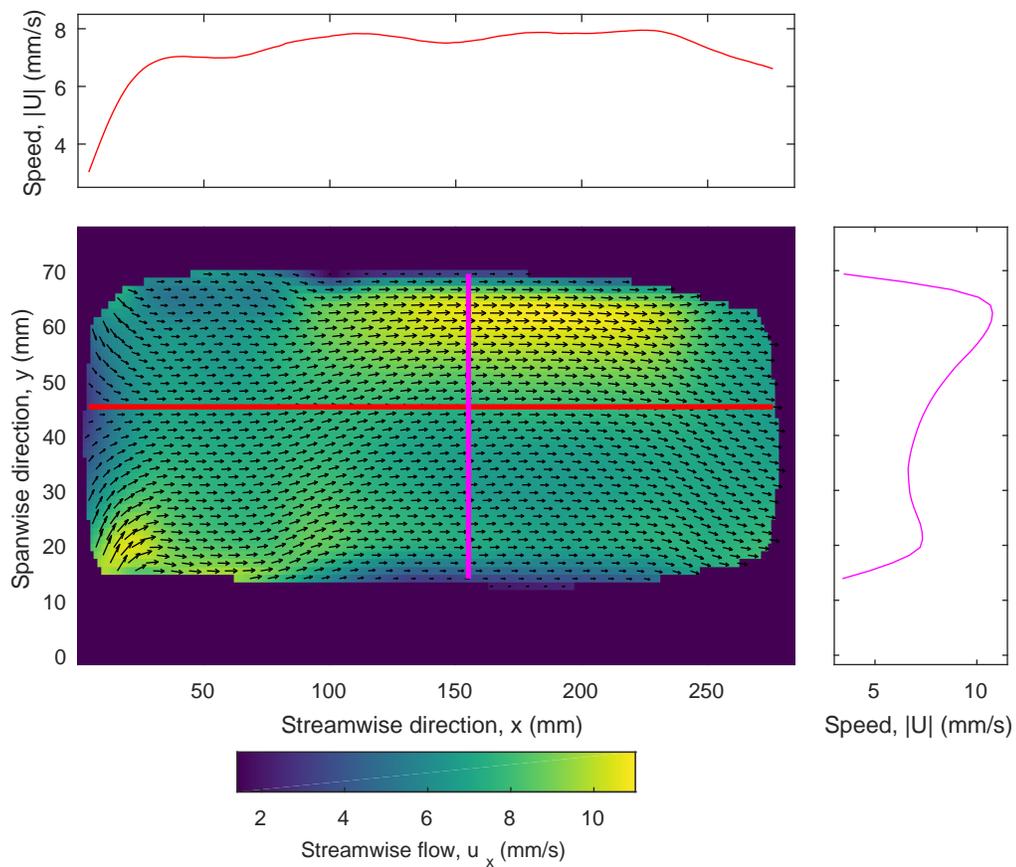


Figure 4.16: Mean flow in a single-layer experiment. Red and magenta lines indicate cross-sections where I measured the speed profiles plotted in the top and right panels, respectively. Work of the author in [101].

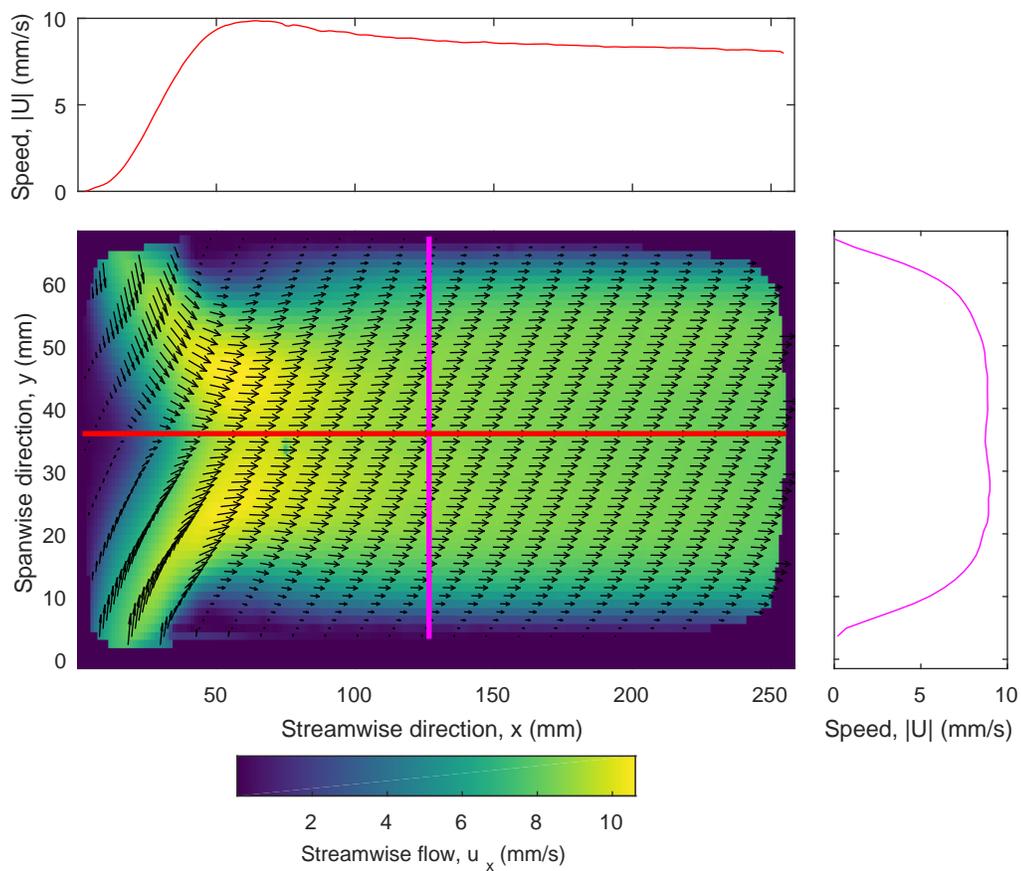


Figure 4.17: Mean flow in a two-layer experiment. Red and magenta lines indicate cross-sections where I measured the speed profiles plotted in the top and right panels, respectively. Work of the author in [101].

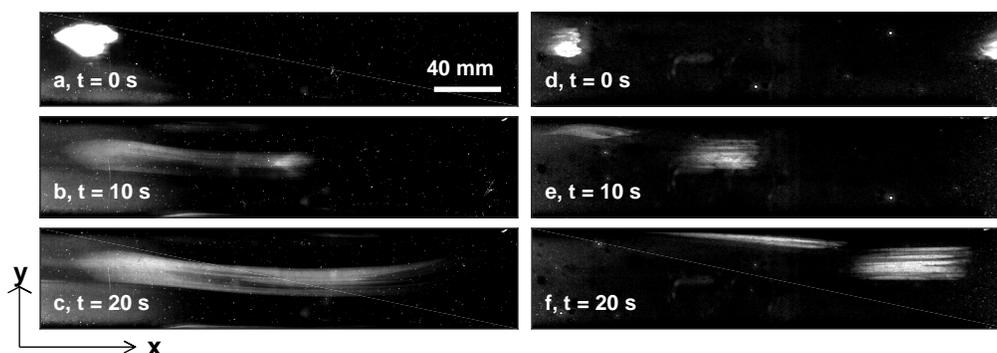


Figure 4.18: Snapshots of the BZ reaction in a single-layer experiment (a-c) and a two-layer experiment (d-f). In both experiments, $U \approx 13$ mm/s and flow goes from left to right, but the reacted regions in (a) evolve differently than in (d). Streamwise smearing is much stronger in the single-layer system because of greater vertical shear. Work of the author in [101].

parameters match those in table 4.1. Fig. 4.18 shows a series of images from each of two experiments, one using the single-layer system, and the other using the two-layer system. Though the flows in these two experiments have comparable speeds, the reacted regions evolve in qualitatively different ways. In the single-layer experiment shown in fig. 4.18(a-c), the reacted region becomes smeared in the streamwise direction over time, therefore dimming. The left end of the region appears to remain nearly stationary. There, the front propagation direction \hat{n} points left, and the rightward flow opposes it, so the observation of a nearly stationary front implies $v_a \approx U \gg v_0$. The right end of the region moves to the right with propagation direction \hat{n} pointing right. The rightward flow supports the front, but measurements show its total speed to be slightly less than the flow speed, so $v_a < 0$. This behaviour was predicted by the simulations, as shown in the lower panel of fig. 4.5: $v_a \approx U$ in opposing flow, and $v_a < 0$ in supporting flow. Interestingly, this behaviour also leads to dilution of the reaction region, which the Eikonal equation could not predict. The two-layer experiment shown in Fig. 4.18(d-f) involves much less smearing, so that the smearing persists only weakly, consistent with the predictions of figure 4.13.

Having shown that the qualitative speeds of experimental fronts in opposing and supporting flows are consistent with trends predicted by simulations, I now make quantitative comparisons. I used front tracking to measure chemical speeds and total speeds at many locations on all fronts in 7 single-layer experiments, each with a different flow speed U . For each experiment, I calculated the mean apparent chemical speed, v_a , at all locations with similar values of $\mathbf{u} \cdot \hat{\mathbf{n}}$ during the first 20 s of flow. Measurements were grouped into bins 0.25 mm/s in extent, and conditional averages were discarded for all bins containing fewer than 100 measurements. Next, I calculated the mean of v_a conditioned on $\mathbf{u} \cdot \hat{\mathbf{n}}$ for all the simulations plotted in fig. 4.5, combining the simulations to produce a single curve, and including error bars sized according to the variation of v_a over the duration of the simulation. The variation in v_a in simulations is due to time dependence of the front speed, which originates from the front's depth profile changing over time.

Fig. 4.19 shows the results. Each experiment gives a different curve, because it represents a variety of $\mathbf{u} \cdot \hat{\mathbf{n}}$ values. Experiments are labeled by the root-mean-square flow velocity measured in that experiment during the 20 s of data collection. Experimental measurements show similar trends to the simulation results: $v_a < 0$ for supporting flows ($\mathbf{u} \cdot \hat{\mathbf{n}} > 0$), and v_a grows with the magnitude of opposing flows ($\mathbf{u} \cdot \hat{\mathbf{n}} < 0$). Five of the seven experimental curves fall within or near the error bars of the prediction from simulations.

Two of the five experimental curves in fig. 4.19 did not fall near the prediction from simulations and deserve further attention. Both experiments involved fast flow that caused severe smearing and poor contrast, making front tracking difficult. When diluted this much the reaction's reacted regions appear to drastically shrink everywhere. The dilution was extreme enough that no reaction appeared again even when the flow was turned off. This front extinction led directly to a lower v_a at all flow values. This indicates a limitation of the Eikonal approximation,

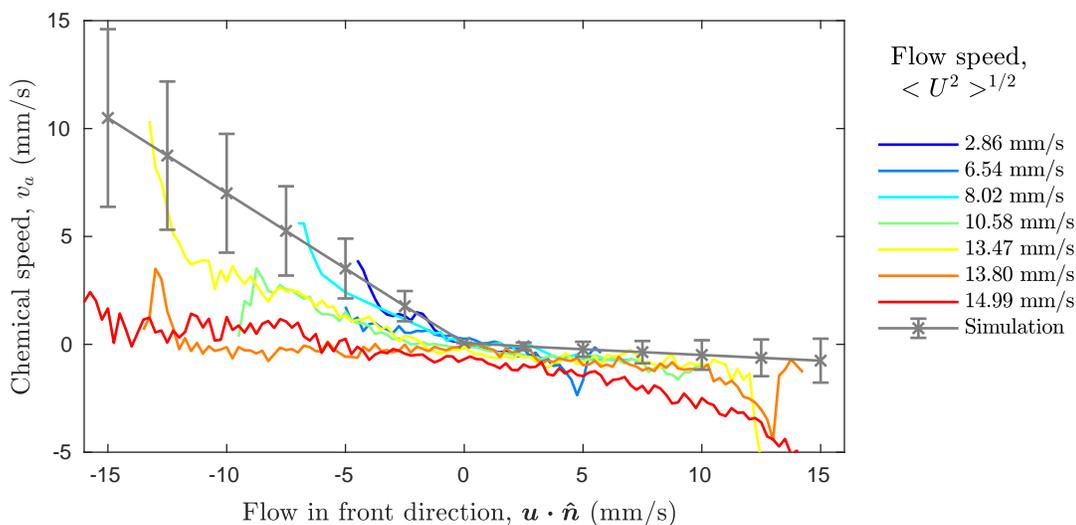


Figure 4.19: Variation of apparent front speed with the front-normal flow speed $\mathbf{u} \cdot \hat{\mathbf{n}}$ in single-layer experiments and simulations. Error bars are developed from data in figure 4.5. Experiments and simulations agree reasonably well. Work of the author in [101].

because in this case the magnitude of the concentration field did matter. These two experiments also correspond to the lowest Da_I . It seems the dilution effect matters beyond some shear level, because slower flows eventually get brighter in spite of the shear. I will discuss dilution effects due to stretching in more detail in chapter 5.

I also considered two-layer experiments. Using the same methods, I calculated the mean of v_a conditioned on $\mathbf{u} \cdot \hat{\mathbf{n}}$ in 16 two-layer experiments and in all the two-layer simulations plotted in fig. 4.13. Fig. 4.20 shows the results. All experimental curves fall within or near the error bars of the prediction from simulations. The shape of the graph is similar to single-layer experiments, but the magnitude of v_a is only about 20% as large, consistent with the fact that $\partial u_x / \partial z$ is only about 20% as large in the two-layer system as in the single-layer system (fig. 4.2 and 4.11). Notice that there is no data set with the large drop like that observed in one layer, even though I reach higher surface speeds. Reduced dilution due to smearing makes the Eikonal approximation hold much better. The extinction of

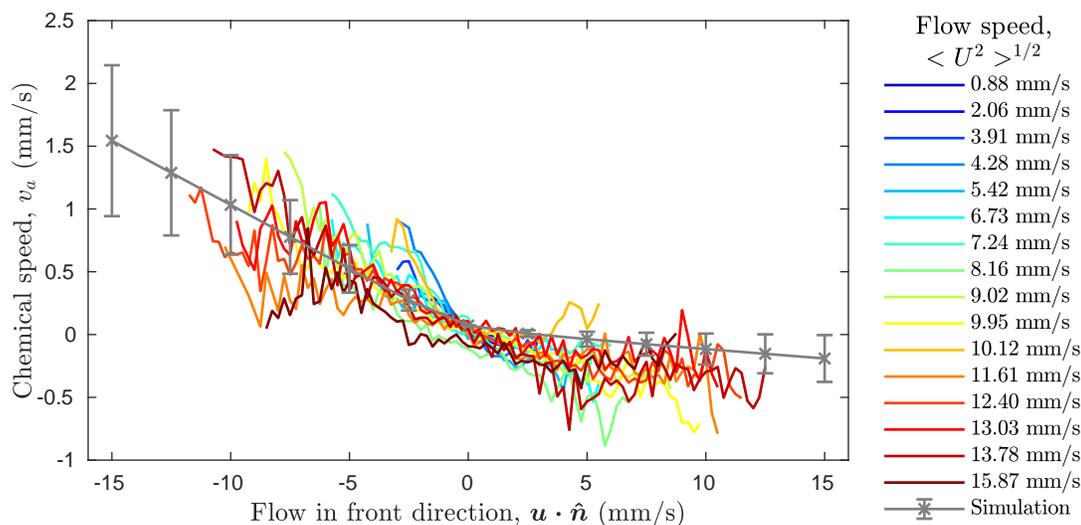


Figure 4.20: Variation of signed front speed with the front-normal flow speed $\mathbf{u} \cdot \hat{\mathbf{n}}$ in two-layer experiments and simulations. Error bars are developed from data in figure 4.13. Experiments and simulations agree reasonably well. Work of the author in [101].

the front in two layers requires a much higher surface flow speed for the same straining and dilution. This is an additional benefit to the predictability and repeatability of two-layer experiments over one-layer.

In experiments, v_a varies with $\mathbf{u} \cdot \hat{\mathbf{n}}$ in a way that is quantitatively consistent with predictions from our simple simulation built on the Eikonal approximation, neglecting curvature and assuming thin fronts moving at constant speed, despite the fact that Da_I was small in some cases. The match holds for single-layer and two-layer experiments.

