IMPERIAL

Imperial College London
Department of Mathematics

Nonlinear dynamics, bifurcations and collective motion of active filaments

Bethany Clarke

November 28, 2025

Submitted in part fulfilment of the requirements for the degree of Doctor of Philosophy in Mathematics of Imperial College London and the Diploma of Imperial College London.

Statement of Originality

I can confirm that the work presented in this thesis is my own. All other works and sources have been appropriately cited in the text. This work has not been submitted for any other degree or purpose.

Copyright Declaration

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Abstract

This thesis investigates the nonlinear dynamics of active filaments submerged in viscous fluids through a combination of initial-value-problem studies and bifurcation analysis. The central focus of this work is the follower force model, in which filament deformation is driven through compressive forces which act along the tangent of the filament centreline. This model provides a powerful framework for understanding biologically relevant systems, such as molecular motor-driven biopolymer filaments and the beating of cilia and flagella.

We begin by examining the behaviour of a single filament, systematically exploring the parameter space of several variants of the follower force model, starting with its simplest formulation. Numerical simulations, complemented with computational dynamical systems techniques, are used to identify new dynamical behaviours and map the bifurcations in the state space. Following studies investigate the influence of the domain, the forcing distribution, and the mechanical properties of the filament, uncovering new solution types and highlighting the differences between modelling approaches.

Building on these results, we extend the study to filament ensembles to examine the stability of coordinated states. We observe multistability between metachronal wave solutions of different wavelengths in arrays of asymmetrically beating filaments, suggesting that while hydrodynamic interactions can facilitate these patterns, additional regulatory mechanisms are required in biological systems to select the desired wavelength solution. Overall, this work provides new insights into the mechanical principles underlying both individual and collective filament dynamics, revealing the rich nonlinear behaviours that emerge from simple forcing mechanisms.

Acknowledgements

I would first like to express my deepest gratitude to my supervisor, Eric Keaveny, for his unwavering support, guidance, and encouragement throughout my PhD. Your insight and patience, both academically and personally, have shaped this journey and guided me through a pivotal (and sometimes stressful) part of my life. I am equally grateful to my unofficial second supervisor, Yongyun Hwang, whose expertise, thoughtful feedback, and enthusiasm have been invaluable. It has been a privilege to work with you both, and I am deeply thankful for the kindness and mentorship you've shown me over the past four years.

I am extremely grateful to my examiners, Raymond Goldstein and Ory Schnitzer, for generously offering their time, thoughts and expertise on this thesis.

I would also like to thank the members of my research group, past and present, for their helpful discussions and support. In particular, I am grateful to Hang for being a sounding board, IT wizard, conference companion, badminton partner, and - most importantly - a friend. Much of my positive experience at Imperial comes down to the luck of being based in Huxley 759. Special thanks to Bo, Hang, Benedetta, Jacob, Alasdair, and Kamal for the laughter, long chats, and for generally being great people. A PhD can be lonely, but having such a strong (not just emotionally but also physically, thanks to the office push-ups) support system in you all made the experience infinitely easier and more enjoyable.

My time at Imperial has also been heavily shaped by life outside the research - editing for Chalkdust, acting as a department rep, teaching, and outreach. Thank you to everyone I met through these activities, and to those who encouraged me when juggling everything felt overwhelming. A special thank you to Alice for keeping me sane and bringing so much fun into life outside of the PhD, and to Tom Bennison for inspiring a love of maths, encouraging me to pursue my dream university, and for being a mentor and a friend.

Finally, I am endlessly grateful to my family for their constant love and endless support. Thank you to my parents for cheering me on and always encouraging me to do whatever makes me happy, and to Josh for making me laugh (and keeping me humble). Thank you also to my second family - Sue, Andy, Shadia, and Vinnie - for their support throughout. And thank you above all to Luke, for being my rock over the last ten years, and for bringing so much joy and colour to my life. I could not have been more privileged to have shared this experience, and our life in London and beyond, with you.

Contents

Abstract **v**

Acknowledgements **vii**

1 Introduction **1**

1.1 Biological background 2

1.1.1 Biopolymer filaments-molecular motor complexes 2

1.1.2 Cilia and flagella 6

1.2 Modelling approaches in the literature 10

1.2.1 Single filament-motor protein systems 10

1.2.2 Modelling cilia and flagella 12

1.3 Thesis aims and outline 15

2 Background - Stokes Flow **17**

2.1 Stokes equations 17

2.1.1 Fundamental solutions of the Stokes equations 20

2.1.2 Boundary integral formulation - single layer potential 23

2.1.3 Multipole expansion . . . 23

2.1.4 Faxéns laws . . . 26

2.1.5 The mobility problem . . . 27

2.2 Filament hydrodynamics . . . 28

2.2.1 Rotne-Prager-Yamakawa (RPY) tensor . . . 30

2.2.2 FCM . . . 32

2.2.3 Fast FCM . . . 34

3 Background - Krylov Subspace Methods 36

3.1 Arnoldi method . . . 37

3.2 Generalised minimal residual method (GMRES) . . . 39

4 Numerical Methods 41

4.1 Problem formulation . . . 41

4.2 Follower force model . . . 42

4.2.1 Model description and continuous equations . . . 42

4.2.2 Numerical discretisation . . . 43

4.2.3 Hydrodynamics . . . 45

4.2.4 Time integration . . . 46

4.2.5 Basal boundary conditions . . . 49

4.3 Jacobian-free Newton Krylov method . . . 49

4.3.1 Choosing a state variable . . . 50

4.3.2 Time-invariant solutions 52

4.3.3 Newton-Raphson method 53

4.3.4 Locally constrained optimal hook 55

4.3.5 Comments on implementation 58

4.4 Linear stability analysis 59

4.5 Bisection 61

5 Individual Dynamics of a Follower Force-Driven Filament 63

5.1 Introduction 63

5.2 Simulation parameters and nondimensionalisation 65

5.3 State space 66

5.4 Two-dimensional buckling 69

5.5 Double Hopf bifurcation 70

5.6 From whirling to beating 75

5.7 From beating to writhing 78

5.8 Discussion 82

6 Modelling Extensions 87

6.1 Effect of twist 88

6.2 Effect of hydrodynamic wall 90

6.2.1 Initial value problems 90

6.2.2 Bifurcation and stability analyses 92

6.3 Effect of force distribution 96

6.3.1 Initial value problems 99

6.3.2 Bifurcation and stability analyses 100

6.4 Discussion 102

7 Breaking Symmetry in a Follower Force-Driven Filament 106

7.1 Introduction 106

7.2 Methods and simulation parameters 108

7.3 Intrinsic curvature in isotropic filaments 110

7.3.1 Planar dynamics 110

7.3.2 Non-planar dynamics - initial value problems 112

7.3.3 Bifurcation and stability analyses 114

7.4 Aniostropic filaments with no intrinsic curvature 122

7.4.1 Initial value problems 122

7.4.2 Bifurcation and stability analyses 123

7.5 Using bending modulus anisotropy to stabilise asymmetric beating 127

7.6 Discussion 129

8 Collective Dynamics of Instability-Driven Filaments 134

8.1 Introduction 134

8.2 Numerical methods 138

8.2.1 Boundary implementation 139

8.2.2 Modified filament model . . . 141
8.2.3 Non-zero mean flow . . . 142
8.3 Discretisation verification - Poiseuille flow . . . 147
8.3.1 Comparison to analytic results . . . 149
8.3.2 Effect of the relative particle radius . . . 151
8.3.3 Effect of the wall density . . . 152
8.4 Collective dynamics of filaments . . . 153
8.4.1 Simulation parameters . . . 153
8.4.2 MCWs stable above a critical wavelength . . . 155
8.4.3 Dependence on inter-filament spacing . . . 157
8.4.4 Multistability between different wavelength solutions . . . 158
8.5 Discussion . . . 164

9 Conclusions and Outlook **167**

Bibliography **172**

A Derivatives of Oseen's Tensor **195**

B Half-Space RPY Tensor **196**

C Approximate Jacobian **198**

D Lie Algebra Relation to Local Basis **202**

E Permissions to Re-use Copyrighted Work **205**

List of Tables

6.1 The critical values of the forcing, f, at which each bifurcation occurs for $\alpha = 44$, both with and without the hydrodynamic wall at $z = 0$. 94

6.2 The critical values of the forcing, $\tilde{f}$, at which each bifurcation occurs when follower forces are distributed along the filament length for $\alpha = 44$, both with and without the hydrodynamic wall. The bifurcation diagrams corresponding to the two domains have the same qualitative properties, with the quantitative differences in the critical forcing summarised in this table. 100

7.1 The critical forcing at which each bifurcation occurs in Figure 7.4 for three values of intrinsic curvature, κ_0. 117

7.2 The critical forcing at which each bifurcation occurs in Figure 7.9 for three values of the bending moduli ratio, β. 125

List of Figures

1.1 Schematic showing the biopolymer filaments described in the text, and the motor proteins which traverse them. (Top) Actin filaments are helical structures with distinguishable ends, given by the 'barbed' or 'pointed' end. Actin filaments are traversed by myosin motor proteins. Different myosin motor proteins can move in either direction along the filament. (Bottom) Microtubules are hollow structures constructed from $\alpha-$ and $\beta-$tubulin. The two ends of the microtubule are distinguishable as the 'plus' end and 'minus' end. Dynein and kinesin are motor proteins which translocate along microtubules, and move in opposite directions. Myosin, dynein and kinesin all have different structures, as shown in the Figure. (Reproduced with permission from Gross et al. [2007]). 3

1.2 (a) An electron microscope image of the cross-section of the two flagella of *Chlamydomonas*, exhibiting the '9+2' axonemal structure in motile cilia. (Image released in the public domain by [Dartmouth Electron Microscope Facility, Dartmouth College]). (b) A schematic of the '9+2' axonemal structure showing the 9 microtubule doublets surrounding the central pair, interconnected by radial spokes, nexin proteins and dynein. In particular, we note the beat plane direction with reference to the orientation of the central pair, and the bridge between microtubule doublets 5 and 6. 6

1.3 Electron microscope images of (a) several *Chlamydomonas* algae, each driven by its two flagella and (b) the ciliated epithelium in the lung trachea. (Both images released in the public domain by [Dartmouth Electron Microscope Facility, Dartmouth College]). (c) A schematic showing a metachronal wave in a 2D array of filaments. Each filament is performing the same periodic beat, with a constant phase shift compared to its neighbour. The individual beat pattern is obtained using the results presented in Chapter 7 for the parameter choice $(f, \kappa_0, \beta) = (40, 3\pi/4, 2)$. 8

1.4 Schematics of some of the modelling choices used in the literature to describe the motion of motors translocating along the length of an elastic filament. (a) Some studies allow the motors to move forwards or backwards along the filament, depending on some pre-determined criteria such as applied load. (b) Other works neglect details of the biomechanics, by considering a collection of motors moving in one direction only. (c) Following this approach, models may incorporate the effect of the motors through the drag the motors experience from carrying cargo (which we call backflow), as well as compressive forces on the filament (follower forces) following Newton's third law. This model can be simplified further, by (d) neglecting the backflow and then (e) considering only one such follower force, corresponding to the effect of a single molecular motor at the filament tip. . . . 11

1.5 (a) A schematic of the axoneme adapted from Figure 1.2. Inset, right shows how two microtubule doublets (represented by black lines) are connected by dynein motors (grey) and nexin proteins (dashed lines). The dynein motors cause relative sliding of the doublets (blue arrows), which leads to bending. In order to generate overall bending of the cilium, there must be some sort of regulation on dynein attachment across the axoneme. (b) Schematics showing three possible regulatory controls for dynein attachment from the literature: (left) the geometric clutch hypothesis, where dynein motors detach if interdoublet distances are too large, (middle) curvature control, where dynein motors respond to the local curvature of the cilium, and (right) sliding control, where the dynein motors detach if forces tangential to the microtubule doublets, leading to sliding, are too large. 13

4.1 (a) A schematic of the follower force model. The filament has cross-sectional radius a, and length L. The bending and twisting moduli are K_B and K_T, respectively. The filament base is clamped to a rigid planar wall at $z = 0$ and surrounded by a fluid with viscosity η. Also shown is the compressive follower force of strength f at the filament's distal end. (b) A schematic describing the numerical method used to compute filament evolution. The filament is discretised into N_p segments of length ΔL. Segment i has position $\boldsymbol{Y}_i$ and local basis $\{\hat{\boldsymbol{t}}_i, \hat{\boldsymbol{\mu}}_i, \hat{\boldsymbol{\nu}}_i\}$ where $\hat{\boldsymbol{t}}_i$ is constrained to be the tangent of segment i. To evaluate the hydrodynamic mobility, the segments are modelled as particles with hydrodynamic radius a. 44

5.1 (a) The tip displacement x, (solid line) as a function of time for $\alpha = 44$. The dashed line $\sim \exp(-t/\tau)$, where τ is the relaxation timescale determined through fitting the displacement data. The schematic depicts how the computation of the relaxation timescale is performed. (b) The relaxation timescale computed numerically for our model as a function of aspect ratio, compared to the timescale arising from resistive force theory. 65

5.2 (a) Figures showing the steady, whirling, planar beating and writhing states for $\alpha = 44$ (from left to right). Views of the $(x, z)-$ plane (side) and $(x, y)-$plane (above) are provided, and the passage of time is indicated by the increasing boldness of the lines. (b) The state space of solutions over a range of α and f when dynamics are restricted to 2D. (c) The state space of solutions over a range of α and f for fully 3D dynamics. In panels (b) and (c), the black lines are drawn to indicate the approximate boundaries of the solutions regions. . . . 67

5.3 Results of 2D stability analysis on the trivial steady state. (a) The real and imaginary components of the dominant eigenvalue, λ_1, from the linear stability analysis of the trivial steady state for $f \in [0, 50]$, $\alpha = 44$. The eigenvalue is purely imaginary at the bifurcation. (b) The real part of the dominant 6 eigenvalues from the same analysis. The corresponding eigenvectors associated to the complex conjugate pair of eigenvalues becoming unstable at the (c) first, (d) second and (e) third buckling events. We note that the wavelength of these eigenvectors decreases as we move through successive bucklings. 69

5.4 (a) The dominant eigenvalue from the linear stability analysis of the trivial steady state for $f \in [0, 60]$ and three values of α. The steady state becomes unstable at $f^* \approx 35.3$ regardless of the aspect ratio. The eigenvalues are purely imaginary at the bifurcation. (b) The dominant four eigenvectors of the steady state at the bifurcation. (c) The norm of the solution grows like $\sqrt{f - f^*}$ near the bifurcation. (d) The difference between the solution and its the projection onto the eigenspace at the bifurcation, $E_{\text{projection}} := \min_{a_x, a_y \in \mathbb{C}} ||\boldsymbol{U} - Re(a_x \boldsymbol{v_x}) - Re(a_y \boldsymbol{v_y})||$, grows like $f - f^*$ close to the bifurcation. (e) The difference between the whirling frequency and the imaginary part of the eigenvalue at the bifurcation, $E_{\text{frequency}} := ||2\pi/T - \omega||$, for $\alpha = 44$. The difference grows like $f - f^*$ near the bifurcation. 71

5.5 (a) The period of the whirling (blue) and planar beating (red) solutions for $f \in [f^*, 600]$ and $\alpha = 44$. Solid (dashed) lines correspond to stable (unstable) solutions. Inset: A close up showing the range $f \in [130, 150]$ indicating that two bifurcations occur. (b) Eigenvalues from the Floquet analysis of the whirling solution for $f \in [80, 200]$. Whirling becomes unstable at $f \approx 137.3$. (c) Eigenvalues from the Floquet analysis of the beating solution for $f \in [80, 200]$. Beating becomes stable at $f \approx 140.0$. 75

5.6 (a) The maximum norm of the Lie algebra elements, $||\boldsymbol{U}||$, for the whirling, beating and quasiperiodic solutions as a function of f. Inset: Close up of the quasiperiodic branch connecting with the whirling and beating branches. The three markers indicate the solutions shown in (b) and (c). (b) Figures of the filament in the quasiperiodic regime for three values of f. The passage of time is indicated by an increasing in the boldness of the lines. The eccentricity of the ellipse traced by the filament tip decreases with increasing f. (c) A Poincaré section of the whirling and beating solutions and QP1 for different f. (d) The (α, f) state space in the region where QP1 exists. The range of f for which QP1 exists increases with α. (e) The filament shapes for (left) the whirling solution, $\boldsymbol{U}_{\text{whirl}}$ for $f = 137.4$ and $\alpha = 44$ and (middle/right) the sum of the whirling solution and a linear combination of the unstable eigenvectors, $\boldsymbol{U}_{\text{whirl}} + \beta(\boldsymbol{\zeta} \pm \bar{\boldsymbol{\zeta}})$, for $\beta = 3/2$. 77

5.7 (a) Eigenvalues from the 2D Floquet analysis of the planar beating state for $f \in [440, 600]$ for different α. Planar beating becomes unstable at $f \in \{578.9, 523.7, 498.1\}$ for $\alpha \in \{44, 66, 88\}$, respectively. (b) Eigenvalues from the 3D Floquet analysis of the planar beating state for $f \in [400, 500]$ for different α. The planar beating becomes unstable at $f \in \{462.8, 434.4, 422.4\}$ for $\alpha \in \{44, 66, 88\}$, respectively. (c) The transition region between the planar beating (red crosses) and QP2 (blue stars) solutions. Solutions in the region are divided into three types: chaotic (green squares), intermittent (green diamonds), and time-periodic (green circles). (d)-(g): Filament motion in the transition region as viewed from the side and above for four of the observed behaviours. The passage of time is indicated by an increase in line boldness. (d) Intermittence between beating with a growing out-of-plane component (left) and a combination of beating and whirling near the base and waving at the tip (right). (e) Periodic solution with beating at the base and waving at the tip. (f) Periodic solution where the filament beats while bent in the direction normal to the beat plane. (g) Chaotic solution corresponding to erratic beating at the filament tip. 79

5.8 (a) Filament motion in the QP2 regime. The passage of time is indicated by an increase in line boldness. (Top) A view from above showing the dynamics projected in the $(x, y)-$plane. The dashed lines indicate the minimum and maximum horizontal distance at $s = L$, and $\Delta r(s) = r_{max}(s) - r_{min}(s)$. (Bottom) A view from the side showing the dynamics projected in the $(x, z)-$plane. The dashed line corresponds to $s/L = s^* \approx 0.3$ from Figure 5.9b and c. (b) The Poincaré section corresponding to the QP2 solution for $f = 560$ and $\alpha = 44$. (c) The time-series of the filament tip displacement for $f = 560$ and $\alpha = 44$ showing periodic oscillations in the z component and irregular oscillations in the $(x, y)-$components. 80

5.9 (a) The frequency associated to the dominant mode, ω_1, for various f and α. We see that ω_1 increases with both f and α. Right: Single sided amplitude spectrum of QP2 for $(f, \alpha) = (540, 44)$. We take ω_1 to be the frequency with the largest amplitude. (b) The variation in radius, $\Delta r(s)/L$, against the arc-length, s/L, for $\alpha = 44$ and various f. (c) The variation in radius, $\Delta r(s)/L$, as a function of s/L for $f = 560$ and various α. 81

6.1 The phase diagram for $f \in (100, 200)$, $\alpha \in [44, 132]$ for (a) $\gamma = 1$ and (b) $\gamma = 0.01$. (c) Floquet analysis on the beating solution for a filament with aspect ratio $\alpha = 44$ for $\gamma \in \{0.01, 1\}$. We note that the bifurcation occurs later for $\gamma = 0.01$, and the eigenvalue is complex at the bifurcation. 89

6.2 (a) A schematic showing the model we use. We consider a single filament with clamped-free boundary conditions in an open domain. (b) The results of initial value problems for $\alpha = 44$, $f \in [0, 600]$ for the case with a hydrodynamic wall at $z = 0$, from Chapter 5, and the case without the wall. (c) The bending energy against time for the whirling solutions, at $f = 100$, and the planar beating solutions, at $f = 200$, with and without a hydrodynamic wall over 5 periods. The time is non-dimensionalised by the period of the solution, obtained via JFNK. (d) The dynamics of the planar beating solutions at $f = 200$ (top row) and whirling solutions at $f = 100$ (bottom row) both with (left column) and without (right column) a hydrodynamic wall at $z = 0$. We plot each dynamic over 1 period. The passage of time is indicated by the increasing boldness of the lines. 91

6.3 (a) The bifurcation diagram for a filament in an unbounded domain, driven by a single follower force applied at its tip, for $\alpha = 44$ and $f \in [0, 450]$. We plot the mean bending energy of each solution, obtained as described in the text. Solid (dashed) lines refer to regions in which the solution is stable (unstable). Inset: zoom at $f \in [145, 160]$, showing the QP1 solution branch and the surrounding bifurcations. We highlight $f = 153.5$ with a dashed line. (b) The bending energies of the planar beating, whirling and QP1 solutions at $f = 153.5$. (c) A plot showing the maximum bending energy of several simulations, where the initial conditions are obtained via the bisection method for $f = 153.5$. As the number of bisections increases, the solutions stay on the QP1 solution for longer before diverging to one of the stable states - planar beating (red) or whirling (yellow). The QP1 solution (blue) is obtained by taking an initial condition much later in the bisection algorithm. 93

6.4 (a) The critical forcing, f^*, at which the whirling solution becoming unstable (yellow) and the planar beating solution becoming stable (red) as we vary the filament aspect ratio, α. The lines intersect at $\alpha \approx 80$ (dashed line). To the left of the dashed line, there is a region where whirling and beating are bistable. Here, the QP1 solution exists as an unstable branch connecting the two solutions. To the right of the dashed line, the QP1 solution is a stable branch which connects the two solutions. (b) The real and imaginary parts of the dominant eigenvalue associated to the planar beating solution for $f \in [320, 450]$. We note that beating briefly becomes unstable for $358.4 \lesssim f \lesssim 396.0$, where we observe bent beating in its place, and then becomes unstable again for a final time at $f \approx 407.6$. . . . 95

6.5 (a) A schematic showing the model we use. We consider a single filament clamped at its base in an open domain. The filament is driven by active forces along its length, $\boldsymbol{F}^a(s,t)$, as defined in the main text. (b) The results of initial value problems for $\alpha = 44$, $\tilde{f} \in [0, 7500]$ for the case (top) with and (bottom) without a hydrodynamic wall at $z = 0$. (c) The dynamics of whirling solutions for $\tilde{f} = 220$ (left), 440 (middle), and 1320 (right). All dynamics shown are plotted over one period. The passage of time is indicated by the increasing boldness of the lines. (d) The dynamics of planar beating solutions for $\tilde{f} = 220$ (left), 880 (middle), and 2200 (right). (e) The dynamics of the S-beating solution for $\tilde{f} = 6600$ (left) and the corresponding time trace of the tip coordinates over four periods (right). (f) The shape traced by the filament tip over one period for, from top left to bottom right, $\tilde{f} \in \{4500, 5060, 5280, 6600\}$. The dynamics change from planar beating to S-beating as the forcing increases. 98

6.6 All diagrams in this figure correspond to a single filament in an unbounded domain, driven by a distribution of follower forces along its length. (a) The bifurcation diagram for $\alpha = 44$ and $\tilde{f} \in [0, 7500]$. We plot the mean bending energy of each solution, obtained as described in the text. Solid (dashed) lines refer to regions in which the solution is stable (unstable). Inset: zoom for $\tilde{f} \in [1100, 1600]$, showing the QP1 solution branch and the surrounding bifurcations. (b) The real and imaginary parts of the dominant eigenvalue corresponding to the steady state for $\alpha = 44$ and $\tilde{f} \in [0, 250]$. We see that the steady state becomes unstable at $\tilde{f} \approx 169.7$. (c) The real and imaginary parts of the dominant eigenvalue associated to the beating solution for $\alpha = 44$, and $\tilde{f} \in [4700, 5200]$. The eigenvalue passes the $y-$axis at $\tilde{f} \approx 4912.3$, and is purely real at the bifurcation. 101

7.1 Schematics of the model. (a) The bending modulus anisotropy and intrinsic curvature are defined through the bending moment, which acts on the filament cross-section. It is expressed in terms of the filament's material frame, $(\hat{\boldsymbol{t}}, \hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\nu}})$. We emphasise that the filament cross-section is still circular. (b) The intrinsic curvature, $\kappa_0 = L\tilde{\kappa_0}$ is defined to be the rest-curvature of the filament under zero forcing. 109

7.2 (a) Results from 2D numerical simulations of a filament with intrinsic curvature ($\kappa_0 \geq 0$, $\beta = 1$). The black line is drawn to indicate the approximate boundaries of the solution regions. For $f < f^* \approx 35.3$, filaments remain in their steady non-trivial state which varies with (b) κ_0 (shown for $f = 0$, $\kappa_0 \in \{0, \pi/4, \pi/2, \pi/4\}$) and (c) f (shown for $\kappa_0 = 3\pi/4$, $f \in \{0, 5, 10, 20\}$). For $f > f^*$ filaments undergo asymmetric oscillations for $\kappa_0 > 0$, as shown for (d) $\kappa_0 = 3\pi/4$. The passage of time is indicated by the increasing boldness of the lines. 111

7.3 (a) Results of numerical simulations of a filament with intrinsic curvature in 3D ($\kappa_0 \geq 0, \beta = 1$). The black lines are drawn to indicate the approximate boundaries of the solution regions. (b) Snapshots of filaments over one period undergoing whirling dynamics for two values of intrinsic curvature close to ($f = 40$) and further from ($f = 100$) the bifurcation. We see asymmetries in whirling are more pronounced close to the bifurcation, and for higher κ_0. (c) Snapshots of filaments over one period for $f = 200$, increasing κ_0 from 0 to $3\pi/4$. The dynamics change from planar beating for $\kappa_0 = 0$ (left) to P1 for $\kappa_0 > 0$ (right). The passage of time is indicated by the increasing boldness of the lines. 113

7.4 Bifurcation diagrams showing the different solutions as we vary f, and the corresponding mean bending energy of these solutions, $\bar{E}_b$, for an isotropic filament ($\beta = 1$) with intrinsic curvature for (a) $\kappa_0 = 0$, (b) $\kappa_0 = \pi/4$, (c) $\kappa_0 = \pi/2$ and (d) $\kappa_0 = 3\pi/4$. Full (dashed) lines refer to stable (unstable) solutions. Square markers are used to indicate the location of bifurcations as we vary f. Inset, left, shows the bifurcation diagram near the initial buckling event and inset, right, shows the bifurcation diagram around the bistable regions, if these are present. As κ_0 increases we see the emergence of the P1 solution branch (b-d), from which whirling bifurcates (b and c, left inset). We observe bistability between P1 and whirling (b and c, right inset) and, for the largest values of κ_0, whirling vanishes completely (d). 115

7.5 The real part of the four largest eigenvalues, λ_i for $i = 1, ..., 4$, associated to the steady state for $\beta = 1$, (a) $\kappa_0 = 0$, and (b) $\kappa_0 = 3\pi/4$. Analysing the corresponding eigenvectors at each bifurcation, obtained as described in the main text and displayed on the right, allows us to identify the branches corresponding to off-plane and in-plane beating. (c) The real part of the largest eigenvalue associated to the steady state for various values of κ_0. We see that the eigenvalue becomes unstable for smaller f as we increase κ_0. (d) The real part of the largest eigenvalue associated to the P1 state for various values of κ_0. For $\kappa_0 > \kappa_0^*$, we see the eigenvalue remains stable. 116

7.6 (a) The real part of the dominant four eigenvalues from a linear stability analysis of the steady state for $\kappa_0 = 3\pi/4$, $\beta = 1$ associated to the off-plane buckling (blue dashed/yellow dotted) or in-plane buckling (red) when the curvature is introduced in either frame. (b) The off-plane component of the direction of the intrinsic curvature for various values of the forcing. These are obtained from initial value problems, as described in the text. We note that if the curvature is fixed in the lab frame, this would be zero for all time. 119

7.7 Plots of filament dynamics for solutions obtained via bisection from two views: a 2D view from above (top) and a 3D view (bottom). The passage of time is indicated by an increase in line boldness. Plots for (a) QP1, $(\kappa_0, \beta, f) = (0, 1, 138)$, (b) QP1*, $(\kappa_0, \beta, f) = (\pi/4, 1, 123)$, (c) P3, $(\kappa_0, \beta, f) = (\pi/2, 1, 80)$, (d) P3*, $(\kappa_0, \beta, f) = (\pi/2, 1, 90)$ 120

7.8 (a) Results of numerical simulations of an anisotropic filament with no intrinsic curvature ($\kappa_0 = 0, \beta \geq 1$). The black lines are drawn to indicate the approximate boundaries of the solution regions. (b) Snapshots of filaments over one period undergoing whirling dynamics in the isotropic case, close to the bifurcation ($\beta = 1, f = 36$), an anisotropic case near the bifurcation ($\beta = 1.2, f = 46$) and an anisotropic case further from the bifurcation ($\beta = 1.2, f = 100$). The passage of time is indicated by the increasing boldness of the lines. The three dynamics are indicated on the phase diagram in (a) using numbered circles. 123

7.9 (a) Bifurcation diagrams showing the different solutions as we vary f, and the corresponding mean bending energy of these solutions, $\bar{E}_b$, for an anisotropic filament with no intrinsic curvature ($\kappa_0 = 0$) for (a) $\beta = 1.2$, (b) $\beta = 1.3$ and (c) $\beta = 1.4$. Full (dashed) lines refer to stable (unstable) solutions. Square markers are used to indicate the location of bifurcations as we vary f. We see the separation of the in-plane and off-plane solution branches as β increases, and observe bistability between in-plane beating and whirling (inset). For larger values of β, i.e. (c), the whirling solution branch vanishes completely. 124

7.10 (a) The real part of the four largest eigenvalues associated to the steady state for $\beta = 1.5$. Analysing the corresponding eigenvectors, obtained as described in the text and visualised on the right, allows us to identify the branches corresponding to in-plane beating (bottom right), and off-plane beating (top right). (b) Plots of P4 dynamics for $(\kappa_0, \beta, f) = (0, 1.2, 120)$ from above and a side view. (c) The real part of the second largest eigenvalue associated to off-plane beating for various values of β. We associate this eigenvalue becoming stable with P4 collapsing onto the off-plane beating branch. (d) The real part of the largest eigenvalue associated to the in-plane beating state for various values of β. For $\beta > \beta^*$, we see the eigenvalue remains stable. 126

7.11 (a) Results of numerical simulations of an anisotropic filament with intrinsic curvature ($\kappa_0 = 3\pi/4, \beta \geq 1$). The black lines are drawn to indicate the approximate boundaries of the solution regions. (b) The eigenvalue associated to off-plane beating for several values of β. For $\beta \lesssim 1.4$, this eigenmode becomes unstable before the in-plane eigenmode (i.e. before $f = f^*$). For $\beta \gtrsim 1.4$, the in-plane eigenmode becomes unstable first. (c) The bifurcation diagram showing the different solutions as we vary f, and the corresponding mean bending energy of these solutions, $\bar{E}_b$, for $\kappa_0 = 3\pi/4$, $\beta = 2$. We see that only asymmetric planar beating (bottom) is stable. 128

8.1 Schematics of (a) the computational domain of size $L_x \times L_y \times L_z$ with $N_f = 1$ filament, and (b) the particles making up the wall at $z = 0$. The particles have radius a_w, and are seeded hexagonally with separations Δ_x^{wall} and Δ_y^{wall} in the $\hat{\boldsymbol{x}}-$ and $\hat{\boldsymbol{y}}-$directions respectively. 138

8.2 Visualisations of the flow field of a filament during part of the effective (recovery) stroke when the net flow is (a) ((c)) neglected or (b) ((d)) included. The thin white lines visualise the streamlines of the flow and the thick white arrow refers to the direction of beating. The filament is undergoing beating oscillations obtained using the parameters $(f, \kappa_0, \beta) = 40, 3\pi/4, 2)$, and is in a domain of size $L \times L \times 2L$ with a wall at $z = 0$ constructed of $N_w = 100$ particles. 146

8.3 Heatmaps showing the (a) analytic solution and (b) numerical solution for the fluid velocity in Poiseuille flow with $G = 8\eta/H^2$. The domain is of size $100a \times 100a \times 100a$. The wall at $z = 0$ (and therefore, also at $z = H$ by periodicity) is discretised into 625 hexagonally seeded particles, separated by distances $\Delta_x^{wall} = \Delta_y^{wall} = 4a$. (c) The velocity profiles of the analytic and numerical solutions for this fluid velocity and (d) the error between the two solutions. 150

8.4 (a) A schematic showing the set-up for simulations where we vary the domain height. We fix the (x, y)-domain and the wall density such that $N_w = 4$, $\Delta_x^{wall} = \Delta_y^{wall} = 4a$, and $L = \Delta_x^{wall}\sqrt{N_w}$. We vary the height of the channel, H. (b) The relative pointwise error between the analytic velocity profile, $\boldsymbol{u}^a(z)$, and the numerical velocity profile, $\boldsymbol{u}^b(z)$, as we vary the height of the box. We see that the error scales like $\epsilon \sim O(a/H) \sim O(1/\sqrt{N_w})$. 151

8.5 The relative pointwise error in the analytic velocity profile, $\boldsymbol{u}^a(z)$, and the numerical velocity profile, $\boldsymbol{u}^n(z)$, as we vary the spacing between the particles. We see that this reaches a minimum around $\Delta_x^{wall} \approx 4.5a$ regardless of the domain height. Inset: Plot of the $z = 0$ plane in the domain, showing the wall particles when the particle separation is $\Delta_x^{wall} = 2.5a$ (left) and $\Delta_x^{wall} = 4.5a$ (right). . . . 152

8.6 (a) A schematic of the domain, with dimensions $L_x \times L_y \times L_z$, and the 3×3 rectangular array of filaments. The beat plane of each filament, defined for $(f, \kappa_0, \beta) = (40, 3\pi/4, 2)$, is aligned with the $\hat{\boldsymbol{x}}-$axis. The distance between neighbouring filaments in the $\hat{\boldsymbol{x}}$ and $\hat{\boldsymbol{y}}$ directions is given by Δ_x^{fil} and Δ_y^{fil} respectively. (b) Schematic of the dynamics of a single filament in an open domain with a no-slip wall at $z = 0$, calculated using the RPY tensor with a no-slip wall and for the parameters $(f, \kappa_0, \beta) = (40, 3\pi/4, 2)$, as found in Chapter 7. The passage of time is indicated by the increasing boldness of the lines. (c) A view of the domain from above. Square markers represent the filament seeding. There are two main types of MCWs. If the MCW direction aligns with the beat plane, we define the wave as symplectic. If the wave direction is perpendicular to the beat plane, we observe a diaplectic wave. 154

8.7 (a) Results of initial value problems of arrays of $1 \times N_f$ and $N_f \times 1$ filaments, corresponding to diaplectic (top) and symplectic (bottom) waves respectively. The initial condition for each simulation is a MCW, based on the dynamics of a single cilium in an infinite domain, bound by a wall at $z = 0$, as found in Chapter 7. The inter-filament spacing is fixed, so increasing N_f also corresponds to increasing the length of the domain. For diaplectic (symplectic) waves, metachronal initial conditions lead to synchronous beating for $N_f \leq 13$ ($N_f \leq 9$). For larger numbers of filaments, the array converges to a MCW. (b) Heat maps of filament phase over 40 periods of the simulation for a 1×10 array of filaments, and visualisations of the dynamics at three points in time. We see that the filaments synchronise their beat. (c) Heat maps of filament phase over 40 periods of the simulation for a 1×15 array of filaments. We see that the solution converges to a metachronal beat within the first 5 periods. 156

8.8 (a) The results of initial value problems with metachronal initial conditions on a 1×15 array of filaments. We fix the spacing of the filaments in the $\hat{\boldsymbol{y}}-$direction to be $\Delta_y^{fil} = L/2$, and vary the inter-filament spacing in the $\hat{\boldsymbol{x}}-$direction, Δ_x^{fil}. We observe MCWs below a critical spacing. Above this spacing, the filaments synchronise. The dynamics of one filament in the array over one period for (b) $\Delta_x^{fil} = L/2$ and (c) $\Delta_x^{fil} = 2L$. The passage of time is indicated by the increasing boldness of the lines. We note that decreasing the separation between filaments changes the dynamics of the filaments at the individual level, adding an aplanar component to the beat. 157

8.9 Heat maps of filament phase at two points in time for 1×60 arrays of filaments with various initial conditions. The left column shows times $t/T \in [10, 15]$, and the right column shows times $t/T \in [100, 105]$. The initial conditions are all MCWs. From top to bottom, the wavelength of the corresponding initial condition is given as $\lambda/L_y = \{1, 1/2, 1/3, 1/4, 1/5\}$. We see that if the wavelength of the solution exceeds the critical value, as found in Figure 8.7a, the MCW is stable. Below this, the dynamics converge to the synchronised solution. 159

8.10 The norm of the mean net velocity, $||\bar{\boldsymbol{V}}^{\infty}||$, for the different stable wavelength solutions found in Figure 8.9. The net velocity is normalised by the norm of the mean for the synchronous beating case. Synchronous beating is categorised by the wavelength $L_y/\lambda = 0$. We observe that the MCW solutions pump more fluid than the synchronous beating solution. Smaller wavelength solutions are more effective at pumping fluid. 160

8.11 Initial conditions for four simulations following the cut-and-stitch approach for MCW generation outlined in Chakrabarti et al. [2022]. (a) Two MCWs spanning half the domain propagating in opposite directions, leading to a total charge of $Q = 0$. (b) Three MCWs, each spanning a third of the domain. Two are propagating in one direction, and the third in the opposite direction, leading to a total charge of $Q = 1$. (c) One MCW spanning a third of the domain. The rest of the filaments are in phase, leading to a total charge of $Q = 1$. (d) Two MCWs, each spanning a third of the domain. The rest of the filaments are in phase, leading to a total charge of $Q = 2$. 162

8.12 Heatmaps showing the phase evolution of each filament starting from the initial conditions outlined in Figure 8.11. We see the solutions converge to (a) a synchronous beating state with charge $Q = 0$, or MCWs with charge (b) $Q = 1$, (c) $Q = 1$ or (d) $Q = 2$. In all cases, we find that this charge is conserved over time. 162

E.1 Proof of permission to reuse Figure 1.1 from Gross et al. [2007]. 206

Chapter 1

Introduction

Hairlike filaments are a universal feature of life, shaping cellular structure and driving motion in fluids. Cilia sweep mucus in our airways, flagella propel microscopic swimmers through liquid, and networks of biopolymer filaments help shape the cell itself. Despite their small size, these filaments exhibit rich dynamics, arising from the interplay of elasticity, hydrodynamic coupling, and applied forcing.

This thesis focuses on understanding the mechanics of filaments driven by internal stresses generated by molecular motors or other biochemical instruments, otherwise known as active filaments. We do this through a minimal yet powerful framework: the follower force model, in which compressive forces act tangentially along the filament centreline. Using numerical simulations and computational dynamical systems techniques, we map out how varying these internal stresses can give rise to buckling into nonlinear behaviours, and explore the transitions and bifurcations between these dynamics. We also investigate the collective dynamics of these filaments, where hydrodynamic coupling between filaments can give rise to non-trivial coordinated states consisting of travelling waves of phase. By studying this simplified model, we aim to capture key features of active filament dynamics which, more generally, provide insight into the mechanisms underlying ciliary actuation and the role of coordination in fluid propulsion at microscopic scales.

We begin by describing manifestations of active filaments in biology, with a focus on biopolymer

filament-motor protein systems including cilia and flagella, before discussing techniques used in the literature to model such problems. With this in hand, we will outline the scope and objectives of this work, as well as the structure of this thesis.

1.1 Biological background

1.1.1 Biopolymer filaments-molecular motor complexes

Biopolymer filaments, such as actin filaments and microtubules, are fundamental structures for transport at the cellular level [Cooper, 2000, Needleman and Shelley, 2019], and are essential in many life processes. These long, slender protein filaments make up the cytoskeleton, with a primary role of maintaining the shape and structure of the cell [Cooper, 2000]. Actin filaments are constructed from helically combined actin proteins, whereas microtubules are constructed of protofilaments: filaments made of pairs (called tubulin dimers) of $\alpha-$ and $\beta-$tubulin, each approximately 8nm in length [Bray, 2000]. Microtubules are typically formed of 13 protofilaments, generating a hollow tube [Bray, 2000]. Schematics of these filaments are shown in Figure 1.1. Both actin filaments and microtubules have a length on the scale of microns, ranging between approximately $0.1\mu\text{m} - 100\mu\text{m}$ for actin filaments [Sept et al., 1999, Schmitz et al., 2005] and $1\mu\text{m} - 100\mu\text{m}$ for microtubules [Bray, 2000, Pampaloni et al., 2006, Ling et al., 2018], and have an approximate diameter of 7nm [Cooper, 2000] and 25nm [Ledbetter and Porter, 1963] respectively. These microfilaments are structured in such a way that they have a polarity, meaning that each end of the filament is distinguishable and denoted either the 'plus' or the 'minus' end (as labelled in Figure 1.1).

In many contexts, microfilaments are driven by molecular motors; proteins which translocate along the filaments, transporting cargo around the cell. There are three categories of cytoplasmic motor proteins which traverse the cytoskeleton, distinguished by their characteristics and the type of filament on which they move; kinesin, myosin and dynein [Schliwa and Woehlke, 2003, Shelley, 2016]. These motors have complex structures, with 'heads' (the part of the motor

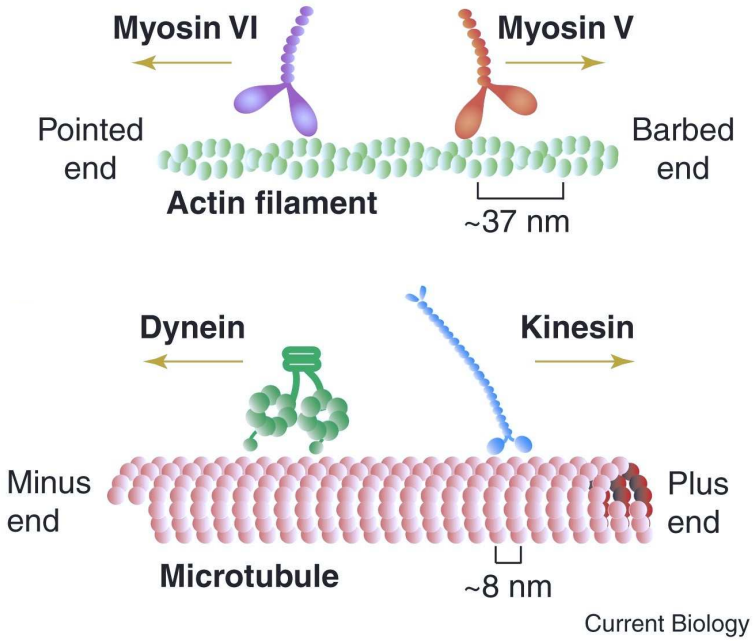


Figure 1.1: Schematic showing the biopolymer filaments described in the text, and the motor proteins which traverse them. (Top) Actin filaments are helical structures with distinguishable ends, given by the 'barbed' or 'pointed' end. Actin filaments are traversed by myosin motor proteins. Different myosin motor proteins can move in either direction along the filament. (Bottom) Microtubules are hollow structures constructed from $\alpha-$ and $\beta-$tubulin. The two ends of the microtubule are distinguishable as the 'plus' end and 'minus' end. Dynein and kinesin are motor proteins which translocate along microtubules, and move in opposite directions. Myosin, dynein and kinesin all have different structures, as shown in the Figure. (Reproduced with permission from Gross et al. [2007]).

which connects to the filament, sometimes called the motor domains, the arms or, in the case of myosin, the actin-binding sites) that may be connected to a cargo-carrying 'tail' through long stalks [Bray, 2000, Cho and Vale, 2012]. Schematics of these motor proteins are shown in Figure 1.1. Motor proteins are responsible for the function of many life processes, often acting cooperatively to carry cargo along filaments [Leduc et al., 2004, Klumpp and Lipowsky, 2005]. Defective motor proteins can lead to infertility [Afzelius, 1976], deafness [Kunishima and Saito, 2010] and many other health issues [Ainsworth, 2007, Maravillas-Montero and Santos-Argumedo, 2012].

These motors are activated chemically, through a process called adenosine triphosphate (ATP) hydrolysis [Schliwa and Woehlke, 2003, Cho and Vale, 2012]; a catabolic reaction which transfers ATP into adenosine diphosphate (ADP) and inorganic phosphate. The energy released from

this process causes a conformational change in the motor head [Schliwa and Woehlke, 2003], which communicates with the motor arms to trigger binding or unbinding from the filament [Cho and Vale, 2012]. This results in a single step of the motor along the filament. The motors take steps of discrete sizes, on the scale of nanometers [Singh et al., 2005, Fehr et al., 2008, Maravillas-Montero and Santos-Argumedo, 2012]. Repeated iterations of this process allows the motors to 'walk' along the filament [Shingyoji et al., 1998], at a speed on the scale of micrometers per second [Kural et al., 2005].

First found in Kühne [1864], around 100 years before dynein was isolated and named for the first time in Gibbons and Rowe [1965], myosin motors are found along actin filaments, taking 36nm steps along the filament in a 'hand-over-hand' fashion [Maravillas-Montero and Santos-Argumedo, 2012]. These myosin-actin systems are ubiquitous in nature; for instance in cytokinesis, the final stage of cell division, where a contractile ring of actin and myosin filaments separates the parent cell into two daughter cells [Cooper, 2000, Hartman and Spudich, 2012, Leite et al., 2019], as well as during muscle contraction, which is generated by the myosin-driven sliding of actin filaments in the muscle cells [Cooper, 2000, Hartman and Spudich, 2012, Squire, 2016]. These complexes are also used in cellular locomotion; actin-rich domains (called lamellipodia) form protrusions constructed of bundles of actin filaments (called filopodia), which a cell can use to 'drag' itself along a substrate [Mattila and Lappalainen, 2008].

Kinesin and dynein, on the other hand, translocate along microtubules [Schliwa and Woehlke, 2003]. Dynein 'walk' towards the minus end of the microtubule (i.e. from the base to the tip) [Sale and Satir, 1977] taking steps of discrete sizes of $8, 16, 24$ or 32nm [Singh et al., 2005], while most kinesin move towards the plus end [Schliwa and Woehlke, 2003], taking uniform steps of size 8nm [Fehr et al., 2008]. Kinesin motors usually move processively in one direction [Gennerich and Vale, 2009], whereas dynein motors have experimentally been observed to perform oscillatory forwards-backwards motion [Shingyoji et al., 1998]. Microtubule-kinesin/dynein systems are also vital in cell division - the motors carry cargo around the cell via microtubules to facilitate timings throughout mitosis [Wojcik et al., 2001], telling the microtubules when to rapidly decrease in length to pull apart the cell's chromosomes, causing the transition between metaphase to anaphase [Goshima and Vale, 2003, Forth and Kapoor, 2017]. Microtubules also

play a key role in intracellular transport at the microscopic scale [Fourriere et al., 2020], for instance in the context of axoplasmic transport [Hirokawa et al., 2010, Holzbaur and Scherer, 2011, Guedes-Dias and Holzbaur, 2019].

The ATP-driven motion of the motors allows them to carry cargo along the filament, hence generating a drag force that the motors experience from the surrounding fluid. By Newton's third law, the motors then exert a compressive force, equal in magnitude, back on the filament. Experimentally, this force has been found to be in the range $1-10$pN [Svoboda and Block, 1994, Schliwa and Woehlke, 2003, Mallik et al., 2004]. In an unbounded domain, this compressive force would cause the filament to slide in the direction of the force. However if the motion of the filament is restricted in some way, for instance if the filament is attached to a cellular body, or if it reaches an obstacle, this forcing can lead to filament buckling. Such buckling has been observed in motility assays where filaments (actin or microtubules) are driven by ATP-fuelled motor proteins (myosin or kinesin/dynein respectively) fixed to rigid surfaces [Allen et al., 1985, Bourdieu et al., 1995, Gittes et al., 1996, Kawamura et al., 2010, Young, 2010, Lam et al., 2014, Isele-Holder et al., 2015]. Undisturbed, the filaments are observed to glide along the surface. However in the presence of an obstacle or surface defect, such as a defective motor protein, part of the filament will become pinned or clamped into place. Then, continued forcing from the motors will cause the filament to buckle.

In another context, it has been observed that collections of motor-driven microtubules can coordinate their motions to generate fluid flows in a process called cytoplasmic streaming [Niwayama et al., 2011, Goldstein and van de Meent, 2015, Klughammer et al., 2018]. An example of this occurs within *Drosophila* oocytes [Gutzeit and Koppa, 1982, Quinlan, 2016, Stein et al., 2021], which are lined internally with kinesin-driven microtubules. Kinesin motors induce compressive forces along each microtubule, which in turn causes each microtubule to buckle. The dynamics of neighbouring microtubules are coupled through the flow within the cell, which allows the filaments to coordinate their motion to generate an overall steady state, in which each microtubule is bent in the same direction. This directs the fluid around the cell, hence allowing for the transportation of important organelles and molecules.

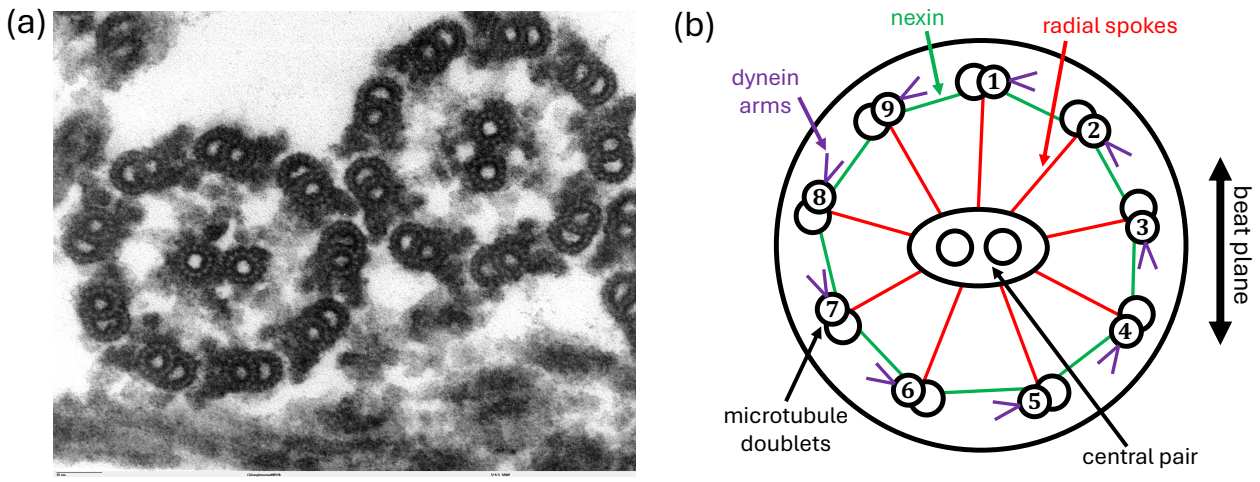


Figure 1.2: (a) An electron microscope image of the cross-section of the two flagella of *Chlamydomonas*, exhibiting the '9+2' axonemal structure in motile cilia. (Image released in the public domain by [Dartmouth Electron Microscope Facility, Dartmouth College]). (b) A schematic of the '9+2' axonemal structure showing the 9 microtubule doublets surrounding the central pair, interconnected by radial spokes, nexin proteins and dynein. In particular, we note the beat plane direction with reference to the orientation of the central pair, and the bridge between microtubule doublets 5 and 6.

1.1.2 Cilia and flagella

Collections of filament-motor protein systems can form more complicated structures, allowing for the transport of fluid flow at a much larger scale. The most notable such example is cilia and flagella: hair-like filaments which act as a fundamental mechanism for cellular propulsion by flagellated cells [Suarez and Pacey, 2006, Bennett and Golestanian, 2015, Pedley et al., 2016, Wan and Goldstein, 2018]. In this thesis we use the terms cilia and flagella interchangeably, though in the literature the term flagellum tends to describe a longer filament, and cilia are associated to shorter filaments, which often exist in patches. The internal structure of cilia is called the axoneme – a superstructure constructed of several microtubule-dynein complexes. In Figure 1.2a, we show an electron microscope image of a cilium cross-section and in Figure 1.2b, we provide a labelled schematic of the axoneme.

Motile cilia are said to have a '9+2' axonemal structure, consisting of a central pair of microtubules, surrounded by nine pairs of microtubules, called microtubule doublets (labelled 1-9 in Figure 1.2b) [Gibbons, 1981, Ishikawa, 2017, Gilpin et al., 2020]. While this structure of the axoneme is common, variations can exist. For instance nodal cilia, which are found in

the embryo in the early stages of its development, have a '9+0' axonemal structure [Nonaka et al., 1998] referring to an absence of the central pair of microtubules. Similarly, some cilia are known to have '9+1', '6+0' or '3+0' structures [Ishikawa, 2017]. It is postulated that these different structures may yield different dynamics [Nonaka et al., 1998, Ishikawa, 2017], though for simplicity we describe only the '9+2' structure further.

The outer microtubule doublets are connected to the inner pair via radial spokes, and adjacent doublets are linked together by nexin proteins and regularly spaced dynein motor proteins. Each motor protein is anchored to one microtubule doublet and can interact with the neighbouring doublet using its heads, commonly referred to as arms in this context. Dynein motor proteins have inner and outer arms; the inner arms are known to control the waveform of the cilium, and the outer arms control the power of the beating [Lin and Nicastro, 2018]. At its base, the axoneme is anchored to a structure called the 'basal body' [Ishikawa et al., 2021], which itself is formed by microtubules.

It was once thought that the motion of cilia was caused by the compression or changing length of the microtubules, where the microtubules on one side of the axoneme would either increase in length or extend, while the microtubules on the opposite side would decrease in length or compress to generate an overall bending of the cilium. However in Satir [1965, 1968], it was shown that this is not the case; rather, motion is caused by interdoublet sliding, which itself is due to the dynein motors. The dynein arms interact with the neighbouring microtubule doublet by walking towards the 'minus' end of the microtubule [Sale and Satir, 1977, Fox and Sale, 1987], causing the doublets to slide adjacent to each other which, due to structural constraints, causes the doublets to bend. Coordination of this bending between doublets across the axoneme allows the cilium to beat.

On an individual level, each cilium undergoes beating oscillations which are able to move fluid. Ciliated or flagellated cells use these beating appendages for propulsion [Brennen and Winet, 1977]. *Chlamydomonas reinhardtii*, for instance, is a single-cell algae which has two flagella (see Figure 1.3a) which are used for this purpose [Bennett and Golestanian, 2015]. *Volvox*, on the other hand, is a multicellular algae with a surface covered by cilia [Pedley et al., 2016].

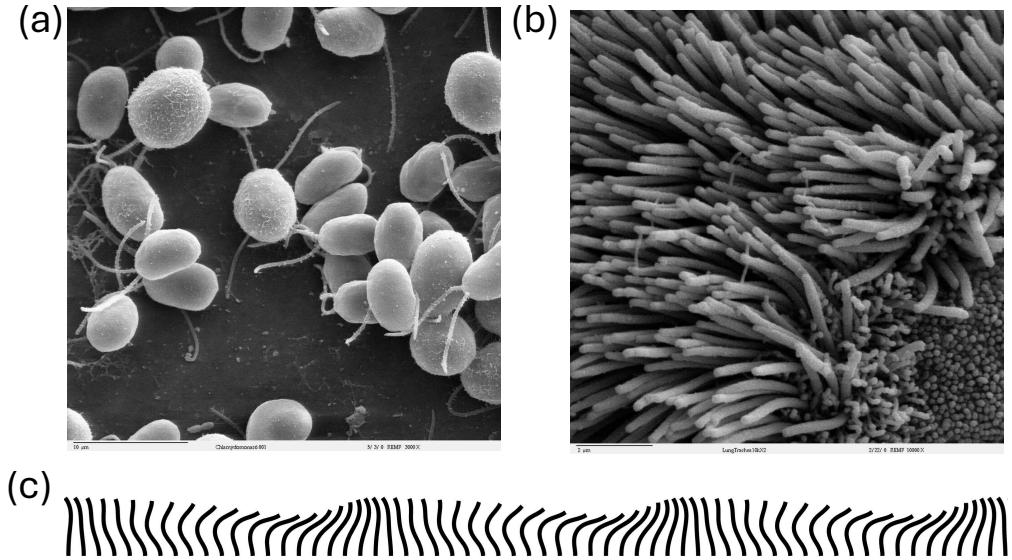


Figure 1.3: Electron microscope images of (a) several *Chlamydomonas* algae, each driven by its two flagella and (b) the ciliated epithelium in the lung trachea. (Both images released in the public domain by [Dartmouth Electron Microscope Facility, Dartmouth College]). (c) A schematic showing a metachronal wave in a 2D array of filaments. Each filament is performing the same periodic beat, with a constant phase shift compared to its neighbour. The individual beat pattern is obtained using the results presented in Chapter 7 for the parameter choice $(f, \kappa_0, \beta) = (40, 3\pi/4, 2)$.

This cilia-controlled locomotion allows microswimmers to capture food [Gilpin et al., 2017], to reorient in response to external stimuli such as light (phototaxis) [Drescher et al., 2010, Bennett and Golestanian, 2015], or the current of the flow (rheotaxis) [Marcos et al., 2012], and to move towards nutrients [Stocker et al., 2008].

Arrays of cilia can coordinate their beating, for instance producing metachronal waves [Han and Peskin, 2018, Chakrabarti et al., 2022, Kanale et al., 2022] whereby a constant phase difference between neighbouring cilia is observed, to control fluid transport at a much larger scale. We show a schematic of a 2D metachronal wave in Figure 1.3c - each individual filament is following the same periodic asymmetric beating pattern, following the methods and results presented in Chapter 7, but at a constant phase difference to its neighbours, generating a travelling wave of phase along the array. This large scale coordination is constantly utilised in the human body, for instance within the brain. The ventricular system is lined with ciliated epithelia called ependyma [Worthington and Cathcart, 1966]. Coordinated beating of this cilia controls the flow of cerebrospinal fluid (CSF) through the four ventricles [Faubel et al., 2016],

allowing for the delivery of essential nutrients and proteins around the brain [Guldbrandsen et al., 2014, Zhang et al., 2015] as well as the disposal of neurotoxins and other waste products in each ventricle [Worthington and Cathcart, 1966]. CSF also transports young neurons from the hypothalamus, where they are made, through the brain and into the central nervous system [Sawamoto et al., 2006], where they play a critical role in the transport of information through the body. The respiratory epithelium which lines our airways is also ciliated (see Figure 1.3b); coordination of ciliary beating is used to transport mucus [Lai et al., 2009], a viscous fluid itself containing passive filaments called mucins [Hwang et al., 1969, Quraishi et al., 1998], to protect the lungs and clean the airways via a process called mucociliary clearance [Sleigh et al., 1988, Smith et al., 2008b].

Cilia and flagella also have several roles within the fertility process; sperm cells have a single flagellum which is used for propulsion [Gaffney et al., 2011]. Inside the uterus, sperm navigate through viscous cervical mucus to reach the egg, where sperm with low motility are filtered out by the mucus [Suarez and Pacey, 2006]. After fertilisation, cilia then have an important role in embryo development [Smith et al., 2019]. Nodal cilia are present in the embryo, and have been found to undergo a bent-over whirling motion [Nonaka et al., 1998], pushing fluid to the left and hence breaking the left-right symmetry in the embryo [Smith et al., 2008a, 2011, 2019]. Defective cilia can lead to a range of diseases called ciliopathies [Tobin and Beales, 2009]. Symptoms of ciliopathies can range from infertility [Gaffney et al., 2011], to rare conditions such as *situs inversus* [Afzelius, 1976], whereby the major organs form on the wrong side of the body, like a mirror image. In Afzelius [1976], it was hypothesised that *situs inversus* could be a result of defective nodal cilia, which resulted in the left-right symmetry of the embryo being broken incorrectly. It was also suggested that the immotile cilia was a result of defective dynein arms in the axoneme, indicating that the nanoscale effect of the motors can thus cause macroscopic health conditions, and highlighting the multiscale nature of ciliary beating.

1.2 Modelling approaches in the literature

Because of their significance in life processes and driving fluid flow at the microscale, there has been much attention to the problem of generating mathematical models for these filament-motor protein complexes. These usually consider an elastic filament with internal active stresses, associated with the action of the motor proteins. Throughout this work we will focus on filaments clamped to a point in space, usually to a no-slip wall, for instance representing cilia anchored onto the epithelial cells lining the respiratory system in the human body. While freely suspended filaments exist in many biological contexts, such as motility assays, the buckling dynamics of interest arise in these systems when the filaments hit a defect or a wall. Therefore we focus on clamped-free end conditions henceforth. Thus, there are two essential components to generating such a model: (i) a model which describes the deformations of a slender elastic filament submerged in viscous fluid, and (ii) incorporating the effect of the molecular motors. We discuss the former in Chapter 2, and focus now on the different forcing mechanisms used in the literature in the two biological contexts discussed earlier.

1.2.1 Single filament-motor protein systems

There are many levels of detail that can be employed when generating a model for motor proteins translocating along the length of a filament, which we summarise through a schematic in Figure 1.4. For instance, the oscillatory forward-backwards motion undertaken by dynein along a microtubule [Shingyoji et al., 1998] suggests the existence of a feedback mechanism involving ATP concentration, which would account for the motor changing direction. Various ideas for regulating dynein motility have been proposed in the literature: from simple models, where the motor switches direction after a predetermined load is reached [Brokaw, 2000], to more complex models, where a low concentration of ATP in the system triggers a second, passive, pathway of diffusion for the dynein [Goedecke and Elston, 2005]. Stochastic Monte Carlo processes have also been used to model the diffusion of dynein along the microtubule, taking into account the different possible step-sizes the dynein motors can take [Singh et al.,

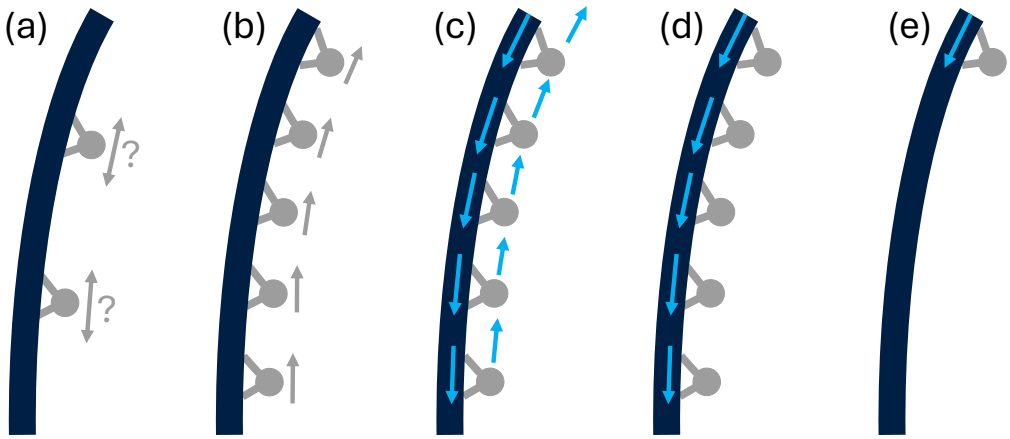


Figure 1.4: Schematics of some of the modelling choices used in the literature to describe the motion of motors translocating along the length of an elastic filament. (a) Some studies allow the motors to move forwards or backwards along the filament, depending on some predetermined criteria such as applied load. (b) Other works neglect details of the biomechanics, by considering a collection of motors moving in one direction only. (c) Following this approach, models may incorporate the effect of the motors through the drag the motors experience from carrying cargo (which we call backflow), as well as compressive forces on the filament (follower forces) following Newton's third law. This model can be simplified further, by (d) neglecting the backflow and then (e) considering only one such follower force, corresponding to the effect of a single molecular motor at the filament tip.

2005].

However, the goal of many mathematical models is to generalise the coordinated effect of these molecular motors, and so do not include explicit regulation mechanisms. Rather, they consider the effect of a motor protein carrying cargo along the filament length. The motor experiences a drag force from the surrounding fluid as it carries cargo along the filament. By Newton's third law, it applies a compressive force on the filament, equal in magnitude but opposite in direction. This idea has motivated a host of forcing mechanisms for active filaments (see Figure 1.4b-e). The simplest incarnation of such a model considers a single 'follower force' at the tip of a filament clamped at its base to a rigid surface. The follower force is a compressive force directed along the filament tangent and thus changes orientation as the filament undergoes deformations [Herrmann and Bungay, 1964, Bayly and Dutcher, 2016, De Canio et al., 2017, Ling et al., 2018, De Canio, 2019, Westwood and Keaveny, 2021, Schnitzer, 2025]. This is representative of a single motor residing at the tip of a filament, and neglects the effect of the drag force from the motor carrying the cargo. When the magnitude of the follower force exceeds a critical value, the filament buckles, giving rise to a time dependent state. If restricted to a

plane, the motion results in planar beating [De Canio et al., 2017], whereas in three dimensions (3D), buckling gives rise to a fully 3D whirling regime [Ling et al., 2018]. Increasing the forcing further, we recover the planar beating solution in both cases. For the highest values of the forcing, we observe a regime of higher order modes, which we call writhing.

The basic follower force model has been extended in several ways to more accurately represent the motor-driven forcing of microfilaments, to consider a distribution of follower forces along the filament length, corresponding to a continuous array of motor proteins [Bayly and Dutcher, 2016, De Canio, 2019], as well as including the back flow associated to the motors acting on the surrounding fluid, either through an opposite force exerted on the fluid [De Canio, 2019, Stein et al., 2021] or through a slip flow on the filament surface [Laskar and Adhikari, 2017, Ozdemir et al., 2025]. Other studies have investigated the follower force model varying the boundary conditions and initial configuration of the filament [Sekimoto et al., 1995, Fatehiboroujeni et al., 2018, Fily et al., 2020, Fatehiboroujeni et al., 2021] to replicate buckling dynamics observed experimentally in motility assays [Sekimoto et al., 1995, Fily et al., 2020]. This model has further been used to generate a model of the microtubules which line the *Drosophila* oocyte, and recovers the collective bending motion that leads to cytoplasmic streaming [Stein et al., 2021], although further ingredients, such as the back-flow, are necessary to recover this non-trivial steady state.

1.2.2 Modelling cilia and flagella

Ciliary dynamics are inherently a multiscale problem [Gilpin et al., 2020]; the overall beat pattern is due to the synchronisation of individual cilia, which itself is determined at the axoneme level. Axoneme bending is generated by interdoublet sliding, which, in turn, is driven by the ATP-fuelled piconewton scale forces of dynein motors (see Figure 1.5a). To generate this large scale beating, it is hypothesised that a regulatory feedback loop must exist to coordinate the dynein motors across the axoneme, to allow for bending in different directions. This has been confirmed experimentally, using mussel gill cilia, where it was found that dynein on one side of the axoneme (say, on microtubule doublets 1-5) will be active during the recovery stroke

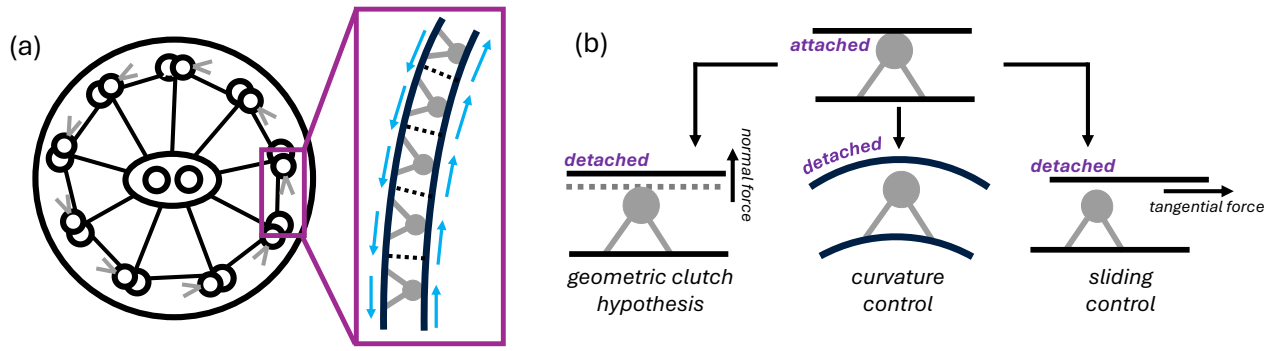


Figure 1.5: (a) A schematic of the axoneme adapted from Figure 1.2. Inset, right shows how two microtubule doublets (represented by black lines) are connected by dynein motors (grey) and nexin proteins (dashed lines). The dynein motors cause relative sliding of the doublets (blue arrows), which leads to bending. In order to generate overall bending of the cilium, there must be some sort of regulation on dynein attachment across the axoneme. (b) Schematics showing three possible regulatory controls for dynein attachment from the literature: (left) the geometric clutch hypothesis, where dynein motors detach if interdoublet distances are too large, (middle) curvature control, where dynein motors respond to the local curvature of the cilium, and (right) sliding control, where the dynein motors detach if forces tangential to the microtubule doublets, leading to sliding, are too large.

while the rest are passive [Satir and Matsuoka, 1989]. During the effective stroke, this switches. In the literature, it is therefore common to incorporate a regulatory feedback system typically based on the idea of a switch-point model whereby the dynein on one half of the axoneme are 'switched on' for the first half of the beat, while the others are 'switched off'. This then flips for the second half of the beat, allowing for oscillations [Satir and Matsuoka, 1989, Satir et al., 2014].

There are many ideas for the mechanism controlling this feedback loop [Woolley, 2010], which we summarise in Figure 1.5b. These include the 'geometric clutch model' [Lindemann, 1994a,b, 2002] which suggests dynein activity is dictated by the local geometry of the axoneme. Specifically, as the axoneme bends the interdoublet distances change. Beyond some critical distance, the dynein on one half of the axoneme must detach - the spacing is too large for the dynein arms to physically attach to the neighbouring microtubule doublet - and hence are 'switched off'. Then, when the interdoublet spacing is small enough, these dynein can reattach and are activated again, allowing the axoneme to bend in the other direction. Other suggestions include that dynein activity depends on the local curvature of the axoneme, either instantaneously or by a delayed curvature [Dillon and Fauci, 2000, Dillon et al., 2003, 2007, Woolley, 2010, Sartori

et al., 2016, Gallagher et al., 2023], or through a sliding control mechanism [Riedel-Kruse et al., 2007, Müller et al., 2008, Chakrabarti and Saintillan, 2019a,b, Chakrabarti et al., 2022], in which dynein are regulated by the relative sliding between doublets, and detach if the forces opposing sliding become too large. Many of these studies are able to recover dynamics similar to those observed experimentally.

Despite much research into these different regulation techniques and their resulting dynamics, the exact mechanism behind dynein regulation in the axoneme is still not clear. In fact, in Bayly and Dutcher [2016] and Woodhams et al. [2022], it was shown that a regulatory feedback system is not actually necessary to achieve buckling and non-trivial oscillations in axoneme models. By implementing the effect of dynein through steady follower forces in opposite directions in multi-filament models of the axoneme, the authors managed to obtain a non-trivial beating state, raising questions regarding the true process behind cilia actuation: while experiments suggest that motor regulation does occur, we are yet to understand the fundamental differences between models with explicit feedback mechanisms, and those without regulation.

With this in mind, many studies choose to coarse-grain the filament forcing and neglect the internal mechanics of the axoneme. Motivated by the planar beating dynamics observed by cilia in nature, some studies take the instability-driven filament approach and utilise the follower force model to model a cilium [Westwood and Keaveny, 2021, Wang et al., 2023]. By choosing the magnitude of the follower force appropriately, one can recover symmetric beating dynamics [De Canio et al., 2017, Ling et al., 2018]. This has been used in a variety of studies, for instance in investigating the effect of fluid rheology on the dynamics and pumping ability of cilia [Wang et al., 2023, Link et al., 2024], and computationally modelling spherical ciliated swimmers [Westwood and Keaveny, 2021]. While the follower force model does not capture the shear-driven actuation we know to be responsible for ciliary beating, it offers a route to self-sustained oscillations that are reminiscent of ciliary dynamics, while remaining relatively easy to interpret and computationally scalable.

1.3 Thesis aims and outline

As discussed in the previous sections, the follower force model has been used in many applications, from understanding filament-motor systems - recreating the cytoplasmic streaming observed in fruit fly oocytes [Stein et al., 2021] and dynamics in motility assays [Sekimoto et al., 1995, Fily et al., 2020] - to investigating ciliary dynamics, through the buckling of the axoneme leading to the onset ciliary beating [Bayly and Dutcher, 2016, Woodhams et al., 2022], assessing the effect of the surrounding fluid on the cilium [Wang et al., 2023, Link et al., 2024], to evaluating the collective behaviours of cilia [Westwood and Keaveny, 2021]. Despite being used to model such complex phenomena, a complete understanding of the simplest case – a single follower force applied at the tip of an individual filament – is not yet known. Furthermore, the follower force model lacks a symmetry-breaking mechanism as required to generate asymmetric cilia-like beating patterns for net fluid flow, and the collective dynamics of these follower force-actuated systems of filaments are not known. While there are several commonly used additions to this model, such as the inclusion of entrainment flow from the motors, or more complex forcing mechanisms that more accurately reflect the biological systems we discuss, it is of interest to first establish a more thorough understanding of the simplest model - a single follower force at the filament tip. These problems outline the aims of this thesis, which is organised as follows.

In Chapters 2 and 3, we outline the relevant mathematical background required to understand how to model filaments submerged in viscous fluids, as well as the essential Krylov subspace algorithms necessary for this thesis. The methodology for modelling a follower force-driven filament is described in Chapter 4, as well as the computational dynamical systems tools and techniques we use to identify steady and time-periodic solutions, and to analyse their stability. We also describe a bisection method, for uncovering unstable solution branches lying in regions of bistability.

In Chapter 5, we combine these methods to explore the state space for a single filament clamped to a no-slip planar wall, driven at its tip by a single follower force. Through the use of initial value problems and bifurcation analysis, we investigate the effect of the strength of the follower

force, as well as the filament aspect ratio, on the resulting dynamics. We extend these results in Chapter 6 by analysing three extensions of the basic follower force model; the effect of the twist modulus, a different forcing, in which follower forces are distributed along the filament length, and a different boundary condition, where now the filament is clamped at its base in an unbounded domain. In doing so, we characterise the differences in the emerging dynamics of these models. A comprehensive extension of this chapter can be found in Ozdemir et al. [2025].

In Chapter 7, we explore how structural anisotropies and asymmetries, motivated by the structure of cilia, can be utilised to recover asymmetric cilia-like beating patterns. We do this through the inclusion of anisotropic bending moduli and an intrinsic curvature within the filament. With these symmetry-breaking mechanisms in hand, in Chapter 8 we consider collective filament dynamics. First we describe how our model is extended to triply periodic boundary conditions, and validate our boundary condition implementation through a small study with Poiseuille flow. We then present results on the coordination of arrays of asymmetrically beating filaments, with a particular focus on metachronal wave formation. Finally, in Chapter 9 we discuss our overall conclusions and outlook.

The results in Chapter 5 and part of Chapter 6 have been published [Clarke et al., 2024]. The other results reported in Chapter 6 are partially based on the dissertation author's contributions to a publication currently under consideration for publication [Ozdemir et al., 2025]. Finally, the results in Chapter 7 have been published [Clarke et al., 2025].

Chapter 2

Background - Stokes Flow

In this chapter, we cover the essential mathematical background surrounding modelling fluid-structure interactions at low Reynolds number. We begin by reviewing Stokes flow, in which viscosity dominates over inertia, by deriving the Stokes equations and describing fundamental solutions of these equations. Next, we discuss how this couples to slender filaments submerged in the flow, to allow us to capture the hydrodynamic effects on the structure.

2.1 Stokes equations

The Navier–Stokes equations are momentum and mass conservation equations which describe the motion of fluids. In general, these are given by

$$\rho\left(\frac{\partial}{\partial t}+\boldsymbol{u}\cdot\nabla\right)\boldsymbol{u}=\nabla\cdot\boldsymbol{\sigma}+\boldsymbol{f}, \tag{2.1}$$

$$\frac{\partial\rho}{\partial t}+\nabla\cdot(\rho\boldsymbol{u})=0, \tag{2.2}$$

where $\boldsymbol{u}(\boldsymbol{x},t)$ is the fluid velocity, $\rho(\boldsymbol{x},t)$ is the fluid density, $\boldsymbol{\sigma}(\boldsymbol{x},t)$ is the stress tensor, and $\boldsymbol{f}(\boldsymbol{x},t)$ is the external applied force (per unit volume) acting on the fluid. For Newtonian fluids,

we assume that the stress tensor takes the constitutive form,

$$\boldsymbol{\sigma} = -p\boldsymbol{I} + \eta\left(\nabla\boldsymbol{u} + (\nabla\boldsymbol{u})^T\right), \tag{2.3}$$

where η is the dynamic viscosity of the fluid, p is the pressure and $\boldsymbol{I}$ is the identity. We assume that the fluid is incompressible, meaning that the density is constant. Therefore, for incompressible Newtonian fluids, these equations read:

$$\rho\left(\frac{\partial}{\partial t} + \boldsymbol{u}\cdot\nabla\right)\boldsymbol{u} = \eta\nabla^2\boldsymbol{u} - \nabla p + \boldsymbol{f}, \tag{2.4}$$

$$\nabla\cdot\boldsymbol{u} = 0, \tag{2.5}$$

We nondimensionalise these equations using typical velocity and length scales U and L, implying time and pressure dimensionalisations L/U and $\eta U/L$ respectively. The corresponding conservation of momentum equation reads

$$Re\left[\frac{\partial}{\partial t} + \boldsymbol{u}\cdot\nabla\right]\boldsymbol{u} = \nabla^2\boldsymbol{u} - \nabla p + \boldsymbol{f}. \tag{2.6}$$

Here, Re is a non-dimensional parameter called the Reynolds number, given by

$$Re = \frac{\rho U L}{\eta}. \tag{2.7}$$

The Reynolds number measures the relative importance of the fluid viscosity to inertia. We call ρ/η the kinematic viscosity, which is approximately $\rho/\eta \approx 10^{-6}m^2s^{-1}$ for water. In the context of cilia and flagella, the Reynolds number is small. For instance, a human sperm cell has characteristic scales $U \approx 200\mu m s^{-1}$ and $L \approx 50\mu m$, which correspond to a Reynolds number $Re \approx 10^{-2}$ [Lauga and Powers, 2009]. As such, we consider the case of low Reynolds number, where viscous stresses dominate and we neglect inertia. Due to the negligible effects of the fluid

inertia in this case, we are left with the Stokes equations:

$$\nabla^2 \boldsymbol{u} - \nabla p + \boldsymbol{f} = \boldsymbol{0}, \tag{2.8a}$$

$$\nabla \cdot \boldsymbol{u} = 0. \tag{2.8b}$$

There are several important properties of Stokes flow that we will utilise in this thesis. The first is a result of the lack of explicit time dependence in Eq (2.8). As such, the fluid velocity depends only on the current configuration of the problem, and thus is determined instantaneously by the boundary conditions. This property is often called instantaneity, and has the particular consequence that a solution with reversed time, i.e. $t \to -t$, is also a solution of the Stokes equations. Therefore, motion in Stokes flow is reversible, which has the pertinent consequence that reciprocal motion (meaning temporally symmetric motion) in Stokes flow cannot induce net flow at zero Reynolds number, as described in Purcell's scallop theorem [Purcell, 1977]

A second result that will be of interest in this chapter is the Lorentz reciprocal theorem [Kim and Karrila, 1991]. This states that for a region of fluid, V, bounded by a surface S, if $\boldsymbol{u}(\boldsymbol{x})$ and $\tilde{\boldsymbol{u}}(\boldsymbol{x})$ are both velocity fields satisfying the Stokes equations, with corresponding stress fields $\boldsymbol{\sigma}$ and $\tilde{\boldsymbol{\sigma}}$ respectively, then

$$\int_S u_i \tilde{\sigma}_{ij} n_j \, dS = \int_S \tilde{u}_i \sigma_{ij} n_j \, dS. \tag{2.9}$$

Finally, we note that in removing the inertia terms in the Navier–Stokes equations, the remaining Stokes equations are linear. As such, superimposed solutions of the Stokes equations are also solutions to the Stokes equations. Boundary conditions are required to identify a particular solution. If the value of the flow on the boundary is prescribed, then the Stokes flow solution is unique; a consequence of the minimum dissipation theorem, which states that Stokes flow minimises the viscous dissipation. Furthermore, this linearity implies the existence of Green's functions for Stokes equations — a fact which has great value when seeking solutions of the flow, as we shall explore further in the next section.

2.1.1 Fundamental solutions of the Stokes equations

Because of the linearity of the Stokes equations, explicit solutions can be found for some example cases. The simplest solution of Stokes equations follows when considering a point force of strength $\boldsymbol{F}$, situated at a position $\boldsymbol{Y}$. In this case, the flow field satisfies:

$$\nabla^2 \boldsymbol{u} - \nabla p = -\boldsymbol{F}\delta(\boldsymbol{x} - \boldsymbol{Y}), \tag{2.10a}$$

$$\nabla \cdot \boldsymbol{u} = 0 \tag{2.10b}$$

with boundary conditions, $|\boldsymbol{u}|, p \to 0$ as $|\boldsymbol{x} - \boldsymbol{Y}| \to \infty$.