Reaction fronts are characterized not only by a chemical speed, but also by a thickness, which is the distance separating reacted and unreacted regions, as represented by a fitting parameter (see chapter 2). Predicting thickness is impossible with my simple simulations, which are built on the thin-front assumption. Simulations built from numerical solutions of equation 4.1 would predict front thickness but are substantially more numerically demanding and lie beyond the scope of this study. My experimental measurements, however, include front thick-

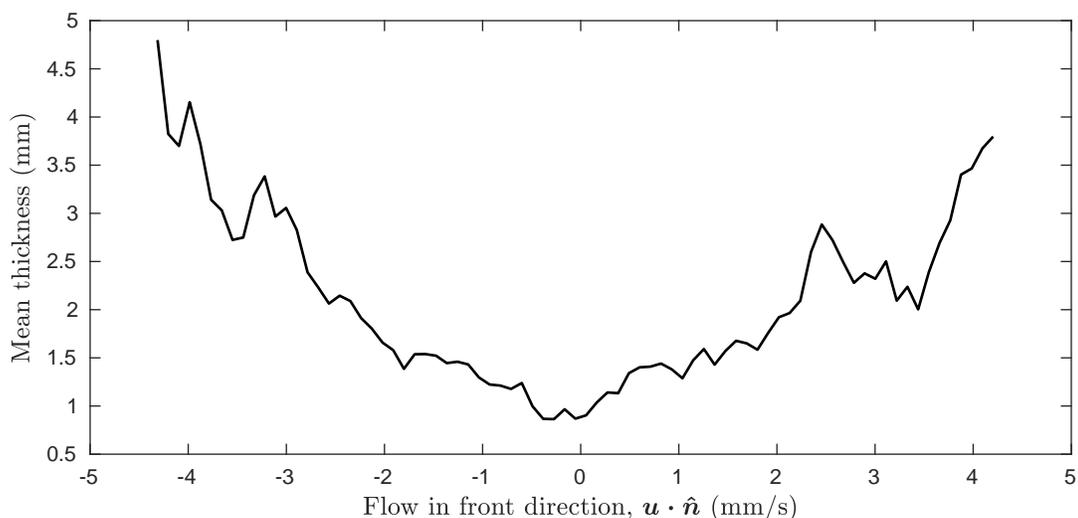


Figure 4.21: Variation of front thickness with the front-normal flow speed $\mathbf{u} \cdot \hat{\mathbf{n}}$ in single-layer experiments. Vertical shear increases front thickness in either supporting or opposing flow. Work of the author in [101].

ness. I calculated the mean thickness, conditioned on $\mathbf{u} \cdot \hat{\mathbf{n}}$, for all fronts in the single-layer experiment with $U = 2.86$ mm/s. As shown in fig. 4.21, fronts appear thicker when $\mathbf{u} \cdot \hat{\mathbf{n}}$ has greater magnitude, whether positive or negative. Greater normal speed causes greater vertical shear perpendicular to the front, which makes fronts thicker, regardless of whether the flow supports or opposes front propagation. Future work might characterize the variation of front thickness with flow speed in more detail, perhaps providing a functional form to explain the variation shown in fig. 4.21.

4.7 Conclusions

In this chapter I provided an explanation for the puzzling prior observation that apparent chemical speed of reaction fronts depends on flow speed in thin-layer experiments from chapter 3. Assuming that the fronts are sharp and that their actual chemical speed is constant, I simulated their motion in vertical cross-sections of

thin-layer flows. I found that even if fronts are initially vertical, shear distorts them over time, so that locating apparent fronts using depth-averaged concentration leads to apparent chemical speeds which differ from the true chemical speed. However, simulations also showed that adding a lubricating layer can reduce shear distortion by 80%. Single-layer experiments designed to match the simulations showed the same distortion by shear, and two-layer experiments showed the same reduction in distortion. The dependence of apparent chemical speed on flow speed agreed closely between simulation and experiment, for both single-layer and two-layer systems. The agreement held even when dimensionless parameters suggested the underlying assumptions to be unfounded. These results resolve the puzzle by showing how the observed variation can arise from the dynamics proposed by [44], via shear. The agreement between simulations and experiments also further validates our front tracking algorithm by confirming that the results it revealed are physical.

Based on these results, the immiscible two-layer configuration for reactive mixing experiments substantially reduces shear and produces apparent chemical speeds much closer to the true chemical speed, making it the ideal setup to use for such experiments. Whereas velocity measurements from particle tracking or particle image velocimetry are relatively unaffected by 3D phenomena until some forcing threshold is exceeded [76, 77], apparent front speed is strongly affected by 3D phenomena, even under gentle forcing. Reactive mixing has been studied in other thin-layer flows, notably Hele-Shaw configurations [33, 106], but the results of section 4.3 indicate that the effect is greater in Hele-Shaw configurations. While I did not perform Hele-Shaw experiments, I hope that these predictions will be useful to future researchers. In this work, I have considered only a single lubricant type, and only a single lubricating layer thickness, $h_d = 5$ mm. Both could be optimized to reduce shear further. In particular, it is known [110] that increasing the viscosity of the electrolyte layer greatly reduces shear. Using the two-layer

configuration does introduce new technical challenges, notably that reaction bubbles are more difficult to eliminate, and stirring the solution after adding ferroin is trickier. Still, they can be managed with careful experimental methods.

Open questions remain and give opportunity for future work. Apparent front speed displays a piecewise-linear variation with flow speed in both the uniform flows considered here and the vortex flows considered in chapter 3 — but the slopes differ. It would be interesting to use analysis similar to that presented here to derive that slope. The non-uniformity of vortex flow may provide the answer, since the local vertical shear depends on the local speed and direction. Initial studies might use Kolmogorov-like flows with simple horizontal speed variation, and correspondingly simple variation of vertical shear. Rotation may have separate effects and could be considered later. Time-dependent flows may involve still more complications.

In future work, it would be interesting to check the agreement between experiments and simulated front dynamics when the Damköhler numbers are even lower than considered here. Also, I have considered only the Belousov-Zhabotinsky reaction, which has $v_0 = 72 \mu\text{m/s}$. Future work might test that these predictions hold for different chemical kinetics and different true chemical speeds. Ultimately both of these questions probe a deeper question: what are the limits of the Eikonal approximation, and how can we model those limitations to expand its applicability.

In the next chapter, I set my sights on a limitation which is visible in fig. 4.19. There dilution of the front led to extinction. Since the Eikonal equation does not consider concentration on either side of the front, this is beyond its ability to predict. Put another way, in a single layer the flow at the base of the domain is no-slip, so the disappearance of the front cannot be attributed to the front being blown downstream, since there is always a slow zone, where $u_x < v_0$. The front transitions from a pinned state to an unreacted state, not a “blown away” state.

Furthermore, previous works in reactions growing in adverse flows have observed a “frozen front” regime that lasts for some distance before the front is blown downstream [10]. In the Eikonal simulations above, there was no speed at which the front could not eventually achieve a frozen state. This eventual detaching from a no-slip location must be beyond the simple Eikonal equation I have used, and may also be beyond the curvature correction.

Since the depinning appears in the single-layer configuration but not the two-layer configuration, shearing is the probable cause. This may be due to another geometric effect of shearing and strain-rate in general. Shearing does not just cause fronts to turn, it can also change front thickness, and change front perimeter. In reactions of limited spatial extent, straining flows can also pull reacted regions into thin filaments, where concentration levels begin to matter. In the next chapter I will present an analysis of strain over a period of time interacting with nonlinear chemical kinetics to cause effects which are beyond the Eikonal equation.

4.8 Acknowledgements

This chapter is a reprint of material in “Vertical shear alters chemical front speed in thin layer flows” which is currently under review for publication[101], and has been modified to fit in this thesis. The thesis author is the primary author and was responsible for conceiving and performing the experiments and simulations presented in the paper, writing the paper, and preparing the figures (except fig. 4.15). Co-author Douglas Kelley provided guidance on the development of the experiments, was responsible for editing the paper and figures, and created fig. 4.15. Experimental design help from R. S. Russell, J. Wang, and J. Tithof was used to design the experiment. An earlier draft of the paper received editorial help from J. Thomas. T. D. Nevins was supported by the Department of Defense through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Pro-

gram throughout the work for this chapter.

5 Optimal Stretching and Blowout in Excitable Reactions

This chapter is based on a combination of published and unpublished results. It reproduces Reference [5] which is the source of all single layer vortex flow experimental results. Sections from Reference [99] regarding flow past a bluff body are also included as this was a later experiment showing the same phenomenon in a new setting. These two previously published sections have been reformatted and edited to fit within the thesis. The thesis author was the primary author on the work from Reference [5], and a co-author on Reference [99]. I have obtained permission to use the materials in Reference [99]. The last section of the chapter is entirely unpublished analytic and simulation results aimed at explaining the results of the previous papers.

5.1 Introduction

So far in this thesis, I have focused my attention on understanding the most basic, and expected behaviors of advection-reaction-diffusion (ARD) systems. The steps I have taken to do this involved expanding on the simplified front tracking model,

and I have built the experimental and analysis tools required to test that model. As such, most of the previous chapters have been focused on validating the tools I created, and exploring some of their immediate implications. In this chapter I will instead focus on phenomena that are *beyond* that simple model. A hint of what I found can be seen in the two unusual experimental results in chapter 4, fig. 4.19. In this chapter I will discuss some unusual behavior that I observed in our experiments which compared how reactions responded to different levels of fluid stretching. Importantly, this interaction between reaction and stretching seems likely to be universal to excitable ARD systems, making it a fundamental concept to understand for predicting and controlling ARD.

In particular, I am interested in studying ARD systems in which the reaction is both excitable, in that it proceeds only if the concentration exceeds some threshold; and autocatalytic, in that it spreads in fronts once triggered, because the reaction product is also a catalyst. While at first this may seem an unusually specific reaction class, such reaction dynamics describe the well studied oscillating Belousov-Zhabotinsky (BZ) chemical reaction [2, 46, 48–50, 80, 114, 115], flame dynamics [2], and phytoplankton blooms [22, 23]. So these reactions are experimentally accessible, and also can have important applications.

The experiments of this chapter are comprised of two different experimental setups over a five year period. One conducted by me, and one by undergraduate Jinge Wang, partially under my supervision. For all these experiments, we drive thin-layer flows of ferroin-catalyzed BZ reaction [2, 92] as in chapter 4. When triggered, the BZ reaction oxidizes ferroin, changing from red to blue [2]. BZ also produces other products which are catalysts for further reaction, making BZ autocatalytic. Diffusion spreads those catalysts, causing reacted regions to grow outward. Because BZ has a large Damköhler number [48], sharp fronts separate reacted and unreacted regions. We measured BZ's kinematic viscosity, $\nu = 1.6 \times 10^{-6}$ m²/s using a TA Instruments DHR-2 rheometer. We trigger

reaction by immersing a silver wire near the vessel's center for 20 s.

Reaction and flow fields were imaged using two hardware synchronized cameras (Emergent HS-4000M) as described in chapters 2-4. One has a dichroic filter to pass blue light matching the color of the LED illumination and the reacted BZ, so reacted regions appear bright when light is reflected from a white background. The other camera has a dichroic filter which blocks blue but passes the color of the fluorescent tracer particles we add to the fluid to visualize advection. We measure advection by tracking each particle [75]. The particle tracking makes a large number of flow measurements and these measurements are distributed over space, giving us a picture of the flow field. Particles range in size $66\ \mu\text{m}$ to $100\ \mu\text{m}$ in diameter and follow fluid motions accurately with small Stokes number ($St < 0.1$). By simultaneously imaging a calibration pattern with both cameras, we determine their resolution (which varies between experiments) and the translational and rotational offsets between their images. Once the cameras are calibrated, we can overlay simultaneous measurements of advection and reaction state. We adjust camera frame rates with flow speed.

I conducted the first experiment in this project in a 2-mm-thick, single-layer experiment, with flow driven by passing current through BZ over an array of magnets, patterned as in chapter 3 for the small length scale vortex flow (fig. 3.5(b)). As such, I will call this the single-layer vortex experiment. The vessel has lateral dimensions of 230×330 mm, and I avoid the vessel edges by imaging only the central 150×200 mm. The magnet spacing is $L = 25.4$ mm and produces flows with Reynolds numbers $0 \leq Re \leq 350$, where $Re = UL/\nu$. Throughout the experiments U is the measured root-mean-square velocity. The Reynolds number compares inertial forces to viscous forces, and in this apparatus I observe a transition from steady flow to time-dependence at $Re \sim 100$.

Jinge Wang and Rony Colon built the second experiment where they attempted to replicate the results of the first experiment, but in the open flow past a bluff

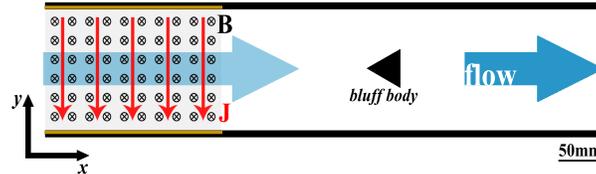


Figure 5.1: Schematic of the channel flow experiment, viewed from above. In the drive section to the left a cross-stream current density \mathbf{J} (indicated by red arrows) in the presence of a downward magnetic field \mathbf{B} produced by nearby permanent magnets, causes Lorentz forces in the x -direction. Reaction is measured in the wake behind a bluff body, located far from the drive section and far from channel ends. Work of Jinge Wang published in [99].

body. This experiment will be referred to as the open flow experiment. This apparatus (shown in fig. 5.1) is a 640 mm long and 139 mm wide channel filled to a 4 mm depth with a single-layer of BZ. In this situation electric current is passed through the BZ in the spanwise direction ($-\hat{y}$), in the presence of a downward ($-\hat{z}$) magnetic field \mathbf{B} , so a Lorentz force drives fluid in the streamwise direction (\hat{x}). The magnetic field is provided by large rectangular ceramic magnets of size 152 mm \times 102 mm \times 12.7 mm. Current is only applied in the area with magnetic field, so a unidirectional flow is obtained. This unidirectional flow using electromagnetic forcing is — as Jack Ryan [116] would say — “a magnetohydrodynamic drive, or caterpillar”. Jinge placed a triangular bluff body (equilateral, with side length $L = 45$ mm and height greater than the fluid depth) downstream from the forcing and symmetric with the channel center line. This creates a recirculating wake. Boundary layers on the leading edges of the triangle separate at the sharp corners, resulting in a pair of vortices behind the bluff body, as is common in wakes. The triangular shape was chosen, so a wake could form with flow speeds comparable to the $v = 72 \mu\text{m/s}$ of BZ. Jinge took high resolution data on one experiment with $Re = 520$.

In this chapter I will investigate the phenomenon shown in fig. 5.2, where I observe that reacted regions primarily occupy vortex edges in low Reynolds num-

ber flows, but occupy vortex centers for high Reynolds number flows in the single layer vortex flow. In the vortex flow, this is comprised of two effects. First in section 5.2, I will show that reaction fronts move towards vortex centers more quickly as the Reynolds number increases, but passive scalar fronts do not. This explains the ease with which reactions reach vortex cores in faster flows. In section 5.3 I find that advective stretching rates [55], related to the finite time Lyapunov exponent (FTLE) [63], correlate with reaction state. The probability of a region being reacted is highest for moderate stretching rates, whereas strong stretching prevents the reaction and weak stretching does little to enhance reaction. Just as a little wind boosts a fire, but a strong gust can extinguish it. Optimal stretching is ubiquitous in our all experiments indicating that it is a flow effect that does not depend on advection details. Since optimal stretching does not depend on flow details, understanding it may explain a huge variety of phenomena in very different systems. For example, regions of fast ocean flow may lack plankton because of strong stretching preventing reproduction, and segregation of species into geographically distinct niches [27] may be explained by optimal stretching values that differ between species. In section 5.4 I explore the probable causes of this optimal stretching value through analysis of the front behaviors in experiments. Finally section 5.5 considers the implications of these findings and possible future directions.

5.2 Radial front growth enhancement

In the single-layer vortex flow I observe that reacted regions spend more time near vortex edges in low- Re flows, but prefer vortex centers in high- Re flows, as shown in fig. 5.2. Supposing a smooth transition between these two distinct states, I would expect that reaction fronts move toward vortex centers with a speed that increases with Re . Figure 5.3 sketches my algorithm for measuring

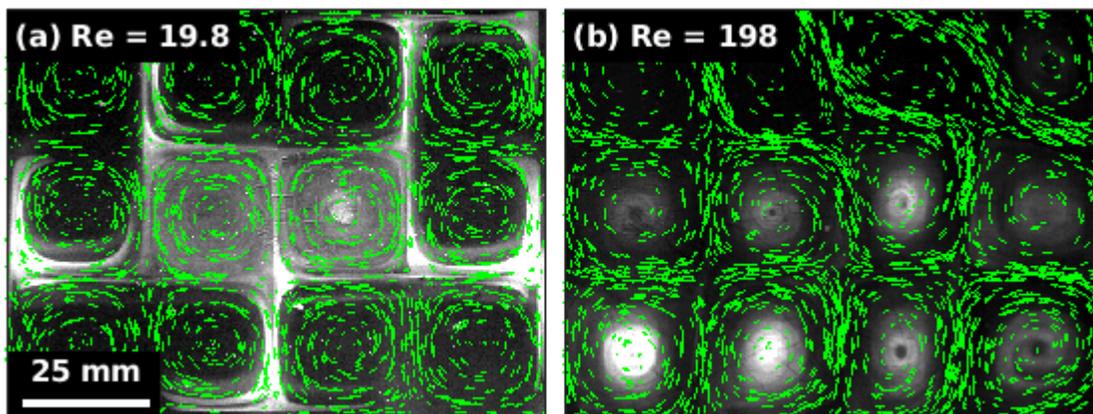


Figure 5.2: Snapshots of ARD experiments with the BZ reaction. Arrows indicate the measured flow velocity. In slow ($Re = 19.8$) flow, reacted regions persist near vortex edges, but in fast ($Re = 198$) flow, reacted regions occupy vortex centers. Work of the author published in [5].

that speed. In each reaction state image I average the brightness of all pixels at a given radial distance from the center of the nearest vortex. The result is a measurement of azimuthally-averaged brightness versus radius for each frame, and I display many subsequent frames on space-time plots. Using the Laplacian of a Gaussian edge finding method [117], I obtain the front location in each frame, and find that the reacted region always moves to smaller radii eventually. To show that the inward propagation of reacted regions is not trivially due to advection, Doug Kelley simulated the motion of a passive scalar in the same flow. In the simulation, he seeded the initially reacted region with tracers, then integrated the measured velocity fields to determine their trajectories. As shown in fig. 5.3(d), he finds little or no *radial* propagation of the passive scalar, consistent with an incompressible flow, but differing starkly from reacting material. (*Azimuthal* speed increases linearly with Re , as expected for passive tracers, or for reactions moving tangent to flow [35, 37].) Thus propagation of reacted regions toward vortex centers depends on coupling between reaction, diffusion, and advection, and cannot be explained by advection alone.

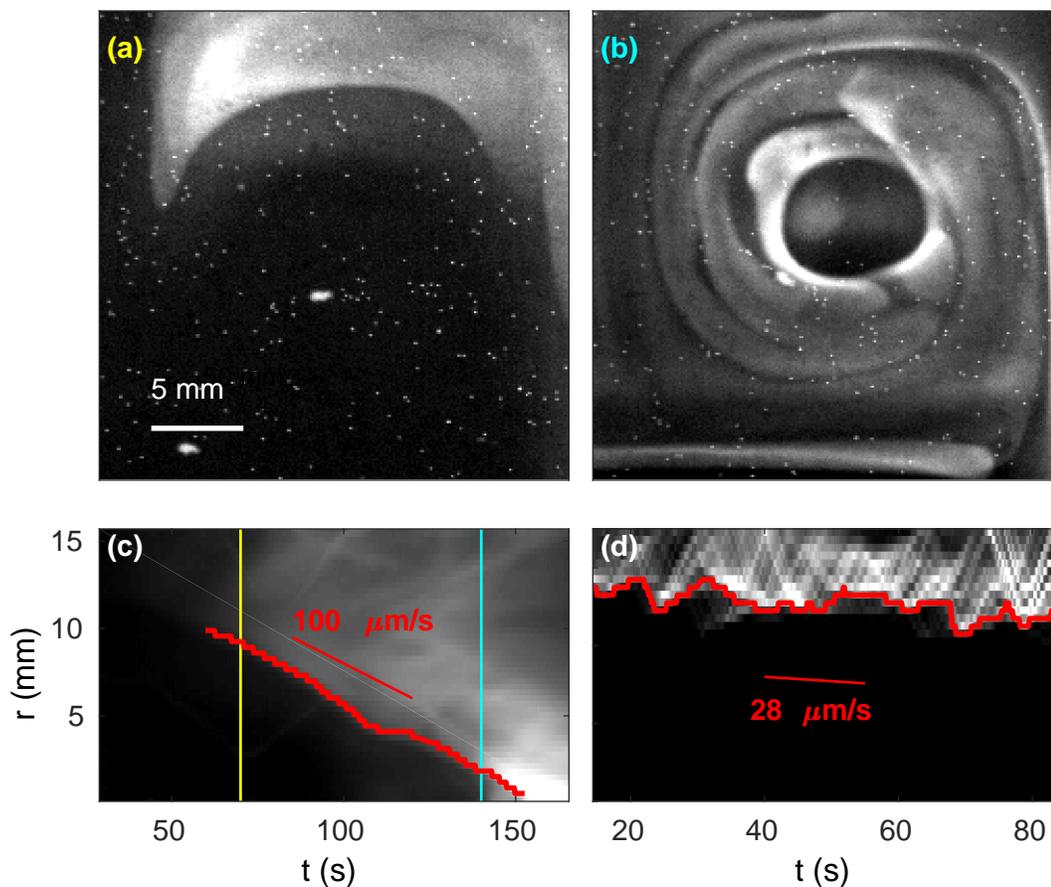


Figure 5.3: Propagation of reacted regions toward vortex centers. (a), A reacted region entering a vortex at time $t = 70$ s. (b), The same vortex at $t = 140$ s, the bright reacted region has propagated inward. (c), Space-time plot showing azimuthally-averaged brightness varying with radial position r and time t , in the vortex shown in (a–b). The thick curve marks radial front position, and its slope measures the inward propagation speed: $100 \mu\text{m/s}$. Vertical lines locate the snapshots in (a–b). Here $Re = 9.3$. (d), Space-time plot generated by simulating a passive scalar in a vortex in a $Re = 18.6$ flow, with the edge and slope indicated as in (c). Inward propagation speed of the passive region is nearly zero (in this experiment, $28 \mu\text{m/s}$). (a–c) Work of the author published in [5]. (d) Work of Douglas Kelley, published in [5].

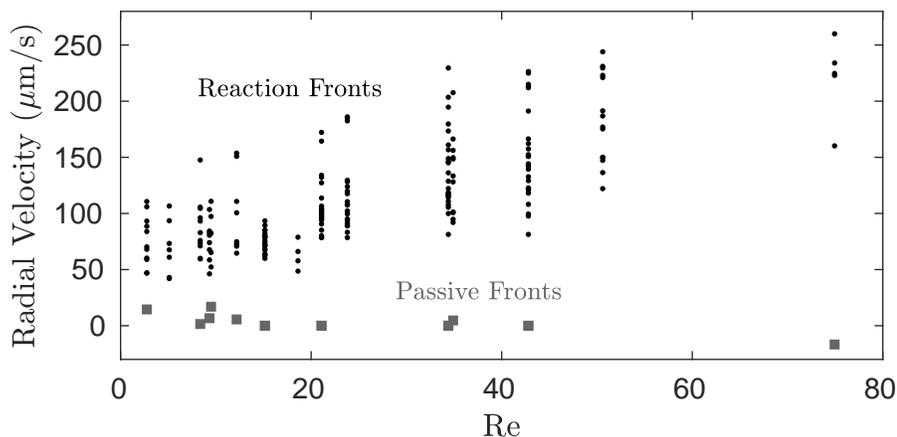


Figure 5.4: Speed toward vortex centers increases with Reynolds number for reaction fronts, but not for a passive scalar. Work of the author published in [5].

With this algorithm I measure *inward radial* front speed at many different flow speeds, and for all reacted vortices not triggered by impurities in the single layer vortex flow, I obtain fig. 5.4. I find that reacted regions do indeed propagate toward vortex centers more quickly as Re increases. The large variation in speed at each value of Re is not an indication of measurement error, nor time-varying advection, as all these experiments involve steady flow. Rather, it indicates variation in initial conditions: reacted regions first reach vortex edges at different times and places for each vortex. Passive scalar fronts show negligible variation with Re , again demonstrating that the propagation of reacted regions toward vortex centers depends on coupling between reaction and advection.

This observation is not consistent with the Eikonal equation. After the initial burn in of a reaction, where the front is wrapped around the vortex, a circular front propagating inward is formed. The front normal direction becomes $\hat{\mathbf{n}} = -\hat{\mathbf{r}}$. The Eikonal equation then predicts that front speed in the $-\hat{\mathbf{r}}$ direction should be $v_0 + (-\hat{\mathbf{r}} \cdot \mathbf{u})$. Since $(-\hat{\mathbf{r}} \cdot \mathbf{u}) \approx 0$, this means the Eikonal equation predicts a constant inward radial speed with Re at $72 \mu\text{m/s}$ for the BZ reaction. This speed is about a third the size of the speed measured in the fastest flows. The

curvature correction [41, 42] cannot fix this, since the radius at which curvature can double front speed is $r = 14 \mu\text{m}$. My resolution is on the order of $100 \mu\text{m}/\text{px}$, so I could not be observing this. Also, flow is perpendicular to the front direction, so the results about depth shear of chapter 4 do not apply. Changing radial speed with increased flow is consistent with Ekman pumping [100], but Ekman pumping would also appear in the simulations of passive fronts, because we used measured flow fields for passive tracers. On top of this, we expect Ekman pumping to resist inward propagating fronts, not support them. This observation is consistent with prior predictions that by spreading reacted regions, advection lengthens fronts and thereby drives diffusion, which supplies reactants more quickly and enhances reaction [52–54], but this effect is beyond the Eikonal equation.

5.3 Observation of optimal stretching

If advection always enhanced reaction, I would expect reacted regions to persist near vortex edges, where advection is fastest, especially for large values of Re . My observations show otherwise (fig. 5.2). A transition occurs after which increased advection no longer enhances reaction, and it is natural to expect the transition to be governed by a quantifiable property of the flow. A property has been identified before in a simplified model that qualitatively fits this behavior. Neufeld [53] reduced a chaotic ARD system by considering the behavior of a single filament. He developed a one-dimensional model whose single flow parameter was the strain rate, which quantifies the rate at which regions of fluid are stretched by advection. Through simulation, he found that the spatially-averaged reactant concentration at large time was maximized at moderate strain rates. Larger strain rates blew out the reaction, and smaller strain rates did little to enhance reaction. Specifically, in Neufeld *et al.* [29] this behavior was identified as a saddle-node bifurcation, where the filament stabilized at progressively thinner widths until the small width could

not be maintained anymore at some critical strain rate. This model is limited, because the strain rate was spatially uniform, and the analysis only covers global reaction state in the steady state.

In more complex flow a reaction can travel through many strain rates through time, and the reaction may experience different strains in different places. Therefore, I would like a different flow quantity that is analogous to Neufeld's strain rate, but which quantifies strain that an element of a reaction has experienced recently. That is, the quantity must be Lagrangian (so it tracks a reaction element), and it must consider a finite time interval. The quantity analogous to strain rate which meets these criterion is the stretching S , which can vary in space and time and can be measured in simulations and experiments. The stretching is a dimensionless quantity defined as the square-root of the largest eigenvalue of the right Cauchy-Green strain tensor [55], and is obtained from ϕ , the function that maps each point \mathbf{x}_0 in the flow at time t to its location \mathbf{x} at time $t + T$. The Cauchy-Green strain tensor $C_T(\mathbf{x})$ is given by [69]

$$\begin{aligned} (C_T(\mathbf{x}_0))_{ij} &= (\nabla\phi_T)_{ki} (\nabla\phi_T)_{kj}, \text{ where} \\ (\nabla\phi_T)_{ij} &= \frac{\partial x_i}{\partial x_{0,j}}. \end{aligned}$$

Which makes the stretching equal to

$$S_T(\mathbf{x}_0, t) = \sqrt{\max \text{eigenvalue}(C_T(\mathbf{x}_0))}. \quad (5.1)$$

Finite time Lyapunov exponents (FTLEs) commonly used for locating barriers to scalar mixing are defined as [55, 62, 63] $\lambda = T^{-1} \log S^{1/2}$. In the results described in this chapter, I consider S with $T = -15$ s, that is, the stretching imposed by the prior 15 s of advection. A consistent stretch time is essential so that in all experiments, the physical meaning of the stretch, as *experienced by the reaction*, is maintained. The value $T = -15$ s corresponds to half an eddy turnover time in the slowest flows measured, and proved long enough for

qualitative convergence of S , which depends only weakly on T if T is sufficiently large [67].

Recent numerical experiments have observed correlations between stretching and reaction rate. In simulations of the double-gyre flow, a two-dimensional prescribed model, the reaction rate was observed to be enhanced in regions of strong stretching [32]. When two competing autocatalytic species were simulated in double-gyre flow, it was observed that the one initiated in a region of stronger stretching almost always dominated long-term [31]. Similarly, in simulations of competing species in sine flow, another two-dimensional prescribed model, it was observed that local FTLEs accurately predict long-term concentrations [8]. Experiments have shown that the global averaged FTLE is correlated with global reaction rate [79]. It seems a global correlation is well established, but local effects are not as clear, and should provide an explanation of where these global correlations come from.

The chaotic advection I observe is not prescribed, but is produced physically using the apparatuses described previously. Furthermore, I measure ongoing reaction state at each location, not only long-term behavior or global behavior. Nonetheless, the observed reaction state correlates with measured stretching, as shown for single-layer vortex-flow in fig. 5.5 and for open flow in fig. 5.6. The correlation changes with advection speed. For small values of Re , reacted regions lie primarily along vortex edges and in narrow filaments where stretching is strongest, consistent with simulations showing that stretching enhances reaction [8, 31, 32]. At large values of Re , however reacted regions are confined where S is weakest. At moderate values of Re , reacted regions lie where stretching is moderate, away from both the large- S corners and the small- S vortex cores. In all cases, reacted regions lie where stretching has similar magnitude — an optimal range of S for enhancing reaction. Optimal stretching differs from the simulations of non-excitable reactions [8, 31, 32], but is consistent with Neufeld's simple model [29, 53].

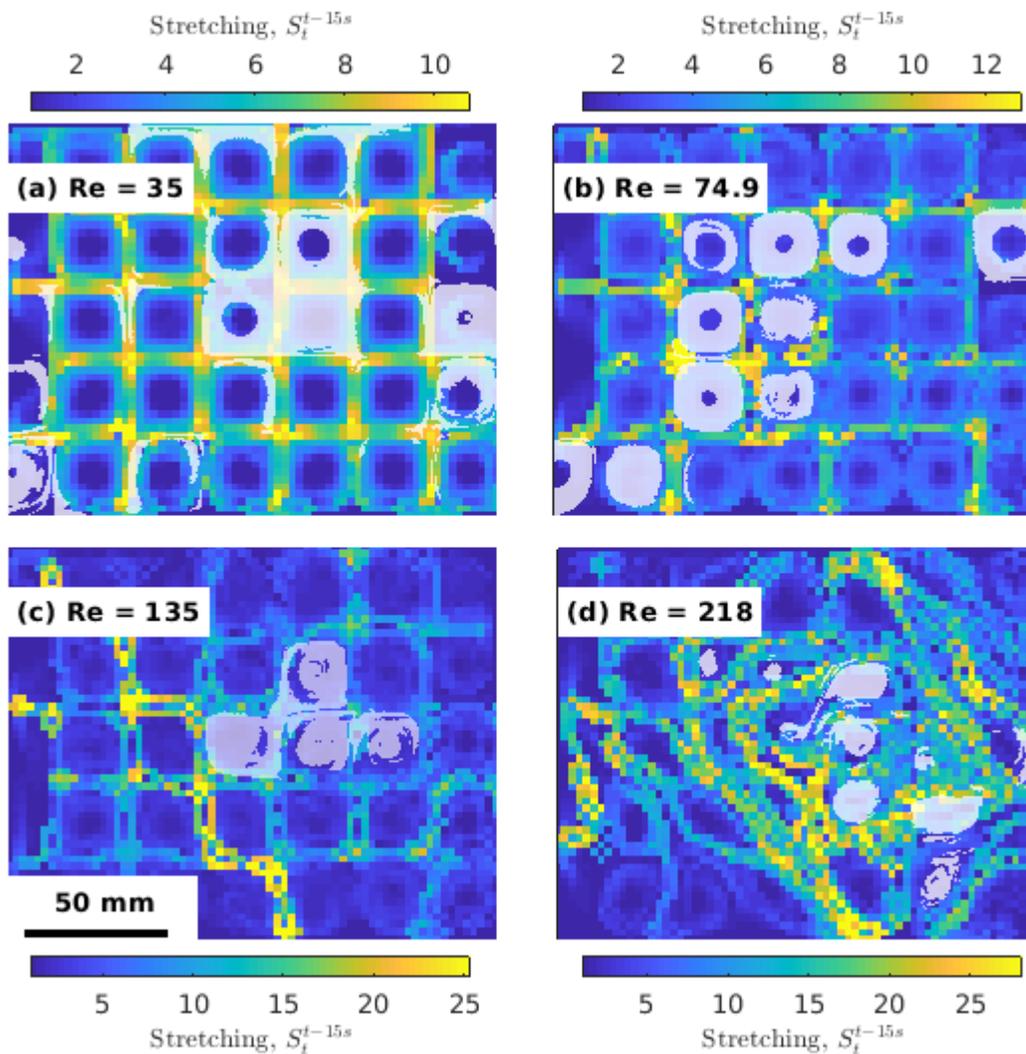


Figure 5.5: Observed reaction state and measured stretching S overlaid, for (a) $Re = 35$, (b) $Re = 74.9$, (c) $Re = 135$, and (d) $Re = 218$. Reacted regions are shaded lighter. The typical stretching in reacted regions is similar in all four experiments, but the range of stretching values increases with Re . Color scales vary. Stretching is calculated on a 2.5 mm grid, below our resolution but allowing sufficient data for stretching measurements at all grid points. Work of the author published in [5].