The corresponding flow field is known as the Stokeslet [Kim and Karrila, 1991, Pozrikidis, 1992, Lauga and Powers, 2009], and can be expressed:

$$u_i(\boldsymbol{x}) = \mathcal{G}_{ij}(\boldsymbol{x} - \boldsymbol{Y})F_j, \tag{2.11}$$

where we sum over pairs of indices, following standard Einstein summation notation. Here $\mathcal{G}_{ij}$ is the Oseen's tensor, given by

$$\mathcal{G}_{ij}(\boldsymbol{r}) = \frac{1}{8\pi\eta}\left(\frac{\delta_{ij}}{r} + \frac{r_i r_j}{r^3}\right) \tag{2.12}$$

where $r = |\boldsymbol{r}|$. Oseen's tensor, also sometimes known as the Stokeslet tensor, is a Green's function for the flow - if, instead, we had a continuous distribution of force with density $\boldsymbol{f}(\boldsymbol{x})$, then the flow field would be given by

$$u_i(\boldsymbol{x}) = \int_V \mathcal{G}_{ij}(\boldsymbol{x} - \boldsymbol{y}) f_j(\boldsymbol{y})\, d^3\boldsymbol{y}. \tag{2.13}$$

We can also calculate the corresponding pressure and stress tensor for the point force,

$$p(\boldsymbol{x}) = \mathcal{P}_i(\boldsymbol{x} - \boldsymbol{Y})F_i, \tag{2.14}$$

$$\sigma_{ij}(\boldsymbol{x}) = \mathcal{K}_{ijk}(\boldsymbol{x} - \boldsymbol{Y})F_k \tag{2.15}$$

where

$$\mathcal{P}_i(\boldsymbol{r}) = \frac{r_i}{4\pi r^3}, \tag{2.16}$$

$$\mathcal{K}_{ijk}(\boldsymbol{r}) = -\frac{3r_i r_j r_k}{4\pi r^5}. \tag{2.17}$$

It follows immediately by the linearity of the Stokes equations, that if we now consider N point forces where the i^{th} point is located at $\boldsymbol{Y}_i$ with force $\boldsymbol{F}_i$ for $i = 1, ..., N$, then the flow field is given by

$$u_i(\boldsymbol{x}) = \sum_{i=1}^{N} \mathcal{G}_{ij}(\boldsymbol{x} - \boldsymbol{Y}_i)F_j. \tag{2.18}$$

As a further consequence of the linearity of the Stokes equations, derivatives of the Green's function are also solutions to the Stokes equations. This has a connection to the multipole expansion, as we discuss in Section 2.1.3. The first derivative leads to the force dipole (also known as the Stokeslet dipole or Stokes doublet),

$$u_i(\boldsymbol{x}) = \frac{\partial}{\partial x_k}\mathcal{G}_{ij}(\boldsymbol{x} - \boldsymbol{Y})D_{jk}. \tag{2.19}$$

Here, D_{jk} is the force dipole tensor. This solution corresponds to the limit as two point forces, of strength $\pm\boldsymbol{F}$, approach each other. Specifically, if the points are located at $\boldsymbol{Y}_\pm = \boldsymbol{Y} \pm \boldsymbol{d}/2$, this is the limit as $|\boldsymbol{d}| \to 0$. In this case, the solution is given by

$$\begin{aligned} u_i(\boldsymbol{x}) &= \mathcal{G}_{ij}(\boldsymbol{x} - \boldsymbol{Y}_+)F_j - \mathcal{G}_{ij}(\boldsymbol{x} - \boldsymbol{Y}_-)F_j \\ &= -d_k\frac{\partial}{\partial x_k}\mathcal{G}_{ij}(\boldsymbol{x} - \boldsymbol{Y})F_j + O(|\boldsymbol{d}|^2) \\ &= [(\boldsymbol{d}\cdot\nabla)\boldsymbol{\mathcal{G}}(\boldsymbol{x} - \boldsymbol{Y})\cdot\boldsymbol{F}]_i + O(|\boldsymbol{d}|^2), \end{aligned} \tag{2.20}$$

i.e. Eq (2.19), where $D_{jk} = -F_j d_k$, and $\boldsymbol{\mathcal{G}}(\boldsymbol{r}) = (\boldsymbol{I} + \boldsymbol{r}\boldsymbol{r}^T/r^2)/8\pi\eta r$ is the tensor form of Oseen's tensor, given in component form in Eq (2.12).

The force dipole tensor can be decomposed into symmetric and antisymmetric parts via

$$D_{jk} = S_{jk} + A_{jk} + \frac{1}{3} D_{ii}\delta_{jk}. \tag{2.21}$$

The trace, $D_{ii}\delta_{jk}/3$, has no contribution to the solution since $\nabla \cdot \boldsymbol{\mathcal{G}} = 0$, as a consequence of the fluid being incompressible. The symmetric, traceless component, S_{jk}, is called the Stresslet. In terms of the dipole connection discussed above, this would correspond to taking $\boldsymbol{F} \propto \boldsymbol{d}$. The stresslet represents a purely straining motion of the fluid, and gives rise to the flow field:

$$u_i(\boldsymbol{x}) = \frac{1}{8\pi\eta r^3}\left(\frac{3(\boldsymbol{F}\cdot\boldsymbol{x})(\boldsymbol{d}\cdot\boldsymbol{x})x_i}{r^2} + (\boldsymbol{d}\cdot\boldsymbol{F})x_i\right). \tag{2.22}$$

The antisymmetric part is called the rotlet, and is related to the total torque on the particle, $\boldsymbol{T}$, via the relation $T_i = \varepsilon_{ijk}D_{jk}$. In the force dipole context, this torque is therefore given by $\boldsymbol{T} = \boldsymbol{d} \times \boldsymbol{F}$. The rotlet corresponds to the solution when a point torque is inserted into the flow, and yields the flow field

$$u_i(\boldsymbol{x}) = \frac{1}{8\pi\eta r^3}\varepsilon_{ijk}T_k x_k. \tag{2.23}$$

Further derivatives of the Greens function give higher order moments, such as the quadrupole and the octopole. However from the definition of the Green's function, it is clear that the solutions decay faster as we take subsequent derivatives; the Stokeslet solution decays like $O(1/r)$, the rotlet, stresslet and force dipole decay like $O(1/r^2)$, and the force quadrupole like $O(1/r^3)$. As such, only the first few derivatives are commonly used to represent flows in the far field.

Solutions involving boundaries can be obtained via superpositions of the fundamental solutions, following the linearity of the Stokes equations. One domain which will be of particular interest in this thesis is the case in which a no-slip planar wall is present. For a point force submerged in this domain, Blake [1971a] shows, using the method of images, that the fundamental solution is constructed using a Stokeslet, force dipole and a quadrupole. This combination of solutions combine to ensure no slip on the boundary.

2.1.2 Boundary integral formulation - single layer potential

If we consider a surface S in the fluid, then Kim and Karrila [1991] and Pozrikidis [1992] show using the reciprocal theorem that the flow field can be expressed in terms of the Green's function:

$$u_i(\boldsymbol{x}) = -\int_S \mathcal{G}_{ij}(\boldsymbol{x}-\boldsymbol{y})\sigma_{jk}n_k \, dS(\boldsymbol{y}) + \int_S u_j(\boldsymbol{y})\mathcal{K}_{jik}(\boldsymbol{x}-\boldsymbol{y})\hat{n}_k(\boldsymbol{y}) \, dS(\boldsymbol{y}), \tag{2.24}$$

where $\boldsymbol{f} = \boldsymbol{\sigma}\cdot\hat{\boldsymbol{n}}$ is the force density exerted by the fluid on the surface, $\hat{\boldsymbol{n}}$ is the unit normal pointing into the fluid and $\mathcal{K}_{ijk}$ is the stress field associated to the Stokeslet, given in Eq (2.17). The first term in Eq (2.24) is called the single layer potential, and the second term is the double-layer potential. When considering a rigid particle, Kim and Karrila [1991] show that the second term integrates to zero, meaning that the flow field is described only by the surface traction on the boundary of the surface.

Boundary integral methods are based on solving this integral representation of the flow in the context of multiple suspended particles in bounded or unbounded domains [Pozrikidis, 1992]. However, integrating the Green's function can be computationally expensive. We can circumvent this using Taylor series, as we discuss in the next section.

2.1.3 Multipole expansion

To avoid directly integrating the Green's function in Eq (2.24), in the far-field flow we can Taylor expand $\boldsymbol{\mathcal{G}}(\boldsymbol{x}-\boldsymbol{y})$ around the particle centre, $\boldsymbol{y}=\boldsymbol{Y}$. As we will see, this leads to some helpful simplifications. Considering a single rigid particle with centre $\boldsymbol{Y}$ and surface S, using the previous section we know that the fluid velocity can be expressed using the single layer potential,

$$u_i(\boldsymbol{x}) = -\int_S \mathcal{G}_{ij}(\boldsymbol{x}-\boldsymbol{y})\sigma_{jk}n_k \, dS(\boldsymbol{y}). \tag{2.25}$$

Defining $\boldsymbol{\zeta} = \boldsymbol{y} - \boldsymbol{Y}$ and Taylor expanding the Green's function about the particle centre gives,

$$\boldsymbol{\mathcal{G}}(\boldsymbol{x} - \boldsymbol{y}) = \boldsymbol{\mathcal{G}}(\boldsymbol{x} - \boldsymbol{Y}) - (\boldsymbol{\zeta} \cdot \nabla)\boldsymbol{\mathcal{G}}(\boldsymbol{x} - \boldsymbol{Y}) + O(|\boldsymbol{\zeta}|^2). \tag{2.26}$$

Substituting this into our equation for the flow field leads us to an expansion called the multipole expansion,

$$u_i(\boldsymbol{x}) = -\mathcal{G}_{ij}(\boldsymbol{x} - \boldsymbol{Y})F_j + \frac{\partial}{\partial x_k}\mathcal{G}_{ij}(\boldsymbol{x} - \boldsymbol{Y})D_{jk} + \cdots . \tag{2.27}$$

where

$$F_j = \int_S (\boldsymbol{\sigma} \cdot \hat{\boldsymbol{n}})_j \, dS(\boldsymbol{y}) \tag{2.28}$$

is the hydrodynamic force on the particle, and

$$D_{jk} = \int_S \zeta_k (\boldsymbol{\sigma} \cdot \hat{\boldsymbol{n}})_j \, dS(\boldsymbol{y}) \tag{2.29}$$

is the dipole tensor introduced in Section 2.1.1. As previously mentioned, the trace of the dipole tensor has no effect on the flow as a consequence of the incompressibility of the flow. Therefore the remaining dipole tensor can be split into symmetric and antisymmetric components:

$$S_{jk} = \frac{1}{2}\left(D_{jk} + D_{kj}\right) - \frac{1}{3}D_{ii}\delta_{jk} \tag{2.30}$$

$$A_{jk} = \frac{1}{2}\left(D_{jk} - D_{kj}\right) \tag{2.31}$$

where S_{jk} is the symmetric component, leading to the stresslet, and A_{jk} is the antisymmetric component, leading to the torque. The hydrodynamic torque of the particle is given by

$$\begin{aligned} T_i &= \int_S \varepsilon_{ijk}\zeta_j (\boldsymbol{\sigma} \cdot \hat{\boldsymbol{n}})_k \, dS(\boldsymbol{y}) \\ &= \int_S \varepsilon_{ikj}\zeta_k (\boldsymbol{\sigma} \cdot \hat{\boldsymbol{n}})_j \, dS(\boldsymbol{y}) \\ &= -\int_S \varepsilon_{ijk}\zeta_k (\boldsymbol{\sigma} \cdot \hat{\boldsymbol{n}})_j \, dS(\boldsymbol{y}) \\ &= -\varepsilon_{ijk}A_{jk}, \end{aligned} \tag{2.32}$$

where the second line follows from swapping the k and j indices, and the third follows by the definition of the Levi–Civita symbol. By a similar calculation, we can also show that

$$A_{jk} = -\frac{1}{2}\varepsilon_{jkl}T_l. \tag{2.33}$$

Therefore, the contribution of the rotlet in the multipole expansion is given by

$$\begin{aligned}\frac{\partial}{\partial x_k}\mathcal{G}_{ij}(\boldsymbol{x}-\boldsymbol{Y})A_{jk} &= -\frac{1}{2}\varepsilon_{jkl}T_l\frac{\partial}{\partial x_k}\mathcal{G}_{ij} \\ &= \frac{1}{2}T_l\left(\varepsilon_{lkj}\frac{\partial}{\partial x_k}\mathcal{G}_{ij}\right) \\ &= \frac{1}{2}\left[\boldsymbol{T}\cdot(\nabla\times\boldsymbol{\mathcal{G}})\right]_i.\end{aligned} \tag{2.34}$$

We can combine these results to formulate the multipole expansion:

$$u_i(\boldsymbol{x}) = \mathcal{G}_{ij}F_j + \left(\frac{1}{2}\varepsilon_{lkj}\frac{\partial}{\partial x_k}\mathcal{G}_{ij}\right)T_l + \frac{1}{2}\left(\frac{\partial}{\partial x_k}\mathcal{G}_{ij} + \frac{\partial}{\partial x_j}\mathcal{G}_{ik}\right)S_{jk} + \cdots, \tag{2.35}$$

where the coefficient of the third term is the symmetric part of the dipole tensor, as a result of the symmetry of the Stresslet.

The expansion continues to higher order derivatives of Oseen's tensor, with the next term corresponding to a force quadrupole. As discussed in Section 2.1.1, higher order modes decay faster in the far field, and so usually it is sufficient to only consider the first few terms. In fact, in some cases, the flow field can be determined completely by a few terms in the multipole expansion. For instance, the flow field corresponding to a single rigid sphere translating in a flow at a velocity $\boldsymbol{U}^\infty$ can be written in terms of a Stokeslet and a degenerate quadrupole [Guazzelli and Morris, 2011],

$$\begin{aligned}\boldsymbol{u}(\boldsymbol{x}) &= \frac{3a}{4}\left(\frac{1}{r}\boldsymbol{I} + \frac{1}{r^3}\boldsymbol{x}\boldsymbol{x}^T\right)\boldsymbol{U}^\infty + \frac{3a^3}{4}\left(\frac{1}{3r^3}\boldsymbol{I} - \frac{1}{r^5}\boldsymbol{x}\boldsymbol{x}^T\right)\boldsymbol{U}^\infty \\ &= 6\pi\eta a\left(1 + \frac{a^2}{6}\nabla^2\right)\boldsymbol{\mathcal{G}}(\boldsymbol{x}-\boldsymbol{Y})\cdot\boldsymbol{U}^\infty,\end{aligned} \tag{2.36}$$

where the $6\pi\eta a\boldsymbol{U}^\infty$ coefficient is to ensure the boundary condition, $\boldsymbol{u}(\boldsymbol{x})|_{r=a} = \boldsymbol{U}^\infty$, is satisfied.

Similarly, for a sphere rotating with angular velocity $\boldsymbol{\Omega}$, the flow field is given by a rotlet [Guazzelli and Morris, 2011]:

$$\boldsymbol{u}(\boldsymbol{x}) = \frac{a^3}{r^3}\,\boldsymbol{\Omega} \times \boldsymbol{x}, \tag{2.37}$$

where the rotlet solution has been multiplied by $8\pi\eta a^3$ to satisfy the boundary condition, $\boldsymbol{u}(\boldsymbol{x})|_{r=a} = \boldsymbol{\Omega} \times \boldsymbol{x}$.

2.1.4 Faxéns laws

By either comparing the solutions in Eqs (2.36) and (2.37) to the multipole expansion, Eq (2.35), or by explicitly calculating the surface traction in the examples of translating and rotating spheres, we can recover the hydrodynamic force for a sphere translating through flow,

$$\boldsymbol{F} = 6\pi\eta a \boldsymbol{U}, \tag{2.38}$$

and the hydrodynamic torque for a rotating sphere,

$$\boldsymbol{T} = 8\pi\eta a^3 \boldsymbol{\Omega}. \tag{2.39}$$

These relationships can be generalised for a rigid spherical particle in a more general flow. By considering the reciprocal theorem with one solution corresponding to a particle of radius a and centre $\boldsymbol{Y}$ moving with velocity $\boldsymbol{U}$ in otherwise quiescent unbounded flow, and the second solution corresponding to a stationary particle of radius a responding to a Stokeslet situated externally to the particle, Kim and Karrila [1991] derive Faxén's equations:

$$\boldsymbol{F} = 6\pi\eta a \left[\left(1 + \frac{a^2}{6}\nabla^2 \right) \boldsymbol{u}(\boldsymbol{x}) \Big|_{\boldsymbol{x}=\boldsymbol{Y}} - \boldsymbol{U} \right], \tag{2.40}$$

$$\boldsymbol{T} = 8\pi\eta a^3 \left[\frac{1}{2}\,\nabla \times \boldsymbol{u}(\boldsymbol{x})|_{\boldsymbol{x}=\boldsymbol{Y}} - \boldsymbol{\Omega} \right]. \tag{2.41}$$

Here, $\boldsymbol{U}$ and $\boldsymbol{\Omega}$ are the translational and rotational velocities of the particle respectively, $\boldsymbol{F}$ and $\boldsymbol{T}$ are the hydrodynamic force and torque exerted by the fluid on the particle, and $\boldsymbol{u}$ is the

disturbance velocity, owing to the disturbance in the flow. These tell us that additional terms are required to relate the drag on the particle to the particle's velocity and angular velocity, corresponding to the curvature of the flow. A similar expression exists for the stresslet, but this will not be used in this thesis and so is omitted for brevity.

2.1.5 The mobility problem

In previous sections, we demonstrated how the translational and angular velocities of particles are linearly related to the hydrodynamic forces and torques exerted by the particles onto the surrounding fluid. For N particles, we can extend this linear relationship. If particle i exerts a force, $\boldsymbol{F}_i$, and torque, $\boldsymbol{T}_i$, on the fluid, the corresponding velocity, $\boldsymbol{V}_i$, and angular velocity, $\boldsymbol{\Omega}_i$, of the particle is given by

$$\begin{bmatrix} \boldsymbol{V} \\ \boldsymbol{\Omega} \end{bmatrix} = \mathcal{M} \begin{bmatrix} \boldsymbol{F} \\ \boldsymbol{T} \end{bmatrix}, \tag{2.42}$$

where here,

$$\boldsymbol{F} = \begin{bmatrix} \boldsymbol{F}_1 \\ \vdots \\ \boldsymbol{F}_N \end{bmatrix}, \quad \boldsymbol{T} = \begin{bmatrix} \boldsymbol{T}_1 \\ \vdots \\ \boldsymbol{T}_N \end{bmatrix}, \quad \boldsymbol{V} = \begin{bmatrix} \boldsymbol{V}_1 \\ \vdots \\ \boldsymbol{V}_N \end{bmatrix}, \quad \boldsymbol{\Omega} = \begin{bmatrix} \boldsymbol{\Omega}_1 \\ \vdots \\ \boldsymbol{\Omega}_N \end{bmatrix}, \tag{2.43}$$

are $3N \times 1$ vectors containing the forces, torques, velocities and angular velocities of each particle respectively, and $\mathcal{M}$ is a $6N \times 6N$ matrix, called the mobility matrix. As a result of the reciprocal theorem, $\mathcal{M}$ is symmetric [Kim and Karrila, 1991]. The mobility matrix is also positive semidefinite, which follows from the fact that the viscous dissipation cannot be negative. The entries of this matrix depend on the domain, as well as the size and position of the particles. Eq (2.42) couples the behaviour of all particles in the domain.

In the case where the particles do not interact with each other (an approximation we will make later in forming an approximate Jacobian), we expect the mobility matrix to be diagonal, with

entries corresponding to Stokes drag. Specifically the mobility matrix will take the block form,

$$\mathcal{M} = \begin{bmatrix} \frac{1}{6\pi\eta a}\boldsymbol{I} & \mathbf{0} \\ \mathbf{0} & \frac{1}{8\pi\eta a^3}\boldsymbol{I} \end{bmatrix} \tag{2.44}$$

where $\boldsymbol{I}$ is the $3N \times 3N$ identity matrix, and $\mathbf{0}$ is a $3N \times 3N$ matrix of zeroes.

2.2 Filament hydrodynamics

There are several levels of detail that have been adopted in the literature to solve for the motion of a slender filament in low Reynolds number flow. Bead and rower models [Vilfan and Jülicher, 2006, Uchida and Golestanian, 2010, Friedrich and Jülicher, 2012, Meng et al., 2021, Chakrabarti et al., 2022, Bennett, 2025] lie at the coarsest scale of filament modelling, and are usually used to model collections of cilia. They exploit the periodic beating of cilia to reduce the filament to a single sphere, modelled as a point force, moving along a prescribed closed path. As such, these models assume that the individual beat pattern is fixed for all time, and only the frequency of the beat can change. While this can be useful for a coarse grained view of collective cilia dynamics, allowing for an investigation into synchronisation in large ciliary arrays, this fails to capture changes in beat patterns we may expect to observe, for instance due to the density of cilia present, or to external forces. Therefore more detail is required to provide insights into filament deformation, particularly at the level of an individual filament.

To capture the full dynamics of a filament, we need to resolve along the whole filament length. A key component of this will be resolving the hydrodynamic coupling between the fluid and the filament. One method to resolve fluid-structure interactions is resistive force theory (RFT) [Gray and Hancock, 1955], which assumes a linear relation between the hydrodynamic force and the fluid velocity perpendicular or parallel to the filament, i.e.

$$\boldsymbol{F}_\perp = \zeta_\perp \boldsymbol{u}_\perp, \quad \boldsymbol{F}_\parallel = \zeta_\parallel \boldsymbol{u}_\parallel, \tag{2.45}$$

where $\zeta_\perp, \zeta_\parallel$ are local drag coefficients depending on the aspect ratio of the filament and the fluid viscosity. As the aspect ratio tends to infinity, the ratio of these coefficients tends to a constant, $\zeta_\perp/\zeta_\parallel \to 2$, reflecting the notion that it is easier to pull a slender object through viscous fluid along its axis, rather than perpendicular to it. RFT neglects non-local hydrodynamic effects, for instance the effect of neighbouring filaments, or even of other parts of the same filament. As such, further details are needed to accurately capture the long-range hydrodynamic interactions between filaments.

Slender body theory (SBT) [Keller and Rubinow, 1976] does exactly this through the inclusion of a non-local operator, which encapsulates the hydrodynamic interactions between different parts of the filament, as well as with other filaments in the flow. SBT is based on the assumption of a large aspect ratio (so filament length is sufficiently larger than filament radius), using the method of matched asymptotics to match the far-field flow (which exploits the large aspect ratio assumption to solve for the flow around a line distribution of singularities) to the flow field near the filaments surface (where we can solve for the flow around a cylinder). A boundary integral approach to SBT is given in Johnson [1980].

Originally used to model the flow around heart valves [Peskin, 1973], the immersed boundary method (IBM) is an alternative method to measure fluid-structure interactions [Peskin, 2002]. The method couples Lagrangian points on the immersed structure to an Eulerian fluid grid, by enforcing that each material point moves with the local fluid velocity interpolated from the grid. In the discrete setting, regularised delta functions are used to transfer forces from the structure to the fluid, and to interpolate velocities from the fluid back to the structure. In this framework, the Stokes equations are solved with a forcing term representing the forces generated by the immersed structure, and the resulting velocity field is used to update the motion of the filament or boundary [Peskin, 2002].

As we discuss further in Section 4.2, the approach we take is to discretise along the filament centreline and then account for the impact of each segment on the flow. It would be possible to model each segment as a point force; in this case, regularisation is important to handle the singularities in the Green's function. The regularised Stokeslet method [Cortez, 2001, Cortez

et al., 2005] achieves this by smoothing the forcing to act over a small region, opposed to a singular point, which removes the computational issues surrounding singular functions. In this thesis, we instead regularise by modelling each segment as a particle with a finite hydrodynamic radius, from which we can extract the particle velocities and angular velocities from the forces and torques these particles exert on the flow using the mobility relationship, Eq (2.42). There are several ways to solve this mobility problem for a collection of spherical particles suspended in the flow; we discuss the methods relevant to this thesis now.

2.2.1 Rotne-Prager-Yamakawa (RPY) tensor

The Rotne-Prager-Yamakawa (RPY) tensor [Rotne and Prager, 1969, Yamakawa, 1970, Wajnryb et al., 2013] provides an explicit expression for the mobility matrix, $\mathcal{M}$, allowing for the pairwise contributions to the particle velocities to be evaluated explicitly. Therefore RPY is applied via matrix-vector multiplication, which can be easily parallelised. Below, we present the RPY matrix formulation for particles submersed in an unbounded domain [Wajnryb et al., 2013] as well as in half-space [Swan and Brady, 2007].

Consider N particles submerged in the flow, where particle i has centre $\boldsymbol{Y}_i$ and radius a_i for $i = 1, ..., N$. For the purpose of this thesis, we can assume these particles do not overlap, and are each of the same radius, i.e. $a_i = a$ for all i, though extensions of the methodology exist to allow for overlapping [Wajnryb et al., 2013], as well as for particles of different radii [Zuk et al., 2014]. RPY finds an explicit relation for the mobility matrix $\mathcal{M}$, by considering the velocity of a particular particle,

$$\begin{bmatrix} \boldsymbol{V}_i \\ \boldsymbol{\Omega}_i \end{bmatrix} = \sum_{j=1}^{N} \begin{bmatrix} M_{ij}^{\boldsymbol{VF}} & M_{ij}^{\boldsymbol{VT}} \\ M_{ij}^{\boldsymbol{\Omega F}} & M_{ij}^{\boldsymbol{\Omega T}} \end{bmatrix} \begin{bmatrix} \boldsymbol{F}_j \\ \boldsymbol{T}_j \end{bmatrix}. \tag{2.46}$$

where $M_{ij}^{\boldsymbol{VF}}, M_{ij}^{\boldsymbol{VT}}, M_{ij}^{\boldsymbol{\Omega F}}$ and $M_{ij}^{\boldsymbol{\Omega T}}$ are 3×3 matrices, encapsulating the contribution of the hydrodynamic force and torque of particle j, $\boldsymbol{F}_j$ and $\boldsymbol{T}_j$ respectively, to the translational velocity and angular velocity of particle i, $\boldsymbol{V}_i$ and $\boldsymbol{\Omega}_i$ respectively.

In order to construct these matrices, we consider the disturbance velocity from particle j.

Expanding the surface tractions gives the following expression for the disturbance velocity for $|\boldsymbol{x} - \boldsymbol{Y}_j| > a$ [Swan and Brady, 2007]:

$$\boldsymbol{u}^d(\boldsymbol{x}; \boldsymbol{Y}_j) = \left(1 + \frac{a^2}{6}\nabla^2\right) \boldsymbol{\mathcal{G}}(\boldsymbol{x} - \boldsymbol{y})|_{\boldsymbol{y}=\boldsymbol{Y}_j} \cdot \boldsymbol{F}_j + \frac{1}{2}\nabla \times \boldsymbol{\mathcal{G}}(\boldsymbol{x} - \boldsymbol{y})|_{\boldsymbol{y}=\boldsymbol{Y}_j} \cdot \boldsymbol{T}_j + \cdots \tag{2.47}$$

where the higher order terms decay faster in the far-field. We note that the degenerate quadrupole term is included, to account for the finite size of the particles.

Then, using Faxén's laws and exploiting the linearity of the Stokes equations, we can recover the translational and angular velocities on particle i, $\boldsymbol{V}_i$ and $\boldsymbol{\Omega}_i$ respectively, as:

$$\boldsymbol{U}_i = \frac{\boldsymbol{F}_i}{6\pi\eta a} + \sum_{j\neq i} \left(1 + \frac{a^2}{6}\nabla^2\right) \boldsymbol{u}^d(\boldsymbol{x}; \boldsymbol{Y}_j)\Big|_{\boldsymbol{x}=\boldsymbol{Y}_i}, \tag{2.48}$$

$$\boldsymbol{\Omega}_i = \frac{\boldsymbol{T}_i}{8\pi\eta a^3} + \frac{1}{2}\sum_{j\neq i} \nabla \times \boldsymbol{u}^d(\boldsymbol{x}; \boldsymbol{Y}_j)\big|_{\boldsymbol{x}=\boldsymbol{Y}_i}. \tag{2.49}$$

By expanding the disturbance velocity, truncating at the torque term, and comparing to Eq (2.46), we can recover the mobility entries:

$$M_{ij}^{\boldsymbol{VF}} = \left(1 + \frac{a^2}{6}\nabla_{\boldsymbol{x}}^2\right)\left(1 + \frac{a^2}{6}\nabla_{\boldsymbol{y}}^2\right) \boldsymbol{\mathcal{G}}(\boldsymbol{x} - \boldsymbol{y})\Big|_{\boldsymbol{x}=\boldsymbol{Y}_i, \boldsymbol{y}=\boldsymbol{Y}_j} \tag{2.50a}$$

$$M_{ij}^{\boldsymbol{VT}} = \left(1 + \frac{a^2}{6}\nabla_{\boldsymbol{x}}^2\right)\left(\frac{1}{2}\nabla_{\boldsymbol{y}} \times \boldsymbol{\mathcal{G}}(\boldsymbol{x} - \boldsymbol{y})\right)\Big|_{\boldsymbol{x}=\boldsymbol{Y}_i, \boldsymbol{y}=\boldsymbol{Y}_j} \tag{2.50b}$$

$$M_{ij}^{\boldsymbol{\Omega F}} = \frac{1}{2}\nabla_{\boldsymbol{x}} \times \left(1 + \frac{a^2}{6}\nabla_{\boldsymbol{y}}^2\right) \boldsymbol{\mathcal{G}}(\boldsymbol{x} - \boldsymbol{y})\Big|_{\boldsymbol{x}=\boldsymbol{Y}_i, \boldsymbol{y}=\boldsymbol{Y}_j} \tag{2.50c}$$

$$M_{ij}^{\boldsymbol{\Omega T}} = \frac{1}{2}\nabla_{\boldsymbol{x}} \times \left(\frac{1}{2}\nabla_{\boldsymbol{y}} \times \boldsymbol{\mathcal{G}}(\boldsymbol{x} - \boldsymbol{y})\right)\Big|_{\boldsymbol{x}=\boldsymbol{Y}_i, \boldsymbol{y}=\boldsymbol{Y}_j} \tag{2.50d}$$

if $i \neq j$, corresponding to the effect of particle j, and

$$M_{ij}^{\boldsymbol{VF}} = \frac{1}{6\pi\eta a}\boldsymbol{I} \tag{2.51a}$$

$$M_{ij}^{\boldsymbol{VT}} = M_{ij}^{\boldsymbol{\Omega F}} = \boldsymbol{0} \tag{2.51b}$$

$$M_{ij}^{\boldsymbol{\Omega T}} = \frac{1}{8\pi\eta a^3}\boldsymbol{I} \tag{2.51c}$$

if $i = j$, corresponding to the self-mobility of the particle. Note that at this point, we have

not specified a particular definition of the Green's function, $\mathcal{G}$. In fact, it is this freedom which allows us to include boundaries, such as a no-slip planar wall, easily. In the case of unbounded flow, the Green's function is Oseen's tensor, given in Eq (2.12). Defining $\boldsymbol{r}_{ij} = \boldsymbol{Y}_i - \boldsymbol{Y}_j$, $r_{ij} = |\boldsymbol{r}_{ij}|$, the above expressions can be expanded to give the corresponding mobility matrices for $i \neq j$:

$$M_{ij}^{\boldsymbol{VF}} = \frac{1}{8\pi\eta|\boldsymbol{r}_{ij}|}\left(\left(1+\frac{2a^2}{3|\boldsymbol{r}_{ij}|^2}\right)\boldsymbol{I} + \left(1-\frac{2a^2}{|\boldsymbol{r}_{ij}|^2}\right)\frac{\boldsymbol{r}_{ij}\boldsymbol{r}_{ij}^T}{|\boldsymbol{r}_{ij}|^2}\right) \tag{2.52a}$$

$$M_{ij}^{\boldsymbol{\Omega T}} = \frac{1}{16\pi\eta|\boldsymbol{r}_{ij}|^3}\left(3\boldsymbol{r}_{ij}\boldsymbol{r}_{ij}^T - \boldsymbol{I}\right), \tag{2.52b}$$

$$M_{ij}^{\boldsymbol{VT}} = M_{ij}^{\boldsymbol{\Omega F}} = \frac{1}{8\pi\eta|\boldsymbol{r}_{ij}|^2}\,\boldsymbol{\epsilon}\cdot\boldsymbol{r}_{ij}, \tag{2.52c}$$

where $\boldsymbol{\epsilon}$ is the 3D Levi–Civita. The self-mobility, on the other hand, is again given by Eq (2.51).

To incorporate the effect of a no-slip planar wall [Swan and Brady, 2007], we instead use the Blakelet solution as introduced in Section 2.1.1. In this case, we include the additional interactions of the 'image' solutions. We show the corresponding entries in the mobility matrix in Appendix B.

2.2.2 FCM

Alternative matrix-free methods exist for solving Eq (2.42). One such method is the Force Coupling Method (FCM) [Maxey and Patel, 2001, Lomholt and Maxey, 2003], which makes use of the Fast Fourier transform to solve the Stokes equations in Fourier space, with the consequence that the resulting domain is triply periodic. As such, FCM will have particular use in Chapter 8, where we consider arrays of filaments, to avoid edge defects.

Consider again the problem of N particles each with identical radius a, and centre $\boldsymbol{Y}_i$, submerged in a viscous fluid filling a domain of size $L_x \times L_y \times L_z$. Say each particle experiences a force, $\boldsymbol{F}_i$, and torque, $\boldsymbol{T}_i$. The collective forces and torques of all particles induce a translational and rotational velocity, $\boldsymbol{V}_i$ and $\boldsymbol{T}_i$ respectively. FCM relates the effect of these forces and torques

to the fluid through the force density

$$\boldsymbol{f}(\boldsymbol{x}) = \sum_{i=1}^{N} \boldsymbol{F}_i \Delta(\boldsymbol{x} - \boldsymbol{Y}_i) - \frac{1}{2} \sum_{i=1}^{N} \boldsymbol{T}_i \times \nabla \tilde{\Delta}(\boldsymbol{x} - \boldsymbol{Y}_i), \tag{2.53}$$

where

$$\Delta(\boldsymbol{x}) = \frac{1}{(2\pi\sigma^2)^{3/2}} \exp\left(-\frac{|\boldsymbol{x}|^2}{2\sigma^2}\right), \tag{2.54}$$

$$\tilde{\Delta}(\boldsymbol{x}) = \frac{1}{(2\pi\tilde{\sigma})^{3/2}} \exp\left(-\frac{|\boldsymbol{x}|^2}{2\tilde{\sigma}^2}\right). \tag{2.55}$$

This can be considered as a regularised force multipole expansion, where we have smoothed the δ-functions with Gaussian functions.

Then, we solve the Stokes equations subject to triply periodic boundary conditions to find the corresponding fluid velocity. This is done in Fourier space; taking the Fourier transform in 3 dimensions, we can rearrange the Stokes equations to yield the fluid velocity and pressure:

$$\hat{\boldsymbol{u}}(\boldsymbol{k}) = \frac{1}{\eta k^2} \left(\boldsymbol{I} - \frac{\boldsymbol{k}\boldsymbol{k}^T}{k^2}\right) \hat{\boldsymbol{f}}, \tag{2.56}$$

$$\hat{p}(\boldsymbol{k}) = -\frac{i(\boldsymbol{k} \cdot \hat{\boldsymbol{f}})}{k^2}, \tag{2.57}$$

where $\boldsymbol{k} = (k_x, k_y, k_z)$ and $k = |\boldsymbol{k}|$.

Once $\hat{\boldsymbol{u}}(\boldsymbol{k})$ has been obtained, we take the reverse Fourier transform to calculate the fluid velocity, $\boldsymbol{u}(\boldsymbol{x})$. Then, we can recover the particle velocities by spatially averaging the fluid flow over the domain V using the kernel,

$$\boldsymbol{V}_i = \int_V \boldsymbol{u}(\boldsymbol{x}) \Delta(\boldsymbol{x} - \boldsymbol{Y}_i)\, d^3\boldsymbol{x}, \tag{2.58}$$

$$\boldsymbol{\Omega}_i = \frac{1}{2} \int_V (\nabla \times \boldsymbol{u}(\boldsymbol{x}))\, \tilde{\Delta}(\boldsymbol{x} - \boldsymbol{Y}_i)\, d^3\boldsymbol{x}. \tag{2.59}$$

In order to match with Stokes drag for a single particle in an unbounded domain, the 'width' of the Gaussians are chosen to be $\sigma = a/\sqrt{\pi}$ and $\tilde{\sigma} = a/(6\sqrt{\pi})^{1/3}$ [Maxey and Patel, 2001,

Lomholt and Maxey, 2003].

While FCM provides a robust method to calculate the particles mobility, a large disadvantage is the number of grid points needed to resolve the Gaussian sufficiently. To circumvent this issue, in larger simulations we use an alternative method, called Fast FCM, which truncates the Gaussian to allow for much larger grid sizes to be used [Su and Keaveny, 2024]. We discuss this method in the next section. Another issue with FCM is related to a net flux condition which arises as a result of the singularity of the flow field in Fourier space, given by Eq (2.56) when $k = 0$ — we discuss this further in Section 8.2.3.

Several extensions of FCM exist. For instance, the force density in Eq (2.53) is based on a multipole expansion of the force, truncated at the dipole term, and neglecting the stresslet. This stresslet could be included [Lomholt and Maxey, 2003], but comes at an additional computational cost and has been found to have little difference in the resulting dynamics [Schöller, 2019, Schöller et al., 2021]. A secondary, particularly relevant, extension of FCM considers adding boundaries to the domain. We discuss this further in Section 8.2.1.

2.2.3 Fast FCM

In order to avoid the small grid sizes we are restricted to using with FCM, for larger simulations later in this thesis we will employ an alternative solver, aptly named Fast FCM [Su and Keaveny, 2024]. In essence, Fast FCM splits the mobility solver into two components,

$$\mathcal{M}_{FCM} = \underbrace{\tilde{\mathcal{M}}_{FCM}}_{\text{mobility, larger } \sigma} + \underbrace{\mathcal{M}_{FCM} - \tilde{\mathcal{M}}_{FCM}}_{\text{correction}}. \tag{2.60}$$

The first component $\tilde{\mathcal{M}}_{FCM}$, is FCM mobility, as discussed in Section 2.2.2. However, in order to reduce the number of grid points we need, we modify the Gaussian and replace the Gaussian widths $\sigma, \tilde{\sigma}$ with $\Sigma > \sigma, \tilde{\Sigma} > \tilde{\sigma}$ respectively. Therefore significantly fewer grid-points are required to resolve these Gaussians sufficiently, allowing for much larger domains to be considered.

The velocities calculated by this coarse-grained mobility will be corrected by the second component of the mobility, $\mathcal{M}_{FCM} - \tilde{\mathcal{M}}_{FCM}$. This correction is applied pairwise to particles within a defined cut-off radius, accounting for the error due to the larger Gaussian width. This correction term, and the modified form of the Gaussians, are chosen such that the correction to the mobility of any two particles vanishes exponentially as their separation tends to infinity [Su and Keaveny, 2024]. Fine-tuning parameters such as the Gaussian widths and the cut-off radius allow us to calculate dynamics within different tolerances [Su and Keaveny, 2024].

Chapter 3

Background - Krylov Subspace Methods

Throughout this thesis, there will be several occasions at which we want to solve linear systems, $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$, or eigenvalue problems, $\boldsymbol{A}\boldsymbol{v} = \lambda\boldsymbol{v}$, for some matrix $\boldsymbol{A}$, vectors $\boldsymbol{x}$, $\boldsymbol{b}$ and $\boldsymbol{v}$, and scalar λ. As such, we will make use of methods which iteratively solve these problems. In this work, we will work in a context in which $\boldsymbol{A}$ is not known explicitly, but rather the action of $\boldsymbol{A}$ on a vector, $\boldsymbol{A}\boldsymbol{x}$ for arbitrary $\boldsymbol{x}$, is known. Thus, it is beneficial to use a framework which bypasses the need for explicit knowledge of $\boldsymbol{A}$, and allows us to construct a solution regardless of this.

With this in mind, we use Krylov subspace methods to tackle these problems [Trefethen and Bau, 1997], which work by projecting our matrix onto a lower dimensional Krylov subspace. The $k-$dimensional Krylov subspace of $\boldsymbol{A} \in \mathbb{C}^{n\times n}$ is defined as

$$\mathcal{K}_k(\boldsymbol{A}, \boldsymbol{b}) = \text{span}\{\boldsymbol{b}, \boldsymbol{A}\boldsymbol{b}, \cdots, \boldsymbol{A}^{k-1}\boldsymbol{b}\}. \tag{3.1}$$

As we increase k, $\boldsymbol{A}^{k-1}\boldsymbol{b}$ converges to the largest dominant eigenvector of $\boldsymbol{A}$, and as such the

matrix

$$\boldsymbol{K}_k = \begin{bmatrix} | & | & & | \\ \boldsymbol{b} & \boldsymbol{A}\boldsymbol{b} & \cdots & \boldsymbol{A}^{k-1}\boldsymbol{b} \\ | & | & & | \end{bmatrix} \tag{3.2}$$

becomes increasingly ill-conditioned. Therefore, many Krylov subspace methods involve generating an orthogonal basis for the Krylov subspace. We first discuss an algorithm that iteratively generates an orthogonal basis, called the Arnoldi method. The Arnoldi method has further applications in approximating eigenvalues, as we discuss in Section 4.4. Next, we discuss the Generalised Minimal Residual method (GMRES); a Krylov-subspace method for solving linear systems, by seeking solutions comprised of the Krylov subspace. Both methods have the advantage that the matrix, $\boldsymbol{A}$, does not need to be known explicitly. Further details of these methods, and more general direct and iterative methods for solving these types of problems, are described in great depth in Trefethen and Bau [1997].