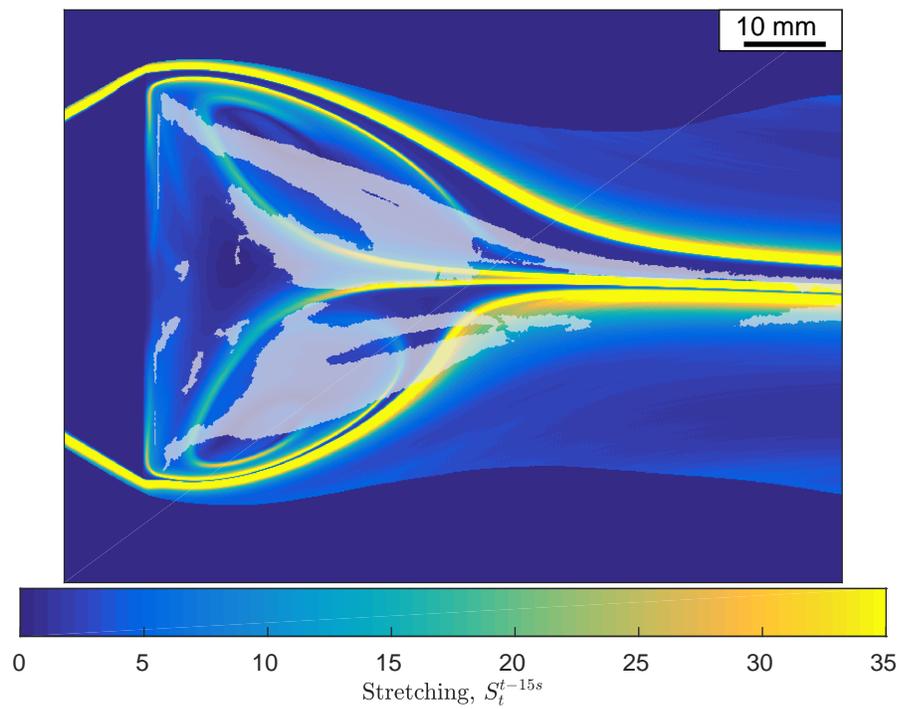


Figure 5.6: Observed reaction state and measured stretching S overlaid, for the open flow experiment at $Re = 586$. Reacted regions are shaded lighter. The typical stretching in the reacted regions is similar to those in fig. 5.5. Work of Jinge Wang published in [99].

To quantify the relationship between stretching and reaction state, I calculate S throughout the observation region at all times after an initial 15 s, and compare the observed reaction state. As in fig. 5.5, regions are classified as reacted or unreacted based on whether their brightness exceeds a chosen threshold. No reacted regions smaller than 5.8 mm^2 are considered. I calculate the reaction probability by normalizing the number of reacted pixels with the given stretching by the total number of points with the given stretching, giving $p(\text{reacted}|S)$. These probabilities estimate a binomial distribution's probability parameter for each bin of S , and as such I can calculate their error bars as being the reliability of this estimator. The probabilities in three single-layer vortex flow experiments are shown in fig. 5.7(a–c), and Jinge's results in open flow are shown in fig. 5.8. Since S is usually small [55], estimation errors are highest for large S . Our analysis shows that for small values of Re , the probability of being reacted increases with stretching. For larger values of Re , larger values of S become available, but the probability there is low. This is the phenomenon I call blowout: regions with stretching stronger than some blowout level S_b allow little or no reaction. Such regions are visible in fig. 5.5(b–d). One striking example is shown in fig. 5.5c. In that experiment, reaction was triggered in the bright vortex just left of the center of the image, but did not spread further left because of a blowout barrier, instead spreading right. I observe $S_b \sim 15$ for BZ with $T = -15 \text{ s}$ in single-layer vortex flow, and Jinge found an optimal range of $8 \leq S_b \leq 20$ in open flow.

Further support for the idea of optimal stretching comes from fig. 5.7(d), where I plot the average stretching $\langle S \rangle$ for reacted and unreacted regions. For large values of Re , $\langle S \rangle$ is larger in unreacted regions; but for small values of Re , $\langle S \rangle$ is larger in reacted regions. Reacted regions persist at highest available S for low Re and at lowest available S for high Re . The transition between these two regimes occurs around $Re \sim 70$, and is not coincident with the transition to time-varying advection at $Re \sim 100$. Alongside the agreement with open flows, this strongly

suggests that this reaction behavior is independent of flow details.

The rough agreement between the single-layer vortex flow and the open flow is truly remarkable. These two systems exist at very different Reynolds numbers, advective timescales, and have different geometries. The pattern even persists between steady and unsteady fluid flows. Also, the single-layer vortex flows only reached flow speeds up to 15 mm/s, while the open flow experiment studied by Jinge was faster yet at 18 mm/s. Characteristic time scales in the vortex flow went from 1.85 s to 11.5 s, while the open flow had a timescale of 2.4 s. Vortex flow had a large number of stagnation points, and the open flow had only a handful in the wake region. Shear is also much higher in the open flow, while vorticity is more substantial in the vortex flow. In fact most global properties are substantially different between the two experiments, but local ARD processes are not different. The reaction kinetics have remained the same, and stretching is calculated at the local level. Assuming stretching is the dominant advection parameter, one would expect that regions of similar stretching would behave similarly regardless of the global flow, and indeed fig. 5.7 and 5.8 support this idea. What's more, in both experiments the regions of $S > S_b$ play a major role in reaction transport, because reactions cannot cross these regions. Since the reaction cannot cross, these zones may act as a kind of "blowout barrier".

The results from the stretching dependence of the two experiments indicates that local flow behavior is the dominant effect, but there are some important caveats that should be discussed. First of all, I have treated all points in a domain as being randomly able to be reacted, and uncorrelated to show the change in ratio of reacted regions with stretching, but reaction growth is not random; reaction could only happen in regions suitably near the reaction front.

Second, these experiments were conducted before the results of chapter 4 were uncovered. As such Jinge and I did not use the two layer system which would more accurately replicate the two-dimensional behavior we measure. Therefore it is not

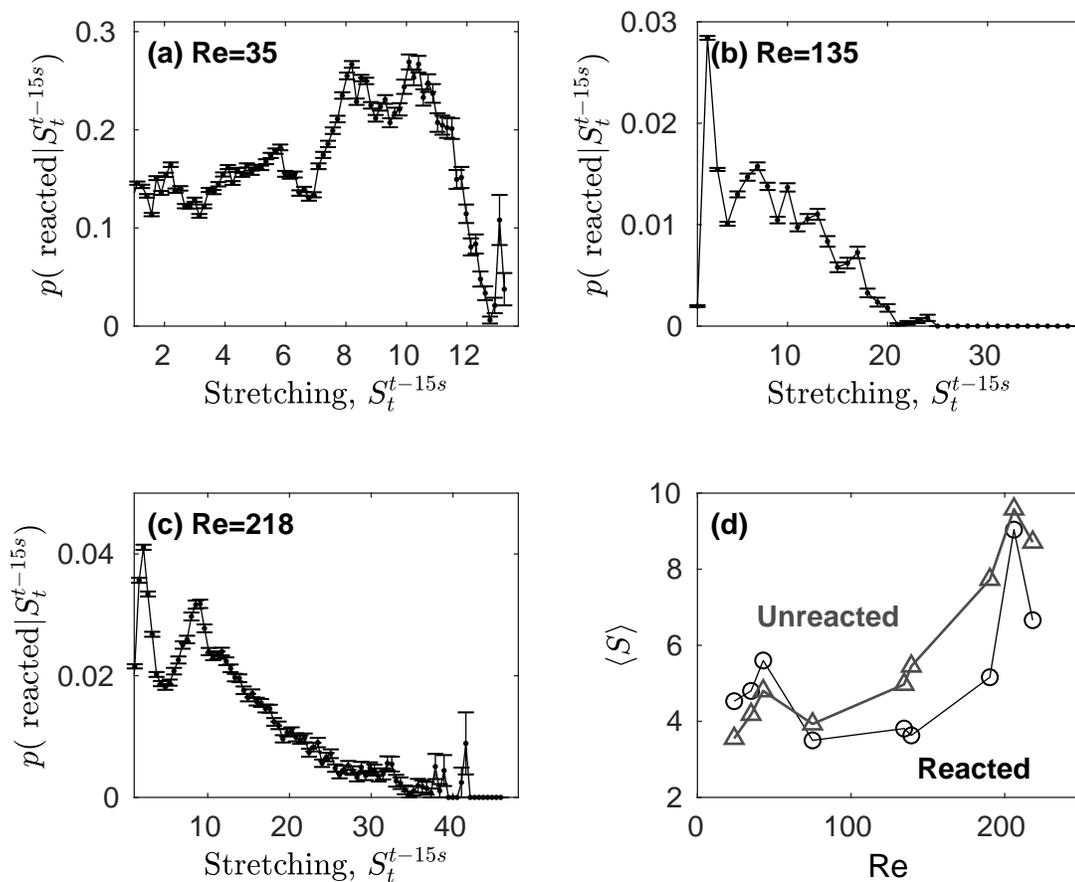


Figure 5.7: Dependence of reaction state on stretching in the single layer vortex flow. (a–c), The probability $p(\text{reacted}|S)$ that a region with a given stretching S will be reacted, at varying Re . Because we plot no data point involving fewer than 100 measurements, the probabilities plotted as zero are well-estimated by a probability of zero. We avoid transients by excluding the first 15 s of each experiment. (d) Average stretching $\langle S \rangle$ in reacted and unreacted regions over nine entire experiments, each at least 500 frames long, at varying Re , error bars are smaller than the plot markers. For $Re < 70$, reacted regions have larger $\langle S \rangle$, but for $Re > 70$, unreacted regions have larger $\langle S \rangle$. Flows with $Re > 100$ vary over time. Work of the author published in [5].

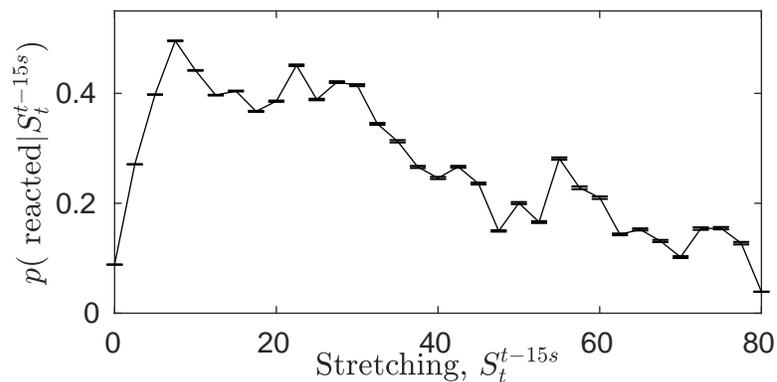


Figure 5.8: Dependence of the reaction state on stretching in the open flow experiment. No data point involving fewer than 100 measurements is plotted. Measurements from the open flow are similar to the single-layer vortex flow. Work of Jinge Wang published in [99].

clear how the depth shear effect will matter. This could be problematic, because strain in the vertical dimension is quite strong, but high vertical straining is correlated with high horizontal straining in these flows. Therefore, the most likely effect of vertical straining is to change the apparent critical stretching, because the vertical straining helps stretch out the reaction. The qualitative behavior should not change though.

Third, the results of Neufeld [53] show that the bifurcation happens when instantaneous strain passes a critical value. Stretching meanwhile has a finite time T over which it is calculated, during which it could be thought of as “adding up” the strain rates experienced by a fluid element. The problem with this can be seen by noting that Neufeld built his simulation with the flow near a hyperbolic point, for which I calculated the stretching analytically to be

$$S_T(\mathbf{x}_0, t) = e^{|T|\gamma} \quad (5.2)$$

where γ is the uniform strain near the hyperbolic point. Now note that no matter the time allowed, Neufeld observed that $\gamma < \gamma_c$ never went extinct while continually getting longer. What this means is that if I were to repeat my analysis on Neufeld’s simulation, I would find that my critical stretching would depend on

the T I chose ($S_b = e^{|T|\gamma_c}$). A different researcher using a larger T would find reaction still present with $S_{T'} > S_b$. S_b is not just dependent on the chemistry the way γ_c is. I have avoided this problem by fixing a constant T , but it does indicate that the true effect driving optimal stretching is an instantaneous effect that needs a certain amount of time to take full effect. As such, deeper mathematical investigation is needed to obtain a better flow quantity.

To handle each of these problems, a deeper understanding of why stretching should matter in the first place is necessary. Previously I have mentioned that Neufeld observed a critical strain rate which caused a transition from a steady reaction filament to extinction, and that stretching could be used as an analogue in more complex systems. It is not clear that squeezing a filament in this manner is the only reason for optimal stretching, and the filament transition does not provide an obvious reason moderate stretching helps a reaction. In the next section, I attempt to put optimal stretching on a more firm theoretical ground by considering three effects that emerge from the spatial derivatives of flow.

5.4 Understanding optimal stretching

To use optimal stretching for prediction or control of ARD systems, it would be useful to have a theoretical understanding of what happens in a stretched reaction that both helps and hurts reaction growth. Our experiments so far indicate that stretching plays a major role. On an intuitive level, this is because stretching relates to diffusion. Stretching relates to how much an initially circular region would be deformed by the fluid flow around it. Thus a large stretching maps a circle to a new shape with a large major axis. In an incompressible flow any material region advected by the flow cannot increase in area due the flow being divergenceless. Since a circle (or sphere in 3D) is the smallest perimeter that can enclose a given area, stretching increases perimeter. This increases availability of

reactants for the reaction, but also increases dilution of catalyst. The blowout can then be understood as when the dilution — via diffusion — outpaces reaction.

While this is a useful qualitative picture, it is not a rigorous explanation of why the Eikonal equation is no longer useful. Put another way: what happens to a reaction front that causes the effect of diffusion to increase, especially to the point where diffusion can reduce concentration below threshold, causing blowout? Understanding the dominant effect behind optimal stretching may someday enable the prediction of the blowout value in any system without having to conduct an experiment. Then from that value, an accurate prediction of the optimal stretching region that will contain most of the reaction growth.

To understand optimal stretching more deeply, notice that perfectly uniform flow cannot have any effect on the total reaction occurring due to Galilean invariance. That is to say, \mathbf{u} cannot change a reaction, but its derivatives in space and time could. Stretching is linked with the spatial derivatives, so we should consider the different ways that the flow gradient tensor, $\nabla\mathbf{u}$ could interact with a reaction's growth. This flow gradient tensor can be decomposed into the symmetric and antisymmetric components,

$$\nabla_i u_j = S_{ij} + \frac{1}{2}R_{ij}$$

where the symmetric part S_{ij} is the strain-rate tensor, and twice the antisymmetric part is the rotation-rate tensor R_{ij} . I will show that the antisymmetric part has no effect on total reaction growth besides displacing the reaction. I will also examine how the symmetric part boosts reaction by both changing reaction geometry, changing the 1D profile of the front, and pulling the reaction into filaments.

5.4.1 Strain-rate drives perimeter change

The first effect of a flow gradient is to change reacted region perimeter. This is a unique effect, because it can change the global reaction rate, without changing

the chemical front speed. To understand why, consider an initial reacted region with a front curve, R , and area A . Within the Eikonal equation, the growth of this reacted region would be

$$\begin{aligned} \frac{dA}{dt} &= \oint_R \mathbf{w} \cdot \hat{\mathbf{n}} dR \\ &= \oint_R (\mathbf{u} \cdot \hat{\mathbf{n}}) dR + \oint_R (v \hat{\mathbf{n}} \cdot \hat{\mathbf{n}}) dR \\ &= \int_A \nabla \cdot \mathbf{u} dR + v \oint_R dR \\ &= vP. \end{aligned}$$

where P is the perimeter, and $\hat{\mathbf{n}}$ is the unit normal to the front. Here using incompressibility I have eliminated the flow component, then shown that the growth is identical to front perimeter times the front speed, assuming the front does not collide with itself. The perimeter P of a material line R in 2D ARD systems changes according to:

$$\frac{dP}{dt} = \oint_R r_i r_j S_{ij} dR. \quad (5.3)$$

Here the r_i and r_j are unit tangent vectors written in Einstein notation. This extends the textbook equation for change in a infinitesimal material line to a finite, closed material line. When fronts are not self-intersecting and have zero front speed, this equation is exact. This result shows that faster global reaction (as given by growth of a total reaction area) does not imply a faster chemical front speed. It also indicates why strain-rate — and by extension stretching — would support reaction growth. A larger perimeter has more access to fresh reactant. This is shown by an increase in P leading to an increase in A . The question still remains though: How might front speed be changed by flow gradient?

5.4.2 Strain-rate modifies reaction front profile

Perimeter change is the simplest of possible flow effects on reaction growth, because it holds as an effect even if chemical front speed v is constant. The second

and third effects of the flow gradient can change v . Consider the one dimensionalization of the reaction diffusion equation in chapter 2. Taking into account the overall front shape in a second and third dimension just provides the curvature effect [41]. I could solve the 1D equation and obtain a front profile that moved at a constant speed. To disrupt the constant front speed, flow must do something to invalidate this solution. One way to do this is for the front to modify the 1D reaction profile of the front, because this alters how catalyst spreads. Flow gradient provides a forcing that changes this profile. To see how flow modifies front profile, consider the equation for the reaction gradient as it is advected, which can be obtained by taking the gradient of eq. 3.1

$$\frac{\partial g_i}{\partial t} + (\nabla_i u_j) g_j + u_j \nabla_j g_i = D \nabla^2 g_i + \alpha g_i \frac{dG(c)}{dc}. \quad (5.4)$$

where $g_i = \nabla_i c$ is the gradient of the reaction field. Now notice that the terms on the right come from diffusion and reaction, so they would create the stagnant profile without flow. The third term is an advection term, but now the quantity being advected is the gradient field. Only the second term here will change the chemical gradient differently away from the stagnant front profile. The second term acts as though the flow gradient squeezes or stretches the front causing the front gradient to get steeper or flatter. This changes how much space is required to go from reacted to unreacted, or in other words, this changes the front thickness.

Since the dynamics of the reaction along the one dimensional cross section lose all information about direction, we really care about what happens to the front gradient's magnitude. Using the fact that $d|\mathbf{x}|/dt = \hat{\mathbf{x}} \cdot d\mathbf{x}/dt$ and substituting eq. 5.4, I can write the gradient of the front over time as

$$\frac{\partial |g|}{\partial t} = \hat{g}_i [- (\nabla_i u_j) g_j + (\dots)]$$

where I have gathered all the terms not of interest at the right. Now consider just the rotation-rate tensor portion of the flow gradient tensor. This term would

appear in the above equation as

$$\frac{\partial|g|}{\partial t} = -\frac{1}{2|g|} (g_i R_{ij} g_j) + (\dots)$$

but this rotation-rate term is always zero because the rotation-rate tensor is anti-symmetric. Therefore, I expect that rotation has no effect on front thickness, but does have an effect on the front angle. Once again I am left with strain-rate as the only term which could affect front growth.

To see that rotation rate has no effect both on reaction perimeter and front profile, see simulation results in fig. 5.9. These show that as rotation rate is increased, there is no corresponding change in global average reaction state. In this simulation I simulate a strain-less vortex with rotation-rate constant Ω where

$$\mathbf{u} = \Omega (-y\hat{\mathbf{x}} + x\hat{\mathbf{y}}). \quad (5.5)$$

using the full ARD equation with an Oregonator driven reaction (eq. 2.16), with Eulerian time stepping and central differencing. An initially reacted circular region slightly offset from the center of the vortex turns in space counterclockwise, but continues to grow outward at the same speed as if it were not moving at all. This is interesting, because this is a case where the system is not a Galilean transformation away from stagnant front growth, but it still acts as though it is.

Strain rate on the other hand can change the magnitude of the reaction front profile. To study its effects, I performed simulations of a reaction growing in a uniform strain,

$$\mathbf{u} = \gamma (-x\hat{\mathbf{x}} + y\hat{\mathbf{y}}) \quad (5.6)$$

where γ is the strain rate. By making the initial condition an infinite vertical strip of reacted region, the system becomes independent of y at all times. This matches how stretching tends to elongate initial reacted regions into filaments which are much narrower than they are long, especially reactions in areas like the ones providing blowout in my experiments. This setup was originally used by Neufeld

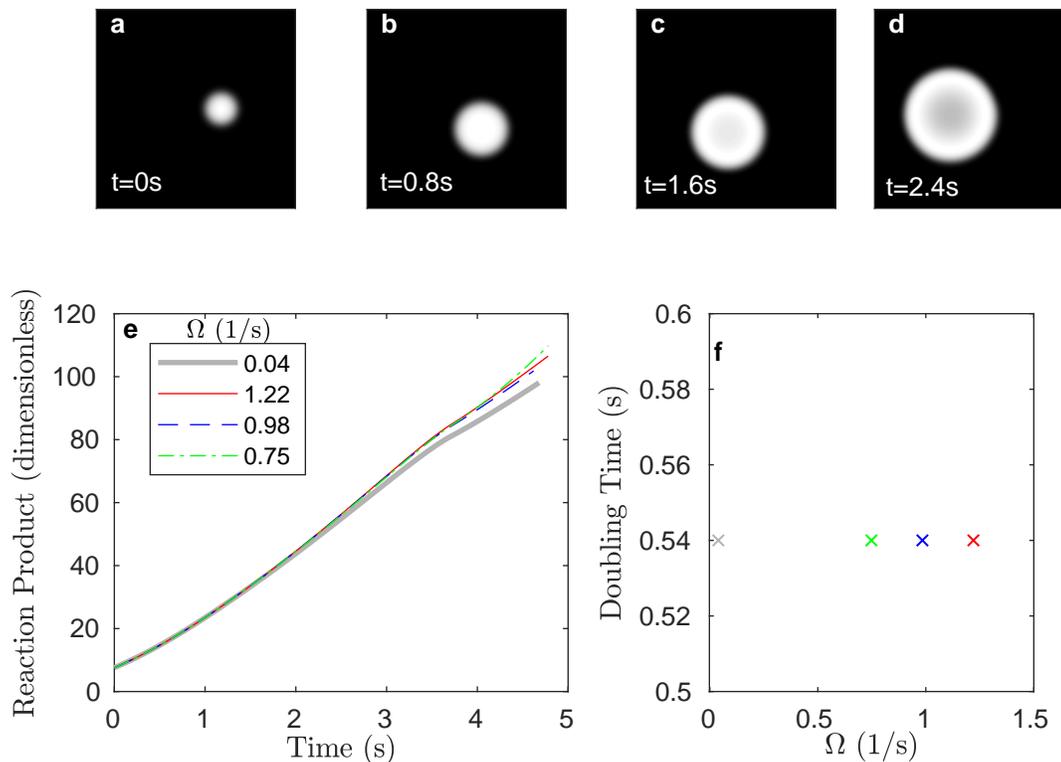


Figure 5.9: Results of simulations of the Oregonator growing in a strainless vortex. Top row (a–d) is snapshots from the highest rotation rate data. I find that to within numeric imprecision of the simulation, the global amount of reaction product is independent of the strength of rotation (e). This can also be seen by the constant time it takes to double reaction product with Ω (f). Unpublished work of the author.

[53]. By simulating this 1D cross section, I can reproduce the dynamics of a real 2D filament. For the Oregonator this has governing equation,

$$\frac{\partial c_1}{\partial t} - \gamma x \frac{\partial c_1}{\partial x} = D \frac{\partial^2 c}{\partial x^2} + \alpha \left[c_1(1 - c_1) - \frac{c_1 - q}{q + c_1} f c_2 \right] \quad (5.7)$$

$$\frac{\partial c_1}{\partial t} - \gamma x \frac{\partial c_1}{\partial x} = D \frac{\partial^2 c}{\partial x^2} + \alpha \epsilon [c_1 - c_2] \quad (5.8)$$

where c_1 and c_2 are the activator and inhibitor species of the Oregonator model respectively, and $q = 9 \times 10^{-5}$, $\epsilon = 0.01$, and $f = 3$ which are constants to match the experiment. This can then be nondimensionalized as

$$\frac{\partial c_1}{\partial \tilde{t}} - \tilde{x} \frac{\partial c_1}{\partial \tilde{x}} = \frac{\partial^2 c}{\partial \tilde{x}^2} + Da_I \left[c_1(1 - c_1) - \frac{c_1 - q}{q + c_1} f c_2 \right] \quad (5.9)$$

$$\frac{\partial c_1}{\partial \tilde{t}} - \tilde{x} \frac{\partial c_1}{\partial \tilde{x}} = \frac{\partial^2 c}{\partial \tilde{x}^2} + Da_I \epsilon [c_1 - c_2] \quad (5.10)$$

where I have taken $\tilde{t} = \gamma t$ and $\tilde{x} = x \sqrt{\gamma/D}$. In this system, strain controls $Da_I = \alpha/\gamma$.

I simulated the dimensionless ARD equation (eq. 5.9) using Euler's method in time, and a central differencing scheme in space at various values of $Da_I = \alpha/\gamma$. Time stepping in the simulation is determined based on the input spatial resolution squared times 0.15 to provide some simulation stability. Half filaments are simulated to shorten computation time, because the system is symmetric across the y axis. I have observed that numeric dispersion and diffusion are reduced using these methods, but that they are still present. The effect is negligible far from a critical $Da_I = Da_{Ic}$. At this critical value the filament transitions from a stable solution to blowout which I will investigate in more detail below. Simulations near this critical point take a very long time to reach their final state, allowing the nonlinear reaction term to grow the numeric dispersion effects. In a given numeric test, I choose an initial condition in dimensional units before converting it to dimensionless units. All the information about concentrations and flow fields are available to me in this simulation, so detailed investigations of transient behaviors are possible.

Reactions in the simulations have a front normal of $\mathbf{g} = -g\hat{\mathbf{x}}$, and a constant strain rate tensor with diagonal elements equal to $\pm\gamma$. Therefore the flow gradient effects from eq. 5.4 on the gradient is

$$\frac{\partial|g|}{\partial t} = |g|\gamma + [(\dots)]. \quad (5.11)$$

The front profile will be steepened by the straining flow. In fig. 5.10 I show simulations of a reaction filament where a large reacted region exists at the start along with an Eikonal front initiated at the same starting position. I find that while the two look qualitatively similar (up to blowout), the actual measurement of front speed (at $c_1 = 0.5$) is very different than the expected constant front speed. This speed difference occurs for any Da_I . However, as strain rate increases, so does the deviation *away from* the stagnant front speed.

The two simulations in fig. 5.10 are identical in the starting *dimensional* profile. I have also simulated front profiles that are steeper or more gradual in the $Da_I = 2.1$ case. I found that the gradual profile started with front speed greater than v_0 , and the steep profile started with front speed less than v_0 . Both quickly converged to the same asymptotic front speed as shown. From this it is clear front profile is closely linked with front speed. BZ has diffusion time much longer than reaction time, so a steep front limits the available area for reaction and slows the front. Wide fronts expand this area, and the front moves faster. In this situation the flow is converging, so it steepens and slows fronts. If I allow the front any orientation in 2D, then $\mathbf{g} = (|g|\cos\theta, |g|\sin\theta)$. θ is the angle the gradient makes with the direction of converging flow (in this case the x axis), then the flow gradient changes according to,

$$\frac{\partial|g|}{\partial t} = \gamma g(\cos^2\theta - \sin^2\theta) + [(\dots)]. \quad (5.12)$$

The minus sign on the sine term here means that fronts tend to get more gradual along the diverging direction (my y axis). Thus I expect that while fronts pointed in a converging flow direction will steepen, those in an expanding direction will

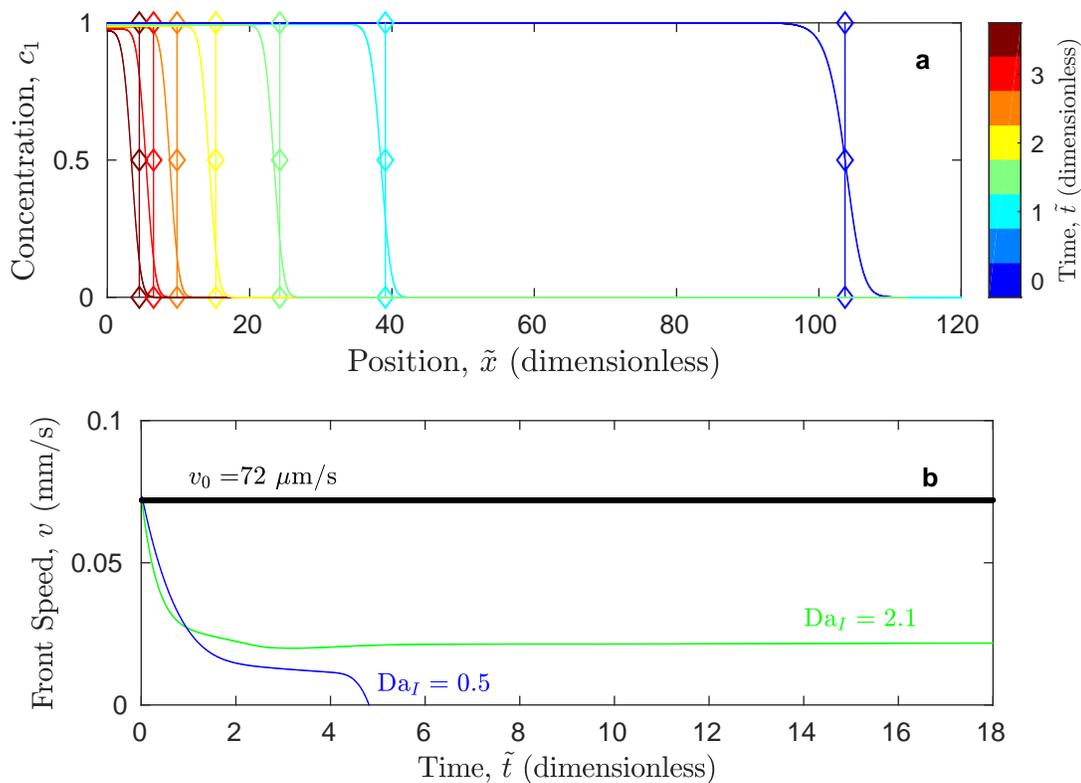


Figure 5.10: Comparison of full ARD and Eikonal equation in the uniform strain situation. (a) Vertical lines with diamonds indicate Eikonal approximation, and curves indicate the ARD result in a $Da_I = 0.5$ situation which eventually blows out. Qualitatively the agreement looks good, but actual calculation of front speed (b) shows significant deviation between the speed of an Eikonal front and the true front speed. The front speed deviation is also present in a $Da_I = 2.1$ simulation, which reaches a nonzero steady state. Blowout transition is shown by a sudden negative spike of front speed. Unpublished work of the author.

widen.

Strain-rate causing a change in front profile does not by itself indicate that shear is changing front speed. In fig. 5.11 I plot a number of simulations that are eventually stable with $1.5 \leq Da_I \leq 3$, and all initiated to match the same dimensional condition. I measure chemical front speed at a threshold of 50% reacted, just as it would be measured in experiments. Front thickness — as defined in chapter 2 — is also measured to provide quantification of the gradient. Front speed and thickness are not constant with strain, but take some finite time to converge to a new stable point. The interesting result occurs when the front speed is plotted versus the thickness. The speed at any given time is linearly related to the thickness at that time. These numerical results indicate the equation for front speed, v , given its thickness, τ , is $v = 1.1(\text{s}^{-1})\tau - 5.4(\mu\text{m/s})$. The plots only deviate from this speed when the filament gets thin enough for boundary effects to matter. Notice that fronts that have very different Da_I can have the same front speed if their front thickness is equivalent. The agreement shows that it is not flow detail that changes front speed, but flow detail changes front profile, which in turn changes front speed.

This result has large potential importance in the dynamic modeling of reaction fronts. Firstly, this shows that fluid flow *can* change chemical front speed, and that a mechanism for this change is strain-rate altering reaction front thickness. Second, strain is normally assumed to be negligible to derive the Eikonal equation in flow [43], but these results indicate a way to model some of strain's effects. For instance, this could be done by having chemical speed change in the system by Mitchell and Mahoney [44] as,

$$v = A\tau + B, \quad (5.13)$$

where A and B are empirical constants which model how a given reaction front's velocity changes due thickness. These constants likely depend on specific chemistry. There remains a problem in predicting reaction front growth from the fluid

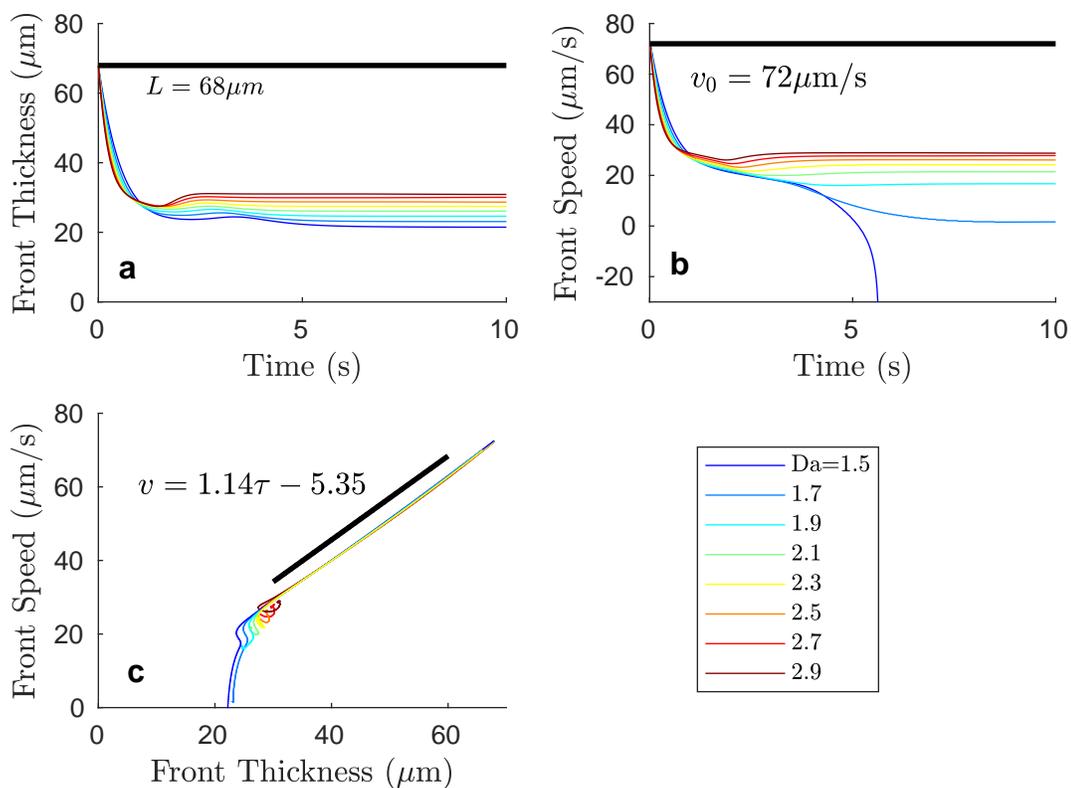


Figure 5.11: Trends of front speed and front thickness in uniform strain simulations at Da_I from 1.5 to 3. (a) Starting from the stagnant profile, strain causes the fronts to narrow, with stronger strain resulting in thinner fronts. (b) Similarly, stronger strain results in slower front speeds. (c) Plotting front speed versus thickness shows a collapse onto a single line of $v = 1.1(\text{s}^{-1})\tau - 5.4(\mu\text{m/s})$ for all Da_I . Unpublished work of the author.

flow. So far I have only found the reaction growth from the reaction front thickness. The equation that controls τ is still needed. These results indicate that such a dynamic equation will depend on strain-rate, but not rotation-rate. Here I also have not explicitly tested shear rate, which also appears in the strain-rate tensor, so future work is required to understand if it behaves differently. Eq. 5.12 gives a calculation for the gradient in this flow, but this is not the whole story, because diffusion and reaction (which is nonlinear) also matter, and much more complex strain fields may be possible. Fronts in flows are typically advected such that the strain they experience is always changing, so steady situations like this serve only as a simplified problem. Therefore it would be interesting for future researchers to determine a dynamical equation for how τ is changed in an arbitrary straining flow over time, and test a time dependent strain. If this could be done, the Eikonal equation could be expanded to greater accuracy in faster mixing flows with only a small increase in the simulation complexity.