3.1 Arnoldi method

Given $\boldsymbol{A} \in \mathbb{C}^{n\times n}$, the Arnoldi method aims to iteratively reduce $\boldsymbol{A}$ to its upper Hessenberg form, $\boldsymbol{H}$. By upper Hessenberg, we mean that the $(i,j)^{th}$ entry of $\boldsymbol{H}$ is zero if $j > i+1$. Specifically, the Arnoldi method aims to find an orthogonal matrix, $\boldsymbol{Q}$, and an upper Hessenberg matrix, $\boldsymbol{H}$, such that

$$\boldsymbol{A} = \boldsymbol{Q}\boldsymbol{H}\boldsymbol{Q}^*. \tag{3.3}$$

This is achieved iteratively, by projecting $\boldsymbol{A}$ onto the Krylov subspace $\mathcal{K}_k(\boldsymbol{A}, \boldsymbol{b})$ for a given $\boldsymbol{b} \in \mathbb{C}^{n\times 1}$, at each iteration k. Then, at iteration k, the Arnoldi method constructs the matrix

$$\boldsymbol{Q}_k = \begin{bmatrix} | & | & & | \\ \boldsymbol{q}_1 & \boldsymbol{q}_2 & \cdots & \boldsymbol{q}_k \\ | & | & & | \end{bmatrix} \in \mathbb{C}^{n\times k} \tag{3.4}$$

such that $\mathcal{K}_k(\boldsymbol{A}, \boldsymbol{b}) = \text{span}\{\boldsymbol{q}_1, \cdots, \boldsymbol{q}_k\}$ and $\boldsymbol{Q}_k\boldsymbol{Q}_k^* = \boldsymbol{I}_n$, $\boldsymbol{Q}_k^*\boldsymbol{Q}_k = \boldsymbol{I}_k$. The method also constructs an upper Hessenberg matrix $\boldsymbol{H}_k$, given by

$$\boldsymbol{H}_k = \begin{bmatrix} h_{11} & h_{12} & \cdots & h_{1k} \\ h_{21} & h_{22} & \cdots & h_{2k} \\ 0 & h_{32} & \ddots & \vdots \\ \vdots & \ddots & \ddots & h_{k,k} \\ 0 & \cdots & 0 & h_{k+1,k} \end{bmatrix} \in \mathbb{C}^{(k+1)\times k} \tag{3.5}$$

such that

$$\boldsymbol{A} = \boldsymbol{Q}_{k+1}\boldsymbol{H}_k\boldsymbol{Q}_k^*. \tag{3.6}$$

The first vector in the Krylov subspace basis is taken to be $\boldsymbol{q}_1 = \boldsymbol{b}/||\boldsymbol{b}||$, so that $\mathcal{K}_1(\boldsymbol{A}, \boldsymbol{b}) = \text{span}\{\boldsymbol{q}_1\}$. Each iteration then proceeds as follows. At iteration k, we calculate $\tilde{\boldsymbol{q}}_{k+1} = \boldsymbol{A}\boldsymbol{q}_k$, and obtain the new row of upper Hessenberg coefficients using the expression

$$h_{jk} = \boldsymbol{q}_j^*\tilde{\boldsymbol{q}}_{k+1} \tag{3.7}$$

for $j = 1, ..., k$. This expression follows from Eqs (3.5) and (3.6). Then the new vector is orthogonalised against the current basis using Gram-Schmidt,

$$\tilde{\boldsymbol{q}}_{k+1} = \tilde{\boldsymbol{q}}_{k+1} - \sum_{j=1}^{k} h_{jk}\boldsymbol{q}_j. \tag{3.8}$$

The final entry of the upper Hessenberg matrix is then found via

$$h_{k+1,k} = ||\tilde{\boldsymbol{q}}_{k+1}||, \tag{3.9}$$

and the new basis vector is given by

$$\boldsymbol{q}_{k+1} = \tilde{\boldsymbol{q}}_{k+1}/||\tilde{\boldsymbol{q}}_{k+1}||. \tag{3.10}$$

This algorithm breaks down if $h_{k+1,k} = 0$. In this event, Eq (3.8) implies that the new vector

is entirely described by the basis at the previous iteration, i.e. $\tilde{\boldsymbol{q}}_{k+1} \in \mathcal{K}_k(\boldsymbol{A}, \boldsymbol{b})$. This implies $\mathcal{K}_{k+1}(\boldsymbol{A}, \boldsymbol{b}) \subseteq \mathcal{K}_k(\boldsymbol{A}, \boldsymbol{b})$, and so the basis is complete at iteration k, and the algorithm will end.

There are several advantages of using the Arnoldi method, including that an explicit form of the matrix, $\boldsymbol{A}$, does not need to be known. Furthermore, the algorithm does not have to run to completion - once a 'satisfactory' reduction into upper Hessenberg form is reached (we will describe one definition of satisfactory in the next section), we can terminate the algorithm. In many applications, particularly for large matrix sizes n, we find that the algorithm can be terminated for $k \ll n$, reducing the computational cost of the Arnoldi method.

3.2 Generalised minimal residual method (GMRES)

GMRES is a method which aims to find the approximate solution of the $n \times n$ linear system, $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$, where $\boldsymbol{A} \in \mathbb{C}^{n\times n}$ and $\boldsymbol{b}, \boldsymbol{x} \in \mathbb{C}^{n\times 1}$. GMRES is an iterative method, where at iteration k, the solution $\boldsymbol{x}$ is approximated by the vector $\boldsymbol{x}_k \in \mathcal{K}_k(\boldsymbol{A}, \boldsymbol{b})$, which satisfies:

$$\boldsymbol{x}_k = \underset{\boldsymbol{x}_k \in \mathcal{K}_k(\boldsymbol{A},\boldsymbol{b})}{\arg\min} ||\boldsymbol{A}\boldsymbol{x}_k - \boldsymbol{b}||. \tag{3.11}$$

In practice, this is achieved by noting that $\boldsymbol{x}_k \in \mathcal{K}_k(\boldsymbol{A}, \boldsymbol{b})$ if and only if $\boldsymbol{x}_k = \boldsymbol{Q}_k\boldsymbol{y}$ for some vector $\boldsymbol{y} \in \mathbb{C}^{n\times 1}$, where $\boldsymbol{Q}_k$ is found via Arnoldi, as in Eq (3.4). Then, the residual we want to be minimise can be rewritten:

$$\begin{aligned}
\min_{\boldsymbol{x}_k \in \mathcal{K}_k} ||\boldsymbol{A}\boldsymbol{x}_k - \boldsymbol{b}|| &= \min_{\boldsymbol{y}} ||\boldsymbol{A}\boldsymbol{Q}_k\boldsymbol{y} - \boldsymbol{b}|| \\
&= \min_{\boldsymbol{y}} ||\boldsymbol{Q}_{k+1}\boldsymbol{H}_k\boldsymbol{y} - \boldsymbol{b}|| \\
&= \min_{\boldsymbol{y}} ||\boldsymbol{H}_k\boldsymbol{y} - \boldsymbol{Q}_{k+1}^*\boldsymbol{b}|| \\
&= \min_{\boldsymbol{y}} ||\boldsymbol{H}_k\boldsymbol{y} - ||\boldsymbol{b}||\boldsymbol{e}_1||,
\end{aligned} \tag{3.12}$$

where the second equality is a result of the k^{th} iteration of the Arnoldi method (specifically, Eq (3.6)), the third a result of multiplying on the left by $||\boldsymbol{Q}_{k+1}^*||(= 1)$, and the final line follows

by the construction of $\boldsymbol{Q}_k$ (specifically, that $\boldsymbol{q}_1 = \boldsymbol{b}/||\boldsymbol{b}||$).

Therefore, GMRES proceeds as follows. At iteration k of Arnoldi, we have expressions for $\boldsymbol{H}_k$ and $\boldsymbol{Q}_{k+1}$. At this point, we solve the least squares problem in Eq (3.12) using a QR factorisation of $\boldsymbol{H}_k$ [Trefethen and Bau, 1997], and the final solution is recovered via $\boldsymbol{x}_k = \boldsymbol{Q}_k\boldsymbol{y}$. As k increases, the subspace $\mathcal{K}_k(\boldsymbol{A}, \boldsymbol{b})$ spans more of $\mathbb{C}^{n\times n}$, and so the approximation becomes closer to the real solution. We continue building the Krylov subspace, and solving this least squares problem for $\boldsymbol{x}_k$, until the relative error,

$$\varepsilon = \frac{||\boldsymbol{A}\boldsymbol{x}_k - \boldsymbol{b}||}{||\boldsymbol{x}_k||} \tag{3.13}$$

is smaller than some predefined tolerance. Oftentimes, this terminates for $k \ll n$.

Chapter 4

Numerical Methods

Note: some of the material in this chapter is derived from [Clarke et al., 2024] and [Clarke et al., 2025]. The dissertation author was the primary researcher and author of these papers.

4.1 Problem formulation

In this thesis we model several filaments submerged in fluid, where each filament is driven by a follower force - that is, a compressive force directed along the filament tangent - at its distal end. We consider N_f filaments, each of the same length L and cross-sectional radius a, such that the filament aspect ratio is given by $\alpha = L/a$. Each filament has bending and twisting moduli K_B and K_T, respectively, and is surrounded by a fluid of viscosity η. Throughout this thesis, the bending modulus is taken to be equal to the twist modulus, $K_B = K_T$, unless otherwise stated. As a result of the small velocities and length scales associated with filament-motor protein systems, the Reynolds number is small. Hence the dynamics can be considered in the overdamped limit, as discussed in Section 2.1, and consequently the effects of fluid and filament inertia are ignored. Motivated by the biological examples of these filament-motor protein systems, we take each filament to be clamped at its base to a no-slip rigid boundary spanning the $(x, y)-$plane, except where otherwise stated.

In this chapter, we discuss the key methods we will utilise throughout the thesis. First, we

outline the numerical model we use to model filament dynamics. For this purpose, we employ the methods outlined in Schöller et al. [2021] which we summarise in Section 4.2. In the subsequent sections, we introduce the computational dynamical systems toolkit we interface with our numerical time-stepping code, which will allow us to track steady and time-periodic solutions, and then to analyse their stability. These techniques will give us an insight into how stability is exchanged between different emergent states.

4.2 Follower force model

4.2.1 Model description and continuous equations

We begin by describing the force and torque balances in the context of one filament, for brevity. Each filament is treated as identical, and so this holds for each filament. The filament centreline position is denoted by $\boldsymbol{Y}(s,t)$ which is a function of arc-length, $s \in [0, L]$, and time, $t \in [0, \infty)$, as depicted in Figure 4.1a. Residing at each point along the filament centreline is the orthonormal local basis $\{\hat{\boldsymbol{t}}(s,t), \hat{\boldsymbol{\mu}}(s,t), \hat{\boldsymbol{\nu}}(s,t)\}$. This is the material frame, which will bend and twist with the filament. We constrain $\hat{\boldsymbol{t}}(s,t)$ to be the unit tangent to the filament centreline through

$$\frac{\partial \boldsymbol{Y}}{\partial s} = \hat{\boldsymbol{t}}. \tag{4.1}$$

Filament deformation is then described by the following force and moment balances,

$$\frac{\partial \boldsymbol{\Lambda}}{\partial s} + \boldsymbol{f} = \boldsymbol{0}, \tag{4.2}$$

$$\frac{\partial \boldsymbol{M}}{\partial s} + \hat{\boldsymbol{t}} \times \boldsymbol{\Lambda} + \boldsymbol{\tau} = \boldsymbol{0}, \tag{4.3}$$

where $\boldsymbol{\Lambda}(s,t)$ and $\boldsymbol{M}(s,t)$ are the internal forces and moments on the filament cross-section, and $\boldsymbol{f}(s,t)$ and $\boldsymbol{\tau}(s,t)$ are the external forces and torques per unit length. For $N_f = 1$, these are the hydrodynamic forces and torques per unit length, but for $N_f > 1$ we may also choose to incorporate steric forces to prevent overlap between neighbouring filaments. The internal

forces enforce the kinematic constraint, Eq (4.1), and the internal moments are given by the constitutive relation [Schöller et al., 2021]:

$$\boldsymbol{M}(s,t) = K_B \left(\hat{\boldsymbol{t}} \times \frac{\partial \hat{\boldsymbol{t}}}{\partial s} \right) + K_T \left(\hat{\boldsymbol{\nu}} \cdot \frac{\partial \hat{\boldsymbol{\mu}}}{\partial s} \right) \hat{\boldsymbol{t}}. \tag{4.4}$$

The filament is clamped at one end to an infinite, no-slip planar surface located at $z = 0$, fixing the position and local basis of the base to be $\boldsymbol{Y}(0,t) = \boldsymbol{Y}_0$ for some constant (and distinct for each filament) position vector $\boldsymbol{Y}_0$, and $(\hat{\boldsymbol{t}}(0,t), \hat{\boldsymbol{\mu}}(0,t), \hat{\boldsymbol{\nu}}(0,t)) = (\hat{\boldsymbol{z}}, \hat{\boldsymbol{y}}, \hat{\boldsymbol{x}})$. We also impose $\boldsymbol{M}(L,t) = \boldsymbol{0}$ such that the distal end of the filament is moment-free. Applying a follower force to the filament tip corresponds to the condition

$$\boldsymbol{\Lambda}(L,t) = -\frac{f K_B}{L^2} \hat{\boldsymbol{t}}(L,t), \tag{4.5}$$

where f is a non-dimensional parameter controlling the magnitude of the follower force.

4.2.2 Numerical discretisation

Following [Schöller et al., 2021], we discretise the filament into N_p segments of length ΔL. Throughout this thesis, we take $\Delta L = 2.2a$ so that the separation between segments is comparable to the filament width, but large enough to avoid overlap between neighbouring segments when the filament bends. As such, the total filament length is given by $L = N_p \Delta L$. For $i = 1, ..., N_p$, segment i has position $\boldsymbol{Y}_i$ and frame $\{\hat{\boldsymbol{t}}_i, \hat{\boldsymbol{\mu}}_i, \hat{\boldsymbol{\nu}}_i\}$, where $\hat{\boldsymbol{t}}_i$ is constrained to be the tangent of the segment, as sketched in Figure 4.1b. After applying central differencing to Eqs (4.1), (4.2) and (4.3) and multiplying by ΔL, we obtain the force and moment balance for each segment i,

$$\boldsymbol{F}_i^C + \boldsymbol{F}_i = \boldsymbol{0}, \tag{4.6}$$

$$\boldsymbol{T}_i^E + \boldsymbol{T}_i^C + \boldsymbol{T}_i = \boldsymbol{0}, \tag{4.7}$$

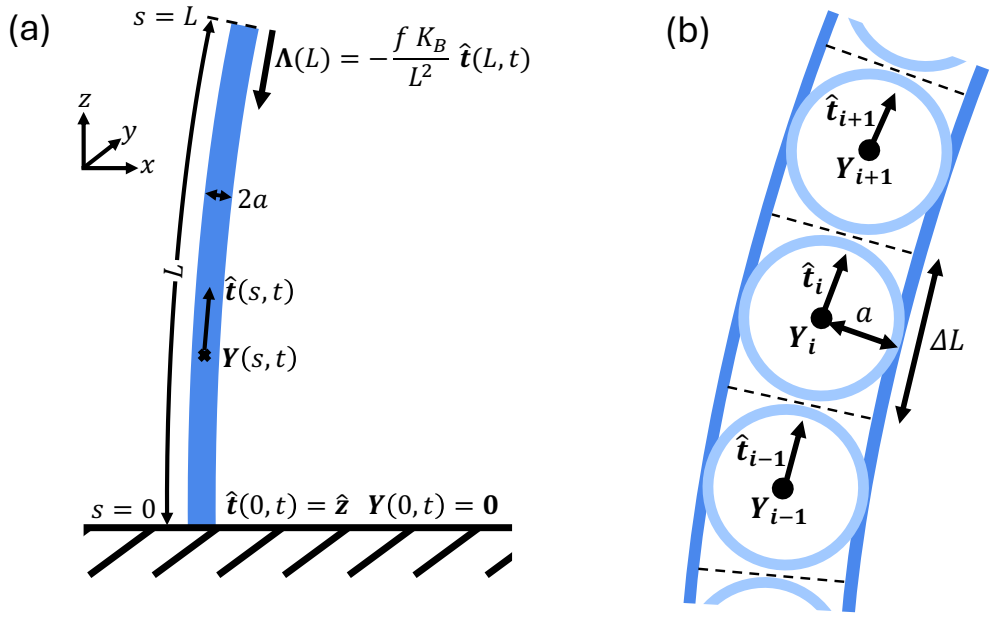


Figure 4.1: (a) A schematic of the follower force model. The filament has cross-sectional radius a, and length L. The bending and twisting moduli are K_B and K_T, respectively. The filament base is clamped to a rigid planar wall at $z = 0$ and surrounded by a fluid with viscosity η. Also shown is the compressive follower force of strength f at the filament's distal end. (b) A schematic describing the numerical method used to compute filament evolution. The filament is discretised into N_p segments of length ΔL. Segment i has position $\boldsymbol{Y}_i$ and local basis $\{\hat{\boldsymbol{t}}_i, \hat{\boldsymbol{\mu}}_i, \hat{\boldsymbol{\nu}}_i\}$ where $\hat{\boldsymbol{t}}_i$ is constrained to be the tangent of segment i. To evaluate the hydrodynamic mobility, the segments are modelled as particles with hydrodynamic radius a.

for $i = 1, ..., N_p$, and the discrete version of the kinematic constraint,

$$\boldsymbol{Y}_{i+1} - \boldsymbol{Y}_i - \frac{\Delta L}{2}\left(\hat{\boldsymbol{t}}_i + \hat{\boldsymbol{t}}_{i+1}\right) = \boldsymbol{0}, \tag{4.8}$$

for $i = 1, ..., N_p - 1$. Here, for segment i, $\boldsymbol{T}_i^E = \boldsymbol{M}_{i+1/2} - \boldsymbol{M}_{i-1/2}$ is the elastic torque, $\boldsymbol{F}_i^C = \boldsymbol{\Lambda}_{i+1/2} - \boldsymbol{\Lambda}_{i-1/2}$ and $\boldsymbol{T}_i^C = \frac{\Delta L}{2}\hat{\boldsymbol{t}}_i \times (\boldsymbol{\Lambda}_{i+1/2} + \boldsymbol{\Lambda}_{i-1/2})$ are the constraint forces and torques respectively, and $\boldsymbol{F}_i = \Delta L \boldsymbol{f}_i$, and $\boldsymbol{T}_i = \Delta L \boldsymbol{\tau}_i$ are the external forces and torques, which we take as

$$\boldsymbol{F}_i = \boldsymbol{F}_i^B + \boldsymbol{F}_i^H, \quad \boldsymbol{T}_i = \boldsymbol{T}_i^H \tag{4.9}$$

where $\boldsymbol{F}_i^B$ are the steric forces (given explicitly in [Schöller et al., 2021]) and $\boldsymbol{F}_i^H, \boldsymbol{T}_i^H$ are the hydrodynamic force and torque the fluid exerts on segment i. We only include steric forces in simulations where $N_f > 1$, to prevent overlap of neighbouring filaments. The discretised

internal moments are defined by

$$\boldsymbol{M}_{i+1/2} = K_B \left(\hat{\boldsymbol{t}}_i \times \frac{(\hat{\boldsymbol{t}}_{i+1} - \hat{\boldsymbol{t}}_i)}{\Delta L} \right) + K_T \left(\hat{\boldsymbol{\nu}}_i \cdot \frac{(\hat{\boldsymbol{\mu}}_{i+1} - \hat{\boldsymbol{\mu}}_i)}{\Delta L} \right) \hat{\boldsymbol{t}}_i \tag{4.10}$$

for $i = 1, ..., N_p - 1$. The internal forces, $\boldsymbol{\Lambda}_{i+1/2}$, on the other hand, act as Lagrange multipliers enforcing the kinematic constraint, Eq (4.8).

4.2.3 Hydrodynamics

For $N_f > 1$ filaments, we can repeat the balances in the previous section to obtain expressions for the forces and torques for each of the $N_p \times N_f$ segments. Thus we now index the segments $i = 1, \cdots, N_p N_f$, such that segment $j \in \{1, ..., N_p\}$ of filament $k \in \{1, ..., N_f\}$ has the index $i = j + N_p(k - 1)$.

Following the previous sections, we have expressions for the hydrodynamic forces and torques which account for the force from the surrounding fluid onto the segments. By Newton's third law, the forces and torques applied to the fluid by the segments is then given by $-\boldsymbol{F}^H, -\boldsymbol{T}^H$. Due to the negligible effects of fluid inertia, the translational and rotational velocities can then be obtained via (compare to Eq (2.42)):

$$\begin{bmatrix} \boldsymbol{V} \\ \boldsymbol{\Omega} \end{bmatrix} = \mathcal{M} \begin{bmatrix} \boldsymbol{F}^C + \boldsymbol{F}^B \\ \boldsymbol{T}^C + \boldsymbol{T}^E \end{bmatrix}, \tag{4.11}$$

where $\boldsymbol{V}$ and $\boldsymbol{\Omega}$ are $3N_pN_f \times 1$ vectors containing the components of the translational and angular velocities of all segments respectively, and $\boldsymbol{F}^C$, $\boldsymbol{F}^B$, $\boldsymbol{T}^C$ and $\boldsymbol{T}^E$ are the corresponding $3N_pN_f \times 1$ vectors containing the constraint forces, steric forces, constraint torques and elastic torques on all segments respectively. This couples the behaviour of all segments in all filaments and, in the case where $N_f > 1$, accounts for the long-range hydrodynamic interactions between the filaments.

As discussed in Section 2.2, there are several choices of hydrodynamic solver we can use. Throughout this thesis, the chapters surrounding a single filament will use RPY tensors, as

described in Section 2.2.1, to handle the hydrodynamics. In Chapters 5 and 7, we will use the adaptation which includes a no-slip planar wall at $z = 0$, and in Chapter 6 we make it clear when we consider the open domain, and when we consider a no-slip planar wall. When we consider periodic domains in Chapter 8, we will instead use FCM and fast FCM, as described in Sections 2.2.2 and 2.2.3.

4.2.4 Time integration

After solving the mobility problem for the velocities and angular velocities of each segment, we integrate the system forward in time and update the segment positions and frame vectors subject to the kinematic constraints. In 2D, we can do this by defining a tangent angle, θ, which allows us to track the bending of the filament over time. However in 3D, our filament can deform by both bending and twisting. To keep a track of the orientation of the local basis as we integrate forward in time, we employ unit quaternions, $\boldsymbol{q} = (q_0, q_1, q_2, q_3)$ with $||\boldsymbol{q}|| = 1$ [Hanson, 2006, Schöller et al., 2021].

Quaternions can be decomposed into scalar and vector parts, q_0 and $(q_1, q_2, q_3) \in \mathbb{R}^3$ respectively. The vector part can be extracted from the quaternion using the operation:

$$[(q_0, q_1, q_2, q_3)]_{\mathbb{R}^3} = (q_1, q_2, q_3). \tag{4.12}$$

We define the associated unit quaternion of segment i, $\boldsymbol{q}_i(t)$, with $||\boldsymbol{q}_i(t)|| = 1$ for all t, such that $\boldsymbol{q}_i(t)$ maps the standard basis, $(\hat{\boldsymbol{x}}, \hat{\boldsymbol{y}}, \hat{\boldsymbol{z}})$, to the segment's local basis at time t, $(\hat{\boldsymbol{t}}_i(t), \hat{\boldsymbol{\mu}}_i(t), \hat{\boldsymbol{\nu}}_i(t))$ through the relation

$$\boldsymbol{R}(\boldsymbol{q}_i(t)) = (\hat{\boldsymbol{t}}_i(t)\ \hat{\boldsymbol{\mu}}_i(t)\ \hat{\boldsymbol{\nu}}_i(t)), \tag{4.13}$$

where $\boldsymbol{R}(\boldsymbol{q})$ is the rotation matrix,

$$\boldsymbol{R}(\boldsymbol{q}) = \begin{bmatrix} 1 - 2q_2^2 - 2q_3^2 & 2(q_1q_2 - q_3q_0) & 2(q_1q_3 + q_2q_0) \\ 2(q_1q_2 + q_3q_0) & 1 - 2q_1^2 - 2q_3^2 & 2(q_3q_2 - q_1q_0) \\ 2(q_1q_3 - q_2q_0) & 2(q_3q_2 + q_1q_0) & 1 - 2q_2^2 - 2q_1^2 \end{bmatrix}. \tag{4.14}$$

for each i.

The evolution of the quaternions, and hence the local basis, is governed by

$$\frac{d\boldsymbol{q}_i}{dt} = \frac{1}{2}(0, \boldsymbol{\Omega}_i) \bullet \boldsymbol{q}_i, \tag{4.15}$$

where $\bullet$ is the quaternion product, defined for two quaternions $\boldsymbol{p} = (p_0, \tilde{\boldsymbol{p}})$ and $\boldsymbol{q} = (q_0, \tilde{\boldsymbol{q}})$ as

$$\boldsymbol{p} \bullet \boldsymbol{q} = (p_0 q_0 - \tilde{\boldsymbol{p}} \cdot \tilde{\boldsymbol{q}},\ p_0 \tilde{\boldsymbol{q}} + q_0 \tilde{\boldsymbol{p}} + \tilde{\boldsymbol{p}} \times \tilde{\boldsymbol{q}}). \tag{4.16}$$

where $\cdot$ and $\times$ refer to the standard dot product and cross product between vectors in $\mathbb{R}^3$.

In practice to keep these quaternions unit length, we utilise Lie algebra elements, $\boldsymbol{u}_i$ for $i = 1, ..., N_p N_f$, which update the quaternions via

$$\boldsymbol{q}_i(t_0 + \Delta t) = \exp(\boldsymbol{u}_i(t_0 + \Delta t, \boldsymbol{q}_i(t_0 + \Delta t)) \bullet \boldsymbol{q}_i(t_0), \tag{4.17}$$

for small timesteps Δt and any time t_0, where again $\bullet$ represents the quaternion product and the exponential is defined by

$$\exp(\boldsymbol{u}) = \left(\cos\left(\frac{||\boldsymbol{u}||}{2}\right), \sin\left(\frac{||\boldsymbol{u}||}{2}\right) \frac{\boldsymbol{u}}{||\boldsymbol{u}||}\right). \tag{4.18}$$

The Lie algebra elements evolve via

$$\frac{d\boldsymbol{u}}{dt} = \mathrm{dexp}^{-1}_{\boldsymbol{u}}(\boldsymbol{\Omega}), \tag{4.19}$$

where the differential of the exponential is defined by

$$\mathrm{dexp}^{-1}_{\boldsymbol{u}}(\boldsymbol{\Omega}) = \boldsymbol{\Omega} - \frac{1}{2}\boldsymbol{u} \times \boldsymbol{\Omega} - \frac{1}{2||\boldsymbol{u}||^2}\left(||\boldsymbol{u}|| \cot\left(\frac{||\boldsymbol{u}||}{2}\right) - 2\right) \boldsymbol{u} \times (\boldsymbol{u} \times \boldsymbol{\Omega}). \tag{4.20}$$

Employing Lie algebra elements by solving Eq (4.19) and updating the quaternions via Eq (4.17), instead of directly updating the quaternions with Eq (4.15), ensures that the quaternions remain unit length. For further details on how this is incorporated into the model, see Schöller

et al. [2021].

Gathering our equations together, the segment positions and orientations evolve according to the differential-algebraic system

$$\frac{d\boldsymbol{Y}_i}{dt} = \boldsymbol{V}_i, \tag{4.21}$$

$$\frac{d\boldsymbol{u}_i}{dt} = \mathrm{dexp}_{\boldsymbol{u}}^{-1}(\boldsymbol{\Omega}_i), \tag{4.22}$$

$$\boldsymbol{Y}_{i+1} - \boldsymbol{Y}_i - \frac{\Delta L}{2}\left(\hat{\boldsymbol{t}}_i + \hat{\boldsymbol{t}}_{i+1}\right) = \boldsymbol{0}, \tag{4.23}$$

where Eqs (4.21) and (4.22) hold for $i = 1, ..., N_p$, and Eq (4.23) holds for $i = 1, ..., N_p - 1$.

We now discretise Eqs (4.21) and (4.22) in time using a second-order backwards differential formula as in Schöller et al. [2021]. This, and the inextensibility constraint given by Eq (4.23), yield a nonlinear system where our unknowns are the Lagrange multipliers and the Lie algebra elements associated to each segment. We substitute the kinematic constraint into our temporally-discretised system to ensure it holds exactly, which results in the nonlinear system

$$\boldsymbol{Y}_1^{j+1} - \frac{4}{3}\boldsymbol{Y}_1^{j} + \frac{1}{3}\boldsymbol{Y}_1^{j-1} - \frac{2}{3}\Delta t \boldsymbol{V}_1^{j+1} = \boldsymbol{0}, \tag{4.24}$$

$$\boldsymbol{u}_i^{j+1} - \frac{1}{3}\boldsymbol{u}_i^{j} - \frac{2}{3}\Delta t \ \mathrm{dexp}_{\boldsymbol{u}_i^{j+1}}^{-1}\left(\boldsymbol{\Omega}_i^{j+1}\right) = \boldsymbol{0}, \tag{4.25}$$

$$\boldsymbol{Y}_1^{j+1} + \frac{\Delta L}{2}\sum_{m=2}^{i}\left(\hat{\boldsymbol{t}}_{m-1}^{j+1} + \hat{\boldsymbol{t}}_m^{j+1}\right) - \frac{4}{3}\boldsymbol{Y}_i^{j} + \frac{1}{3}\boldsymbol{Y}_i^{j-1} - \frac{2}{3}\Delta t \boldsymbol{V}_i^{j+1} = \boldsymbol{0}. \tag{4.26}$$

We can solve this via Broyden's method [Broyden, 1965]: a quasi-Newton method that uses an approximation of the Jacobian, and updates this Jacobian iteratively. We choose a block diagonal initial approximation, which we define using a local Stokes drag model to relate velocities to forces and torques. The entries of the initial Jacobian approximation are explicitly calculated in Schöller et al. [2021], and briefly discussed in Appendix C. Broyden's method then yields the Lagrange multipliers and the Lie algebra elements for each segment, as well as the position of the first segment of each filament, from which we can calculate each quaternion, and hence the filament configuration using the constraints. This process then repeats iteratively for the subsequent timesteps. For further detail on Broyden's method, including comparisons

of Broyden's method performance against other nonlinear solvers, see Schöller et al. [2021].

4.2.5 Basal boundary conditions

The above methods do not fix the degrees of freedom of the base segment of each filament. As the filaments are clamped at their base in our case, the position and orientation of the first segment, $\boldsymbol{Y}_{1+N_p(k-1)}$ and $\boldsymbol{u}_{1+N_p(k-1)}$ for $k = 1, ..., N_f$ respectively, are known and fixed. The constraints corresponding to tethering the filament base, i.e. fixing the position, give rise to Lagrange multipliers, $\boldsymbol{\lambda}^k_{tether}$, which act on the first segment,

$$\boldsymbol{F}^C_{1+N_p(k-1)} = \boldsymbol{\lambda}^k_{tether}, \tag{4.27}$$

to hold the filament in place. The constraints corresponding to clamping the filament base, i.e. setting the Lie algebra element of the first segment to zero, gives rise to Lagrange multipliers, $\boldsymbol{\lambda}^k_{clamp}$, which induce a torque on the first segment

$$\boldsymbol{T}^C_{1+N_p(k-1)} = \mathrm{dexp}^{-1}_{\boldsymbol{u}_{1+N_p(k-1)}}(\boldsymbol{\lambda}^k_{clamp}). \tag{4.28}$$

For further details on why the constraints lead to forces and torques of this form, see Westwood [2020]. These forces and torques will appear in the force and torque balances given in Eqs (4.6) and (4.7) for $i = 1 + N_p(k-1)$ where $k = 1, ..., N_f$. The new unknowns, $\boldsymbol{\lambda}^k_{tether}$ and $\boldsymbol{\lambda}^k_{clamp}$ will be solved for by Broyden's method, in place of $\boldsymbol{Y}_{1+N_p(k-1)}$ and $\boldsymbol{u}_{1+N_p(k-1)}$, at each timestep.

4.3 Jacobian-free Newton Krylov method

The methods discussed in the previous section provide an initial value problem; we prescribe an initial configuration for the filament, and solve the nonlinear system given by Eqs (4.24), (4.25) and (4.26) successively to let the simulation evolve over time. This allows us an insight into the emergent behaviours, but has a few disadvantages. First, that the resulting dynamics can take

a long time to settle into equilibrium, particularly close to bifurcations. Furthermore, the initial value problem can only be used to find stable solutions, but gaining a complete understanding of the state space and bifurcation diagram requires us to follow solution branches into regions where the solution is unstable.

In this section we outline how we can go beyond the initial value problem; discussing how we can use an adapted Newton-Raphson approach as described in [Viswanath, 2007, Willis, 2019], called the Jacobian Free Newton-Krylov subspace method (JFNK), to track steady-state and time-periodic solutions into regions of instability. In subsequent sections, we outline how linear stability analysis and bisection methods can be interfaced with our filament model, to classify the observed bifurcations and identify further dynamics invisible by initial value problems alone. As these methods require a specific state variable which describes the filament configuration entirely, we first discuss how we define this variable in the context of our filament model.

4.3.1 Choosing a state variable

We require a state variable which describes the filament configuration entirely. A naive choice would be the segment positions, $\boldsymbol{Y}_1, ..., \boldsymbol{Y}_{N_p N_f}$. This choice is desirable as stability analysis on this vector would be intuitive - perturbations directly correspond to changes in the filament shape - and this vector would be immediately compatible with algorithms like GMRES. However this would be a poor choice of state variable, as we would not be able to recover the local basis of the filament from the segment positions[1], and so the filament orientations would not be known. As such, this vector does not fully describe the filament configuration.

This motivates an alternative choice; a vector consisting of the quaternions of each segment, $\boldsymbol{q}_1, ..., \boldsymbol{q}_{N_p N_f}$. This encapsulates the local basis of the filament and therefore, in the case where the base of the filament is clamped at a known position, recovers the segment positions via Eq (4.23). However these are also an undesirable choice as quaternions are not commutative, and so do not immediately interface with algorithms such as GMRES; adapted algorithms would

[1] In fact we can recover the tangent vector associated to each segment through Eq (4.23), provided the tangent of the first segment is known. However, this is not sufficient to recover the normal and binormal local basis vectors, or the quaternion of each segment.

need to be sought to take this into account. The Lie algebra updates, $\boldsymbol{u}_1, \cdots, \boldsymbol{u}_{N_p N_f}$, are another contender, but again do not describe the full filament configuration. These variables represent an update with respect to the previous configuration, and so require knowledge of the previous configuration.

As such, we instead utilise the Lie algebra elements, $\boldsymbol{v}_1, ..., \boldsymbol{v}_{N_p N_f}$, that generate the quaternions of each segment with respect to the vertical equilibrium. Accordingly, these are defined such that for each segment i, $\boldsymbol{q}_i$ is related to $\boldsymbol{v}_i$ through the expression,

$$\boldsymbol{q}_i = \exp(\boldsymbol{v}_i) \bullet \boldsymbol{q}_{VE}, \tag{4.29}$$

where $\boldsymbol{q}_{VE} = \frac{1}{\sqrt{2}}(1, 0, -1, 0)$ is the quaternion corresponding to a vertically upright filament, and the exponential map is defined by Eq (4.18). From here, we can use this quaternion to recover the local basis of each segment, and consequently the position through Eq (4.23), if the position of the base segment is known. Therefore, these vectors fully describe the filament configuration at a point in time.

In practice, we will use the state variable as an initial condition, from which we can obtain the position and orientation of each segment. We then progress the simulation forwards using the methods in Section 4.2. At each timestep, the Broyden's solver will output the Lie algebra updates corresponding to the new filament configuration, which are then used to update the filament quaternions and positions. After integrating the simulation forward by N_t timesteps, we can then extract the Lie algebra element for segment i, $\boldsymbol{v}_i = ||\boldsymbol{v}_i|| \hat{\boldsymbol{v}}_i$, from the corresponding quaternion, $\boldsymbol{q}_i$, through

$$\hat{\boldsymbol{v}}_i = \frac{[\boldsymbol{q}_i \bullet \boldsymbol{q}_{VE}^*]_{\mathbb{R}^3}}{||[\boldsymbol{q}_i \bullet \boldsymbol{q}_{VE}^*]_{\mathbb{R}^3}||}, \tag{4.30}$$

$$||\boldsymbol{v}_i|| = 2 \arccos\left((\boldsymbol{q}_i \bullet \boldsymbol{q}_{VE}^*) \cdot \hat{\boldsymbol{e}}_1\right), \tag{4.31}$$

where $\hat{\boldsymbol{e}}_1 = [1, 0, 0, 0]$ and $[\boldsymbol{q}]_{\mathbb{R}^3} = (q_1, q_2, q_3)$.

The Lie algebra elements relative to the vertical equilibrium have a simple physical intuition.

One of these Lie algebra elements, $\boldsymbol{v}$, represents a rotation of the local basis vectors by an angle of $||\boldsymbol{v}||$ around the vector $\boldsymbol{v}/||\boldsymbol{v}||$. As such, the Lie algebra element $\boldsymbol{v} = (v, 0, 0)$ for some $v > 0$ corresponds to the tangent $\hat{\boldsymbol{t}} = (0, -\sin v, \cos v)$. Conversely, the Lie algebra element $\boldsymbol{v} = (0, v, 0)$ for some $v > 0$ corresponds to the tangent $\hat{\boldsymbol{t}} = (\sin v, 0, \cos v)$. We derive the local basis corresponding to a general Lie algebra element in Appendix D.

In this set-up, the trivial state, $\boldsymbol{v}_1 = ... = \boldsymbol{v}_{N_p N_f} = \boldsymbol{0}$, corresponds to the steady, vertically upright filament. We note also that due to the clamped end condition, we have $\boldsymbol{v}_{1+N_p(k-1)}(t) = \boldsymbol{0}$ for all time for $k = 1, ..., N_f$. As a result, any admissible filament configuration can be reached using the state variable:

$$\boldsymbol{U}(t) = (\boldsymbol{v}_2(t)^T, \cdots, \boldsymbol{v}_{N_p}(t)^T, \boldsymbol{v}_{2+N_p}(t)^T, \cdots, \boldsymbol{v}_{2N_p}(t)^T, \cdots, \boldsymbol{v}_{2+N_p(N_f-1)}(t)^T, \cdots, \boldsymbol{v}_{N_p N_f}(t)^T)^T. \tag{4.32}$$

4.3.2 Time-invariant solutions

With the state variable in hand, we now define how it evolves in time. At the temporally continuous level, the time evolution can be summarised as follows. We first transform the $\boldsymbol{v}_i(t)$ contained in $\boldsymbol{U}(t)$ into their corresponding $\boldsymbol{q}_i(t)$ via Eq (4.29). From the quaternions, we determine the filament configuration at time t and solve Eqs (4.21), (4.22) and (4.23) to advance the $\boldsymbol{q}_i(t)$ to time $t + t'$. Finally, we compute $\boldsymbol{v}_i(t + t')$, and hence $\boldsymbol{U}(t + t')$ using Eqs (4.30) and (4.31). This process corresponds to applying the continuous flow-map, $\boldsymbol{\varphi}(\boldsymbol{U}, t)$, which describes the evolution of the state variable over an interval of time. Specifically, we have

$$\boldsymbol{U}(t + t') = \boldsymbol{\varphi}(\boldsymbol{U}(t), t'), \tag{4.33}$$

for all times t, t'.

Then, steady-state solutions satisfy

$$\boldsymbol{G}(\boldsymbol{U}) = \boldsymbol{\varphi}(\boldsymbol{U}, t) - \boldsymbol{U} = \boldsymbol{0}, \tag{4.34}$$

where $t > 0$ is an arbitrary time. Time-periodic solutions, on the other hand, satisfy

$$\boldsymbol{G}(\boldsymbol{U}, T) = \boldsymbol{\varphi}(\boldsymbol{U}, T) - \boldsymbol{U} = \boldsymbol{0} \tag{4.35}$$

for some nonzero period, T. Eqs (4.34) and (4.35) define a nonlinear system, which we wish to solve for $\boldsymbol{U}$ and, for the time-periodic case, T.

4.3.3 Newton-Raphson method

Time-periodic solutions

For brevity, we first discuss the method for time-periodic solutions, before outlining how this is simplified for steady-state solutions. In this case, the nonlinear system given by Eq (4.35) for $\boldsymbol{U}$ and T provides $3(N_p - 1)N_f$ equations for $3(N_p - 1)N_f + 1$ unknowns and is therefore underdetermined. An additional equation is necessary and following [Viswanath, 2007, 2009, Willis, 2017, 2019], we will include it as part of the Newton iteration.

For this purpose, we define a new variable $\boldsymbol{X} = (\boldsymbol{U}, T)$. We denote the initial guess for the solution as $\boldsymbol{X}_0 = (\boldsymbol{U}_0, T_0)$ and seek to determine its update $\boldsymbol{X}_0 + \tilde{\boldsymbol{X}}$, where $\tilde{\boldsymbol{X}} = (\tilde{\boldsymbol{U}}, \tilde{T})$. Linearising Eq (4.35) about the initial guess yields

$$\boldsymbol{G}(\boldsymbol{X}_0) + \boldsymbol{J}_{\boldsymbol{G}}(\boldsymbol{X}_0)\tilde{\boldsymbol{X}} + O(|\tilde{\boldsymbol{X}}|^2) = \boldsymbol{0}, \tag{4.36}$$

where

$$\boldsymbol{J}_{\boldsymbol{G}}(\boldsymbol{X}_0) = \left.\frac{\partial \boldsymbol{G}}{\partial \boldsymbol{X}}\right|_{\boldsymbol{X}_0} \tag{4.37}$$

is the Jacobian of $\boldsymbol{G}(\boldsymbol{X})$ evaluated at $\boldsymbol{X}_0$. Following [Viswanath, 2007], we incorporate the additional equation by insisting that $\tilde{\boldsymbol{U}}$ is perpendicular to its velocity at T, preventing any update of the solution along its periodic orbit. Specifically, the additional equation is

$$\tilde{\boldsymbol{U}} \cdot \dot{\boldsymbol{U}}_0 = 0, \tag{4.38}$$

where $\dot{} = d/dt$. Eqs (4.36) and (4.38) combine to provide the linear system,

$$\boldsymbol{A}\tilde{\boldsymbol{X}} = \boldsymbol{b} \tag{4.39}$$

that we need to determine $\tilde{\boldsymbol{X}}$ at iteration k, where

$$\boldsymbol{A}\tilde{\boldsymbol{X}} = \begin{bmatrix} \boldsymbol{J_G}(\boldsymbol{X}_k)\tilde{\boldsymbol{X}} \\ \dot{\boldsymbol{U}}_k \cdot \tilde{\boldsymbol{U}} \end{bmatrix}, \tag{4.40}$$

and

$$\boldsymbol{b} = \begin{bmatrix} -\boldsymbol{G}(\boldsymbol{X}_k) \\ 0 \end{bmatrix}. \tag{4.41}$$

We solve the linear system defined in Eq (4.39) using GMRES. As we don't know the Jacobian explicitly, the application of the Jacobian is done using the matrix-free difference formula,

$$\boldsymbol{J_G}(\boldsymbol{X}_k)\tilde{\boldsymbol{X}} \approx \frac{1}{\epsilon}\left(\boldsymbol{G}(\boldsymbol{X}_k + \epsilon\tilde{\boldsymbol{X}}) - \boldsymbol{G}(\boldsymbol{X}_k)\right), \tag{4.42}$$

which holds to $O(\epsilon)$, where $\epsilon \ll 1$ is a small parameter and $\dot{\boldsymbol{U}}_k$ is evaluated using a first-order forward differencing. Once GMRES converges, the solution is updated via $\boldsymbol{X}_{k+1} = \boldsymbol{X}_k + \tilde{\boldsymbol{X}}$. We continue the Newton iterations until the relative error,

$$\varepsilon_{newton} = \frac{||\boldsymbol{G}(\boldsymbol{X})||}{||\boldsymbol{X}||}, \tag{4.43}$$

is less than a predefined tolerance. Once the Newton method converges and we obtain a time-periodic solution for a specific set of parameters, we use continuation by using this solution as an initial condition and now seeking the solution at $f + \delta f$, where δf is small in magnitude. As δf is small, we expect the initial condition to be a good approximation of the solution, thus leading to quick convergence. We continue this process for different values of f to track the solution along the solution branch, including values of f for which the solution may be unstable.

Steady-state solutions

For steady-state solutions, the nonlinear system given by Eq (4.34) holds for arbitrary times $T > 0$. As such the system now describes $3(N_p - 1)N_f$ equations for the $3(N_p - 1)N_f$ unknowns in $\boldsymbol{U}$, and so the solution is no longer underdetermined. Then we can proceed as in the previous section with $\boldsymbol{X} = \boldsymbol{U}$. In particular, the linear system is still given by Eq (4.39), where now

$$\boldsymbol{A}\tilde{\boldsymbol{X}} = \boldsymbol{J}_{\boldsymbol{G}}(\boldsymbol{X}_k)\tilde{\boldsymbol{X}}, \tag{4.44}$$

$$\boldsymbol{b} = -\boldsymbol{G}(\boldsymbol{X}_k). \tag{4.45}$$

We again solve this system using GMRES to obtain the update, $\tilde{\boldsymbol{X}}$, which we use to update our solution.

4.3.4 Locally constrained optimal hook

Following the methods above, at each Newton iteration we will be given an update $\tilde{\boldsymbol{X}}$. However if the initial guess is relatively far from the real solution, $\tilde{\boldsymbol{X}}$ may be large, in which case the linearisation in Eq (4.36) is no longer valid. In this case, although the predicted reduction in error,

$$\varepsilon_p = ||\boldsymbol{G}(\boldsymbol{X}_k)|| - ||\boldsymbol{A}\tilde{\boldsymbol{X}} - \boldsymbol{b}||, \tag{4.46}$$

may be satisfactory, the actual reduction in the error

$$\varepsilon_a = ||\boldsymbol{G}(\boldsymbol{X}_k)|| - ||\boldsymbol{G}(\boldsymbol{X}_k + \tilde{\boldsymbol{X}})||, \tag{4.47}$$

may be smaller than desired.

We can circumvent this issue by limiting the size of the Newton update. This will improve the accuracy of the linearisation, while also allowing us to use initial conditions further from the real solution, therefore improving the domain of convergence of Newton's method. This is commonly achieved by damping the update, i.e. updating the solution by $\boldsymbol{X}_{k+1} = \boldsymbol{X}_k + \alpha\tilde{\boldsymbol{X}}$,

where $\alpha \in (0, 1]$ is either fixed *a priori*, or found via a line search during each iteration. In our case, we instead utilise a hookstep as discussed in [Dennis and Schnabel, 1996, Viswanath, 2007, 2009, Willis, 2019], which instead limits the size of the update to be within a certain radius, $||\tilde{\boldsymbol{X}}|| \leq \sigma$. Specifically, we instead aim to solve

$$\min_{\tilde{\boldsymbol{X}}:\, ||\tilde{\boldsymbol{X}}|| \leq \sigma} ||\boldsymbol{A}\tilde{\boldsymbol{X}} - \boldsymbol{b}||. \tag{4.48}$$

In this case, σ is called the trust region radius, and describes 'how well' the linearisation utilised in Eq (4.36) holds, seeing as the error is $O(\sigma^2)$. The solution to Eq (4.48), $\tilde{\boldsymbol{X}}$, is known as the optimal hookstep, so-called as updates can lead to hook-shaped trajectories in valleys. Because we derive the hookstep while minimising the error, this approach should perform better than a straightforward damping of the same size.

A further benefit is that Eq (4.48) can be solved without much additional computational work. In practice, this is implemented during the GMRES iterations. Suppose at iteration k, we have an $\tilde{\boldsymbol{X}}$ such that the relative error

$$\varepsilon = \frac{||\boldsymbol{A}\tilde{\boldsymbol{X}} - \boldsymbol{b}||}{||\tilde{\boldsymbol{X}}||} \tag{4.49}$$

is sufficiently small. Then, at this iteration, we will have orthonormal matrices $\boldsymbol{Q}_k$, $\boldsymbol{Q}_{k+1}$ and an upper Hessenberg matrix, $\boldsymbol{H}_k$ such that

$$\boldsymbol{A}\boldsymbol{Q}_k = \boldsymbol{Q}_{k+1}\boldsymbol{H}_k. \tag{4.50}$$

We can then take a singular value decomposition of $\boldsymbol{H}_k = \boldsymbol{U}\boldsymbol{D}\boldsymbol{V}^T$, where $\boldsymbol{U}$ and $\boldsymbol{V}$ are orthonormal matrices and $\boldsymbol{D}$ is diagonal. This allows us to rearrange our residual to

$$\begin{aligned} ||\boldsymbol{A}\tilde{\boldsymbol{X}} - \boldsymbol{b}|| &= ||\boldsymbol{H}_k\boldsymbol{Q}_k\tilde{\boldsymbol{X}} - \boldsymbol{Q}_{k+1}^T\boldsymbol{b}|| \\ &= ||\boldsymbol{D}\boldsymbol{V}^T\boldsymbol{Q}_k\tilde{\boldsymbol{X}} - \boldsymbol{U}^T\boldsymbol{Q}_{k+1}^T\boldsymbol{b}|| \\ &= ||\boldsymbol{D}\boldsymbol{x} - \hat{\boldsymbol{b}}|| \end{aligned} \tag{4.51}$$

where $\boldsymbol{x} = \boldsymbol{V}^T\boldsymbol{Q}_k\tilde{\boldsymbol{X}}$ and $\hat{\boldsymbol{b}} = \boldsymbol{U}^T\boldsymbol{Q}_{k+1}^T\boldsymbol{b}$. The first line is achieved by multiplying on the left

by $\boldsymbol{Q}_{k+1}^T$, and the second by multiplying on the left by $\boldsymbol{U}^T$. Thus our problem now reduces to minimising $||\boldsymbol{D}\boldsymbol{x} - \hat{\boldsymbol{b}}||$ for $\boldsymbol{x}$, subject to the constraint $||\boldsymbol{x}|| \leq \sigma$, which can be enforced using Lagrange multipliers. As $\boldsymbol{D}$ is diagonal, this gives an explicit expression for the elements of $\boldsymbol{x}$:

$$x_i = \frac{\hat{b}_i D_{ii}}{D_{ii}^2 + \lambda} \tag{4.52}$$

where λ can be found by finding the root of the expression $f(\lambda) = ||\boldsymbol{x}(\lambda)||^2 - \sigma^2$. The optimal hookstep is then recovered via $\tilde{\boldsymbol{X}} = \boldsymbol{Q}_m^T \boldsymbol{V} \boldsymbol{x}$.

At each step, after the optimal hookstep has been calculated, the trust radius is updated. The residual of the linearisation provides a prediction of the error. We can compare the predicted reduction in error, ε_p, to the actual reduction in error, ε_a. If the actual reduction is much smaller than the prediction, this means the linearisation is inaccurate, and therefore a smaller trust radius must be used. This process of finding the optimal hookstep with the new trust radius, and checking the respective errors, is then repeated until either an admissible solution is found, or until the trust region is too small (in which case, no solution is found).

However the trust radius can also be increased. We may find that the actual reduction may be very similar to the prediction. This scenario indicates that the linearisation is a good approximation of the nonlinear system, and therefore the trust radius could perhaps be larger. A larger trust radius allows us to take larger updates, which can accelerate the convergence of Newton's method. Additionally, there is a computational advantage to recalculating the optimal hookstep with twice the trust radius vs running two Newton iterations with the same trust radius; the largest computational expense in each Newton iteration arises from applying the approximate Jacobian, which would be repeated multiple times if running two Newton iterations. Recalculating the optimal hook based on an increased trust radius, on the other hand, involves performing a one-dimensional Newton search for λ, and then performing two matrix-vector multiplications for $\tilde{\boldsymbol{X}}$, and so is much cheaper. Therefore we repeat this process of calculating the optimal hookstep for the new trust radius and checking the respective errors, until it has been found that increasing the trust radius no longer increases our reduction in error. At this stage we update the trust radius to the previous choice, and accept our previous

optimal hookstep as the solution to Eq (4.48), before continuing with the next Newton iteration.

4.3.5 Comments on implementation

In the earlier chapters of this thesis, we use a Matlab implementation of the filament model described in Section 4.2, interfaced with a Matlab implementation of JFNK following Willis [2017] (and described further in Willis [2019]). In the later chapters, we use a CUDA/C++ implementation of the filament model, where now JFNK is implemented in python. The results of both implementations have been compared to ensure good agreement. In both cases, the approximate Jacobian in Eq (4.42) is implemented by calculating $\boldsymbol{G}(\boldsymbol{y})$ for two different initial conditions, $\boldsymbol{y} = \boldsymbol{X}_k$ and $\boldsymbol{y} = \boldsymbol{X}_k + \epsilon\tilde{\boldsymbol{X}}$, by progressing the simulation forward for a time T and then subtracting the initial condition in each case. We compute the difference between these two vectors, and divide by the magnitude of the perturbation, ϵ, to obtain an approximation of the Jacobian.

Newton's method requires an initial 'guess' of the solution to begin. Poor initial guesses can lead to poor convergence, or even no convergence. As such, initial conditions for a new branch are obtained from the initial value problem, where the solution is stable, if possible. We can obtain estimates for the period of the solution from the numerical simulations. Then, continuation allows us to track along the branch.

Finally, there are several tolerances and small parameters to be considered in this problem. We outline the key tolerances we use in the table below.

Error/Perturbation	System	Typical Ranges
Broyden's Error	Solving Eqs (4.24), (4.25), (4.26)	$10^{-8} - 10^{-12}$
Jacobian Approximation Perturbation	ϵ in Eq (4.42)	$10^{-5} - 10^{-6}$
Newton's Error	Eq (4.43)	$10^{-6} - 10^{-8}$
GMRES Error	Eq (4.49)	$10^{-3} - 10^{-4}$

4.4 Linear stability analysis

After identifying steady-state and time-periodic solutions in the system, we wish to analyse their stability. To derive the linearised equations of motion, we first express the continuous-time system given by Eqs (4.21), (4.22) and (4.23) for the state variable as

$$\frac{d\boldsymbol{U}(t)}{dt} = \boldsymbol{f}(\boldsymbol{U}(t)). \tag{4.53}$$

Consider a small perturbation $\delta\boldsymbol{U}$ about a base state $\boldsymbol{U}_0$, such that $\boldsymbol{U} = \boldsymbol{U}_0 + \delta\boldsymbol{U}$. The linearised dynamics of Eq (4.53) for the small perturbation are given by

$$\frac{d\delta\boldsymbol{U}(t)}{dt} = \boldsymbol{A}\delta\boldsymbol{U}, \tag{4.54}$$

where $\boldsymbol{A} = \partial\boldsymbol{f}/\partial\boldsymbol{U}|_{\boldsymbol{U}_0}$. The general solution to Eq (4.54) is given by $\delta\boldsymbol{U}(t) = \boldsymbol{\Phi}(t;\boldsymbol{U}_0)\delta\boldsymbol{U}(0)$, where $\boldsymbol{\Phi}(t;\boldsymbol{U}_0)$ is the state transition operator (or impulse response). When $\boldsymbol{U}_0$ is a steady solution,

$$\boldsymbol{\Phi}(t;\boldsymbol{U}_0) = e^{t\boldsymbol{A}}. \tag{4.55}$$

When $\boldsymbol{U}_0$ is time-periodic with period T, i.e. satisfying Eq (4.35), Floquet's theorem gives

$$\boldsymbol{\Phi}(t;\boldsymbol{U}_0) = \boldsymbol{P}(t)e^{t\boldsymbol{B}}, \tag{4.56}$$

where $\boldsymbol{P}(t) = \boldsymbol{P}(t+T)$ with $\boldsymbol{P}(0) = \boldsymbol{I}$ and $\boldsymbol{B}$ is the Floquet matrix exponent.

The stability of steady and time-periodic solutions in this study is examined by computing the eigenvalues of $\boldsymbol{\Phi}(T;\boldsymbol{U}_0)$, where T is an arbitrary short time interval for a steady solution, or the time period for a periodic solution.

To establish these eigenvalues, we need a method that does not require an explicit form of the transition operator. The power method is a contender, but has the disadvantage of only converging in the case where there is a single dominant eigenvalue, not for instance when there is a complex conjugate pair [Wilkinson, 1965]. Instead, we use the Arnoldi method for

finding eigenvalues, which has the additional benefit of deriving multiple dominant eigenvalues. Therefore, we combine our numerical solver, which advances the filament configuration in time, with the Arnoldi method. The effect of multiplying a vector $\delta\boldsymbol{U}(0)$, where $||\delta\boldsymbol{U}(0)|| = \epsilon \ll 1$, by the state transition operator for the time interval T, is approximately given by

$$\begin{aligned}\boldsymbol{\Phi}(T;\boldsymbol{U}_0)\delta\boldsymbol{U}(0) &= \delta\boldsymbol{U}(T)\\ &= \boldsymbol{\varphi}(\boldsymbol{U}_0(0)+\delta\boldsymbol{U}(0),T)-\boldsymbol{U}_0(T)\\ &= \boldsymbol{\varphi}(\boldsymbol{U}_0(0)+\delta\boldsymbol{U}(0),T)-\boldsymbol{U}_0(0),\end{aligned} \tag{4.57}$$

where the first equality is the definition of $\boldsymbol{\Phi}(T;\boldsymbol{U}_0)$, and the second equality follows from the definition of $\boldsymbol{\varphi}$ given by Eq (4.33) for $\boldsymbol{U} = \boldsymbol{U}_0 + \delta\boldsymbol{U}$. The final equality follows as $\boldsymbol{U}_0$ is either steady, or time-periodic with period T. Defining $\hat{\boldsymbol{x}} = \delta\boldsymbol{U}/\epsilon$, we can apply the Arnoldi method using the approximation of the matrix-vector multiplication between the transition operator and a unit vector, $\hat{\boldsymbol{x}}$,

$$\boldsymbol{\Phi}(T;\boldsymbol{U}_0)\hat{\boldsymbol{x}} \approx \frac{1}{\epsilon}\left(\varphi(\boldsymbol{U}_0+\epsilon\hat{\boldsymbol{x}},T)-\boldsymbol{U}_0\right). \tag{4.58}$$

This is equivalent to solving the initial value problem with initial condition $\boldsymbol{U}_0+\epsilon\hat{\boldsymbol{x}}$, extracting the configuration after a time T, subtracting the steady or time-periodic configuration, $\boldsymbol{U}_0$, and dividing by the magnitude of the perturbation, ϵ.

At iteration k, the Arnoldi method decomposes the transition operator into

$$\boldsymbol{\Phi}(T;\boldsymbol{U}_0) \approx \boldsymbol{Q}_{k+1}\boldsymbol{H}_k\boldsymbol{Q}_k^*, \tag{4.59}$$

as described in Section 3.1. For large k, the eigenvalues of $\tilde{\boldsymbol{H}}_k$, the first $k\times k$ sub-block of $\boldsymbol{H}_k$, approximate those of $\boldsymbol{\Phi}(T;\boldsymbol{U}_0)$. As the extremal eigenvalues of $\tilde{\boldsymbol{H}}_k$ will be the first to converge, the most unstable mode can be determined directly from $\tilde{\boldsymbol{H}}_k$ after several iterations when the matrix is typically small in size and standard eigenvalue solvers can be used. The eigenvalues of $\tilde{\boldsymbol{H}}_k$ (say μ) approximate those of $\boldsymbol{\Phi}(T;\boldsymbol{U}_0)$. Then, we can determine the eigenvalues, say λ, of $\boldsymbol{A}$ or $\boldsymbol{B}$ through

$$\lambda = \frac{1}{T}\log(\mu), \tag{4.60}$$

following from Eqs (4.55) and (4.56). The corresponding eigenmode can then be computed by applying $\boldsymbol{Q}_k$ to the eigenvector of $\tilde{\boldsymbol{H}}_k$.

4.5 Bisection

The tools described thus far in this chapter allow us to identify filament dynamics, whether it be through initial value problems (as described in Section 4.2), or through numerical computation (as described for steady-state and time-periodic solutions in Section 4.3). In this thesis, we may encounter a scenario where we find that two solutions are stable for the same parameter values. If this is the case, we expect for there to be an unstable state which lies on the boundary of the two stable solutions in the state space - the so-called 'edge state' [Skufca et al., 2006].

We can identify the unstable solution by using the bisection algorithm outlined in [Skufca et al., 2006], which iteratively bisects between initial conditions leading to either stable state. Specifically, suppose $\boldsymbol{U}_1, \boldsymbol{U}_2 \in \mathbb{R}^{3(N_p-1)N_f \times 1}$ are the Lie algebra vectors corresponding to the first and second stable state. At each iteration, we bisect these to generate a new initial condition,

$$\boldsymbol{U}_{bisection} = \frac{1}{2}(\boldsymbol{U}_1 + \boldsymbol{U}_2), \tag{4.61}$$

and progress this forwards in time by solving the system discussed in Section 4.2, until one of the stable states is achieved. Suppose the dynamics converge to the first stable state; we then update $\boldsymbol{U}_1 = \boldsymbol{U}_{bisection}$, and repeat. We continue the process until $||\boldsymbol{U}_1 - \boldsymbol{U}_2|| < 10^{-8}$. After a few iterations, we will begin to observe the unstable solution transiently, before the dynamics converge to one of the stable states. As the number of iterations increases, we will find that the unstable solution exists for longer times, allowing us to observe the unstable dynamic.

While this method allows us to identify an unstable solution lying on the boundary surface (separatrix) between the stable states in the state space, we only observe this transiently, before the filament converges to one of the stable states. Therefore characterising the unstable solutions can be difficult, and temporal properties (such as the average bending modulus) can

be difficult to obtain. Furthermore, if we expect multiple unstable solutions to exist for a set of parameters, bisection can only compute a solution stable (or attractive) along the directions tangential to the boundary surface between the two stable states. Any subsequent branches are unable to be sought using bisection alone.

Chapter 5

Individual Dynamics of a Follower Force-Driven Filament

Note: the material in this chapter is derived from Clarke et al. [2024]. The dissertation author was the primary researcher and author of this paper.

5.1 Introduction

In this chapter, we focus on the follower force model; a popular model to drive oscillations in instability-driven filaments, motivated by the effect of a single molecular motor, such as kinesin or dynein, acting on the tip of a filament, such as a microtubule. We will consider a single follower force applied at the filament tip, as described in Section 4.2. When the strength of this compressive force exceeds a critical value, the filament buckles and undergoes self-sustained oscillations. While the different states corresponding to filament whirling, beating, and writhing have been identified in previous studies [De Canio et al., 2017, Ling et al., 2018], these studies have relied primarily on solving initial value problems for different follower force values to determine when these states might arise. As a result, the precise nature of the bifurcations producing these states and a clear understanding of state stability remain outstanding. While stability of the trivial, straight filament has been studied quantitatively by assessing the eigenvalues of

the linear operator at the bifurcation [De Canio et al., 2017, Ling et al., 2018], the differences in emergent states given by the 2D and 3D analyses have yet to be reconciled. The transition between whirling and beating has not yet been explored, and a characterisation of the complex writhing behaviour observed at higher forcing has not been performed. Furthermore, previous studies have focused on how the whirling, beating, or writhing vary with the follower force, leaving the dependence of the emergent state on the filament aspect ratio, a key parameter related to the balance of the viscous and elastic forces, largely unexplored.

As such, in this chapter we combine the techniques from computational dynamical systems (described in Section 4.3 and 4.4) with our filament model (described in Section 4.2) to examine, in detail, the bifurcations that give rise to the beating, whirling and writhing states, exploring how the solutions change as both the strength of the follower force and filament aspect ratio are varied. First, we perform a series of simulations to ascertain the final state achieved for different follower force strengths and filament aspect ratios, by allowing the filament to evolve from a straight configuration with a small perturbation. The resulting state space is provided in Section 5.3. We then ascertain the stability of the steady, trivial state to characterise the initial buckling bifurcation in both 2D and 3D in Sections 5.4 and 5.5. Using this approach, we confirm the 2D buckling of the trivial state into planar beating occurs through a Hopf bifurcation. In 3D, we show that the buckling of the trivial state is a double Hopf bifurcation, and link the 2D buckling dynamics to the whirling dynamic that emerges in fully 3D dynamics. In Section 5.6, we determine directly the time-periodic branches corresponding to whirling and beating solutions and analyse the stability of these branches to explore the transition between whirling and beating. In doing so, we identify a new, quasiperiodic state that connects these two solution branches. Finally, in Section 5.7, we study the writhing regime. We demonstrate that it can be decomposed further into a transitional regime, where a menagerie of quasiperiodic, chaotic and time-periodic solutions are observed, and a regime where only a quasiperiodic solution arises. In this quasiperiodic dynamic, filament motion can be characterised by a whirling base and a beating tip. In performing this study, we present a thorough overview of the state space in this fundamental model.

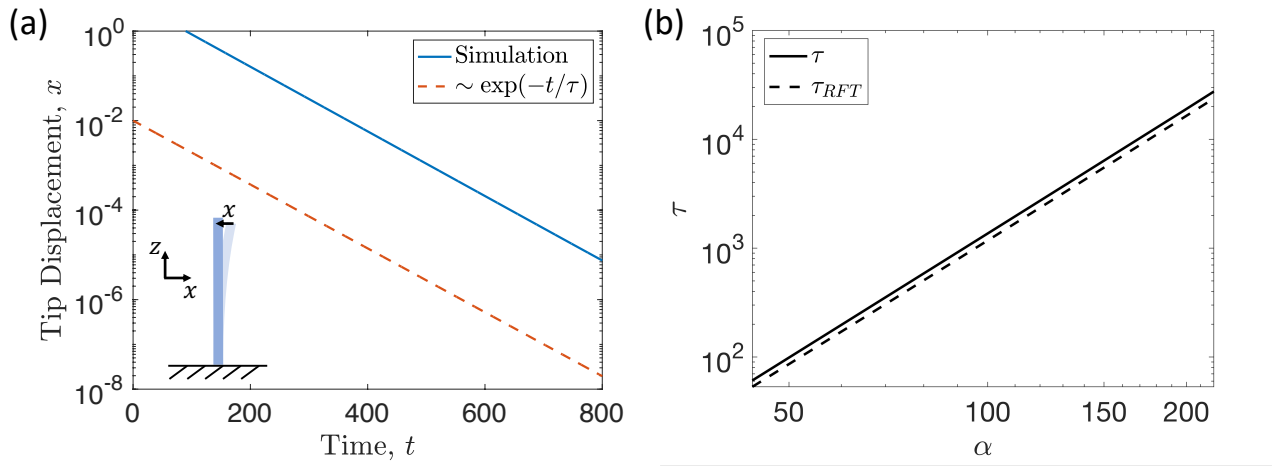


Figure 5.1: (a) The tip displacement x, (solid line) as a function of time for $\alpha = 44$. The dashed line $\sim \exp(-t/\tau)$, where τ is the relaxation timescale determined through fitting the displacement data. The schematic depicts how the computation of the relaxation timescale is performed. (b) The relaxation timescale computed numerically for our model as a function of aspect ratio, compared to the timescale arising from resistive force theory.

5.2 Simulation parameters and nondimensionalisation

In this chapter, we focus on the case of a single filament clamped at its base to a no-slip planar wall at $z = 0$. As such, we use RPY to calculate the hydrodynamic interactions between segments in the mobility relation, adapted to include the no-slip wall, as outlined in Appendix B. We investigate the effect of varying two non-dimensional parameters; the nondimensional follower force, f, as defined in Eq (4.5), and the aspect ratio, $\alpha = L/a$. The aspect ratio is adjusted while keeping the bending modulus of the filament fixed, corresponding to varying the filament length. This results in changing the hydrodynamic response to the filament, while keeping its elastic response fixed.

All timescales in this chapter are reported with respect to the filament's relaxation time, τ, the characteristic timescale associated with a bent filament returning to its undeformed configuration. Due to the dependence of the relaxation time on the hydrodynamic model and also the boundary conditions on the fluid domain, we calculate the relaxation time, τ, numerically. To do this, we measure the tip-displacement of a force-free filament as it relaxes back to the vertical equilibrium from a deformed state. As the initial deformation is small, we are in the linear regime and so expect the tip displacement to decay like $\sim \exp(-t/\tau)$. After using this

to fit the decay data, we can extract τ (see Figure 5.1). As this timescale depends on filament aspect ratio, α, this process must be repeated for filaments with different aspect ratio.

From Wiggins and Goldstein [1998], the relaxation time for a filament is

$$\tau_{RFT} = \frac{\zeta_\perp L^4}{(1.875)^4 K_B} \tag{5.1}$$

where $\zeta_\perp$ is the perpendicular drag per unit length given by resistive force theory [Keller and Rubinow, 1976],

$$\zeta_\perp = \frac{4\pi\eta}{\log(2L/a) - 1/2}. \tag{5.2}$$

Comparing with τ given by our model, we find a similar dependence on the filament aspect ratio with the approximate relation $\tau \approx 1.1\tau_{RFT}$ (see Figure 5.1).

5.3 State space

Before performing a detailed bifurcation and stability analysis, we first solve a series of initial value problems using the methods outlined in Section 4.2 to provide a general overview of the (f, α) state space, and to demonstrate that our code reproduces the results found in previous studies. Unless otherwise stated, in each simulation the filament is initially straight and untwisted, i.e. $(\hat{\boldsymbol{t}}_i(0), \hat{\boldsymbol{\mu}}_i(0), \hat{\boldsymbol{\nu}}_i(0)) = (\hat{\boldsymbol{e}}_z, \hat{\boldsymbol{e}}_x, \hat{\boldsymbol{e}}_y)$ for all i, and a perturbation is introduced in the form of a small, random force applied to three of the segments at the first timestep. If the random forces are co-planar, the resulting dynamics will be in that plane. Otherwise, the dynamics can be fully 3D.

With the filament aspect ratio fixed at $\alpha = 44$, we obtain results consistent with those from previous studies of planar [De Canio et al., 2017] and fully 3D [Ling et al., 2018] filament motion. For planar perturbations, we observe that below a critical follower force value, f^*, the perturbation decays with time and the filament returns to its vertical equilibrium. Above the critical value, however, we observe the onset of self-sustained symmetric oscillations which we call planar beating. When motion is fully 3D, we observe buckling at the same critical

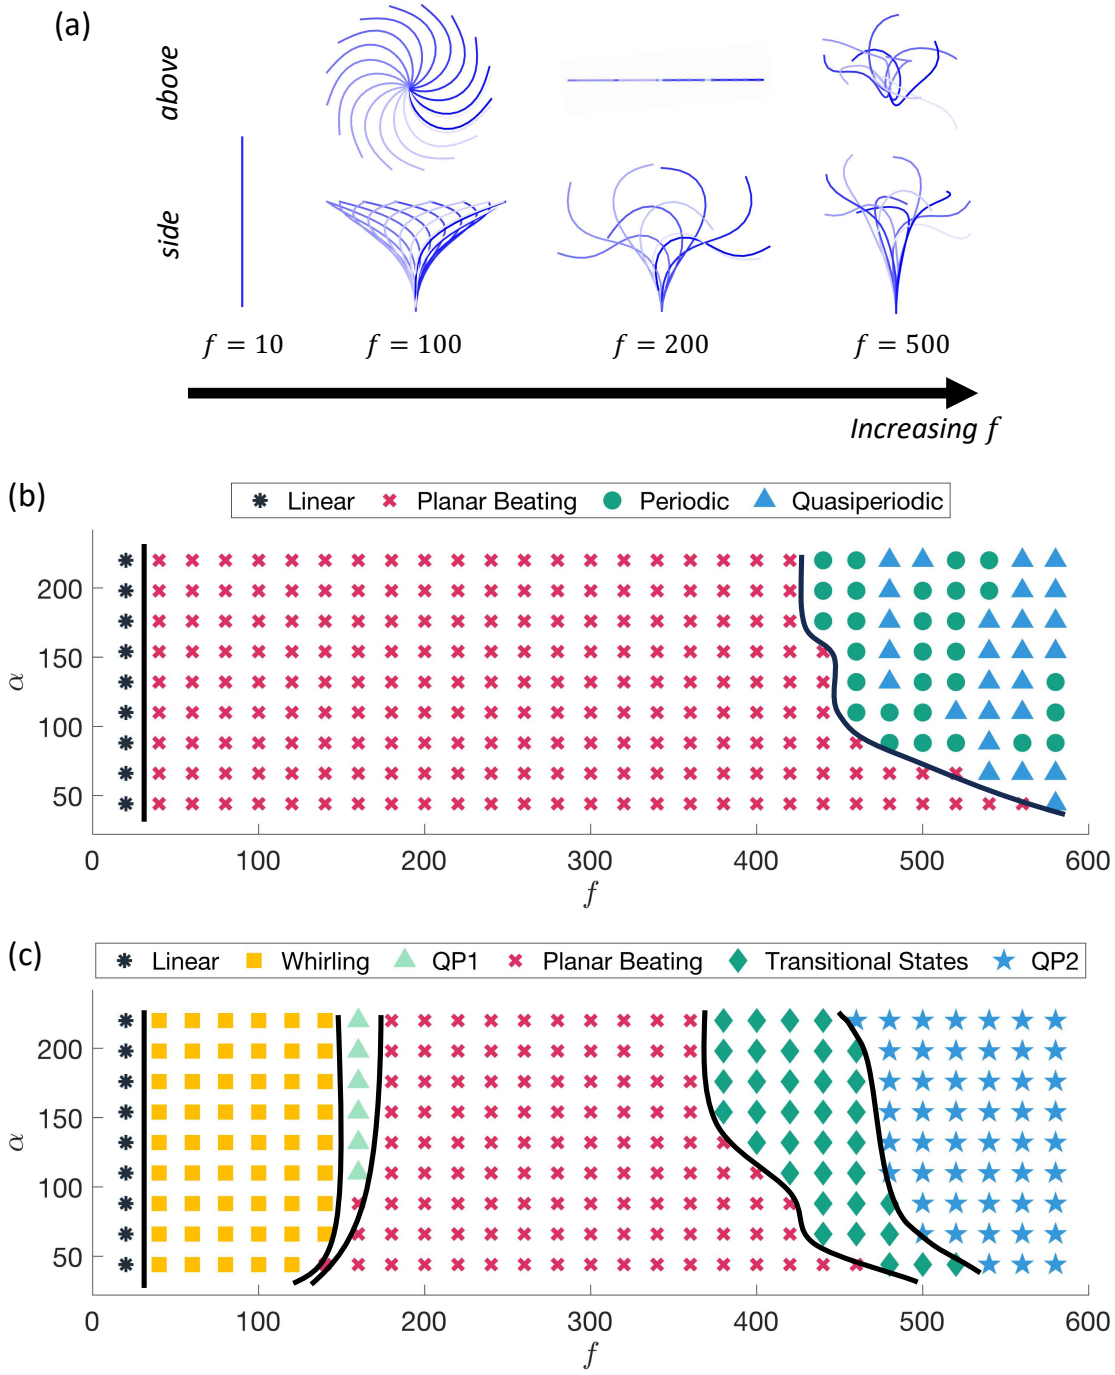


Figure 5.2: (a) Figures showing the steady, whirling, planar beating and writhing states for $\alpha = 44$ (from left to right). Views of the $(x, z)-$ plane (side) and $(x, y)-$plane (above) are provided, and the passage of time is indicated by the increasing boldness of the lines. (b) The state space of solutions over a range of α and f when dynamics are restricted to 2D. (c) The state space of solutions over a range of α and f for fully 3D dynamics. In panels (b) and (c), the black lines are drawn to indicate the approximate boundaries of the solutions regions.

value, f^*, but instead of planar beating, a fully 3D whirling motion emerges. Whirling can be characterised as a rigid body rotation of the filament centreline, with the filament tip tracing out a circle in the $(x, y)-$plane. Increasing the forcing, we find evidence of a second bifurcation as the filament instead performs the same planar beating observed in the 2D model. At very high force values, filament motion enters a writhing regime where it undergoes non-periodic oscillations, frequently changing direction and beat shape. These four distinct behaviours are depicted in Figure 5.2a.

Exploring the state space further for higher values of α, we find that new solutions emerge. The state space when dynamics are restricted to 2D is shown in Figure 5.2b. Here, we see that for $f > 400$, planar beating transitions to a regime where one encounters a mix of quasiperiodic solutions and time-periodic solutions more complicated than planar beating. The force value where this transition occurs decreases as α increases. The state space for fully 3D dynamics is shown in Figure 5.2c. Here, we find the existence of a new, quasiperiodic solution (QP1) that emerges for a small range of f between the whirling and the beating solutions. Furthermore, we see that the region broadly described as writhing can be further divided into two distinct domains of behaviour. The first is a transition region, where beating becomes unstable and gives rise to a diverse array of behaviours highly dependent on α. In the second, occurring at higher f, we find a second quasiperiodic solution (QP2 in Figure 5.2c).

In the subsequent sections, we explore these solutions and study the bifurcations that produce them. In particular, we use JFNK to track the time-periodic solutions of beating and whirling into regions where they are unstable, and perform a linear stability analysis to classify observed bifurcations. We also analyse the new quasiperiodic solutions in depth, using Poincaré sections to elucidate their time-dependence and further illustrate how the solution space changes with filament aspect ratio.

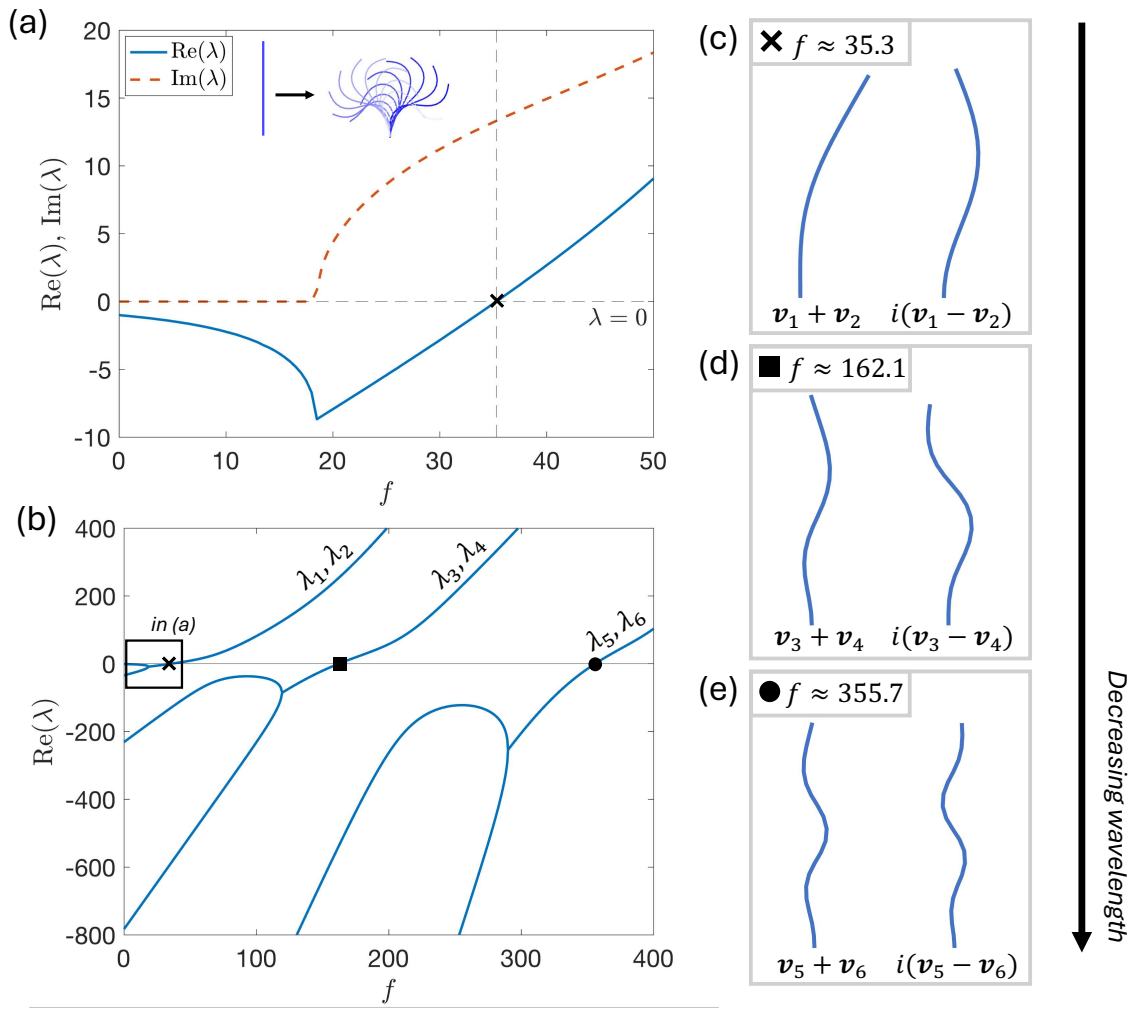


Figure 5.3: Results of 2D stability analysis on the trivial steady state. (a) The real and imaginary components of the dominant eigenvalue, λ_1, from the linear stability analysis of the trivial steady state for $f \in [0, 50]$, $\alpha = 44$. The eigenvalue is purely imaginary at the bifurcation. (b) The real part of the dominant 6 eigenvalues from the same analysis. The corresponding eigenvectors associated to the complex conjugate pair of eigenvalues becoming unstable at the (c) first, (d) second and (e) third buckling events. We note that the wavelength of these eigenvectors decreases as we move through successive bucklings.

5.4 Two-dimensional buckling

We first focus on the initial buckling bifurcation. By solving the initial value problem, we have seen that above a critical forcing, f^*, the filament buckles. In 2D buckling leads to beating, while in fully 3D simulations whirling is observed. To reconcile these observations, we perform a linear stability analysis of the trivial state (the vertically upright filament) using the Arnoldi method described in Section 4.4. This computation yields the dominant eigenmodes which correspond to the least stable states. We extract the dominant eigenvalue, i.e. the eigenvalue

with the largest real part and hence the fastest growing mode, and its associated eigenvector. The eigenvector provides the filament shape that will grow in amplitude with time. We denote each eigenvalue-eigenvector pair as $(\lambda_i, \boldsymbol{v}_i)$.

We first find the 2D unstable modes using a perturbation $\delta\boldsymbol{U}(0)$ that restricts motion to the (x, z)-plane. The real and imaginary components of the dominant eigenvalue, λ_1, are plotted for $\alpha = 44$ in Figure 5.3a. We see that the dominant eigenvalue becomes complex at $f \approx 18.1$, indicating that perturbations to the steady state will still decay, but now oscillate, around the steady state. We observe a bifurcation at $f^* \approx 35.3$, in which a complex conjugate pair of eigenvalues becomes unstable, indicating this to be a Hopf bifurcation. At the bifurcation, these eigenvalues are purely imaginary, and so can be expressed as $\lambda_\pm = \pm i\omega$. Their associated eigenvectors are the complex conjugate pair, say $\boldsymbol{v}_1$ and $\boldsymbol{v}_2 = \bar{\boldsymbol{v}}_1$, and are shown in Figure 5.3c; linear combinations of the eigenvectors are plotted to extract their real and imaginary components. We note that by rotational symmetry, an orthogonal set of eigenvectors, $\boldsymbol{v}_1^\perp$, $\bar{\boldsymbol{v}}_1^\perp$, which correspond to beating in the transverse plane, exist with the same eigenvalues.