In the following section I will investigate how strain rate also alters the usual solution to the one dimensional profile by changing the boundary conditions. Specifically, strain can pull an initial reacted region into a thin filament, where the distance across the filament is thin enough that the reaction no longer reaches fully reacted and strain sweeps away growth due to diffusion faster than it can be replenished leading to the blowout transition.

5.4.3 Straining reaction into thin filaments drives reaction blowout

The third and final effect of flow gradient emerges, because I don't always have the reaction being reacted out to infinity on one side, and unreacted to infinity on the other. Which is to say, the boundary conditions can change. In real situations, a reaction could be spatially isolated, or a reacted region could encounter another

reacted region. I will focus on an isolated regions and note that since the boundary has changed, the 1D solution can change, and therefore front speed can change.

Fluid flow is capable of compressing a region into a thin filament to bring these boundary effects into play, as shown in the simulations by Neufeld [53]. Due to flow incompressibility, any elongation must be accompanied by this compression, so the front elongation in section 5.4.1 is not without cost. During perimeter elongation, the compressed direction gets smaller, and reaction product always diffuses outward to reduce the concentration of the filament. This is related to KiSS theory for the minimum habitat-size problem in phytoplankton populations [18, 24, 25, 52]. Kierstead and Slobodkin [24], Skellam [25] observed that if an initial patch of plankton growing linearly without limit is constrained by some habitat of lateral size l outside which the plankton immediately disappear, then if $l < l_c = \pi\sqrt{D/\alpha}$ the population tends to zero at infinite time. Martin [18] expanded this idea for when the plankton are stretched by a straining flow. In flow plankton outside the initial zone are swept away faster than those at the center, so that the population growth along the 1D profile must compete with the straining flow. In this case, the population trended to zero at infinite time if strain rate exceeds reaction rate, $\gamma > \alpha$. Strain which pulls reacted regions into filaments can cause extinction if the strain is able to outcompete reaction production.

The problem with this theoretical development is that in my experiments I have neither infinite time or infinite space to stretch the filament. In a closed flow domain, flow cannot remove a reaction product, only relocate it. A reaction of first, second, or third order, or any reaction where the unreacted state is unstable, will take infinite time to reduce to zero, so the reaction filament will expand outside the stretching zone before extinction. Since neither advection, nor diffusion, nor reaction can decrease global reaction product in an enclosed volume, the global total reaction product cannot decrease. Blowout would be impossible. But a

dish of BZ with a small reacted region stirred vigorously enough will return to unreacted nonetheless. This indicates that BZ must have a stable unreacted state and a threshold reaction level above which the reaction grows to fully reacted, and below which the reaction returns to unreacted. Other types of reaction with this quality include other excitable reactions like BZ, bistable reactions, combustion reactions (where often the catalyst's role is played by temperature), and more. In the case of BZ, this threshold is very slightly above the unreacted stable state 1.2, but it is present. Since the unreacted state is stable, the reaction term of the ARD equation can decrease global concentration. This tiny threshold allows BZ to go through blowout even though it is contained in a dish, and in finite time, so that I can measure it.

To investigate the relevance of the filament model of reaction blowout to my experiments, I repeated the simulations of Neufeld [53]. These are the same simulations as shown for strain rate in section 5.4.2, but now I will look for the critical strain γ_c and characterize the behavior of filaments which have $\gamma > \gamma_c$ to determine what stretching level should lead to reaction blowout. For these situations I use the same simulation methods, but find that the physical dimensionless initial conditions are usually too large to simulate well. Thus I treat $D = 1$ unit of diffusion distance squared per second. Thus the results are dimensional, but no longer mm. With BZ's diffusivity this works out to about 1 diffusion unit = 31.6 μm . That this is much smaller should not matter since strain quickly shrinks any reacted region anyways.

In fig. 5.12 I show simulations at various times to demonstrate the effect. Straining flow is inward from the right, so the concentration is pushed left over time. If $Da_I < Da_{Ic}$ (in other words, $\gamma > \gamma_c$) then we have the situation in fig. 5.12a. The maximum reaction intensity drops below the reaction threshold in a finite amount of time. Even though the filament is getting longer, the entire filament will die out since it is zero across the width, this is what I call blowout.

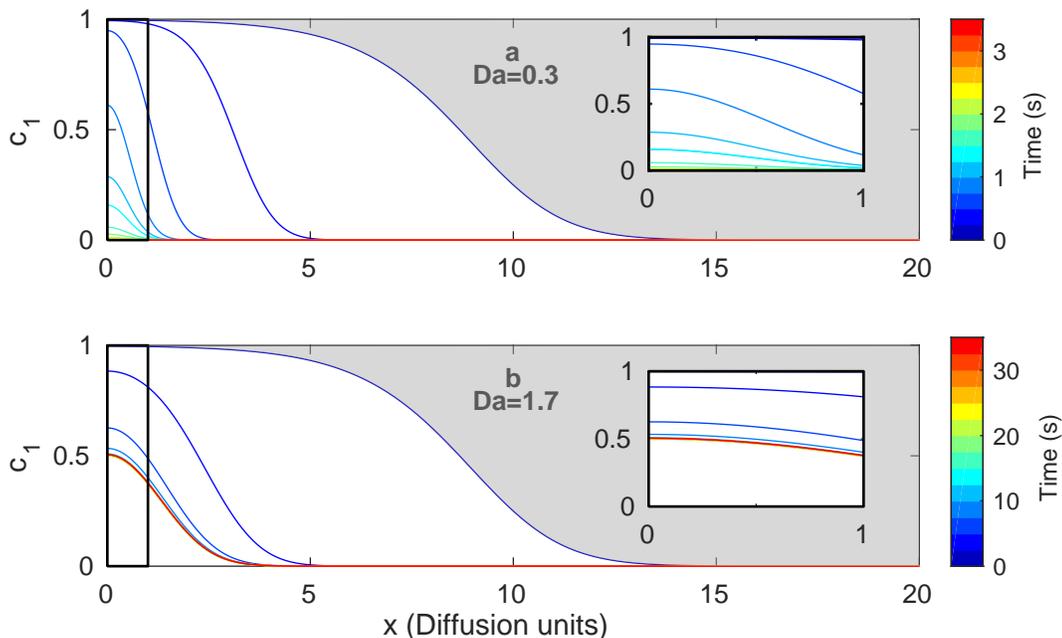


Figure 5.12: Results of two simulations of the Oregonator growing in a hyperbolic strain in 1D. Both simulations start with the same initial condition in dimensional units, but experience different strain. In (a) $Da_{Ic} > Da_I = 0.3$ so the reaction undergoes blowout in a finite time of about 3.25 s. In (b) $Da_{Ic} < Da_I = 1.7$ so the reaction stabilizes. Insets provide a closer look at the time behavior near the origin. Unpublished work of the author.

If $Da_I > Da_{Ic}$ (or $\gamma < \gamma_c$) the solution is stable as in fig. 5.12b, and therefore the reaction filament will grow over time, because it reaches a stable state that does not get thinner while it keeps getting longer.

According to KiSS theory a linear reaction is blown out by $\gamma > \alpha$ [18, 52]. The growth rate of the Oregonator $F_1(c_1, c_2)$ can be shown to be always less than a linear reaction with the same reaction rate, α . Therefore, the Oregonator can't have $Da_{Ic} < 1$, since that would imply the Oregonator resisted strain flow better than linear. Fig. 5.13(a) shows the trends of Oregonator reaction filaments with Da_I by showing the reaction level at $x = 0$ over time. At a certain Da_{Ic} this maximum reaction state switches from being stable over time to decaying to zero,

indicating reaction blowout. From these simulations, I estimate that $Da_{Ic} \approx 1$, or $\gamma_c = \alpha$. The apparent resilience of the Oregonator to strain is very different from the behavior of the Fitz-Hugh Nagumo studied in Neufeld [53], which had a higher Da_{Ic} . This may be because the Oregonator is approximately second order, while Fitz-Hugh Nagumo is third order, and so the Oregonator is stronger near the threshold. That I can identify a transition between optimal stretching and blowout in this filament model based on strain is good evidence that this is the process I am seeing in the experiments when I compare reaction state to stretching. Even better evidence is the fact that if I assume the depth shear is the dominant source of strain-rate in these experiments, then I get $\gamma = \alpha = 1.3$ 1/s when I experiment at $Re = 41$. This is of the same order of magnitude as results of fig. 5.7(d). Higher strain rates may be required in experiments due to the presence of layer boundaries which the reaction can't diffuse through.

Fig. 5.13(b) shows a problem with using stretching to predict blowout. From the experiments I expect that reaction filaments will blow out when they reach some specific S_b . From eq. 5.2 the time being stretched and γ uniquely determine stretching S , so I expect that the time to blowout should change such that as γ increases the blowout time decreases, so the product stays constant. Therefore, I have proven that in this simple situation there is no single critical stretching S_b , but there is a critical strain γ_c , which causes blowout after some time. The constant stretching time used in experiments however would lead to an apparent critical stretching, where all $\gamma > \gamma^* > \gamma_c$ are strong enough to blow out the reaction in less than the stretching time T . Therefore S_b is a good half measure: it tracks the Lagrangian strain effects which do result in blowout, and it has a quantifiable transition that depends only on the reaction, not the flow structure, but the actual value of S_b is dependent on user choice of T . So stretching is not the whole story to explain blowout, but it is a reasonable way to estimate blowout locations, and is very close to the flow quantity that is truly most relevant.

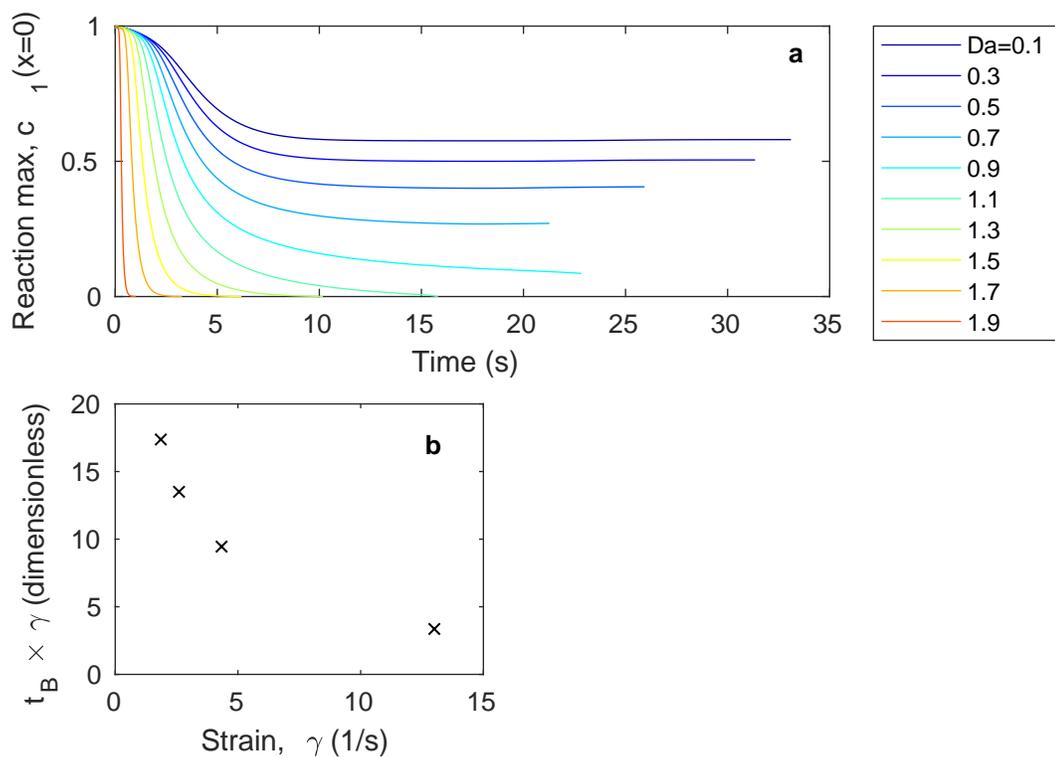


Figure 5.13: Plots of simulation results near Da_{Ic} . (a) Each color is a different Da_I , showing the trends of the maximum reaction over time. This shows a sudden switch at $Da_{Ic} \approx 1$. (b) Time it takes to reach the KiSS radius $l = \pi\sqrt{D/\alpha}$ times γ is not constant as expected if a specific S is needed to cause blowout. Unpublished work of the author.

An open question then is: how could a researcher develop a better Lagrangian quantity which could track blowout more accurately, without an arbitrary stretching time? According to Martin [18] time behavior of linear reactions will go roughly as $\exp(t(\alpha - \gamma))$. Thus a system with strain below α could have an infinitely high stretching with enough time, and yet be below the critical strain. Therefore any improvement to optimal stretching should incorporate the reaction rate constant α .

While these simulations are very simple, being only 1D and in uniform flow they are powerful. Because they are 1D, they focus on the local behavior of the reaction, which is universal for reactive mixing. A 1D profile can always be drawn from a reaction level surface. Growth of a reaction front perimeter is due to strain-rate, so reactions are often spread through filaments. An initial reaction is stretched in filaments which are then folded over each other until they are close enough to interact. However, as strain is increased the narrow dimension of the filament loses front speed. This may be because our reaction is diffusion limited, and therefore the front thickness determines the chemical front speed. A more complex system modeling the development of front thickness τ is required, but nonetheless, a modification of the Eikonal equation should be possible to account for this effect. The strained front which moves slower than normal is more easily squeezed into a thin filament and it can then blow out if the strain rate exceeds the reaction rate. Since this can all be developed from a local picture of a reaction front, this could happen in any ARD system. All three effects are rooted in the interaction of reaction, diffusion, and strain-rate over a finite period of time (stretching). Therefore these are the causes of optimal stretching I observed across different experimental setups. A little stretching helps growth through an increase in perimeter and boosted front speed along spreading directions, but larger stretching slows the front speed in the compressed direction until the filament grows thin enough that the front dies due to diffusive losses

faster than reaction can replenish.

5.5 Summary and Implications

I have observed that reacted regions primarily occupy vortex edges for small values of Re , but primarily occupy vortex cores for large values of Re . I found that as Re increases, reaction fronts move more quickly toward vortex cores. This is consistent with ideas about flow enhancing reaction growth by enhancing access to reactant. I then focused on the disappearance of reaction from vortex edges, and explained this with stretching of material volumes, specifically that for the BZ, there exists an optimal range of S for enhancing reaction. Just as a flame grows in moderate winds but is blown out by high winds, moderate stretching promotes reaction by supplying fresh reactants, but strong stretching inhibits excitable reactions by diluting catalyst. To determine the root of this I performed a series of analytic and numerical results to study different ways the flow gradient tensor could influence the growth of a reaction. Perimeter growth due to strain-rate drives enhancements to reaction growth, which causes moderate stretching to help reaction growth more than small stretching. Front thickness is modified by strain-rate, but not rotation-rate, and front thickness in turn directly modifies reaction front speed. Finally, I examined how reaction compression into filaments is causing the blowout effect, and made predictions about how much straining is needed to cause blowout.

Though the regions of strong stretching, where $S > S_B$, account for only a small fraction of the total observation region, they make large regions inaccessible to the reaction by blocking its spread. In large- Re flows, I observe that the reaction is often limited to a few vortices near the trigger point. A barrier region with $S > S_B$ forms across several vortex cells, and may even surround a few vortices. In small- Re flows, however the barriers are missing, and the initial reaction

spreads throughout. Blowout could therefore represent a new and important type of barrier to reactive mixing, in addition to the barriers predicted by theories of burning invariant manifolds (BIMs) and burning Lagrangian Coherent Structures (bLCS) recently developed by Mitchell, Mahoney, and Solomon [44, 46, 49–51]. Both BIMs and bLCS are calculated using a constant reaction front speed and the flow \mathbf{u} . Those quantities alone are sufficient for locating barriers arising from interactions between front speed and advection speed — roughly, where front speed is matched or exceeded by headwinds. Their power and simplicity are elegant. However, stagnant front speed is the only information about chemical kinetics encoded in BIMs and bLCS; the excitation threshold is absent. Thus blowout would represent a new sort of barrier.

Blowout barriers likely occur in natural and industrial ARD systems outside the laboratory. For example, phytoplankton blooms may be blown out by fast ocean currents, and blowout regions may form barriers that segment the ocean. While plankton blooms depend on many conditions, they have been shown to be well-modeled as excitable ARD systems [22, 23] and are therefore likely to experience blowout. Blowout may play a role in forming the dynamical niches recently observed [27] to promote phytoplankton diversity in Earth’s oceans. Flames in many kinds of combustion are known to experience extinction under high strain [107]. These results open up the possibility of a simplified model of reaction extinction, which need not rely on empirical testing. Stretching is present in all of these situations, and the mechanisms I have outlined combine to cause optimal stretching which appears to be a very important property of these chemical systems independent of the flow details.

Based on the analysis in section 5.4, it is clear that strain-rate of a flow is important to the growth of a chemical reaction. Since it takes some finite amount of time for this strain rate to have an effect, and in experiments the reaction is swept along by the flow, Lagrangian stretching is a reasonable quantity to understand

eventual structures of a reactive scalar. The strain rate a reaction experiences may vary as it is carried, so determining how to measure a more relevant Lagrangian stretching for reactions, and determining front profile dependence on strain-rate remain open questions. Time dependence of the strain-rate can be especially noticeable when flows become unsteady, and vortices start to combine. Understanding all of these effects will be key to understanding reactions in turbulent systems. The results of section 5.4 also indicate that a key feature of blowout in real systems is the presence of an excitation threshold, above which the reaction takes a long excursion through phase space, but below which the system quickly returns to equilibrium. The thresholding of the reaction enables blowout to happen in a finite time, and at local positions. Higher stretching is correlated with increasing perimeter of a reacted region, but because stretching makes thinner structures diffusion can deplete the bulk below the threshold. With too much thinning and dilution, the excitable threshold is crossed in reverse, and blowout occurs. Determining how critical strain relates to critical stretching still requires some work in both theory and experiment, but once blowout is well understood in one system, it will apply in general because it is not a function of flow details. Understanding optimal stretching may someday allow researchers to predict likely reaction structures from only reaction kinetics, regardless of flow details.

5.6 Acknowledgements

This chapter is a reprint of material in “Optimal Stretching in Advection-Reaction-Diffusion Systems” which was printed in the journal *Physical Review Letters*[5], with inclusions reproduced from J. Wang, J. Tithof, T.D. Nevins, R.O. Colon, and D.H. Kelley, “Front tracking velocimetry in advection-reaction-diffusion systems”, *Chaos* **27**, 123109 (2017), with the permission of AIP Publishing, and some unpublished work. These papers have been modified to fit in this thesis, and in-

termixed to fill in some gaps between each. The thesis author is the primary author on material in this chapter, and he was responsible for the experiments in vortex flows, developing the idea of optimal stretching and blowout, all the simulation and analytic results in section 5.4, writing most paragraphs in this chapter, and most preparation of most of the figures. Co-author Jinge Wang was the primary author on the *Chaos* paper, and was responsible for the experiments and paragraphs describing the flow past a bluff body experiments. Co-author Douglas Kelley provided guidance during the experiments, provided guidance on figure formation, and edited the paper and figures. Others who assisted in gathering these results include J. Tithof, R. Colon, and B. Knisely. Insightful conversations with T. H. Solomon, J. Gore, R. F. Ashour, and K. Vats also contributed to this chapter. T. D. Nevins was supported by the Department of Defense through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program for the results of section 5.4.

6 Conclusion

Research into advection-reaction-diffusion (ARD) systems has a long history and a large body of existing literature, enough to fill several books [2, 7, 52, 118]. But number of results does not indicate how well a field is studied; number of remaining questions does. Prediction of a reactive mixing system tends to involve expensive direct numerical simulations [8, 31, 32, 38], which can be difficult to interpret. Control of ARD systems is an even farther goal, because these direct simulations have not provided insight into the dominant effects on the reaction. Even assuming that the flow is given and unaffected by product concentration, intuition about the transient dynamics of an ARD system remains elusive. In this thesis I have attempted to enable simplified modeling methods for ARD systems. For a simplified model to be useful it must (1) be simpler than the full system, (2) provide insight about what is most important to the problem, and (3) provide accurate predictions of real world behavior.

In chapters 2 thru 4 I specifically investigated a reduced dimensional modeling. This reduced dimensional modeling is given by replacing the full partial differential ARD equation with an ordinary differential equation, known as the Eikonal equation [41, 43, 44]. The Eikonal equation models a field of reaction state undergoing diffusion, nonlinear reaction, and advection (which is in general infinite dimensional) as a material line with an outward growth term (using motion in

only 3 dimensions: 2 space, 1 orientation). In this way, the Eikonal equation is clearly simpler, and it provides insight about reactions in flows. The Eikonal equation assumes that all regions are irrelevant except for a small transition region known as the front, and even within that region the reaction and diffusion terms act separately from the advection. This leads to a constant front velocity which is then added to the local flow speed to get the overall dynamics of the reaction front. In this model the primary effect of advection is to carry the reaction around and change its shape.

What the Eikonal equation was lacking was a test of its accuracy in real world situations. In particular, measurements which could compare a real experiment to the Eikonal equation were necessary. Interestingly, since front speed does not appear in the ARD equation anywhere, direct simulations of ARD systems also lack a way to easily measure local chemical front speed, so even in simulations local front speed is not measured. In chapter 2, I began developing the methods to measure a local front speed, by first working without advection. To measure front speed, only the perpendicular direction matters, because tangent to the front is along the level curve. In 1D without flow, the ARD equation has known solutions for second and third order reactions [2, 39]. These solutions indicate that there are only two parameters needed to describe the front: front speed and thickness. These two parameters describe the system well enough that measuring front speed and thickness can be used to back out reaction rate and diffusion. I showed that automated computer measurements of front speed and thickness were possible by use of image thresholding to find reaction fronts, and the combination of curve fitting and interpolation to find chemical speed and thickness. Front speed and thickness measurements using this algorithm matched expected values in simulations and experiments, and the diffusion and reaction rates backed out also matched expected values.

In chapter 3, I expanded this method to flows. An initial front is advected

forward using a measured flow field before the algorithm of chapter 2 was applied. This works off the assumption that the original front's material line is going to be carried to a new position by the fluid flow. In this way, the front tracking with flow mimics the behavior of the Eikonal equation, except that it measures chemical speed instead of assuming that it is constant. A measurement of constant chemical reaction speed with flow speed then indicates an agreement with the Eikonal equation. The method showed good agreement with simulations, indicating that it works as intended. However, the Eikonal equation did not have good agreement with experiments, instead showing that as flow speed increased so did front speed. This trend occurred at all values of Re , so they could not be avoided by restricting experiments to some parameter range.

In chapter 4, I identified why the experiments in chapter 3 disagreed with the Eikonal equation. Getting a good understanding of real world systems through simulation or experiment does not just require good analysis, it also requires good simulation or experiment. The problems experienced in chapter 3 are due to my experimental setup. While flow velocity is well modeled as quasi-2D in thin layer flows, this is partly because flow is tracked by tracer particles at the layer surface. Reactions on the other hand must be depth averaged to measure in a quasi-2D situation. This wouldn't matter if the reaction field were independent of depth. Shear within the reaction layer ensures that this is not the case, and the difference between the apparent front location and the surface advection yields an apparent chemical speed far above the usual chemical speed. I studied the dependency of apparent chemical speed on flow profile for three different experiment types: single-layer magnetohydrodynamically driven, Hele-Shaw, and two-layer magnetohydrodynamically driven. Simulations assuming the Eikonal equation was true in unidirectional flow revealed that the front shape in each case converges to a final shape which moves at the same speed as the leading edge of the layer. Even though these simulations were very simple, and relied only on the Eikonal

equation, they matched very closely with experimental results as long as the flow experienced by the reaction was treated as the flow component along the front normal ($\mathbf{u} \cdot \hat{\mathbf{n}}$). The two-layer system which drastically lowers depth shear also drastically lowered the apparent front speed. Apparent front profile thickness was also found to vary with flow speed due to depth shear, but I have not investigated this in detail.

Based on these three results I was able to conclude that the Eikonal equation could provide a reasonable model of reactions growing in fluid flows, that the algorithm we designed for tracking reaction fronts was working, and that typical quasi-2D reactive mixing experiments were subject to depth shear and therefore should be done using a two-layer system. However, there remained some phenomena that still went beyond the Eikonal equation.

In chapter 5 I investigated these optimal stretching phenomena which were ubiquitous, being present in a variety of flows, and which could be related to behaviors like quenching in combustion [107], and extinction of plankton [19, 23]. In my experiments I observed that the Belousov-Zhabotinsky reaction tended to gather at vortex edges in low Re flows, and at vortex cores in high Re flows. In fact, it avoided reaction edges so much in fast Re flows that the vortex edges became reaction barriers. I was able to identify the cause of these behaviors as being a combination of an enhanced inward radial speed in faster fluid flows, and the optimal stretching for reaction growth. Optimal stretching meant that small stretching did help the reaction grow, but a stretching that was too high led to reaction blowout, which the Eikonal equation could not predict. The optimal stretching range was independent of the flow, so in faster flows the optimal stretching sat in the vortex cores, while in slow flows it was on vortex edges. The same optimal stretching value appeared when Jinge studied the completely different flow, the wake of a bluff body [99]. This all indicates that the optimal stretching is a property of the chemical reaction, and it will be the same for the

reaction in any setting. While previous authors had identified optimal flow for global reaction growth [53, 107] the identification of the Lagrangian stretching field had never been used to identify the same optimal flow at a local level. Optimal stretching appears to serve as a way to simplify reactive mixing systems, which accurately reproduces experimental observations. Estimating the stretching gives a good feel for whether or not a reaction will blowout. Since understanding how stretching changed reaction growth, I conducted an analytic and numerical analysis of reactions exposed to different flow gradients. Stretching is the Lagrangian finite time analogue of the instantaneous flow gradient, so it acts as the combined effect of all flow gradients.

I characterized three possible flow gradient effects. The first is that any material surface in a flow changes shape according to the strain-rate experienced. This causes reactions to grow *globally* faster than they would without flow, since the global reacted region growth is approximately vP . Even if v is constant, P need not be. The second effect changes v , because strain rate can alter the usual 1D solution by adding an extra effect that changes reaction front shape. I found that in simulations the reaction front speed was linearly related to the thickness of a reaction front. This reveals a possible modification to the Eikonal equation: allowing v to change in a way that relates to the current front thickness, and then allowing that thickness to change as a function of strain-rate. Finally, I showed that the results of KiSS theory [18, 24, 25] which predict that if strain-rate exceeds reaction-rate, then for linear chemical reactions a filament will be stretched until it reaches 0 concentration, applied reasonably well to blowout in the Oregonator. Since the Oregonator is everywhere a slower reaction than a linear reaction, the linear approximation serves as an upper bound for γ_c . In simulations I demonstrated that the transition from a stable filament solution to a blown out solution did happen at around $\gamma = \gamma_c$. When I estimated what Reynolds number first has a strain rate exceeding γ_c , I found it a little below the observed transition from

high stretching being reacted to low stretching regions. While blowout happens in infinite time for reactions without a threshold, for the Oregonator it happens in a finite time, which is why blowout has such a big effect on the eventual reaction field. This blowout transition will be entirely chemical dependent, and it may serve to quickly explain the transient structure of reaction fields when mixed.

6.1 Future Directions

These results all serve to help understand the growth of ARD systems through a dynamical systems perspective. I can now measure the growth of reacted regions, study how flow modifies the stagnant behavior of a reaction front, and calculate the importance of strain-rate and stretching to ARD systems. The biggest value provided by this thesis however is the work that it enables in the future. Both in terms of the next basic science steps to develop these ideas, and also in terms of the application of these results to important problems beyond the lab.

Front tracking is largely complete in its formulation. There could be slight improvements to time stepping methods, less arbitrary thresholding, three-dimensional systems, or advancements to track more of the reaction field's information, but each of these would be an incremental improvement over its current state. It would be more interesting to see if other researchers could expand the simple ideas of front tracking to tracking regions besides reactive scalar concentration regions. The goal would be to quantify how a bounding region moves over time, but that region could be any volume that undergoes a sudden transition from one state to another. This more general problem of "edge tracking" exists for a variety of systems, each requiring their own theory. Numerous opportunities exist for study in this area, from the dynamics of cloud formation, motion of interfaces due to surface tension, quantifying growth or shrinkage of flow structures like coherent vortices, or tracking phase transitions to name a few.

The research on depth shear in quasi-2D reactive mixing experiments was intended to enable future research in reactive mixing. While predictions are easier in 2D and ARD systems such as phytoplankton in the ocean can be modeled as 2D ARD systems [18, 19, 22], depth shear means that experiments are rarely comparable. For all the difficulty posed by a quasi-2D experiment, a 3D experiment is even more challenging. With the understanding that depth shear is the cause of the discrepancy, I hope that future experiments will find ways to suppress the shear further and replicate 2D ARD dynamics more accurately. The first step is to develop experimental procedures that reduce depth shear even farther than the reduction I achieved with a two-layer system. For example by tuning the viscosity [110]. On the subject of improving reactive mixing experiments, it would also be good to find new chemical reactions that can be worked with. BZ has the benefit of being very colorful, but its complex kinetics can make it difficult to isolate the causes of different behaviors in experiments. It might also be interesting to think about modifying the Eikonal equation to account for the depth shear effect and see if it better predicts one layer experiments.

Optimal stretching was helpful for understanding my results, but an improved measure is needed. This is because optimal stretching treats strain-rate and time as interchangeable. If more strain rate is needed, the same result can be obtained by increasing time. The finite-time-Lyapunov exponent — which is the logarithm of stretching divided by the stretching time — may be part of the answer, but it remains unclear if this is the only problem. For now, Lagrangian stretching with a fixed stretching time serves the purpose of predicting reaction structures. Another big open problem in optimal stretching is a method for accurately measuring a blowout level S_b . Conceivably, since blowout appears regardless of flow structure, being able to measure blowout level in a lab should provide valuable information on reactions outside the lab. Theory could also be developed on the reaction kinetics dependence of blowout. For instance, how does γ_c depend on

q which determines the threshold, ϵ which sets the strength of interaction with the inhibitor, or even why different order reaction terms behave differently? For that matter how does the blowout time t_b depend on those things? Combined, measuring S_b in experiments and determining its relationship to reaction kinetics could allow enhanced prediction capabilities and someday even control capabilities. For instance, blowout barriers may serve as a major organizer of reaction state in turbulent flows, but predicting where these barriers would be is tricky unless S_b can be predicted from the kinetics, or measured in a lab.

Finally the big remaining basic science questions from this thesis is what will happen when the results are combined. If an experiment is run in a two-layer system, the reaction will experience strain-rate similar to that measured through particle tracking velocimetry. Such a system is driven more easily into unsteady or turbulent flow too, so I would be able to further probe the applicability of these results to turbulence. I could reasonably expect the fronts to be more 2D, and I could therefore use front tracking to measure the various strain-rate effects I have identified. In such an experiment there is no guarantee I would not see even more new behavior, but at the least I could better test the ideas I have presented.

Applications for front tracking are numerous. Even ignoring “edge tracking”, it can already be used to track the movement of any scalar edge. What this means is it provides new information about reactive mixing systems with important applications such as fires, phytoplankton blooms, and stirred chemical reactions. More importantly, the new information matches the information that we are most interested in: Where is the fire? How fast is it advancing? What’s the effect of the wind? By knowing how reaction front speed is changed by different environmental situations, we could enable a real time prediction of how a flame will grow, or a plankton bloom will travel. Within industrial mixing, front tracking provides an easily digestible diagnostic tool for how well a chemical reaction is growing. If it takes a long time for a chemical reaction front to encapsulate an entire vat,

then the mixing is not good. By measuring a reaction front speed and location throughout time, industry can track what they care most about in their mixers. With these applications in mind, I would suggest one of the next steps would be to try front tracking some of these different systems, such as phytoplankton in the ocean, or a flame. In fact, I have recently worked on a collaborative project in Plog *et al.* [119], in which front tracking was used to quantify the inflow of cerebrospinal fluid into the brains of mice. The results of front tracking helped to show that hypertonic saline or mannitol could cause an influx of cerebrospinal fluid. From this the paper was able to conclude that these chemicals could possibly enable drug delivery into the brain through the cerebrospinal fluid.