We can continue this analysis for larger values of the forcing, beyond the initial bifurcation. We plot the real part of the six dominant eigenvalues from the stability analysis on the steady state in Figure 5.3b for $f \in [0, 400]$. We notice that successive buckling events occur, where further complex conjugate pairs of eigenvalues become unstable. We plot the corresponding eigenvectors in Figures 5.3c-e. We note that the later buckling events are associated to smaller wavelength solutions becoming unstable, similar to the standard Euler buckling modes for a clamped-free steady beam.

5.5 Double Hopf bifurcation

Turning our attention now to the full 3D problem by removing any restrictions on $\delta\boldsymbol{U}(0)$, we perform this stability analysis on the steady state for different α, as shown in Figure 5.4a. The analysis indicates modest increases in the real part of the dominant eigenvalue as α increases. We note, however, that the real part of the eigenvalue again becomes positive at f^* for each α,

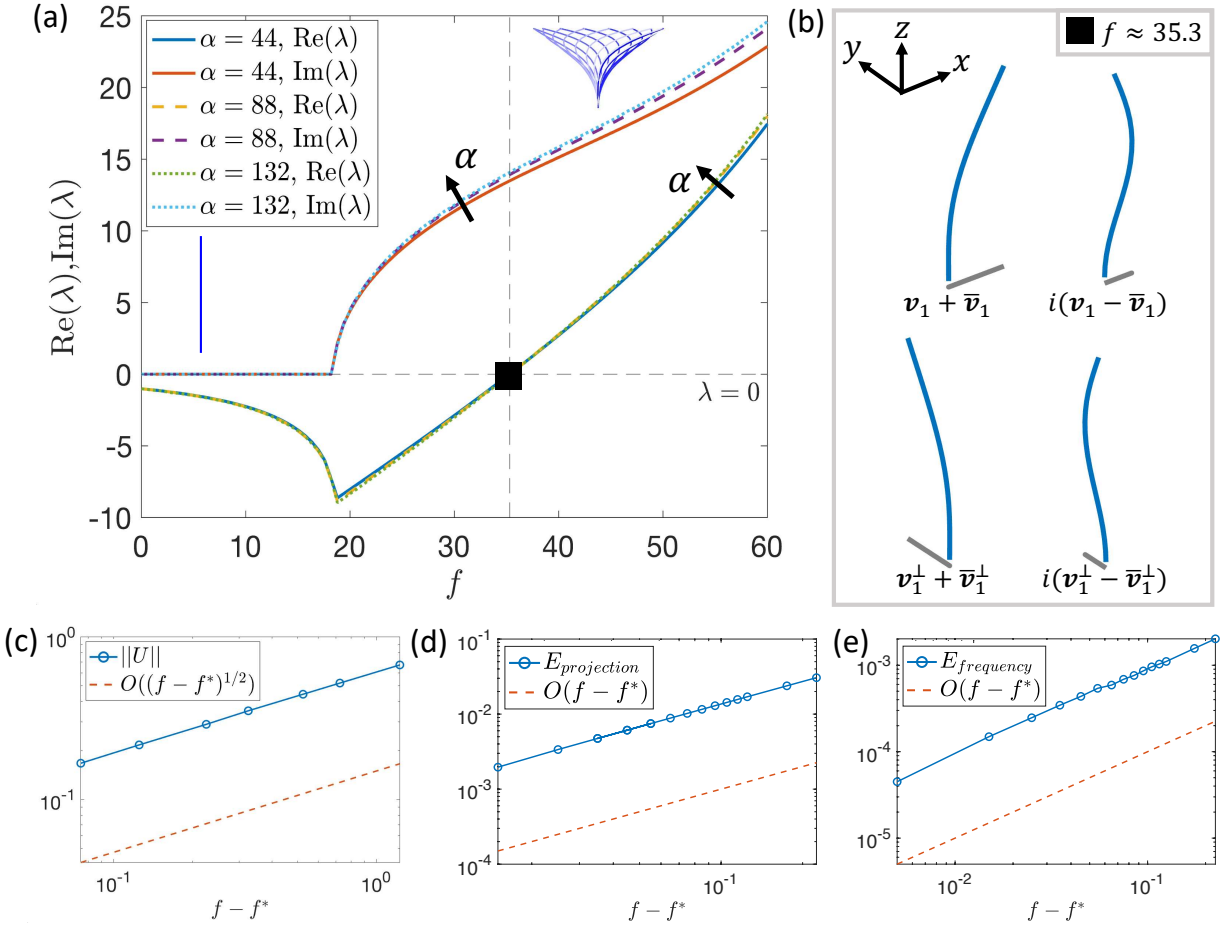


Figure 5.4: (a) The dominant eigenvalue from the linear stability analysis of the trivial steady state for $f \in [0, 60]$ and three values of α. The steady state becomes unstable at $f^* \approx 35.3$ regardless of the aspect ratio. The eigenvalues are purely imaginary at the bifurcation. (b) The dominant four eigenvectors of the steady state at the bifurcation. (c) The norm of the solution grows like $\sqrt{f - f^*}$ near the bifurcation. (d) The difference between the solution and its the projection onto the eigenspace at the bifurcation, $E_{\text{projection}} := \min_{a_x, a_y \in \mathbb{C}} ||\boldsymbol{U} - Re(a_x \boldsymbol{v_x}) - Re(a_y \boldsymbol{v_y})||$, grows like $f - f^*$ close to the bifurcation. (e) The difference between the whirling frequency and the imaginary part of the eigenvalue at the bifurcation, $E_{\text{frequency}} := ||2\pi/T - \omega||$, for $\alpha = 44$. The difference grows like $f - f^*$ near the bifurcation.

demonstrating that the critical value for buckling is independent of α. The imaginary part of the eigenvalue that provides the frequency of the unstable solution near the bifurcation increases slightly with the aspect ratio.

To connect the differences in the 2D and 3D solutions at the bifurcation, we first note that the dominant eigenvalues are identical to those found in 2D, $\lambda_{\pm} = \pm i\omega$. However, we now find that there is not only one, but rather two, complex conjugate pairs of unstable eigenvectors. While we do find that the eigenvectors depend on the initial perturbation used to compute them, they

can always be written as linear combinations of the two sets of planar eigenvectors. Hence we can express the unstable eigenmodes at the bifurcation as

$$\begin{aligned}(i\omega, \boldsymbol{v}_1), (-i\omega, \bar{\boldsymbol{v}}_1),\\ (i\omega, \boldsymbol{v}_1^{\perp}), (-i\omega, \bar{\boldsymbol{v}}_1^{\perp}).\end{aligned} \tag{5.3}$$

A bifurcation with two pairs of imaginary eigenvalues is defined as a double Hopf, or Hopf-Hopf, bifurcation [Kuznetsov, 1998], as opposed to a Hopf bifurcation where only a single pair is present. While in general for a double Hopf bifurcation the eigenvalue pairs can be different, in our specific problem, we obtain repeated pairs due to the rotational symmetry that is present. Bifurcations in the presence of symmetry have been analysed previously using group theoretic techniques [Golubitsky and Stewart, 1985]. We also note that a double Hopf bifurcation occurs in systems with two parameters. The fact that we encounter a double Hopf bifurcation here while varying only the follower force suggests the presence of another hidden parameter already at its critical value. Variation of the hidden parameter must break the rotational symmetry of the two unstable modes, or in other words, the isotropy of unstable motion in planes normal to the filament tangent. Given that the instability is related to buckling, such a hidden parameter may be bending modulus anisotropy, for example. Anisotropy would then allow buckling to occur in one direction at a lower follower force value compared to the other, resulting in Hopf bifurcations at different critical values of the follower force. Only when the bending moduli in both directions are equal would we see the double Hopf bifurcation. We explore the impact of symmetry-breaking on the initial buckling event in Chapter 7.

To better understand the system near the bifurcation, we propose a form of the solution just after it occurs following standard weakly nonlinear analysis [Stuart, 1960, Kuramoto, 1984, Schmid and Henningson, 2001] (for the details in the context of fluid mechanics, see Eq (5.48) in [Schmid and Henningson, 2001]). The solution near the bifurcation is spanned by the unstable eigenvectors with amplitudes that vary over a slow time scale associated with the small growth rate of the instability. When the amplitudes have steady values, as in the beating and whirling

states in this study, the solution may be written as

$$\boldsymbol{U} \approx \boldsymbol{U}^* + A_1\sqrt{f-f^*}\boldsymbol{v}_1 e^{i(\omega+O(f-f^*))t} + A_2\sqrt{f-f^*}\boldsymbol{v}_1^{\perp} e^{i(\omega+O(f-f^*))t} + \text{c.c.} + O(f-f^*), \quad (5.4)$$

for $|f - f^*| \ll 1$, where $\boldsymbol{U}^*$ is the unstable steady state solution and A_1 and A_2 are constant complex amplitudes. Indeed, we find that the error in the projection of the numerical whirling solution onto the eigenspace scales like $\sim f - f^*$, as shown in Figure 5.4d. Eq (5.4) is also consistent with the observation that the growth of the solution scales like $\sqrt{f - f^*}$ (see Figure 5.4c) and that the error in the frequency scales like $f - f^*$ (see Figure 5.4e). In recent work by Schnitzer [2025], this weakly nonlinear analysis on the follower force problem has been carried out analytically, calculating the amplitudes and frequency shifts near the threshold in Eq (5.4), and confirming our numerical analysis. We note that for a standard Hopf bifurcation, the solution would only contain a single mode and its associated amplitude (along with its complex conjugate). Here, for the double Hopf bifurcation that we encounter, we have two modes. The additional dimensions in which solutions live allow for both beating and whirling to emerge at a single bifurcation. From the rotational symmetry of the basic state (i.e. the straight filament), once the filament is saturated into the whirling state we have $|A_1| = |A_2|$, as confirmed by our numerical simulation. The precise values of the complex amplitudes are, however, to be determined from the initial conditions as well as from nonlinear interactions between the two modes.

The double Hopf bifurcation with a solution of the form of Eq (5.4) implies that for the 2D case where the dynamics are confined to the (x, z) or (y, z) plane, either A_1 or A_2 must be zero and the filament would beat in the corresponding plane with angular frequency ω. For dynamics confined to a general plane, the same result follows through a rotation of the axis to align with the beating plane. Based on the stability analysis, we expect beating to be unstable to orthogonal perturbations. However, in 3D, the filament may execute rotational motion about z-axis, as both A_1 and A_2 are expected to be non-zero and incorporate the phase difference needed to achieve the whirling state. Saturation to any state would require a nonlinear mechanism to balance the growth in the unstable modes, however the lack of preferential direction in

the problem suggest that the solutions that do emerge, in this case circular whirling, must reflect this symmetry. Since the eigenvectors are geometrically orthogonal and their size may be normalised arbitrarily, suppose the location of the filament tip in (x, y)-plane is given by two sets of eigenvectors, $(1, 0)$ and $(0, 1)$, respectively. Then, Eq (5.4) implies that that tip location in time would be

$$(x, y)_{\text{tip}} \approx (B_1, B_2)\sqrt{f - f^*}e^{i\omega t} + \text{c.c.}, \tag{5.5}$$

where B_1 and B_2 are complex constants with $|B_1| = |B_2|$. Eq (5.5) indicates that the tip location is, in general, an ellipse with its characteristic angular frequency ω. The whirling state observed in Section 5.3 is admitted when B_1 and B_2 have a particular phase difference: $|\text{Arg}(B_1) - \text{Arg}(B_2)| = \pi/2$. This suggests that there is a nonlinear interaction between the two linear instability modes in Eq (5.4) (e.g. resonance) to form the rotationally invariant whirling state.

The whirling state can be viewed as a steady state in a rotating frame, where separate components of the follower force are balanced by the other forces that are present. The vertical (surface normal) component of the follower force is resolved mostly by the internal stresses in the filament, resulting in the filament's non-trivial centreline shape without changing in time. The horizontal component, however, is instead balanced mostly by the viscous stresses from the fluid on the filament, which may further deform the filament, but more importantly establishes the constant angular velocity associated with the whirling solution. In other words, for the whirling state, the follower force is weak enough, such that its horizontal component does not require large internal stresses to achieve force and moment balances, resulting in a rotational motion with a constant angular velocity. However, if the follower force is increased further, the viscous stresses alone are not able to resolve it. In this case, relatively large amounts of internal stresses along the filament would be required to form a balance with the follower force. This presumably generates an instability, which ultimately leads to an unsteady motion with larger bending along the filament, i.e. the beating state, as we discuss below.

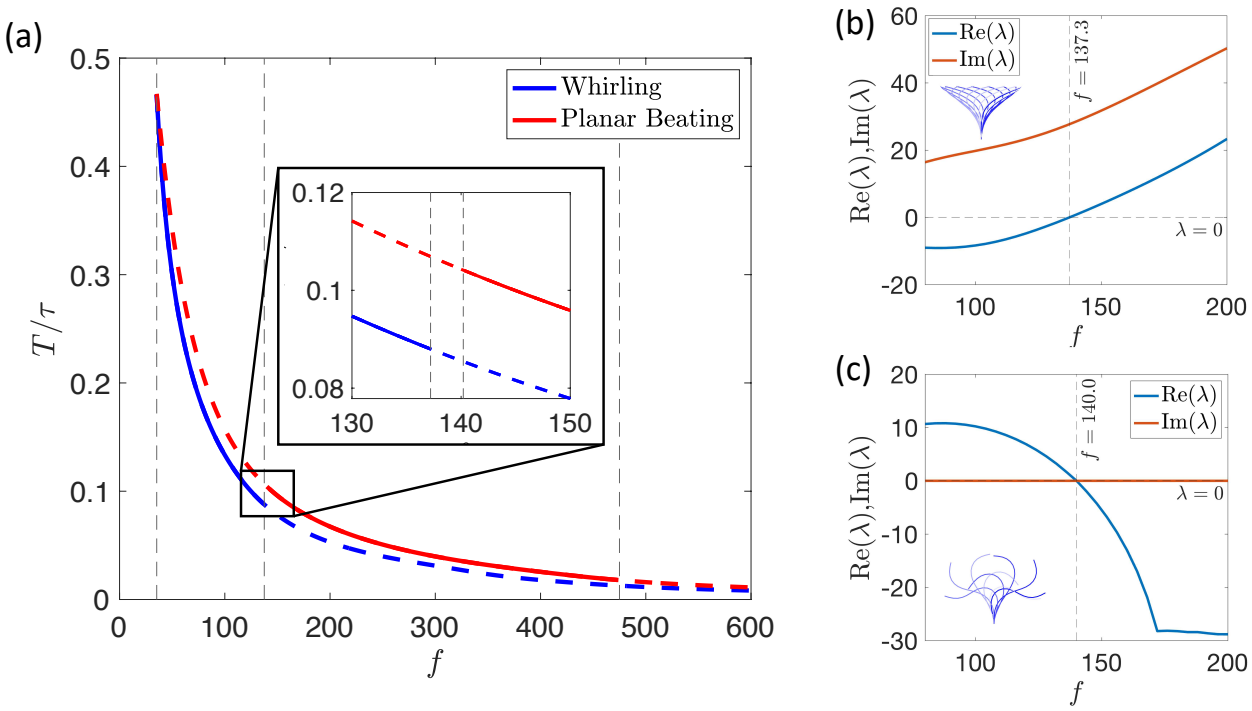


Figure 5.5: (a) The period of the whirling (blue) and planar beating (red) solutions for $f \in [f^*, 600]$ and $\alpha = 44$. Solid (dashed) lines correspond to stable (unstable) solutions. Inset: A close up showing the range $f \in [130, 150]$ indicating that two bifurcations occur. (b) Eigenvalues from the Floquet analysis of the whirling solution for $f \in [80, 200]$. Whirling becomes unstable at $f \approx 137.3$. (c) Eigenvalues from the Floquet analysis of the beating solution for $f \in [80, 200]$. Beating becomes stable at $f \approx 140.0$.

5.6 From whirling to beating

We now consider the next bifurcation in the 3D problem, where the whirling state becomes unstable. Previous works [Ling et al., 2018, Westwood and Keaveny, 2021] based on solving initial value problems, point to a sharp transition between the whirling and beating behaviours at $f \approx 137.5$ for $\alpha = 44$. In order to elucidate this bifurcation, we track the beating and whirling branches as f increases using the JFNK approach described in Section 4.3. We track the whirling and beating branches for $f \in [f^*, 600]$. Figure 5.5a shows how the period of these solutions varies with f for $\alpha = 44$. Changing α produces the same qualitative behaviour, and the quantitative differences that arise are due to differences in the bifurcation points, and that the period decreases with α. In performing this continuation, we see that the planar beating branch can be tracked back to the initial buckling. We can assess the stability of these time-periodic solutions by computing their corresponding Floquet exponents for $f \in [f^*, 600]$ using

the methods from Section 4.4. The Floquet exponents shown in Figure 5.5 reveal that for $\alpha = 44$, whirling becomes unstable at $f \approx 137.3$ (see Figure 5.5b), while beating becomes stable at $f \approx 140.0$ (see Figure 5.5c). Thus there is a region of f where both behaviours are unstable.

Indeed, there is a stable solution that connects the beating and whirling solutions (see Figure 5.6a) and thus there are, in fact, two bifurcations here. As this solution is stable, we track this branch by solving initial value problems. We observe that this stable solution appears as a mixture of the whirling and beating. Two distinct frequencies, one for beating and one for whirling, are associated with this solution, and, in the correct rotating frame, the filament tip traces out an ellipse with an eccentricity that increases continuously with f (see Figure 5.6b). The two periods are not integer divisors, implying that this solution is quasiperiodic. Accordingly, we refer to this solution as QP1.

To study the QP1 solution in more detail, we construct Poincaré sections using a time series of the $x-$coordinate of the filament tip. We label successive local maxima of this signal as $M_1, M_2, ...$, and then plot consecutive maxima against each other. We note that when the filament dynamics are planar, the plane of beating is arbitrary and hence, without loss of generality, we take the beat plane to be the $(x, z)-$plane. Time periodic solutions, such as whirling and beating, appear as single points on the Poincaré section, as shown in Figure 5.6c. Quasiperiodic solutions, on the other hand, appear as curves. In Figure 5.6c, we see that QP1 yield closed curves, consistent with its quasiperiodic behaviour and also characterising the solution as a 2-torus. As the forcing increases, the diameter of the 2-torus trace increases due to the increasing eccentricity of the ellipse traced by the tip.

Using Floquet analysis further, we can classify the two bifurcations that occur. Figure 5.5b reveals that the eigenvalue from the whirling branch is imaginary at the bifurcation, and is distinct from the period of the whirling solution. Thus, the emergence of the quasiperiodic branch corresponds to a supercritical Hopf bifurcation. On the other hand, we observe that the eigenvalue from the planar beating branch is purely real (see Figure 5.5c). This indicates that the second bifurcation is a pitchfork bifurcation, whereby the two stable solutions (corresponding

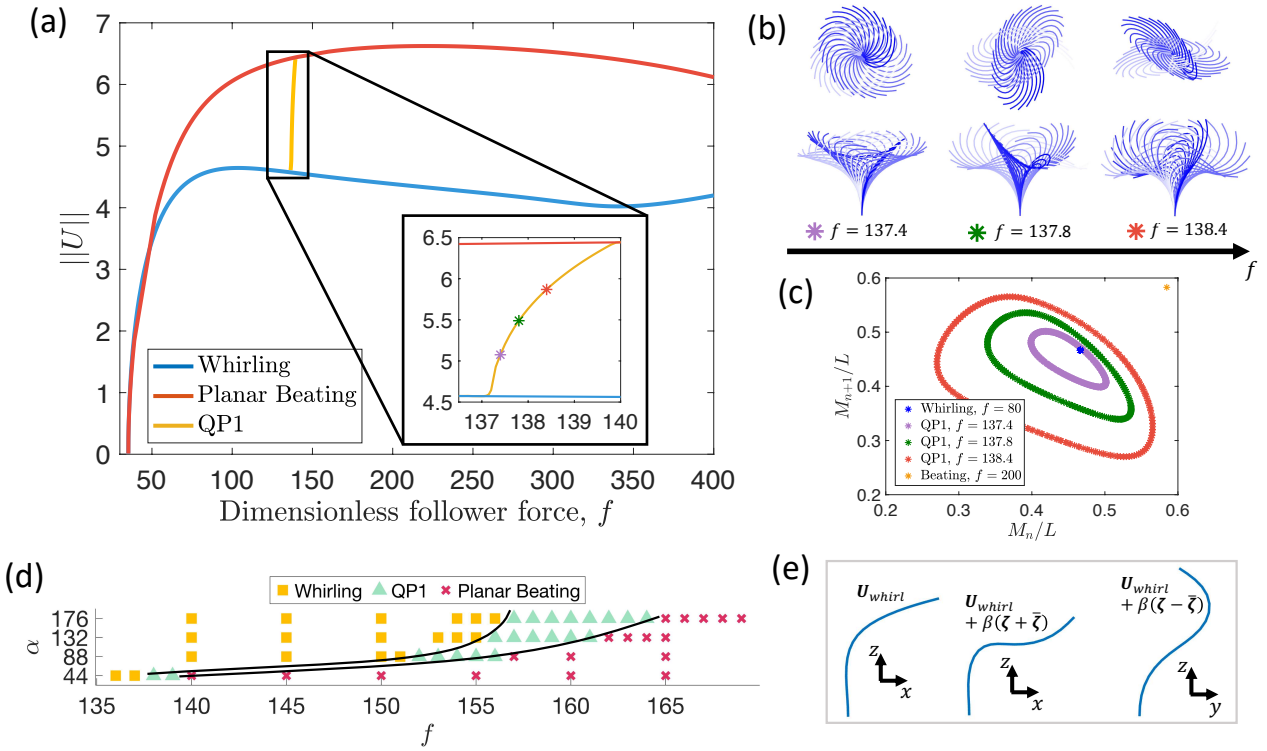


Figure 5.6: (a) The maximum norm of the Lie algebra elements, $||\boldsymbol{U}||$, for the whirling, beating and quasiperiodic solutions as a function of f. Inset: Close up of the quasiperiodic branch connecting with the whirling and beating branches. The three markers indicate the solutions shown in (b) and (c). (b) Figures of the filament in the quasiperiodic regime for three values of f. The passage of time is indicated by an increasing in the boldness of the lines. The eccentricity of the ellipse traced by the filament tip decreases with increasing f. (c) A Poincaré section of the whirling and beating solutions and QP1 for different f. (d) The (α, f) state space in the region where QP1 exists. The range of f for which QP1 exists increases with α. (e) The filament shapes for (left) the whirling solution, $\boldsymbol{U}_{\text{whirl}}$ for $f = 137.4$ and $\alpha = 44$ and (middle/right) the sum of the whirling solution and a linear combination of the unstable eigenvectors, $\boldsymbol{U}_{\text{whirl}} + \beta(\boldsymbol{\zeta} \pm \bar{\boldsymbol{\zeta}})$, for $\beta = 3/2$.

to the quasiperiodic solution rotating clockwise or anticlockwise) cease to exist after beating becomes stable. To investigate how α affects these bifurcations we run additional IVPs close to the whirling and beating bifurcations. Increasing α delays the onset of the first bifurcation and broadens the domain of f for which QP1 exists (see Figure 5.6d).

While ascertaining the precise reasons for the onset of QP1 requires a more detailed analysis, we can gain some understanding of this bifurcation by examining the unstable eigenmode just after the whirling solution becomes unstable. We plot the eigenmodes for $f = 137.4$, $\alpha = 44$ in Figure 5.6e. We note that we plot linear combinations of the complex conjugate pair of eigenvectors, to extract their real and imaginary parts. In particular, we find that this mode

has an approximate wavelength, $l_2 \simeq l_1/2$, where $l_1(= L/2)$ is the wavelength of the unstable mode at buckling of the vertical equilibrium. Given that the critical force for the classical buckling of an Euler beam is inversely proportional to square of the buckling wavelength, it is reasonable that the dimensionless critical follower force for the instability of whirling may be estimated as $f^{**} \approx 4f^*(= 141.2)$. This value of f^{**} is fairly close to the measured values of f^{**} across all α (see Figure 5.5b), in particular $f^{**} = 137.3$ for $\alpha = 44$. This suggests that secondary buckling is a potential physical mechanism for this bifurcation.

5.7 From beating to writhing

By solving initial value problems, we have shown that planar beating is unstable at high follower force values, with the transition occurring at smaller f as α increases. This occurs both when the dynamics are restricted to 2D (Figure 5.2b), as well as in 3D (Figure 5.2c). This picture is confirmed by the Floquet analysis of the beating solution, shown in Figure 5.7a, where we see that the critical follower force in 2D decreases from $f \approx 578.9$ for $\alpha = 44$ to $f \approx 498.1$ for $\alpha = 88$. In 3D, we have $f \approx 462.8$ for $\alpha = 44$ decreasing to $f \approx 422.4$ for $\alpha = 88$. (see Figure 5.7b). The Floquet analysis further reveals that at this bifurcation, in both 2D and 3D, the eigenvalues are imaginary. Focusing on the fully 3D dynamics and revisiting the initial value problem just above the critical follower force, we observe that the solution enters a transition region that contains a whole host of behaviours including intermittency, periodic motion and chaotic beating as shown in Figure 5.7c. The intermittent dynamics are a subset of the chaotic solutions, exhibiting irregular bursts of nearly periodic motion interrupted by chaotic beating (see Figure 5.7d). At $(f, \alpha) = (410, 132)$, we observe a transient period of intermittency that ultimately returns to the planar beating solution.

These numerical experiments reveal that the observed dynamics are highly sensitive to f and α. Immediately after the bifurcation we observe intermittency for lower α. For high α, however, we find a periodic solution where the filament is beating while bent in the direction perpendicular to the beat plane (shown in Figure 5.7f). Additionally, we find that filaments with a smaller α

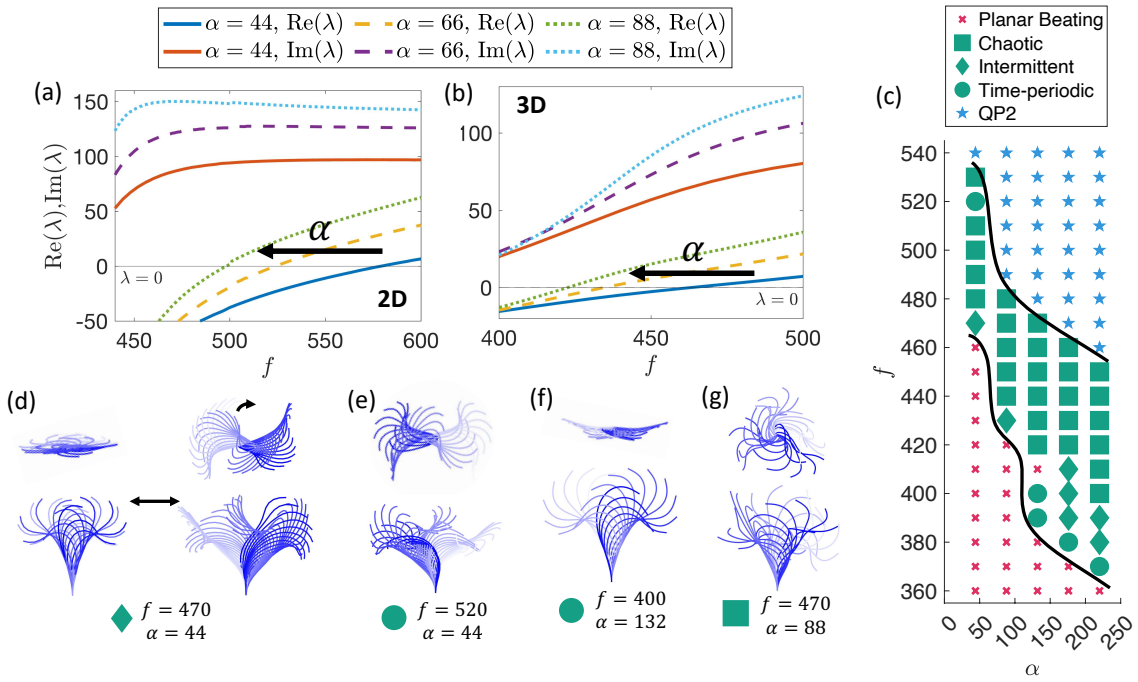


Figure 5.7: (a) Eigenvalues from the 2D Floquet analysis of the planar beating state for $f \in [440, 600]$ for different α. Planar beating becomes unstable at $f \in \{578.9, 523.7, 498.1\}$ for $\alpha \in \{44, 66, 88\}$, respectively. (b) Eigenvalues from the 3D Floquet analysis of the planar beating state for $f \in [400, 500]$ for different α. The planar beating becomes unstable at $f \in \{462.8, 434.4, 422.4\}$ for $\alpha \in \{44, 66, 88\}$, respectively. (c) The transition region between the planar beating (red crosses) and QP2 (blue stars) solutions. Solutions in the region are divided into three types: chaotic (green squares), intermittent (green diamonds), and time-periodic (green circles). (d)-(g): Filament motion in the transition region as viewed from the side and above for four of the observed behaviours. The passage of time is indicated by an increase in line boldness. (d) Intermittence between beating with a growing out-of-plane component (left) and a combination of beating and whirling near the base and waving at the tip (right). (e) Periodic solution with beating at the base and waving at the tip. (f) Periodic solution where the filament beats while bent in the direction normal to the beat plane. (g) Chaotic solution corresponding to erratic beating at the filament tip.

can produce solutions, such as the periodic state shown in Figure 5.7e, that are not observed by filaments with a larger α, and vice versa. In general, at higher forces values, a greater variety of periodic solutions are observed. For instance, as shown in Figure 5.7e, the filament base, broadly speaking, beats in one plane, while the tip oscillates in the transverse plane. Near the final bifurcation (the upper black line in Figure 5.7c), we observe chaotic beating where the filament is rotating near the base while the tip waves aperiodically (Figure 5.7g).

We find that this transition region exists for a finite range of the forcing values. Regardless of filament aspect ratio all filaments undergo the same quasiperiodic behaviour at the highest force

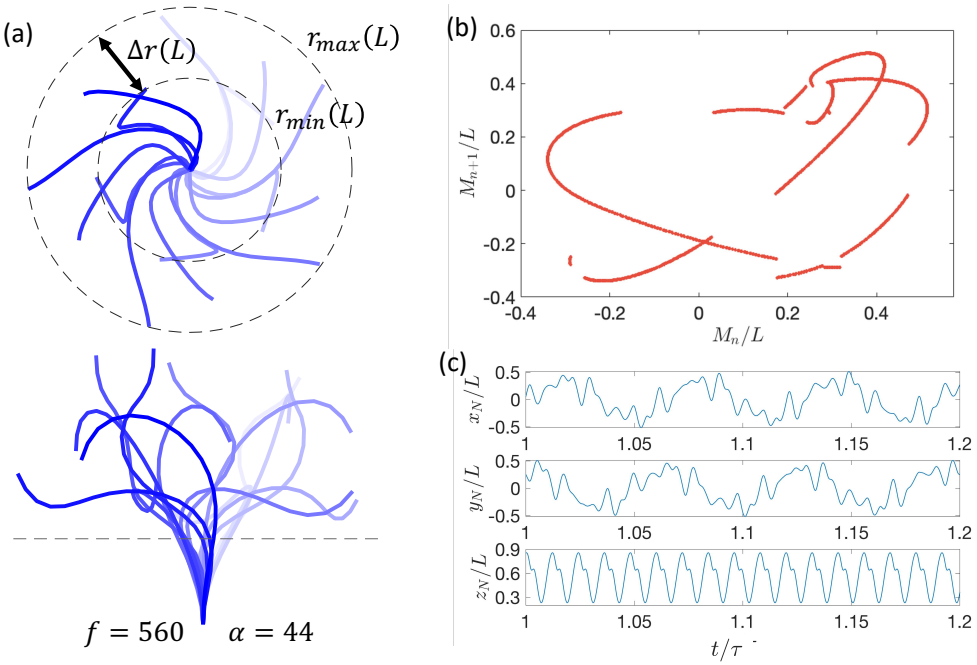


Figure 5.8: (a) Filament motion in the QP2 regime. The passage of time is indicated by an increase in line boldness. (Top) A view from above showing the dynamics projected in the $(x, y)-$plane. The dashed lines indicate the minimum and maximum horizontal distance at $s = L$, and $\Delta r(s) = r_{max}(s) - r_{min}(s)$. (Bottom) A view from the side showing the dynamics projected in the $(x, z)-$plane. The dashed line corresponds to $s/L = s^* \approx 0.3$ from Figure 5.9b and c. (b) The Poincaré section corresponding to the QP2 solution for $f = 560$ and $\alpha = 44$. (c) The time-series of the filament tip displacement for $f = 560$ and $\alpha = 44$ showing periodic oscillations in the z component and irregular oscillations in the $(x, y)-$ components.

values. We refer to this solution as QP2. The transition to QP2 occurs at lower force values as α increases. For the QP2 solutions, the bottom half of the filament appears to be whirling while the upper half beats with a different frequency (see Figure 5.8a). We find that QP2 arises for the variety of initial conditions that we have explored, including those generated by perturbing the trivial equilibrium, the asymptotically stable state in 2D, and the unstable whirling state. This implies that that QP2 is presumably a stable attractor robust to perturbations and initial configurations. To analyse the nature of this solution, we compute its Poincaré section, which reveals several disconnected curves (see Figure 5.8b), confirming the quasiperiodicity of the solution.

Although QP2 is quasiperiodic, meaning the dynamics are a combination of multiple behaviours with periods that are not integer divisors, we can analyse how varying the forcing and aspect ratio affect the dynamics temporally by comparing dominant frequencies using the fast Fourier

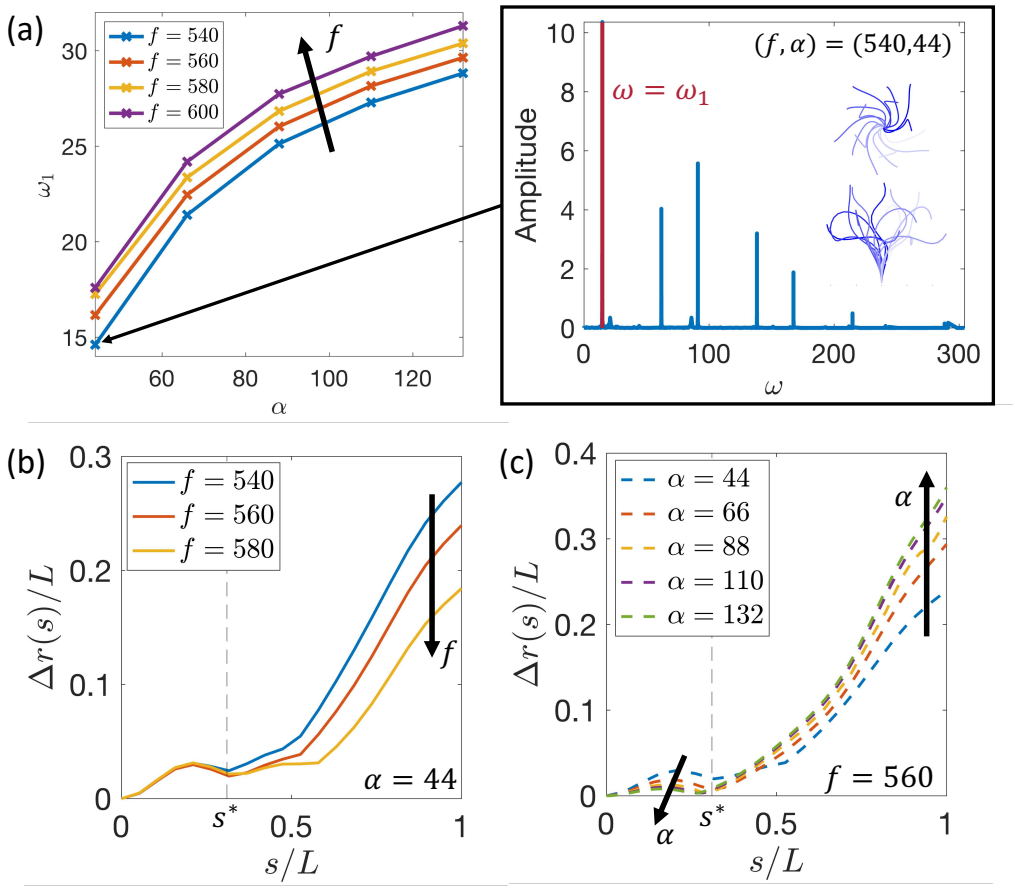


Figure 5.9: (a) The frequency associated to the dominant mode, ω_1, for various f and α. We see that ω_1 increases with both f and α. Right: Single sided amplitude spectrum of QP2 for $(f, \alpha) = (540, 44)$. We take ω_1 to be the frequency with the largest amplitude. (b) The variation in radius, $\Delta r(s)/L$, against the arc-length, s/L, for $\alpha = 44$ and various f. (c) The variation in radius, $\Delta r(s)/L$, as a function of s/L for $f = 560$ and various α.

transform (FFT). In particular, we FFT the $x-$coordinate of the tip displacement to decompose our dynamics into their corresponding frequencies (see Figure 5.9a, right). We then extract the frequency with the most dominant amplitude (i.e. the largest spike) and plot these for various f and α, as shown in Figure 5.9a. As with the planar beating and whirling solutions (as shown in Figure 5.5a), we observe that this frequency increases with both forcing and aspect ratio. Although these results are obtained by decomposing the time-trace of the tip $x-$coordinate, we note that the dominant frequencies are found to be the same when analysing the displacement of any other segment, indicating that this frequency is likely associated to the dynamics of the lower portion of the filament. Recall that in QP2, the base of the filament is whirling whereas the upper portion appears to be beating, while being 'dragged' around by the whirling base. It appears as though the mode corresponding to this whirling is the dominant contribution to

the dynamics along the filament length.

To further investigate the effect of the forcing and aspect ratio on QP2, we compute the horizontal distance from the vertical equilibrium,

$$r(s,t) = ||(\boldsymbol{I} - \hat{\boldsymbol{z}}\hat{\boldsymbol{z}}^T)\boldsymbol{Y}(s,t)||. \tag{5.6}$$

After allowing the solution to reach its asymptotically stable state, we determine the minimum and maximum horizontal distance over time, $r_{min}(s) = \min_{t\geq 0} r(s,t)$ and $r_{max}(s) = \max_{t\geq 0} r(s,t)$, respectively, and plot their difference $\Delta r(s) = r_{max}(s) - r_{min}(s)$. The results are shown for fixed aspect ratio and forcing in Figures 5.9b and 5.9c, respectively. For small values of the arc-length, i.e. for s/L below s^* (the local minimum of $\Delta r(s)$ for $(f,\alpha) = (560, 44)$), the difference in the horizontal distance remains small, corresponding to whirling. Above s^*, the difference rapidly increases as the beating amplitude increases with s. Increasing the follower force decreases the difference in the radius near the filament tip (Figure 5.9b), while increasing the aspect ratio has the opposite effect (Figure 5.9c). Additionally, the local maximum in Δr appearing at $s/L < s^*$ decreases in value as the aspect ratio increases. This implies that as the filament becomes more slender, the horizontal distance approaches a constant value. Hence, the motion near the filament base more closely resembles whirling as α increases, where each segment trajectory lies on a circle of constant radius.

5.8 Discussion

In this chapter, we performed a thorough computational bifurcation analysis of the follower force model for an active filament clamped to a no-slip wall. Along with verifying the findings of previous works [De Canio et al., 2017, Ling et al., 2018], we have extended these studies to pinpoint and classify bifurcations, and identify new solutions. By performing a linear stability analysis of the trivial state, we have established the presence of a double Hopf bifurcation at buckling, reconciling previous results from 2D and 3D studies of the model and establishing the connection between the eigenmodes in 2D with those in 3D. Turning our attention to the

transition between whirling and planar beating occurring at higher force values, we show that there is not only one, but rather two, bifurcations that occur, with a new, stable quasiperiodic solution (QP1) that exists for a small range of f between the two bifurcations. Finally, we examine the solutions at high force values. Here, we show there are two distinct regions: a transition region where a range of periodic and chaotic behaviours are observed, and a second quasiperiodic regime (QP2) where the filament exhibits rotation at the base and waving at the tip. In performing these analyses, we have also established the role of α, the filament aspect ratio, providing a full picture of the parameter space governing the solutions that may arise. In particular, we have shown how increasing α expands the QP1 region, decreases the force where beating becomes unstable, and produces different solutions observed in the transition regime.

Previous studies of the follower force model rely primarily on the use of initial value problems to map out the state space, with stability analysis studies focusing on the initial buckling bifurcation. By interfacing our filament time-stepping algorithm with the numerical tracking and stability analysis tools, we are able to not only explore the first bifurcation in detail - obtaining several of the dominant eigenmodes in the process to analyse further instabilities of the trivial state - but also to examine secondary and tertiary bifurcations in the system. In doing so, we are able to identify previously missed solutions such as QP1, which only exist for small ranges of the forcing, and to characterise these transitions. This is extremely helpful in garnering a more thorough understanding of the state space and, as we will see later, will allow us to identify the differences in the state space more precisely as we make modifications to this model. However a limitation of these tools, at least in their current form, lies in the fact that they apply only to steady and time-periodic solutions. As such, our knowledge surrounding the region where the transitional states begin, and beyond to QP2, is limited; as we are unable to track or evaluate the stability of quasiperiodic solutions, we cannot be sure as to why or how these transitions occur. In particular, it is unclear why the dynamics change from being so sensitive to the parameters in the transitionary region, to more 'regular' at the largest values of the forcing that were tested.

Revisiting the original motivation for the follower force model, recall that the single follower force at the tip is representative of a single motor protein residing at the tip of a filament, such

as a microtubule. A natural question that follows is whether the notion of a motor fixed at a point on the microtubule is accurate; specifically, how the timescale corresponding to motor diffusion along the microtubule compares to the timescale of filament motion. We can generate estimates of both timescales using the results presented in this chapter and experimental values from the literature. The length of microtubules are variable, but can be between $1\mu m - 100\mu m$ [Bray, 2000, Pampaloni et al., 2006, Ling et al., 2018]. Taking the typical microtubule length to be $L \approx 20\mu m$, the microtubule bending modulus as $K_B \approx 10^{-23} Nm^2$ [Gittes et al., 1993], and the viscosity of water to be $\eta = 1$cP, we can estimate the relaxation time, using Eqs (5.1) and (5.2), to be $\tau_{RFT} \approx 2.37$s. Using Figure 5.5a, we see that the period of the whirling and beating solutions are very similar, but reduce drastically as we increase the forcing. Taking a characteristic period to be $T_{fil} \approx 0.05\tau$, taken from the period of the planar beating solution for a moderate value of the forcing, we can use the relaxation time provided by RFT to recover a dimensional estimate of the period: $T_{fil} \approx 0.12$s. We can compare this to the time it takes a single kinesin motor to move along a microtubule, T_{motor}. Taking the characteristic speed of the motor to be $v = 1.5 \times 10^{-6}\text{ms}^{-1}$ [Kural et al., 2005] and the same characteristic microtubule length, we calculate $T_{motor} \approx 13.3$s. Therefore the rate of filament oscillations is much faster than the rate of motor motion - over a single period, the motor will move less than 1% of the filament length - supporting our choice to neglect the motion of the motor along the filament in the follower force model.

It is also interesting to note that while it is reasonable to assume microtubules are actuated by more than one motor protein at any given time, the follower force model predicts that only a single motor protein is required to induce buckling, producing the filament motion explored in this work. We again take the typical microtubule length to be $L \approx 20\mu m$ and the microtubule bending modulus as $K_B \approx 10^{-23} Nm^2$, and now take the motor forcing as $6pN$ (molecular motors exert forces ranging between $1 - 10pN$ [Schliwa and Woehlke, 2003, Svoboda and Block, 1994, Mallik et al., 2004]). These values lead to the corresponding non-dimensional follower force is $f \approx 240$. This value lies squarely in the planar beating regime, firmly suggesting that a single dynein provides sufficient force to cause the onset of self-sustained microtubule oscillations. Due to the large distribution of parameter values reported experimentally, it is

likely that even higher values of the non-dimensional forcing could also be attained.

The follower force model has also been extended to describe more organised collections of microtubules, such as the axoneme within cilia. While the follower force does not replicate the shear-driven actuation in cilia, the dynamics exhibited by the basic follower force model studied in this paper are qualitatively similar to that observed for cilia. For instance, nodal cilia in the embryo undergo a titled whirling, to pump fluid unidirectionally and break symmetry during embryogenesis [Nonaka et al., 1998, Smith et al., 2008a, 2011]. Motile cilia elsewhere, such as respiratory cilia which use this motion to transport mucus in the lungs [Chilvers and O'Callaghan, 2000], typically perform planar beating. Consequently, some studies employ a follower force model as a surrogate model to examine related ciliary motions [Westwood and Keaveny, 2021, Wang et al., 2023, Link et al., 2024]. Furthermore, recent studies have incorporated the follower force model into an axoneme-like configuration of filaments [Bayly and Dutcher, 2016, Woodhams et al., 2022], demonstrating that oscillations can arise by instabilities similar to those studied in the chapter, and hence providing an additional mechanism for cilia actuation.

Returning once more to the particular biological motivation of motor-driven filaments, while the follower force model considered in this paper provides the simplest representation of such a structure, alternative models that provide more features of the biophysical system have been proposed. These include models with follower forces distributed along the entire filament length, or further, the direct effect of the motors on the surrounding fluid, either through an equal and opposite force on the surrounding fluid [De Canio et al., 2017, De Canio, 2019, Stein et al., 2021], or through a slip flow on the filament surface [Laskar and Adhikari, 2017, Ozdemir et al., 2025]. Other variants of the model discussed in this work consider how changing the boundary conditions enforced on the filament can affect the resulting dynamics, for instance clamping both ends ($s = 0$ and $s = L$) a distance $l < L$ apart [Fatehiboroujeni et al., 2018, 2021]. In these cases buckling still occurs, but the resulting dynamics are different to those observed for the clamped and free conditions we consider throughout this paper. Previous studies [Ling et al., 2018, De Canio, 2019] indicate that the clamped-free distributed follower force case has qualitatively similar dynamics to the single follower force at the tip. Moreover, the initial

buckling from the vertical equilibrium, to whirling in the 3D case and beating in 2D, could follow the same mechanism, and hence also be a double Hopf bifurcation. The same argument from Section 5.5 could then apply in this case to explain the connection between 2D and 3D buckling, suggesting that this initial buckling bifurcation may be generic to follower-force driven filament systems. However it is unclear the precise differences these models will make to the state space, in particular to the bifurcations after the initial buckling event. Furthermore, the impact of the hydrodynamic no-slip wall we consider is unclear; previous works have investigated the 2D dynamics and initial buckling through initial value problems and stability analysis [De Canio, 2019], but the precise changes to the state space in three dimensions have not yet been explored. A comparison between the bifurcations that arise in these different models could provide insight into their fundamental differences. These studies will be the focus of the next chapter.

Chapter 6

Modelling Extensions

Note: the first section of this chapter is derived from Clarke et al. [2024]. The dissertation author was the primary researcher and author of this paper. Some of the material in the subsequent sections is derived from Ozdemir et al. [2025]. The dissertation author performed the numerical simulations and stability analysis of the follower force-driven filaments.

In the previous chapter, we explored the dynamics of the simplest incarnation of the follower force model in which a single compressive force is applied to the tip of a filament, which is clamped at its base to a no-slip planar wall at $z = 0$. In this chapter, we turn our attention to several extensions of this model. First, we note that the analysis performed so far has been under the assumption that the bending and twisting moduli are equal. The choice of twist modulus is non-trivial, and it is not immediately clear what impact this will have on the resulting dynamics. Therefore in Section 6.1, we calculate a rough estimate for the range of values the twist modulus may take with respect to the bending modulus in the context of a microtubule, and perform simulations and stability analysis to elucidate how increasing or decreasing the twist modulus can affect the overall state space.

Next in Section 6.2, we investigate the impact of the hydrodynamic wall by modelling a filament in an unbounded domain, driven by a single follower force at its tip. While the initial buckling is a result of the elasticity of the filament rather than a hydrodynamic response, and so we expect nominal changes in the low-actuation regions, we anticipate that the wall may have

larger effects on later bifurcations. Through the use of initial value problems and stability analysis, we outline the state space completely for $\alpha = 44$, and compare to the results of Chapter 5. Finally, in Section 6.3 we investigate the forcing mechanism by now considering a distribution of follower forces along the filament length, motivated by the biological example of several molecular motors translocating along the length of a biopolymer filament, and analyse the resulting changes to the system.

6.1 Effect of twist

The appropriate choice of $\gamma = K_T/K_B$ in our simulations is non-trivial, as these parameters are difficult to measure experimentally for microtubules and cilia, and their values appearing in the literature can vary significantly. Even the parameters which can be used to estimate these moduli, such as the Young's modulus and shear modulus, have experimental estimates spanning several orders of magnitude [Kis et al., 2002, Feng and Mitran, 2018].

For the purpose of this work, we can generate a rough estimate of these parameters for a microtubule by making the assumption that the microtubule is isotropic. Then, the shear modulus can be expressed in terms of the Young's modulus, E, and Poisson's ratio, ν, by $\mu = E/2(1+\nu)$. Experimentally, the Young's modulus of microtubules is suggested to lie in the range 3.1 MPa $-$ 1.4 GPa [Feng and Mitran, 2018], and, following Sirenko et al. [1996], we can estimate Poisson's ratio as $\nu = 0.3$. This gives an estimated range of the shear modulus: $\mu \in (1.2, 540)$ MPa. By approximating the microtubule as a cylindrical shell with internal and external diameters $2r_i = 15$nm and $2r_e = 25$nm respectively [Safinya et al., 2019], the bending and twisting moduli are then given by [Landau and Lifshitz, 1986]:

$$K_B = \frac{\pi E}{4}\left(r_e^4 - r_i^4\right), \quad K_T = \frac{\pi \mu}{2}\left(r_e^4 - r_i^4\right), \tag{6.1}$$

respectively. Using our parameters, this gives broad estimates in the range $K_B \in (5.2 \times 10^{-26}, 2.3 \times 10^{-23})\text{Nm}^2$ and $K_T \in (4.0 \times 10^{-26}, 1.8 \times 10^{-23})\text{Nm}^2$. Experimental evidence suggests the bending modulus of microtubules could vary between around $2 \times 10^{-24}\text{Nm}^2$ [Felgner

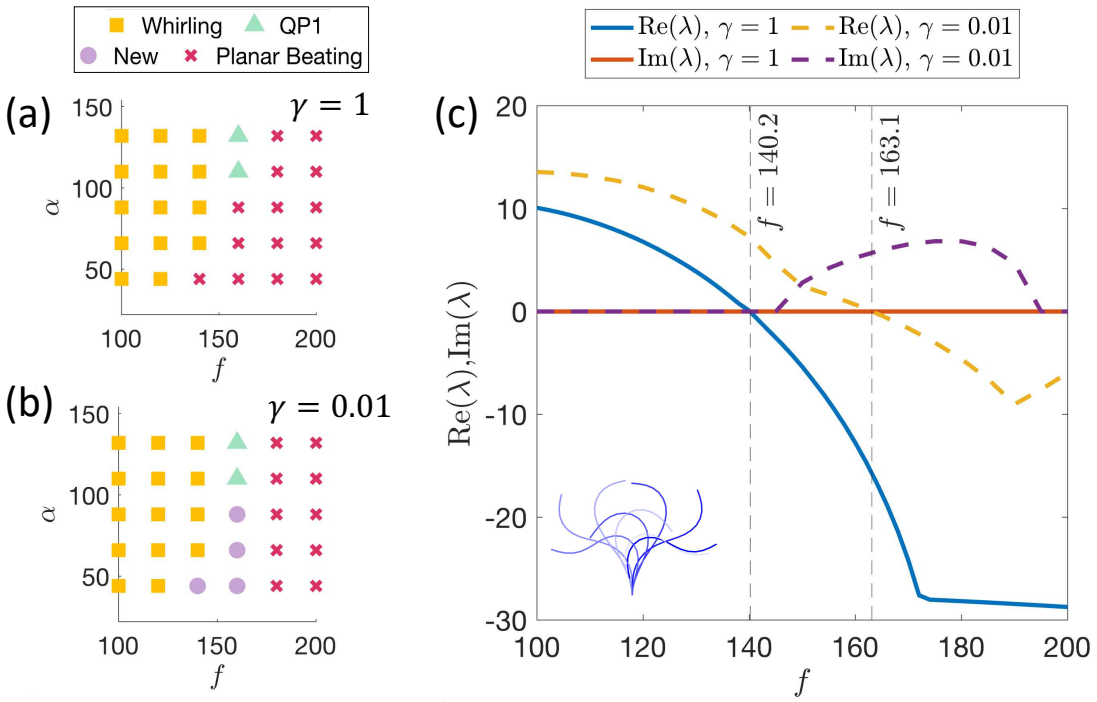


Figure 6.1: The phase diagram for $f \in (100, 200)$, $\alpha \in [44, 132]$ for (a) $\gamma = 1$ and (b) $\gamma = 0.01$. (c) Floquet analysis on the beating solution for a filament with aspect ratio $\alpha = 44$ for $\gamma \in \{0.01, 1\}$. We note that the bifurcation occurs later for $\gamma = 0.01$, and the eigenvalue is complex at the bifurcation.

et al., 1996] and $2 \times 10^{-23}\mathrm{Nm}^2$ [Gittes et al., 1993]. This implies that γ can range between approximately 1.7×10^{-3} and 350 for microtubules; over five orders of magnitude. In [Chełminiak et al., 2010], calculations that take into account the anisotropy of microtubules suggest this ratio to be around $\gamma \approx O(0.01)$.

To investigate the effects of the twisting modulus on our model, we compare results for $\gamma \in \{0.01, 1, 100\}$. We run initial value problems for $f \in \{0, 300\}$ and observe that the state space does not change when increasing from $\gamma = 1$ to $\gamma = 100$. However for $\gamma = 0.01$, we observe differences for filaments with smaller aspect ratios – planar beating becomes stable at a higher value of the forcing, as shown in Figure 6.1. In its place, we see a range of distinct dynamics to those observed for $\gamma = 1$. For instance, for $(f, \alpha) = (140, 44)$ the filament appears as a mixture of whirling, QP1 and planar beating. The filament begins by whirling, and then transitions to planar beating through QP1 type-beating, before returning to whirling in the opposite direction. For $(f, \alpha) = (160, 88)$ on the other hand, the filament undergoes periodic planar deformations.

By performing Floquet analysis on the planar beating behaviours for $\gamma = 0.01$, we can confirm the beating becomes stable for higher forcing values, as indicated by the simulations. Although the dynamics are different, for $\alpha \in \{66, 88\}$, the dominant eigenmode remains purely real at this bifurcation, as for $\gamma = 1$. However, for $\alpha = 44$, the eigenvalue is complex as shown in Figure 6.1c, indicating a change in frequency between the new dynamics and the planar beating at this transition.

These results are not entirely unexpected. The torsional drag scales like $\sim O(\alpha^{-3})$, and so increasing the aspect ratio will decrease the amount of torsion generated. Therefore it is expected that changing the twist modulus should not change the results significantly, and any changes will occur for smaller aspect ratio, as observed and discussed above.

6.2 Effect of hydrodynamic wall

Next, we consider a single filament driven by a single follower force at its tip, again clamped at its base to a point on the $z = 0$ plane as shown in Figure 6.2a. As before, the filament has length L and radius a. Fixing the aspect ratio to $\alpha = 44$ corresponds to discretising the filament into $N_p = 20$ segments. We focus on 3D results only and henceforth fix $K_B = K_T$. In this section, however, we now model the filament in an unbounded domain. In our previous modelling, the no-slip wall at $z = 0$ was incorporated through the hydrodynamic mobility solver. As such, to consider unbounded flow, we instead utilise the RPY tensor without the additional image terms [Wajnryb et al., 2013], as defined in Section 4.2.3.

6.2.1 Initial value problems

We run initial value problems for $0 \leq f \leq 600$, in order to compare our results to the previous chapter, and show the results in Figure 6.2b. The dynamics are qualitatively very similar to those found in Chapter 5; we observe the same steady state below a critical forcing value, whereby the filament remains vertically upright. Beyond this critical forcing value, the filament

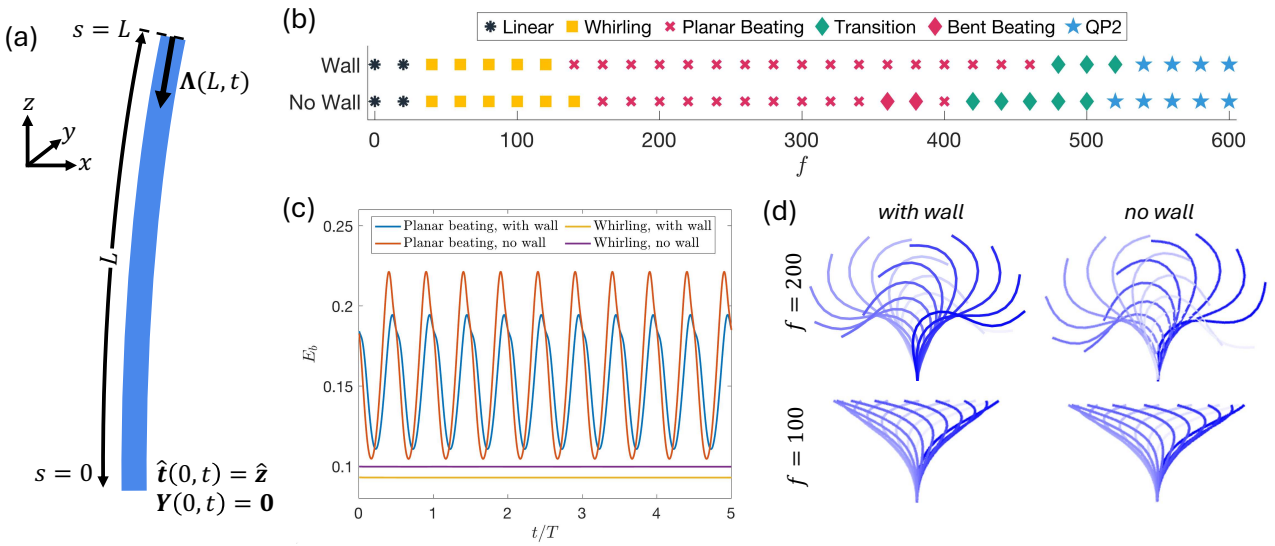


Figure 6.2: (a) A schematic showing the model we use. We consider a single filament with clamped-free boundary conditions in an open domain. (b) The results of initial value problems for $\alpha = 44$, $f \in [0, 600]$ for the case with a hydrodynamic wall at $z = 0$, from Chapter 5, and the case without the wall. (c) The bending energy against time for the whirling solutions, at $f = 100$, and the planar beating solutions, at $f = 200$, with and without a hydrodynamic wall over 5 periods. The time is non-dimensionalised by the period of the solution, obtained via JFNK. (d) The dynamics of the planar beating solutions at $f = 200$ (top row) and whirling solutions at $f = 100$ (bottom row) both with (left column) and without (right column) a hydrodynamic wall at $z = 0$. We plot each dynamic over 1 period. The passage of time is indicated by the increasing boldness of the lines.

buckles and exhibits whirling. For higher forcings, whirling becomes unstable for planar beating. Increasing the forcing further, we observe a transitionary region, followed by QP2 for the highest values of the forcing.

One qualitative difference, we note, is in the transition from planar beating to the transition region. In Figure 6.2b, there is an ambiguity in where to classify the transition region; we distinguish 'bent beating', as the planar beating solution appears to be briefly stable again before the transition region begins, at $f = 400$. However, this is very similar to the numerical simulations for a filament with the wall present for $\alpha = 132$ (see Figure 5.7c), where we observed that planar beating gave way to bent beating, and then we briefly observed planar beating again before the chaotic dynamics prevailed. As such, it could be argued that this is part of the transition region, in which case the transition region extends over even smaller values of the forcing.

In Figure 6.2d, we plot the stable whirling behaviour at $f = 100$, and the stable beating

behaviour at $f = 200$ for the cases both with and without the wall. Visually, these appear very similar, but the time trace reveals slight differences in the curvature over time; in Figure 6.2c, we plot the corresponding bending energies, and see that the presence of the wall leads to lower mean bending energies. We attribute this to be a result of the different flow fields produced in either case; the presence of the wall leads to faster decaying flow fields as a result of the additional image terms which enforce zero slip on the wall (the flow decays like $O(r^{-2})$ for a Blakelet, in comparison to $O(r^{-1})$ for a Stokeslet in unbounded flow). Deformations from the filament yield flow fields which reinforce the bending. Therefore if the flow decays slower, further bending is encouraged, and higher curvature solutions are achieved.

We also observe a quantitative difference in the bifurcations; the whirling solution becomes unstable earlier when there is a wall present, and planar beating remains stable for larger values of the forcing. The latter is likely related to the stabilising effect of the wall, as the wall causes the flow field to decay faster. We also note that we do not observe QP1 solutions through initial value problems for $\alpha = 44$.

6.2.2 Bifurcation and stability analyses

To shine further light on the dynamics we observe from initial value problems, we perform a bifurcation analysis. We can construct the bifurcation diagram by first tracking the time periodic whirling, beating and bent beating solutions, using Section 4.3. Once we have these solution branches, we can calculate a nondimensionalised version of the mean elastic bending energy,

$$\bar{E}_b = \frac{1}{T}\int_0^T \int_0^L \frac{1}{2}\kappa(s,t)^2 \, ds \, dt, \tag{6.2}$$

where $\kappa = |\partial \hat{\boldsymbol{t}}/\partial s|$ is the local curvature of the filament evaluated at arc-length s and time t, and T is a unit of time. For periodic solutions, we take T to be the solution period. For steady states, $\kappa = \kappa(s)$ and so the bending energy no longer needs temporal averaging. For other solutions, T is taken to be a long period of time. We discretise Eq (6.2) both temporally and spatially using finite differencing, and use this as an observable to differentiate between the

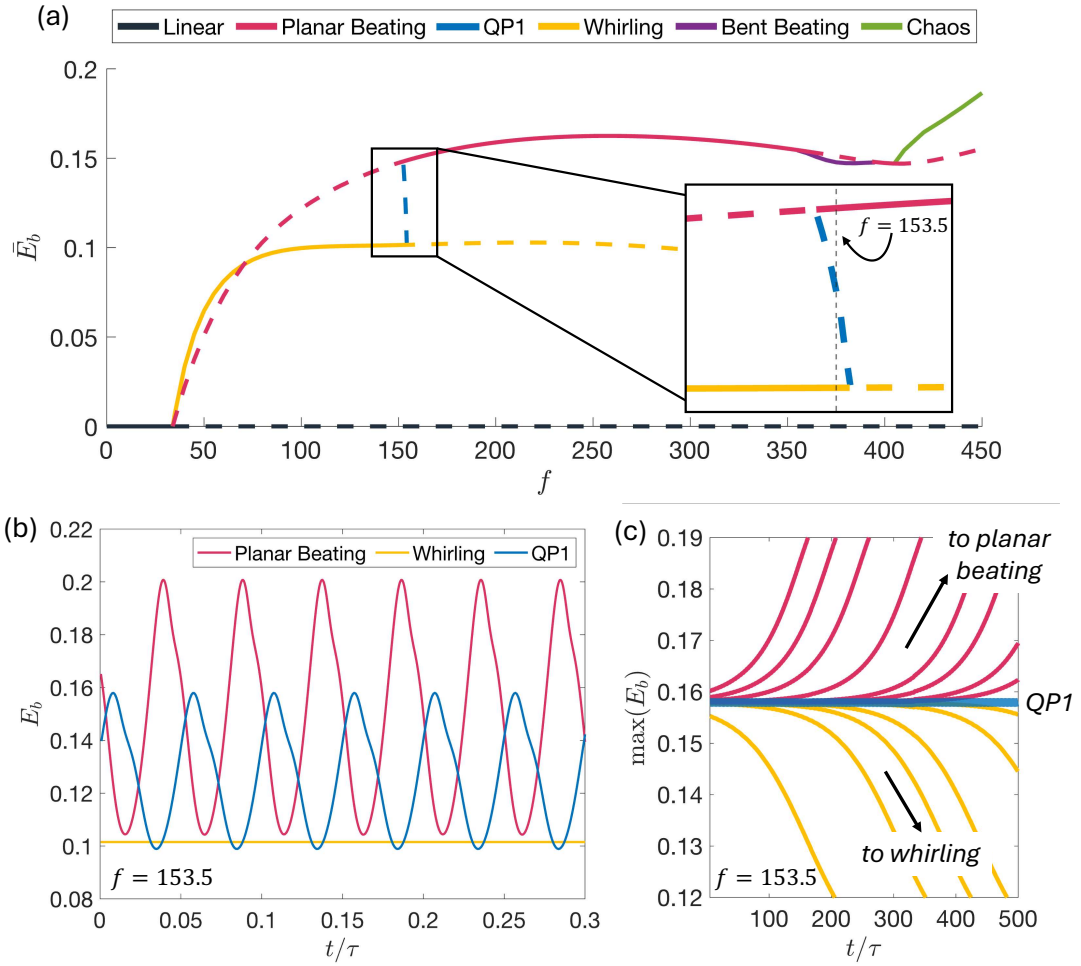


Figure 6.3: (a) The bifurcation diagram for a filament in an unbounded domain, driven by a single follower force applied at its tip, for $\alpha = 44$ and $f \in [0, 450]$. We plot the mean bending energy of each solution, obtained as described in the text. Solid (dashed) lines refer to regions in which the solution is stable (unstable). Inset: zoom at $f \in [145, 160]$, showing the QP1 solution branch and the surrounding bifurcations. We highlight $f = 153.5$ with a dashed line. (b) The bending energies of the planar beating, whirling and QP1 solutions at $f = 153.5$. (c) A plot showing the maximum bending energy of several simulations, where the initial conditions are obtained via the bisection method for $f = 153.5$. As the number of bisections increases, the solutions stay on the QP1 solution for longer before diverging to one of the stable states - planar beating (red) or whirling (yellow). The QP1 solution (blue) is obtained by taking an initial condition much later in the bisection algorithm.

different solution branches. Once we have the solution branches and the corresponding mean bending energies along the branches, as a function of f, we can then analyse the stability of these states, as well as the steady state, using the methods described in Section 4.4 to ascertain the location and nature of the bifurcations. The corresponding bifurcation diagram for the unbounded domain when $\alpha = 44$ is shown in Figure 6.3a. We outline the key critical forcing parameters corresponding to the different bifurcations, both with and without a wall, in Table

Event	Wall	No Wall
Steady state → Whirling	$f \approx 35.3$	$f \approx 33.9$
Whirling → QP1	$f \approx 137.3$	$f \approx 154.2$
QP1 → Planar Beating	$f \approx 140.0$	$f \approx 152.4$
Planar Beating → Transition	$f \approx 462.8$	$f \approx 358.4$

Table 6.1: The critical values of the forcing, f, at which each bifurcation occurs for $\alpha = 44$, both with and without the hydrodynamic wall at $z = 0$.

6.1. We now discuss each of the bifurcations in turn.

The initial buckling event is an elastic response to the follower force; as we increase the forcing from 0, the follower force is balanced by internal stresses in the filament. This continues until the forcing reaches a value such that the internal stresses are not sufficient to resolve the force, and the filament buckles. As such, this occurs at the same value of the force regardless of whether the wall is present, through the same mechanism - a double Hopf bifurcation, from which the beating and whirling solutions are born. An identical argument to Section 5.5 holds to connect the 3D buckling into whirling to the 2D buckling observed in De Canio [2019].

After the initial buckling bifurcation, whirling is stable while planar beating is unstable. The whirling solution becomes unstable at $f \approx 154.2$, while the beating solution becomes stable at $f \approx 152.4$, agreeing with observations from initial value problems (see Figure 6.2b). Accordingly, we observe a region of bistability for $152.4 \lesssim f \lesssim 154.2$. This implies the existence of an unstable state between the two stable states, which we can investigate using bisection, as described in Section 4.5. For each value of the forcing, we successively bisect between initial conditions which lead to beating or whirling solutions respectively. We calculate the bending energy of the dynamics associated with this new, bisected initial condition, to observe when we leave the unstable state and converge to one of the stable states. As we increase the number of bisections, we find that the initial conditions will transiently undergo undulations of the unstable state, allowing us to draw insight into the unstable dynamics. In Figure 6.3b, we plot the bending energy of the two stable states, whirling and planar beating, as well as the unstable state, which we find to be QP1, obtained via bisection for $f = 153.5$. Plotting only the local maxima of the bending energy time trace for various bisection iterations in Figure 6.3c highlights the idea behind bisection; progressive iterations remains on the unstable solution

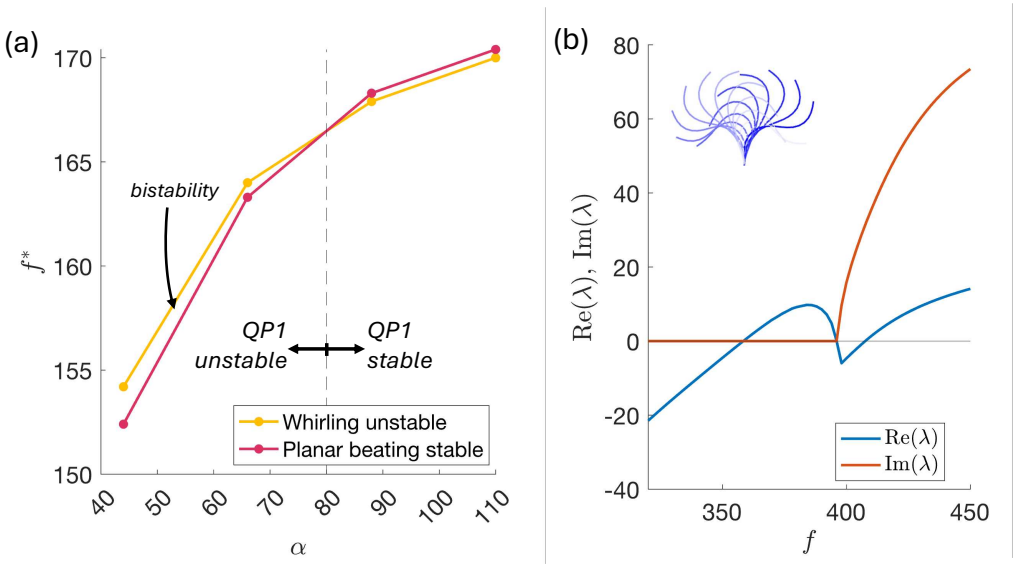


Figure 6.4: (a) The critical forcing, f^*, at which the whirling solution becoming unstable (yellow) and the planar beating solution becoming stable (red) as we vary the filament aspect ratio, α. The lines intersect at $\alpha \approx 80$ (dashed line). To the left of the dashed line, there is a region where whirling and beating are bistable. Here, the QP1 solution exists as an unstable branch connecting the two solutions. To the right of the dashed line, the QP1 solution is a stable branch which connects the two solutions. (b) The real and imaginary parts of the dominant eigenvalue associated to the planar beating solution for $f \in [320, 450]$. We note that beating briefly becomes unstable for $358.4 \lesssim f \lesssim 396.0$, where we observe bent beating in its place, and then becomes unstable again for a final time at $f \approx 407.6$.

for longer periods of time. When the difference between initial conditions goes below a given tolerance, we can use the time signal of the unstable state to infer and classify the dynamics. Each red and yellow curve in Figure 6.3c shows the results from a single iteration of bisection, with the colour of these lines chosen to emphasise the end state of the simulation. The blue line, which we label as the signal of the unstable state, is obtained from a later bisection iteration.

While the mechanism leading to the QP1 branch transitioning from stable, in Chapter 5 where a wall is present, to unstable, in the present case in the absence of a wall, is unclear, we find that the bistability between whirling and beating vanishes as we increase the aspect ratio. In Figure 6.4a, we plot the critical forcing values at which the whirling solution becomes unstable and the planar beating solution becomes stable as we vary the aspect ratio, α. As observed already, for $\alpha = 44$ there is a region of bistability between whirling and beating, as indicated in the figure by the red line lying below the yellow line. While we find that the critical forcing at which whirling becomes unstable, and beating becomes stable, both increase with filament aspect

ratio, the rate at which these change is not the same. At around $\alpha \approx 80$, we observe that the two lines cross. Beyond this value, we no longer see bistability between the two solutions, and the QP1 solution again exists as a stable branch connecting the whirling and beating solutions. To that end, the bistability for smaller aspect ratios appears to be a result of either the aspect ratio or of the spatial discretisation, and thus might not be attributed to the choice of domain.

We can classify the bifurcations on either side of the QP1 branch by analysing the Floquet exponents associated to the periodic whirling and beating solutions. As in Chapter 5, the dominant eigenvalue of the monodromy matrix associated to the whirling solution is purely imaginary at the bifurcation, indicating this to be a Hopf bifurcation. The dominant eigenvalue associated to the beating solution is purely real at the bifurcation, suggesting this to be a pitchfork bifurcation.

The observation that the planar beating solution reappears after a region of bent beating for high forcings can also be seen through Floquet analysis. In Figure 6.4b, we show the dominant eigenvalue associated to the planar beating solution for $f \in [320, 450]$. We see the planar beating solution destabilising at $f \approx 358.4$ through a pitchfork bifurcation, in which we observe the bent beating dynamic, and restabilising at $f \approx 396.0$ as the bent beating solution collapses onto the planar beating branch. The beating solution remains stable again until $f \approx 407.6$, matching observations from initial value problems, where chaotic solutions are now observed in its place. As before, the initial value problems indicate that after this brief transitionary region of chaotic and quasiperiodic solutions, for large enough values of the forcing we again observe QP2.

6.3 Effect of force distribution

Realistically, several motor proteins will be translocating along the length of each microtubule. These motors will exert an additional force on the surrounding fluid, due to the drag from carrying cargo along the filament. Consequently, it is of interest to investigate the impact these different forcing mechanisms can have on the resulting dynamics.

Previous studies have investigated the impact of varying the force distribution on a follower-force driven filament [Ling et al., 2018, De Canio, 2019] through the use of initial value problems, as well as stability analysis of the original buckling event. These studies indicate that the distributed follower force case has qualitatively similar dynamics to the single follower force at the tip, however still leave many questions unanswered about the precise differences in the state space when the forcing mechanism is changed. In particular, focusing on the smaller forcing values, it is not immediately clear whether the mechanism in the initial buckling bifurcation is the same as the single follower force case. While it has been shown that the whirling and beating solution branches emerge again in the distributed forcing case, it is not known whether the exchange of stability between these solutions again occurs through a transitionary stable quasiperiodic solution branch. Furthermore, the three dimensional dynamics at large forcing values are largely unexplored.

To bridge this gap, we now investigate the effect of changing the force distribution along the filament. We consider follower forces distributed along the filament length, as shown in Figure 6.5a. We first define a new parameter,

$$\tilde{f} = \frac{FL^3}{K_B}, \tag{6.3}$$

where again F is the dimensional follower force magnitude, K_B is the bending modulus, and L is the filament length. We note that this is related to the non-dimensional follower force via the relation $\tilde{f} = Lf$.

The other filament parameters we use, such as the filament length, aspect ratio and bending and twisting moduli, are the same as in the previous section. However in order to describe the new forcing mechanism, it is beneficial to rethink how to express the model of a single follower force at the filament tip. Rather than implementing this as a boundary condition applied at the distal end of the filament, we can consider the follower force as an active force we apply to the filament tip. Specifically, we can enforce a force-free end condition, $\mathbf{\Lambda}(L,t) = \mathbf{0}$, and define

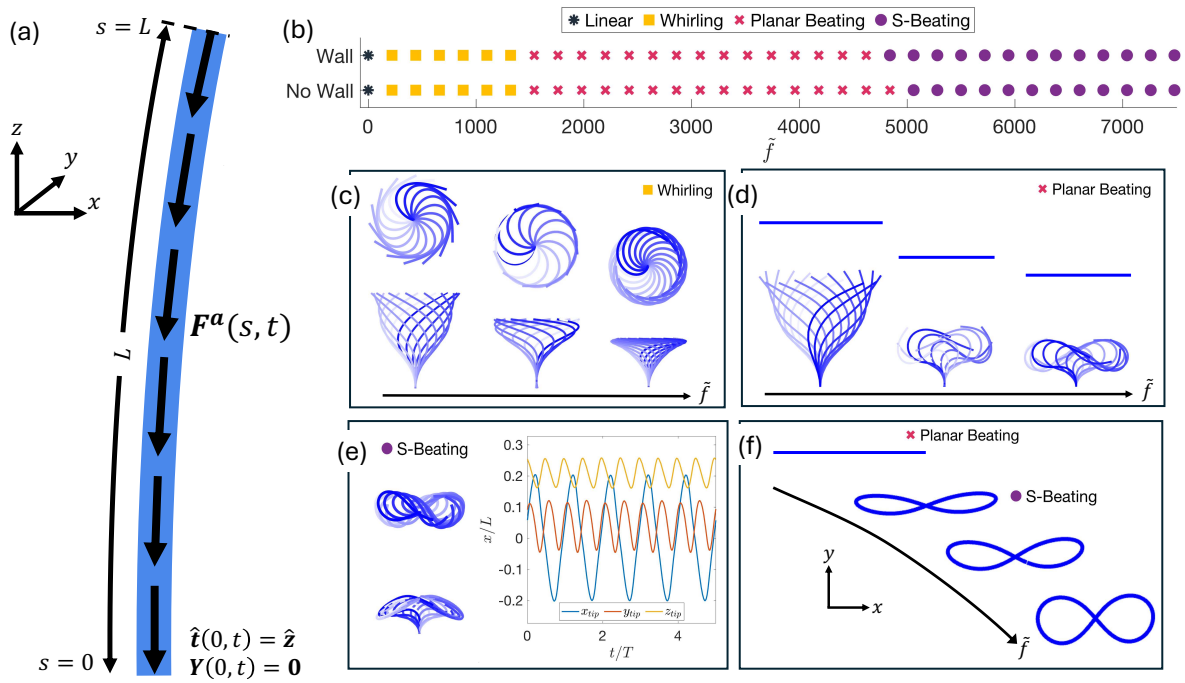


Figure 6.5: (a) A schematic showing the model we use. We consider a single filament clamped at its base in an open domain. The filament is driven by active forces along its length, $\boldsymbol{F}^a(s,t)$, as defined in the main text. (b) The results of initial value problems for $\alpha = 44$, $\tilde{f} \in [0, 7500]$ for the case (top) with and (bottom) without a hydrodynamic wall at $z = 0$. (c) The dynamics of whirling solutions for $\tilde{f} = 220$ (left), 440 (middle), and 1320 (right). All dynamics shown are plotted over one period. The passage of time is indicated by the increasing boldness of the lines. (d) The dynamics of planar beating solutions for $\tilde{f} = 220$ (left), 880 (middle), and 2200 (right). (e) The dynamics of the S-beating solution for $\tilde{f} = 6600$ (left) and the corresponding time trace of the tip coordinates over four periods (right). (f) The shape traced by the filament tip over one period for, from top left to bottom right, $\tilde{f} \in \{4500, 5060, 5280, 6600\}$. The dynamics change from planar beating to S-beating as the forcing increases.

the follower force as an active force applied to segment i,

$$\boldsymbol{F}_i^a = -\delta_{i,N_p} \frac{f K_B}{L^2} \hat{\boldsymbol{t}}_{N_p}, \tag{6.4}$$

which modifies the force balance in Eq (4.6) to

$$\boldsymbol{F}_i^C + \boldsymbol{F}_i^H + \boldsymbol{F}_i^a = \boldsymbol{0}. \tag{6.5}$$

As such, the distributed follower force case will correspond to instead applying the active force:

$$\boldsymbol{F}_i^a = -\frac{\tilde{f} K_B}{L^3} \hat{\boldsymbol{t}}_i \tag{6.6}$$

to each segment. To investigate how this will impact the dynamics, we first run initial value problems for different forcing values. Numerical simulations will be performed for cases in which the filament is in unbounded flow, as well as in the presence of a no-slip planar wall. Following this, we will again perform a bifurcation analysis on the system for $\alpha = 44$, which will primarily focus on the unbounded case, due to the qualitative similarities of the two.

6.3.1 Initial value problems

The results of initial value problems for $0 \leq \tilde{f} \leq 7500$ both with and without a hydrodynamic wall are shown in Figure 6.5b. This range of f is chosen as we find that it encapsulates the key set of dynamics in the state space, in particular the first few bifurcations and the dynamics beyond planar beating, allowing for a comparison with the previous chapter. We observe many similarities with the single follower force case for low to moderate forcing values; below a critical forcing value we observe a steady state, in which the filament aligns with the $\hat{\boldsymbol{z}}-$axis. Beyond this critical forcing, the steady state becomes unstable and gives rise to whirling. For higher forcing values, whirling becomes unstable and planar beating is observed in its place.

The dynamics reflect the strongly compressive forcing acting along the filament length. In Figure 6.5c, we show the whirling behaviour for three values of the forcing; $\tilde{f} = 220, 440$ and 1320. As we get further from the bifurcation, we see that the whirling solution becomes drastically more compressed. Similarly, we show the planar beating dynamic at $\tilde{f} = 220, 880$ and 2200 in Figure 6.5d. Again we observe beats with significantly higher curvature as we increase the forcing, leading to dynamics which are visually significantly more compressed. We also note that, as in the previous section, we don't observe QP1 through initial value problems close to the transition between whirling and beating.

Along the beating solution branch, it appears that the follower force is balanced by the internal stresses of the filament, leading to higher-curvature beat patterns as shown in Figure 6.5d. However, increasing the forcing further, as in the previous models we again reach a critical value of the forcing at which planar beating is no longer observed. At the critical value of the forcing, the filament is no longer able to resolve these forces in the plane, and so favours

Event	Wall	No Wall
Steady state → Whirling	$\tilde{f} \approx 169.7$	$\tilde{f} \approx 169.7$
Whirling → QP1	$\tilde{f} \approx 1426.8$	$\tilde{f} \approx 1500.5$
QP1 → Planar Beating	$\tilde{f} \approx 1151.8$	$\tilde{f} \approx 1216.3$
Planar Beating → Transition	$\tilde{f} \approx 4773.2$	$\tilde{f} \approx 4912.3$

Table 6.2: The critical values of the forcing, $\tilde{f}$, at which each bifurcation occurs when follower forces are distributed along the filament length for $\alpha = 44$, both with and without the hydrodynamic wall. The bifurcation diagrams corresponding to the two domains have the same qualitative properties, with the quantitative differences in the critical forcing summarised in this table.

a solution in which the filament can also deform out of the plane. Consequently, we observe a periodic dynamic which we call S-beating, shown in Figure 6.5e. This is similar to the planar beating solution, but with an off-plane component that grows as we get further from the bifurcation. We plot the trajectory traced by the filament tip, viewed from the $(x, y)-$plane for different values of the forcing in Figure 6.5f. From this figure, we see the filament tip traces a figure of 8 over each period. As we increase the strength of the forcing, the amplitude in the $\hat{\boldsymbol{y}}-$direction increases. We plot the time trace of the solution for $\tilde{f} = 6600$ in Figure 6.5e, right.