Optimal stretching may have practical impacts of its own in all these areas as well. Understanding that strain-rate is the next most dominant effect on global reaction caused by mixing was very useful for me in my research, because it provided an intuition about what a reaction was going to do before doing any experiments. In terms of the growth of phytoplankton in the ocean, it is possible that each phytoplankton species has a different optimal stretching, due to differences in predator avoidance, nutritional needs, and reproduction rates. If different species have different optimal stretching ranges, that means each is best suited to the area of the ocean at that stretching. This means stretching could be selecting the species in a given area of the ocean. Optimal stretching may be the cause of “the paradox of the plankton” [28], but measurements of stretching and plankton in the ocean would be required to confirm this. Optimal stretching regions for wildfires due to wind could be important to track in order to better predict the motion of fires. Since I have shown strain-rate is the primary reason reactions grow faster in flows, the three different strain rate effects should be carefully considered in the design of reaction mixers. It may be interesting to study how to get a given reaction to take place more quickly without increasing the total flow kinetic energy input to the mixer. Industrial mixers have a huge number of possible variables from the

complexity of the machinery, the driving mechanism, the spacing of mixing, and so on, so having an understanding of what promotes reaction mixing may prove invaluable.

Between the simplified modeling enabled by front tracking, the results surrounding strain-rate, and optimal stretching, the results of this thesis enable numerous future projects within the study of ARD. These results serve as a way to understand the evolution of ARD systems, and front tracking in particular may be a useful tool for analyzing reaction dynamics. In the future these methods may ease the difficulties of predicting ARD systems, even in turbulence. From there even control may be reachable. Questions like: “How can I mix these reactions with the least possible energy input?”, or “How fast is the combustion front moving in this motor?”, or even “Why are the plankton here different than over there?” show that there is still much to study. Studying these ARD problems through the lens of dynamical systems will bring new insights, answers, and even more questions.

Bibliography

- [1] G. K. Batchelor, *Journal of Fluid Mechanics* **5**, 113 (1959).
- [2] S. K. Scott, *Oscillations, Waves, and Chaos in Chemical Kinetics* (Oxford University Press Inc., New York, 1994).
- [3] I. R. Epstein, *Nature* **374**, 321 (1995).
- [4] I. Nagypal and I. R. Epstein, *The Journal of Physical Chemistry* **90**, 6285 (1986).
- [5] T. D. Nevins and D. H. Kelley, *Phys. Rev. Lett.* **117** (2016), 10.1103/PhysRevLett.117.164502.
- [6] S. M. Kresta, V. A. Atiemo-Obeng, E. L. Paul, and V. Atiemo-Obeng, *Handbook of Industrial Mixing: Science and Practice* (John Wiley & Sons, Incorporated, Hoboken, United States, 2004).
- [7] N. Harnby, M. F. Edwards, and A. W. Nienow, *Mixing in the Process Industries: Second Edition* (Butterworth-Heinemann, 1997) google-Books-ID: LxdFnHQcXhgC bibtex: nienow1997.
- [8] C. P. Schlick, P. B. Umbanhowar, J. M. Ottino, and R. M. Lueptow, *Chaos* **24**, 013109 (2014).
- [9] K. S. Elvira, X. C. i Solvas, R. C. R. Wootton, and A. J. deMello, *Nat. Chem.* **5**, 905 (2013).

- [10] S. Atis, S. Saha, H. Auradou, D. Salin, and L. Talon, *Phys. Rev. Lett.* **110** (2013), 10.1103/PhysRevLett.110.148301.
- [11] S. Saha, S. Atis, D. Salin, and L. Talon, *EPL* **101**, 38003 (2013).
- [12] W. G. Yao, H. Sato, K. Takahashi, and K. Koyama, *Chemical Engineering Science* **53**, 3031 (1998).
- [13] N. Di Patrizio, M. Bagnaro, A. Gaunand, J.-F. Hochepped, D. Horbez, and P. Pitiot, *Chemical Engineering Journal* **283**, 375 (2016).
- [14] G. Ascanio, *Chinese Journal of Chemical Engineering* **23**, 1065 (2015).
- [15] B. Coriton, J. H. Frank, and A. Gomez, *Combust. Flame* **170**, 37 (2016).
- [16] W. W. Hargrove, R. H. Gardner, M. G. Turner, W. H. Romme, and D. G. Despain, *Ecological Modelling* **135**, 243 (2000).
- [17] C. Punckt, P. S. Bodega, P. Kaira, and H. H. Rotermund, *J. Chem. Educ.* **92**, 1330 (2015).
- [18] A. P. Martin, *J Plankton Res* **22**, 597 (2000).
- [19] A. P. Martin, *Prog. Oceanogr.* **57**, 125 (2003).
- [20] M. Sandulescu, C. Lpez, E. Hernndez-Garca, and U. Feudel, *Nonlinear Processes in Geophysics* **14**, 442 (2007).
- [21] M. Sandulescu, C. Lpez, E. Hernndez-Garca, and U. Feudel, *Ecological Complexity* **5**, 228 (2008).
- [22] J. E. Truscott and J. Brindley, *Bulletin of Mathematical Biology* **56**, 981 (1994).
- [23] Z. Neufeld, *Chaos* **22**, 037102 (2012).

- [24] H. Kierstead and B. Slobodkin, *J. Mar. Res.* **12**, 141 (1953).
- [25] J. G. Skellam, *Biometrika* **38**, 196 (1951).
- [26] W. H. Schlesinger, *Biogeochemistry: an analysis of global change* (Academic Press, San Diego, 1991).
- [27] F. dOvidio, S. De Monte, S. Alvain, Y. Dandonneau, and M. Lvy, *Proc. Natl. Acad. Sci. U.S.A.* **107**, 18366 (2010).
- [28] G. E. Hutchinson, *The American Naturalist* **95**, 137 (1961).
- [29] Z. Neufeld, C. Lopez, E. Hernandez-Garcia, and O. Piro, *Phys. Rev. E* **66** (2002), 10.1103/PhysRevE.66.066208.
- [30] Z. Neufeld, P. H. Haynes, and T. Tl, *Chaos* **12**, 426 (2002).
- [31] W. Tang and A. Dhumuntarao, *Phys. Fluids* **27**, 076601 (2015).
- [32] W. Tang and C. Luna, *Phys. Fluids* **25**, 106602 (2013).
- [33] T. Chevalier, D. Salin, and L. Talon, *Phys. Rev. Fluids* **2** (2017), 10.1103/PhysRevFluids.2.043302.
- [34] D. Bargteil and T. Solomon, *Chaos* **22**, 037103 (2012).
- [35] B. F. Edwards, *Phys. Rev. Lett.* **89** (2002), 10.1103/PhysRevLett.89.104501.
- [36] B. F. Edwards, *Chaos* **16**, 043106 (2006).
- [37] M. Leconte, J. Martin, N. Rakotomalala, and D. Salin, *Phys. Rev. Lett.* **90** (2003), 10.1103/PhysRevLett.90.128302.
- [38] C. O. Mehrvarzi and M. R. Paul, *Phys. Rev. E* **90** (2014), 10.1103/PhysRevE.90.012905.
- [39] R. A. Fisher, *Ann. Eugen.* **7**, 355 (1937).

- [40] A. Kolmogoroff, I. Petrovsky, and N. Piscounoff, in *Dynamics of Curved Fronts* (Elsevier, 1988) pp. 105–130.
- [41] J. P. Keener, *SIAM Journal on Applied Mathematics; Philadelphia* **46**, 18 (1986).
- [42] P. Foerster, S. C. Muller, and B. Hess, *Science* **241**, 685 (1988).
- [43] R. S. Spangler and B. F. Edwards, *J. Chem. Phys.* **118**, 5911 (2003).
- [44] K. A. Mitchell and J. R. Mahoney, *Chaos* **22**, 037104 (2012).
- [45] J. R. Mahoney, J. Li, C. Boyer, T. Solomon, and K. A. Mitchell, *Phys. Rev. E* **92** (2015), 10.1103/PhysRevE.92.063005.
- [46] J. R. Mahoney and K. A. Mitchell, *Chaos* **23**, 043106 (2013).
- [47] M. Doan, J. J. Simons, K. Lilienthal, T. Solomon, and K. A. Mitchell, *Phys. Rev. E* **97**, 033111 (2018).
- [48] S. Gowen and T. Solomon, *Chaos* **25**, 087403 (2015).
- [49] P. W. Megson, M. L. Najarian, K. E. Lilienthal, and T. H. Solomon, *Phys. Fluids* **27**, 023601 (2015).
- [50] J. Mahoney, D. Bargteil, M. Kingsbury, K. Mitchell, and T. Solomon, *EPL* **98**, 44005 (2012).
- [51] J. R. Mahoney and K. A. Mitchell, *Chaos* **25**, 087404 (2015).
- [52] Z. Neufeld and E. Hernandez-Garcia, *Chemical and Biological Processes in Fluid Flows: A Dynamical Systems Approach* (World Scientific, 2009).
- [53] Z. Neufeld, *Phys. Rev. Lett.* **87** (2001), 10.1103/PhysRevLett.87.108301.
- [54] Z. Neufeld, I. Z. Kiss, C. Zhou, and J. Kurths, *Phys. Rev. Lett.* **91** (2003), 10.1103/PhysRevLett.91.084101.

- [55] G. A. Voth, G. Haller, and J. P. Gollub, *Phys. Rev. Lett.* **88** (2002), 10.1103/PhysRevLett.88.254501.
- [56] O. Reynolds, *Proc. R. Soc. Lond.* **35**, 84 (1883).
- [57] A. Thess, *Phys. Fluids A* **4**, 1385 (1992).
- [58] A. Thess, *Physics of Fluids A: Fluid Dynamics* **4**, 1396 (1992).
- [59] J. Tithof, B. Suri, R. K. Pallantla, R. O. Grigoriev, and M. F. Schatz, *Journal of Fluid Mechanics* **828**, 837 (2017).
- [60] A. Hadjighasem and G. Haller, *SIAM Rev.* **58**, 69 (2016).
- [61] G. Froyland, C. Horenkamp, V. Rossi, N. Santitissadeekorn, and A. S. Gupta, *Ocean Modelling* **5253**, 69 (2012).
- [62] G. Haller, *Annual Review of Fluid Mechanics* **47**, 137 (2015).
- [63] G. Haller and G. Yuan, *Physica D: Nonlinear Phenomena* **147**, 352 (2000).
- [64] G. Haller, *Physica D: Nonlinear Phenomena* **149**, 248 (2001).
- [65] G. Haller, *Phys. Fluids* **14**, 1851 (2002).
- [66] M. J. Olascoaga and G. Haller, *Proc Natl Acad Sci U S A* **109**, 4738 (2012).
- [67] S. C. Shadden, F. Lekien, and J. E. Marsden, *Physica D: Nonlinear Phenomena* **212**, 271 (2005).
- [68] M. R. Allshouse and T. Peacock, *Chaos* **25**, 087410 (2015).
- [69] M. R. Allshouse and T. Peacock, *Chaos* **25**, 097617 (2015).
- [70] Y.-K. Tsang, *Phys. Rev. E* **80**, 026305 (2009).
- [71] R. J. Field and R. M. Noyes, *J. Chem. Phys.* **60**, 1877 (1974).

- [72] R. J. Field and R. M. Noyes, *Journal of the American Chemical Society* **96**, 2001 (1974).
- [73] A. T. Winfree, *Journal of Biosciences* **27**, 465 (2002).
- [74] I. Farkas, D. Helbing, and T. Vicsek, *Nature* **419**, 131 (2002).
- [75] N. T. Ouellette, H. Xu, and E. Bodenschatz, *Exp. Fluids* **40**, 301 (2006).
- [76] D. H. Kelley and N. T. Ouellette, *Physics of Fluids* **23**, 045103 (2011).
- [77] J. Tithof, B. C. Martell, and D. H. Kelley, *Physical Review Fluids* **3** (2018), 10.1103/PhysRevFluids.3.064602.
- [78] T. D. Nevins and D. H. Kelley, *Chaos* **27**, 043105 (2017).
- [79] P. E. Arratia and J. P. Gollub, *Phys. Rev. Lett.* **96** (2006), 10.1103/PhysRevLett.96.024501.
- [80] V. Pérez-Muñuzuri and G. Fernández-García, *Physical Review E* **75** (2007), 10.1103/PhysRevE.75.046209.
- [81] J. A. Pojman, R. Craven, and D. C. Leard, *Journal of Chemical Education* **71**, 84 (1994).
- [82] K. Kovacs, M. Leda, V. K. Vanag, and I. R. Epstein, *Physica D: Nonlinear Phenomena* **239**, 757 (2010).
- [83] C. Gendrin, Y. Roggo, and C. Collet, *Journal of Pharmaceutical and Biomedical Analysis* **48**, 533 (2008).
- [84] E. Ray, P. Bunton, and J. A. Pojman, *American Journal of Physics* **75**, 903 (2007).
- [85] J. Glimm, J. W. Grove, X. L. Li, K.-m. Shyue, Y. Zeng, and Q. Zhang, *SIAM J. Sci. Comput.* **19**, 703 (1998).

- [86] I.-L. Chern, J. Glimm, O. McBryan, B. Plohr, and S. Yaniv, *J. Comput. Phys.* **62**, 83 (1986).
- [87] A. Buschmann, F. Dinkelacker, T. Schafer, M. Schafer, and J. Wolfrum (TECFLAM, Stuttgart, 1987) pp. 437–445.
- [88] L. P. H. de Goey, T. Plessing, R. T. E. Hermanns, and N. Peters, *Proceedings of the Combustion Institute* **30**, 859 (2005).
- [89] S. C. Miller, T. Plesser, and B. Hess, *Physica D: Nonlinear Phenomena* **24**, 71 (1987).
- [90] M. J. Ablowitz and A. Zeppetella, *Bull. Math. Biol.* **41**, 835 (1979).
- [91] V. K. Vanag and I. R. Epstein, *J. Chem. Phys.* **117**, 8508 (2002).
- [92] P. M. Wood and J. Ross, *J. Chem. Phys.* **82**, 1924 (1985).
- [93] I. R. Epstein, *Chem. Eng. News Archive* **65**, 24 (1987).
- [94] A. N. Zaikin and A. M. Zhabotinsky, *Nature* **225**, 535 (1970).
- [95] L. Kuhnert, H. J. Krug, and L. Pohlmann, *The Journal of Physical Chemistry* **89**, 2022 (1985).
- [96] T. D. Nevins and D. H. Kelley, *Chaos* **28**, 043122 (2018).
- [97] T. Chevalier, A. K. Dubey, S. Atis, A. Rosso, D. Salin, and L. Talon, *Phys. Rev. E* **95** (2017), 10.1103/PhysRevE.95.042210.
- [98] M. E. Schwartz and T. H. Solomon, *Phys. Rev. Lett.* **100**, 028302 (2008).
- [99] J. Wang, J. Tithof, T. D. Nevins, R. O. Colon, and D. H. Kelley, *Chaos* **27**, 123109 (2017).
- [100] T. H. Solomon and I. Mezi, *Nature* **425**, 376 (2003).

- [101] T. D. Nevins and D. H. Kelley, Under Review (2019).
- [102] J. Xin, SIAM Rev. **42**, 161 (2000).
- [103] W. Vansaarloos, Physics Reports **386**, 29 (2003).
- [104] A. Kiselev and L. Ryzhik, Annales de l'Institut Henri Poincare (C) Non Linear Analysis **18**, 309 (2001).
- [105] M. Abel, A. Celani, D. Vergni, and A. Vulpiani, Phys. Rev. E **64**, 046307 (2001).
- [106] E. Beauvier, S. Bodea, and A. Pocheau, Phys. Rev. E **96**, 053109 (2017).
- [107] B. Coriton, J. H. Frank, and A. Gomez, Combust. Flame **160**, 2442 (2013).
- [108] H. Sevčiková and M. Marek, Physica D: Nonlinear Phenomena **9**, 140 (1983).
- [109] R. Feeney, S. L. Schmidt, and P. Ortoleva, Physica D: Nonlinear Phenomena **2**, 536 (1981).
- [110] B. Suri, J. Tithof, R. Mitchell, R. O. Grigoriev, and M. F. Schatz, Phys. Fluids **26**, 053601 (2014).
- [111] A. Figueroa, F. Demiaux, S. Cuevas, and E. Ramos, Journal of Fluid Mechanics **641**, 245 (2009).
- [112] F. V. Dolzhanskii, V. A. Krymov, and D. Y. Manin, Journal of Fluid Mechanics **241**, 705 (1992).
- [113] J. Sharif, M. Abid, and P. D. Ronney, in *1st Joint US Sections Combustion Institute Meeting* (1999) p. 4.
- [114] A. von Kameke, F. Huhn, A. P. Muuzuri, and V. Prez-Muuzuri, Phys. Rev. Lett. **110** (2013), 10.1103/PhysRevLett.110.088302.

- [115] A. von Kameke, F. Huhn, G. Fernndez-Garca, A. P. Muuzuri, and V. Prez-Muuzuri, *Phys. Rev. E* **81** (2010), 10.1103/PhysRevE.81.066211.
- [116] Clancy, Tom, *The Hunt for Red October* (Naval Institute Press, 1984).
- [117] J. R. Parker, *Algorithms for Image Processing and Computer Vision* (Wiley, Hoboken, US, 2010).
- [118] P. Pelce, *Dynamics of Curved Fronts* (Elsevier, 2012).
- [119] B. A. Plog, H. Mestre, G. E. Olveda, A. M. Sweeney, H. M. Kenney, A. Cove, K. Y. Dholakia, J. Tithof, T. D. Nevins, I. Lundgaard, T. Du, D. H. Kelley, and M. Nedergaard, *JCI Insight* **3** (2018), 10.1172/jci.insight.120922.