The exchange of stability from planar beating onwards is a much simpler transition than in the single follower force case; recall that in this case we instead entered a transitionary region, in which we observed a host of periodic, quasiperiodic and chaotic solutions. Beyond this, the dynamics settled into a quasiperiodic beating, which we called QP2. However now, this transition into a single periodic solution indicates that the distribution of the forcing has a stabilising effect on the dynamics.

6.3.2 Bifurcation and stability analyses

We now turn our attention to performing a bifurcation analysis on the distributed follower force dynamics. We perform a similar analysis on both the bounded and unbounded domain cases, which indicates the wall has no qualitative impact on the dynamics and bifurcations, agreeing with observations from initial value problems (see Figure 6.5b). Therefore we summarise the critical forcing values of each bifurcation for both cases in Table 6.2, to allow for a comparison,

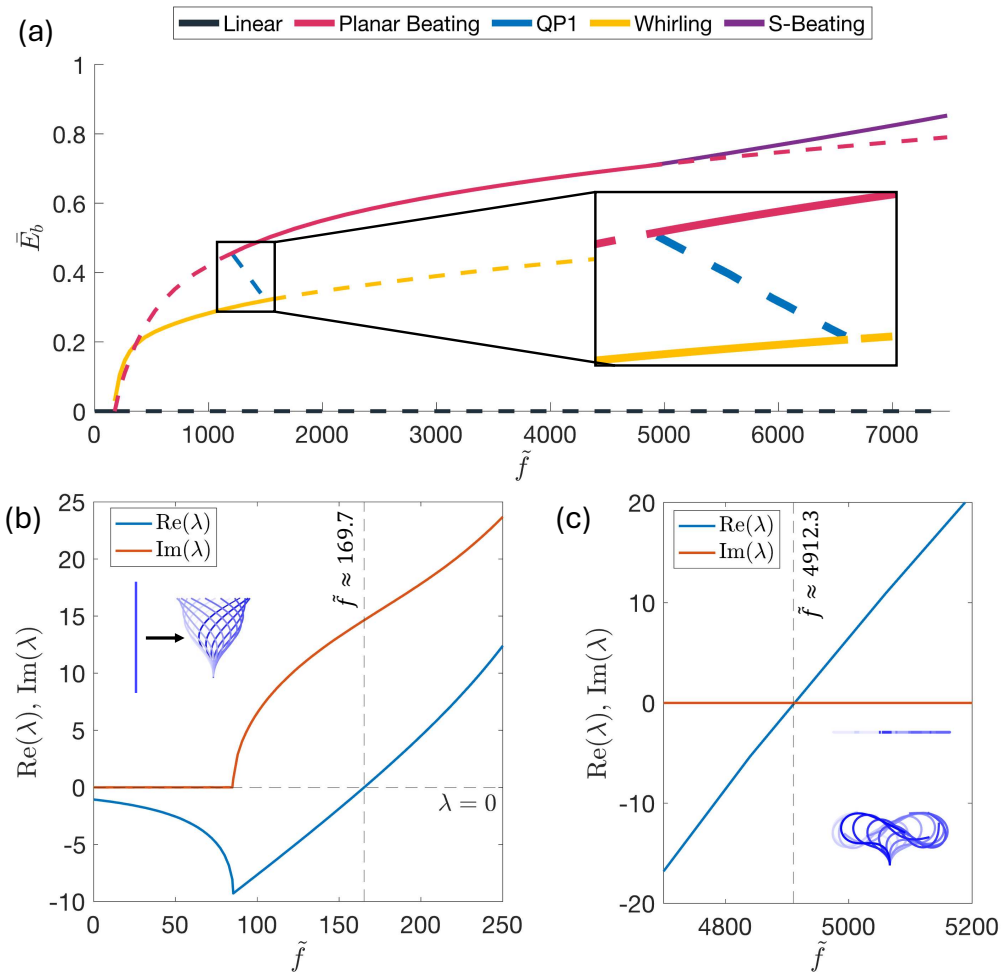


Figure 6.6: All diagrams in this figure correspond to a single filament in an unbounded domain, driven by a distribution of follower forces along its length. (a) The bifurcation diagram for $\alpha = 44$ and $\tilde{f} \in [0, 7500]$. We plot the mean bending energy of each solution, obtained as described in the text. Solid (dashed) lines refer to regions in which the solution is stable (unstable). Inset: zoom for $\tilde{f} \in [1100, 1600]$, showing the QP1 solution branch and the surrounding bifurcations. (b) The real and imaginary parts of the dominant eigenvalue corresponding to the steady state for $\alpha = 44$ and $\tilde{f} \in [0, 250]$. We see that the steady state becomes unstable at $\tilde{f} \approx 169.7$. (c) The real and imaginary parts of the dominant eigenvalue associated to the beating solution for $\alpha = 44$, and $\tilde{f} \in [4700, 5200]$. The eigenvalue passes the $y-$axis at $\tilde{f} \approx 4912.3$, and is purely real at the bifurcation.

and henceforth we focus our discussion of the bifurcation analysis on the unbounded case. We plot the bifurcation diagram for $\alpha = 44$ in the unbounded domain case, obtained in the same way as described in Section 6.2.2, in Figure 6.6a. We now discuss the key features of the bifurcation diagram in turn.

We begin by performing a linear stability analysis of the steady state; we plot the dominant eigenvalue in Figure 6.6b. This analysis indicates that the steady state becomes unstable at

$\tilde{f} \approx 169.7$, again through a double Hopf bifurcation, in which two complex conjugate eigenpairs become unstable at the same forcing value. The eigenvectors are again linear combinations of planar eigenvectors in perpendicular planes, as in the previous models. We find that the critical buckling value is the same, regardless of the aspect ratio.

We next track the periodic whirling and beating branches using JFNK, and assess their stability. We find that the whirling solution becomes unstable at $\tilde{f} \approx 1500.5$, and the beating solution becomes stable at $\tilde{f} \approx 1216.3$. Therefore, we again find bistability between whirling and planar beating for $1216.3 \lesssim f \lesssim 1500.5$. Bisection in this range of $\tilde{f}$ reveals that the usual QP1 solution is the unstable branch which connects the two behaviours (see Figure 6.6a, inset). While we found in Section 6.2 that this bistability would vanish as we increase the aspect ratio, the same does not hold here, in this case. We find that bistability exists regardless of the filament aspect ratio, suggesting that this can be attributed to the modelling, opposed to the spatial resolution. By analysing the Floquet exponents of the periodic solutions, we again find the transition from whirling to QP1 occurs through a Hopf bifurcation, and QP1 to beating through a pitchfork bifurcation.

We can also explore the transition of planar beating to S-beating by analysing the dominant floquet exponents associated to planar beating. We plot these in Figure 6.6c. We find that the beating becomes unstable at $\tilde{f} \approx 4912.3$. At the bifurcation, the dominant eigenvalue is real and zero. This is indicative of a supercritical pitchfork bifurcation. We find that S-beating is stable for the remaining values of the investigated forcing values and, through initial value problems, we find that this solution is robust against different initial conditions.

6.4 Discussion

In this chapter, we have presented results on several extensions of the follower force model introduced in Chapter 5. First, we examined the impact of the choice of twist modulus in Section 6.1. We generated rough estimates of the ratio between the bending modulus and twist modulus for an idealised version of a microtubule, which were used to inform the possible order

of magnitude of this ratio. With this in hand, we ran initial value problems varying the twist modulus relative to the bending modulus for low to moderate forcing values, to distinguish changes to the state space. Small changes at the region between the beating and whirling solutions were observed if the twist modulus is sufficiently smaller than the bending modulus, though this effect vanishes as we increase the aspect ratio. We observe no differences in the state space when the twist modulus is larger than the bending modulus, or when the aspect ratio is large, indicating that the effect of the twist modulus is small.

In Section 6.2, we investigated the impact of removing the hydrodynamic wall in the case of a single follower force driven filament. We found that the wall has a predominantly quantitative effect on the dynamics. The first few bifurcations, in particular the initial buckling into whirling, and the transition between whirling and beating, are primarily dictated by the elasticity of the filament, and so are largely unchanged when the wall is removed. We note a slight difference in the nature of the QP1 branch, which we attribute to the small aspect ratio and spatial resolution of the filament. The role of the wall becomes more prominent for larger values of the forcing, where more energy is being put into the system through the follower force, which results in larger deformations and additional dissipation into the fluid. When the wall is present the flow decays faster, meaning the wall has a stabilising effect on the dynamics. One way in which this manifests is that the beating solution remains stable for a significantly larger region of f when the wall is present, delaying the transition to the later observed dynamics.

In Section 6.3, we turned our attention to the forcing mechanism, investigating the impact of distributing the force over the length of the filament. The bifurcations and dynamics for low to moderate forcing values remain qualitatively similar to the other models considered. We again observe buckling through a double Hopf bifurcation into whirling, which gives way to beating for higher forcing values. We do see a change in the nature of the QP1 branch, which exists as an unstable branch in a region of bistability between whirling and beating, as well as a visual change in the dynamics as the filament responds to the strong compressional forces along its length. In fact in Ozdemir et al. [2025], it was shown that this bistability appears to be a feature of distributing the forcing along the filament length; confining the forcing to the tip leads to stable QP1 branches, at least in the limit of large α. The predominant change in

distributing the forcing, however, is in the dynamics at higher forcing values; distributing the forcing along the length appears to stabilise the dynamics, avoiding the messy transition region that occurs in the tip-driven cases and instead exhibits a single, periodic solution after planar beating destabilises. A more comprehensive discussion of these results, alongside comparisons to a forcing mechanism which includes a slip flow, can be found in Ozdemir et al. [2025].

Previous studies have investigated the impact of varying the force distribution on a follower-force driven filament with these clamped-free end conditions, through varying the distribution of the forcing along the length in 2D [De Canio, 2019], or in 3D [Ling et al., 2018], or by incorporating the drag force from the cargo the motor protein carries along the filament through an equal and opposite force [De Canio et al., 2017, De Canio, 2019, Stein et al., 2021], or through a slip flow on the filament surface [Laskar and Adhikari, 2017, Ozdemir et al., 2025]. Our results agree with the findings of [Ling et al., 2018, De Canio, 2019], that we see the same states for low to moderate forcing values regardless of the forcing distributions. However these studies primarily focus on initial value problems, with any stability analysis surrounding only the first buckling bifurcation. The stability analysis we perform allows us to delve deeper into the differences in the phase space as these parameters are varied.

In [Laskar and Adhikari, 2017, Ozdemir et al., 2025], the backflow from the cargo-carrying motor is incorporated through the use of the squirmer model [Lighthill, 1952, Lauga, 2020], in which each segment of the filament has a prescribed surface slip velocity. Regardless of the actuation mechanism, it seems the buckling into whirling, and subsequent transition of stability to the beating solution, is generic to these systems. The specific mechanisms, such as the stability of the QP1 branch and the existence of bistability between whirling and beating, depend on the details of the forcing mechanism and domain. The key differences in dynamics arise for higher forcing values, beyond the beating solution.

While we have not observed drastic changes in the emerging dynamics from these different forcing mechanisms and domains, we do expect there to be a more prominent distinction when considering collective dynamics, when the hydrodynamic effects become more important. In such a case, the presence of the wall, for instance, will impact the hydrodynamic coupling

between neighbouring filaments, and may affect synchronisation behaviours. Investigating the impact of these modelling choices on arrays of filaments would allow for an effective comparison of these models with actual microtubule dynamics and hence, better ascertain the correct models for predicting more complicated microtubule-motor protein related phenomena. However instead considering these results in the context of ciliary beating, the key conclusion from the results presented in this chapter differs; the states that were most interesting from Chapter 5 - that is, the whirling and planar beating states - appear to exist regardless of the forcing mechanism. Therefore in order to encapsulate cilia-like beating dynamics, the first few bifurcations of the simplest case (a single follower force applied at the filament tip) will suffice. This will be the focus of the next section.

Chapter 7

Breaking Symmetry in a Follower Force-Driven Filament

Note: the material in this chapter is derived from Clarke et al. [2025]. The dissertation author was the primary researcher and author of this paper.

7.1 Introduction

Asymmetry is ubiquitous in nature across multiple scales, from the chirality of molecules [Garbacz, 2024] to the left-right asymmetry in the placement of internal organs in the human body [Hirokawa et al., 2009]. Many animals utilise asymmetry for mating or survival, such as flatfish like flounders and halibut, which have both eyes on one side of the head to allow them to lie flat against the ocean floor and avoid predators [Palmer, 2009]. Anisotropy, here referring to the direction dependence of mechanical response of a material, is also common in nature, a simple example being the human hand, which is easier to bend in one direction than in any other direction. Fibrous materials such as wood are anisotropic due to the natural grain formed by fibre alignment, meaning the strength and flexibility of these materials are direction-dependent [Kretschmann, 2010].

As suggested in Chapter 1, asymmetries and anisotropies also exist within the structure of cilia, and are reflected within the resulting dynamics. Symmetry breaking is essential for fluid transport at these scales, following Purcell's so-called 'scallop theorem' [Purcell, 1977], which states that reciprocal motion does not generate net fluid flow in a Stokes regime, as a consequence of the time-reversibility of the Stokes equations. Accordingly, cilia often undergo asymmetric beating oscillations, to ensure net fluid motion and cellular transport. Examples include the distinct fast, low-curvature effective strokes and slower recovery strokes performed by respiratory cilia [Chilvers and O'Callaghan, 2000, Jorissen and Jaspers, 2023], or the left-tilted three-dimensional whirling motions undergone by nodal cilia in the embryo [Nonaka et al., 1998, Smith et al., 2008a, 2011].

Such asymmetries in ciliary dynamics are thought to be the result of the internal structure of cilia. Recall the '9+2' axonemal structure, common to motile cilia [Gibbons, 1981, Marshall and Nonaka, 2006, Gilpin et al., 2020], consisting of a central pair of microtubules, surrounded by 9 microtubule doublets (see Figure 1.2). The doublets are held together by structural proteins and molecular motors. The bridge between one pair of doublets, labelled doublets 5 and 6 in Figure 1.2, is known to be much stiffer than the others [Afzelius, 1959, Gibbons, 1981], indicating that this pair have limited motion relative to each other. The orientation of the central pair and the stiffer bridge suggest a plane in which it is easier for the axoneme to bend. It has been observed that the beat plane for cilia aligns with this plane [Gibbons, 1981], suggesting that the direction of the cilium beat could be a result of anisotropy in axonemal structure. In previous numerical studies modelling cilia, this has been included explicitly by considering a stiffer bridge between this pair of doublets [Han and Peskin, 2018], or implicitly by increasing the perpendicular bending modulus [Rallabandi et al., 2022]. It has also often been used as motivation for restricting dynamics to be entirely planar [Bayly and Dutcher, 2016].

While these modelling choices can reproduce planar beating, the resulting dynamics are often symmetric and do not reproduce differences in the effective and recovery strokes observed experimentally. Additional asymmetry must be considered to generate more accurate simulations which match experiments. Some works achieve this by incorporating constraints on material

properties of the filament, such as curvature-dependent bending moduli [Han and Peskin, 2018] in which the bending modulus can take different values depending on the local curvature of the filament's centreline, or through biased kinetics of dynein motors [Chakrabarti and Saintillan, 2019a,b, Chakrabarti et al., 2022]. A widely accepted approach, however, includes introducing an intrinsic curvature to the filament [Sartori et al., 2016, Chakrabarti and Saintillan, 2019a,b, Cass and Bloomfield-Gadêlha, 2023, Wang et al., 2023], in which it is assumed that the rest state of the cilium has finite curvature. These numerical studies are successful in demonstrating realistic cilia beats, with a close resemblance to the beats observed in nature, but are often restricted to two dimensional dynamics. Further studies are necessary to determine whether these dynamics persist when the filament is free to deform in any direction.

In this chapter we explore the effect of asymmetries and anisotropies in a follower force-driven filament. While the follower force model does not represent the sliding-mechanism known to generate ciliary beating, as we have shown in Chapter 5 this model has been shown to exhibit beating and whirling dynamics for sufficient forcing values, which are reminiscent of cilia beating patterns observed in nature. We begin by outlining how the modelling techniques will be modified to incorporate asymmetry, through a non-trivial intrinsic curvature, and anisotropy, through anisotropic bending. Next, we investigate the impact of these parameters, beginning with a non-trivial intrinsic curvature to an otherwise-isotropic filament, and establish the dynamics in both 2D, matching those observed in the literature, and 3D. Following this, we discuss the effect of direction-dependent bending moduli in the absence of intrinsic curvature, before investigating how this anisotropy can be exploited on filaments with intrinsic curvature to recover and stabilise asymmetric beating dynamics in 3D.

7.2 Methods and simulation parameters

In this chapter, we model a single filament clamped at it's base to a no-slip planar wall and forced at its tip by a single follower force. We largely follow the model described in Section 4.2. The key difference is that we now allow for non-zero intrinsic curvature as well as differences in

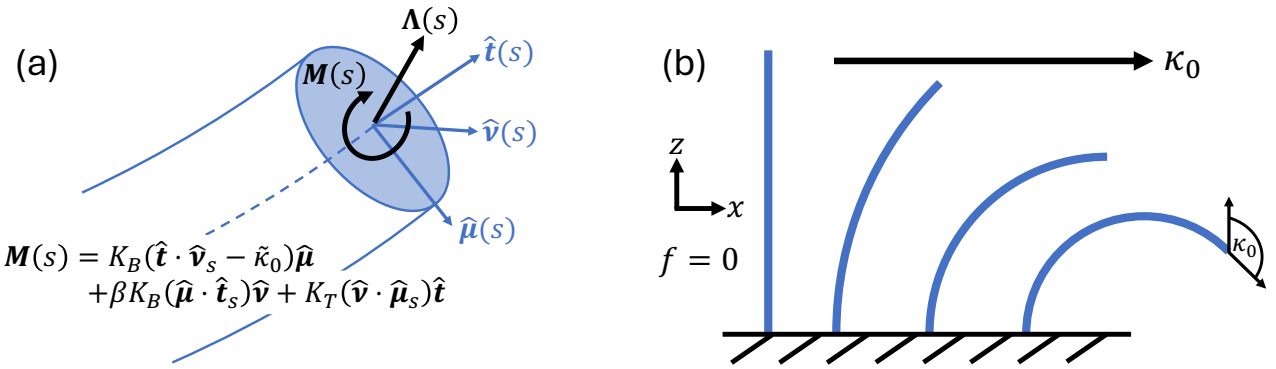


Figure 7.1: Schematics of the model. (a) The bending modulus anisotropy and intrinsic curvature are defined through the bending moment, which acts on the filament cross-section. It is expressed in terms of the filament's material frame, $(\hat{\boldsymbol{t}}, \hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\nu}})$. We emphasise that the filament cross-section is still circular. (b) The intrinsic curvature, $\kappa_0 = L\tilde{\kappa}_0$ is defined to be the rest-curvature of the filament under zero forcing.

the bending modulus, which are incorporated through the constitutive relation for the bending moment along the filament centreline given by Eq (4.4). Specifically, we first use the expression [Schöller et al., 2021]:

$$\hat{\boldsymbol{t}} \times \frac{\partial \hat{\boldsymbol{t}}}{\partial s} = \left(\hat{\boldsymbol{t}} \cdot \frac{\partial \hat{\boldsymbol{\nu}}}{\partial s}\right) \hat{\boldsymbol{\mu}} + \left(\hat{\boldsymbol{\mu}} \cdot \frac{\partial \hat{\boldsymbol{t}}}{\partial s}\right) \hat{\boldsymbol{\nu}}, \tag{7.1}$$

in Eq (4.4) to decompose the bending moment into the components about each local basis vector, and instead use the adapted constitutive relationship:

$$\boldsymbol{M}(s,t) = K_B \left(\hat{\boldsymbol{t}} \cdot \frac{\partial \hat{\boldsymbol{\nu}}}{\partial s} - \tilde{\kappa_0}\right) \hat{\boldsymbol{\mu}} + \beta K_B \left(\hat{\boldsymbol{\mu}} \cdot \frac{\partial \hat{\boldsymbol{t}}}{\partial s}\right) \hat{\boldsymbol{\nu}} + K_T \left(\hat{\boldsymbol{\nu}} \cdot \frac{\partial \hat{\boldsymbol{\mu}}}{\partial s}\right) \hat{\boldsymbol{t}}. \tag{7.2}$$

Here, $\tilde{\kappa}_0 = \kappa_0/L$ defines the intrinsic curvature, such that the rest configuration of the filament is a planar arc with constant curvature κ_0, as sketched in Figure 7.1b. Therefore, larger values of κ_0 correspond to a higher curvature rest configuration for the filament. Also defined in Eq (7.2) are K_B, the filament bending modulus in the direction of intrinsic curvature, βK_B, the bending modulus in the perpendicular plane, and K_T, the twisting modulus. Choosing $\beta > 1$ leads to preferred bending about the $\hat{\boldsymbol{\mu}}-$axis, and $\beta < 1$ leads to preferred bending about the $\hat{\boldsymbol{\nu}}-$axis. Correspondingly, setting $\beta \to \infty$ should restrict the filament dynamics to the $(\hat{\boldsymbol{\nu}}, \hat{\boldsymbol{t}})$-plane entirely. Choosing $\beta = 1, \kappa_0 = 0$ corresponds to the isotropic, symmetric case considered in Chapter 5.

In this chapter, we fix the filament aspect ratio to be $\alpha = 44$. The twist modulus is taken to be

equal to the bending modulus in the $\hat{\boldsymbol{\mu}}-$direction, $K_B = K_T$. In the isotropic case, where $\beta = 1$ and $\kappa_0 = 0$, the magnitude of the twist has little effect on the dynamics (see Section 6.1), and so we anticipate the effect of the magnitude of the twist modulus to be small in this case also. The filament is clamped at its base to a no-slip planar wall, and so the hydrodynamic solver we use is RPY, adapted to include the contributions of the wall (outlined in Appendix B). We investigate the effect of varying three non-dimensional parameters; the intrinsic curvature, κ_0, as in Eq (7.2); the bending anisotropy, quantified by β, as in Eq (7.2); and the nondimensional follower force, f, as defined in Eq (4.5). We focus on ranges of κ_0 and β that follow values explored in other studies [Sartori et al., 2016, Lindemann and Lesich, 2016, Chakrabarti and Saintillan, 2019a,b, Chakrabarti et al., 2022, Rallabandi et al., 2022, Cass and Bloomfield-Gadêlha, 2023, Wang et al., 2023], and a range of f which allows us to access and explore the whirling and beating states. All initial value problems use the initial condition of a vertically upright filament, with random perturbations to the force on each segment. All timescales are non-dimensionalised by the filament's relaxation time, τ, as described in Section 5.2.

7.3 Intrinsic curvature in isotropic filaments

7.3.1 Planar dynamics

We begin by considering an isotropic filament with finite intrinsic curvature, i.e. $\beta = 1$, and $\kappa_0 \neq 0$, with its motion restricted to the plane in which it is deformed by its intrinsic curvature. The corresponding phase diagram of the planar filament states in $\kappa_0 - f$ space is obtained by running one initial value problem for each set of parameters. We summarise the results in Figure 7.2a. For this study we consider $\kappa_0 \in [0, 3\pi/4]$. This range of values is appropriate as it includes values used in the literature in numerical studies [Wang et al., 2023], in particular to recover dynamics similar to those observed experimentally [Sartori et al., 2016, Chakrabarti and Saintillan, 2019a]. These initial value problems reveal that we observe two types of behaviour; steady states below a critical value of the forcing, or self-sustained beating oscillations when this critical value is exceeded, matching observations in [Wang et al., 2023].

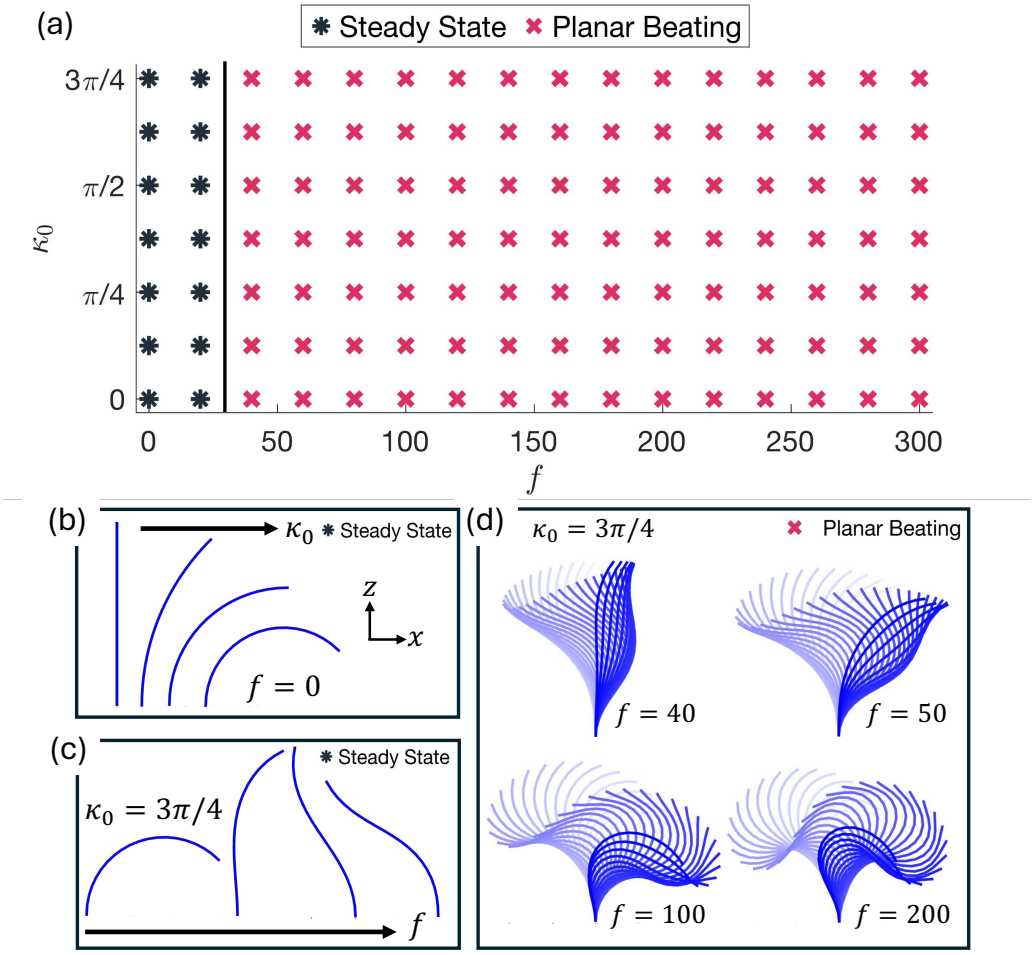


Figure 7.2: (a) Results from 2D numerical simulations of a filament with intrinsic curvature ($\kappa_0 \geq 0$, $\beta = 1$). The black line is drawn to indicate the approximate boundaries of the solution regions. For $f < f^* \approx 35.3$, filaments remain in their steady non-trivial state which varies with (b) κ_0 (shown for $f = 0$, $\kappa_0 \in \{0, \pi/4, \pi/2, \pi/4\}$) and (c) f (shown for $\kappa_0 = 3\pi/4$, $f \in \{0, 5, 10, 20\}$). For $f > f^*$ filaments undergo asymmetric oscillations for $\kappa_0 > 0$, as shown for (d) $\kappa_0 = 3\pi/4$. The passage of time is indicated by the increasing boldness of the lines.

For $\kappa_0 = 0$, the steady state is a vertically upright filament. However increasing $\kappa_0 \neq 0$, the steady state becomes non-trivial. For $f = 0$, it is given by the prescribed intrinsic curvature, such that

$$\hat{\boldsymbol{t}}(s) = (\sin(\kappa_0 s/L), 0, \cos(\kappa_0 s/L)), \quad s \in [0, L], \tag{7.3}$$

as shown in Figure 7.2b. When $f > 0$, this steady state changes; the tip of the filament is pushed by the follower force, until a new steady state is reached in which the elastic forces balance the follower force. This gives a non-trivial steady state, as shown for $\kappa_0 = 3\pi/4$ in Figure 7.2c.

When the force exceeds a critical value, the filament buckles and we observe planar oscillations. To find this critical value for each κ_0, we track the steady state using JFNK and use stability analysis. For $\kappa_0 = 0$, this reveals that the buckling occurs at $f = f^* \approx 35.3$, through a Hopf bifurcation. For $\kappa_0 > 0$, the bifurcation occurs through the same mechanism, but we observe a small increase in the critical forcing of approximately 0.5% from $\kappa_0 = 0$ to $\kappa_0 = 3\pi/4$. For $\kappa_0 = 0$, numerical simulations after the bifurcation show that the planar beating oscillations are symmetric. However, this symmetry is broken by intrinsic curvature; for $\kappa_0 > 0$, simulations reveal asymmetric beating patterns as shown in Figure 7.2d, and observed in [Wang et al., 2023]. Close to the bifurcation this asymmetry is strongest, but further away from the bifurcation, where the follower force is strong, the effect of intrinsic curvature is less pronounced and the beating becomes more symmetric.

7.3.2 Non-planar dynamics - initial value problems

When the filament is free to deform in any direction, we observe dramatic changes in filament dynamics as we vary κ_0 and f. We summarise the filament states in Figure 7.3a. For the smallest values of the forcing, we observe the same steady states as in the planar case. However, we now find that the filament exhibits fully 3D dynamics after buckling and the critical value of the forcing has a dependence on κ_0.

The case where $\kappa_0 = 0$ was investigated in Chapter 5. We summarise the findings from this chapter here, for clarity. Immediately after the buckling, we observe a symmetric whirling state, where the filament tip traces a circle over one period. Increasing the forcing further, whirling is replaced by QP1, which exists for a small region of f, before being replaced with planar beating. Planar beating is also observed in the 2D simulations, but is now symmetric (as $\kappa_0 = 0$) and can occur in any plane (due to the rotational symmetry of the problem).

Increasing $\kappa_0 > 0$, the initial buckling occurs at smaller values of the forcing. The whirling behaviour still exists but is no longer symmetric, as shown in Figure 7.3b. In fact, the asymmetry is strongest as we increase the intrinsic curvature, and choose forcing values closer to

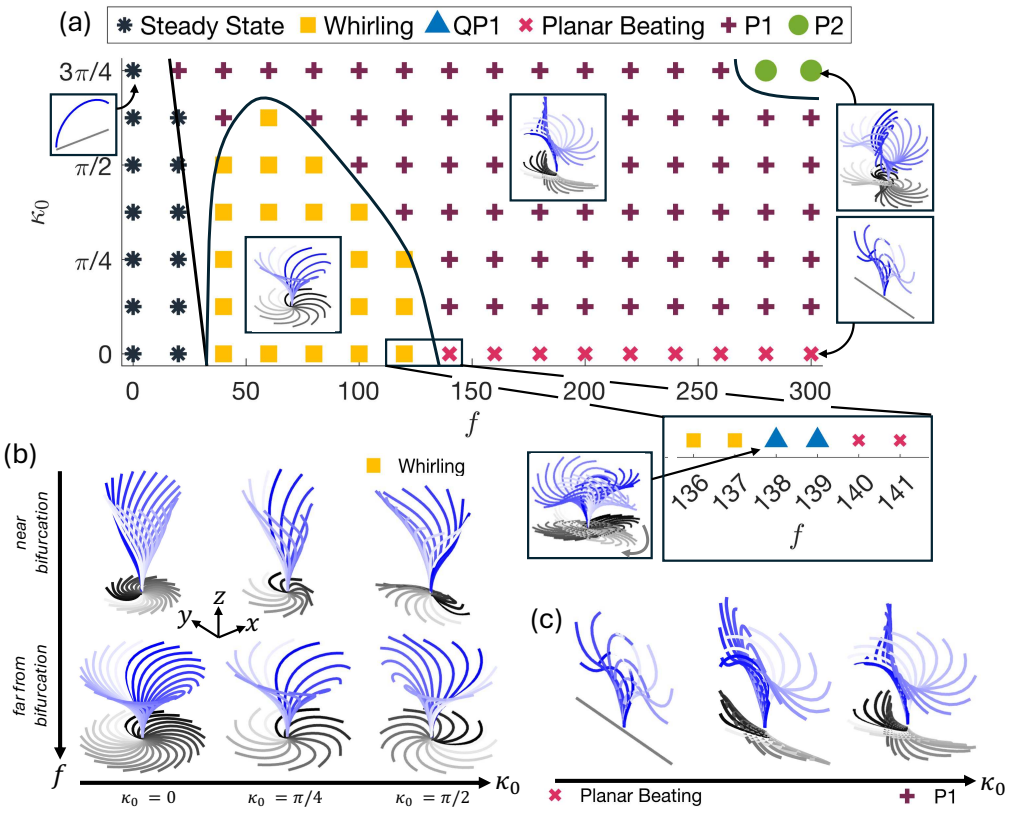


Figure 7.3: (a) Results of numerical simulations of a filament with intrinsic curvature in 3D ($\kappa_0 \geq 0, \beta = 1$). The black lines are drawn to indicate the approximate boundaries of the solution regions. (b) Snapshots of filaments over one period undergoing whirling dynamics for two values of intrinsic curvature close to ($f = 40$) and further from ($f = 100$) the bifurcation. We see asymmetries in whirling are more pronounced close to the bifurcation, and for higher κ_0. (c) Snapshots of filaments over one period for $f = 200$, increasing κ_0 from 0 to $3\pi/4$. The dynamics change from planar beating for $\kappa_0 = 0$ (left) to P1 for $\kappa_0 > 0$ (right). The passage of time is indicated by the increasing boldness of the lines.

the initial bifurcation. Increasing the forcing has the effect of smoothing out the asymmetry, as the follower force dominates over the torques maintaining the intrinsic curvature.

Increasing f further, we observe that whirling gives way to a new periodic solution for $\kappa_0 > 0$, which we call P1. This consists of a bent-over beating, where the filament bends in the plane it would deform in the absence of the follower force, and beats in the plane perpendicular to it, as shown in Figure 7.3a inset and Figure 7.3c. The intrinsic curvature breaks left-right symmetry, in such a way that the filament beats in the left half of the plane. In Figure 7.3c, we fix $f = 200$ and increase κ_0 from 0 to $3\pi/4$ to show how P1 continues from the planar beating solution. We see that increasing κ_0 increases the maximum perpendicular distance of the tip from the beat plane. Varying $\kappa_0 > 0$, planar beating is no longer visible with initial value problems alone - it

appears that the 2D asymmetric beating patterns observed in the planar case are unstable in the 3D model.

For larger values of κ_0, we note that the whirling state is also no longer visible using initial value problems, regardless of the initial filament configuration. Furthermore, for large κ_0 and f, we see a new periodic mode appearing, which we call P2. P2 beats similarly to P1, but traces a distinct shape of 8 with its tip over one period (see Figure 7.3a, inset).

7.3.3 Bifurcation and stability analyses

To explore in more detail how filament dynamics change with κ_0, we track the steady and time-periodic states with JFNK and perform a stability analysis to establish bifurcations within the system. An overview of the resulting bifurcation diagrams for $f \in [0, 200]$ is shown in Figure 7.4, where we plot the non-dimensionalised mean bending energy given by Eq (6.2) for each solution branch. While the bending energy is usually defined with respect to the intrinsic curvature, i.e. replacing $\kappa(s)$ with $(\kappa(s) - \kappa_0)$ in Eq (6.2), we instead use the standard bending energy (with $\kappa_0 = 0$) for ease of comparison as we vary κ_0. We summarise the critical forcing value at which each bifurcation in these diagrams occurs in Table 7.1. We now discuss the key changes in the bifurcation diagrams as we increase κ_0.

Isotropic case

We begin by recapping the isotropic case discussed in Chapter 5, where $\kappa_0 = 0$. The bifurcation diagram is shown in Figure 7.4a. The initial buckling event was found to be a double Hopf bifurcation at $f = f^*$, in which two complex conjugate pairs of eigenmodes become unstable. To understand the buckling bifurcation further, which will be particularly helpful for $\kappa_0 > 0$, we can visualise the eigenvectors. Defining each eigenvalue-eigenvector pair as $(\lambda_i, \boldsymbol{x}_i)$, we note that the $\boldsymbol{x}_i$ are complex for $i = 1, ..., 4$. Therefore, to interpret their physical meaning, we take linear combinations of the eigenvectors, $\boldsymbol{v}_1 = C(\boldsymbol{x}_1 + \boldsymbol{x}_2)$, $\boldsymbol{v}_2 = C(\boldsymbol{x_1} - \boldsymbol{x}_2)$ and similar for $\boldsymbol{v}_3$ and $\boldsymbol{v}_4$ using $\boldsymbol{x}_3$ and $\boldsymbol{x}_4$, where C is an arbitrary constant chosen to emphasise the structure of the

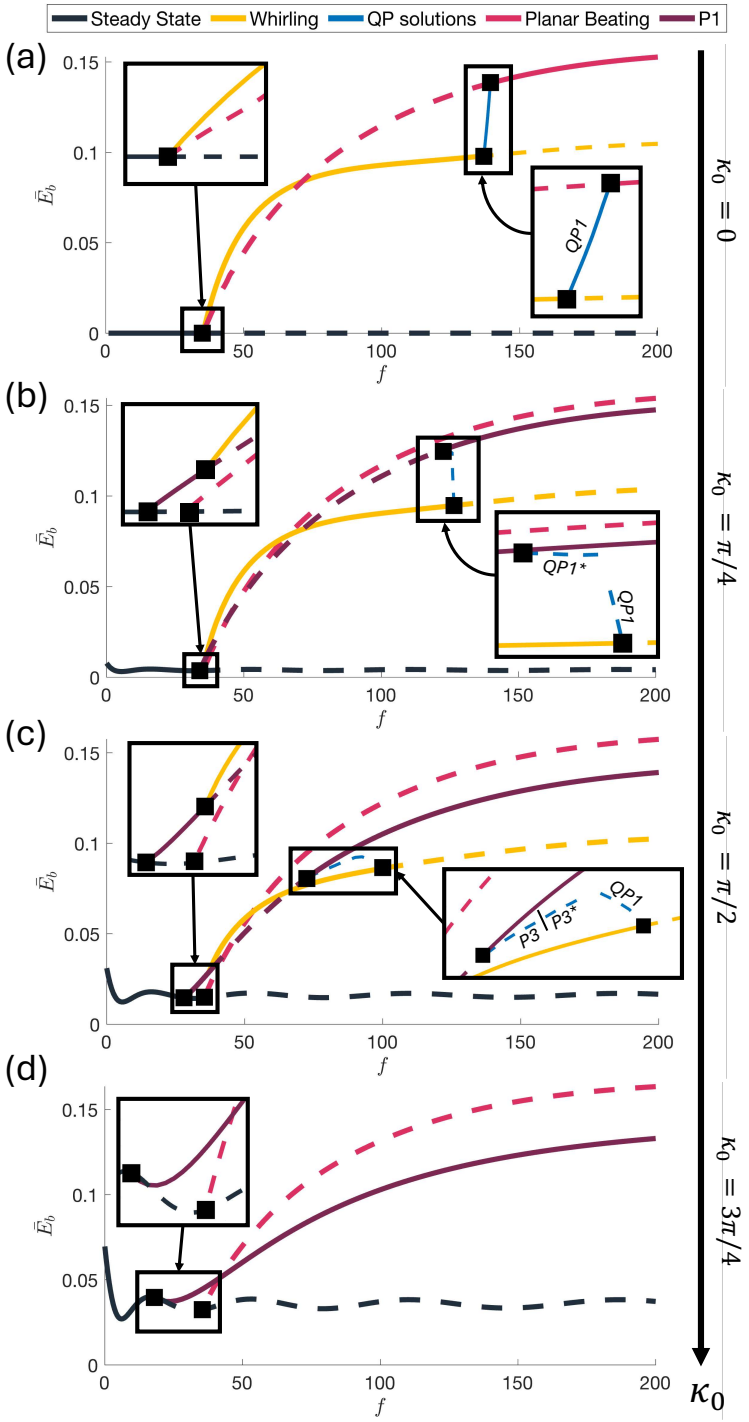


Figure 7.4: Bifurcation diagrams showing the different solutions as we vary f, and the corresponding mean bending energy of these solutions, $\bar{E}_b$, for an isotropic filament ($\beta = 1$) with intrinsic curvature for (a) $\kappa_0 = 0$, (b) $\kappa_0 = \pi/4$, (c) $\kappa_0 = \pi/2$ and (d) $\kappa_0 = 3\pi/4$. Full (dashed) lines refer to stable (unstable) solutions. Square markers are used to indicate the location of bifurcations as we vary f. Inset, left, shows the bifurcation diagram near the initial buckling event and inset, right, shows the bifurcation diagram around the bistable regions, if these are present. As κ_0 increases we see the emergence of the P1 solution branch (b-d), from which whirling bifurcates (b and c, left inset). We observe bistability between P1 and whirling (b and c, right inset) and, for the largest values of κ_0, whirling vanishes completely (d).

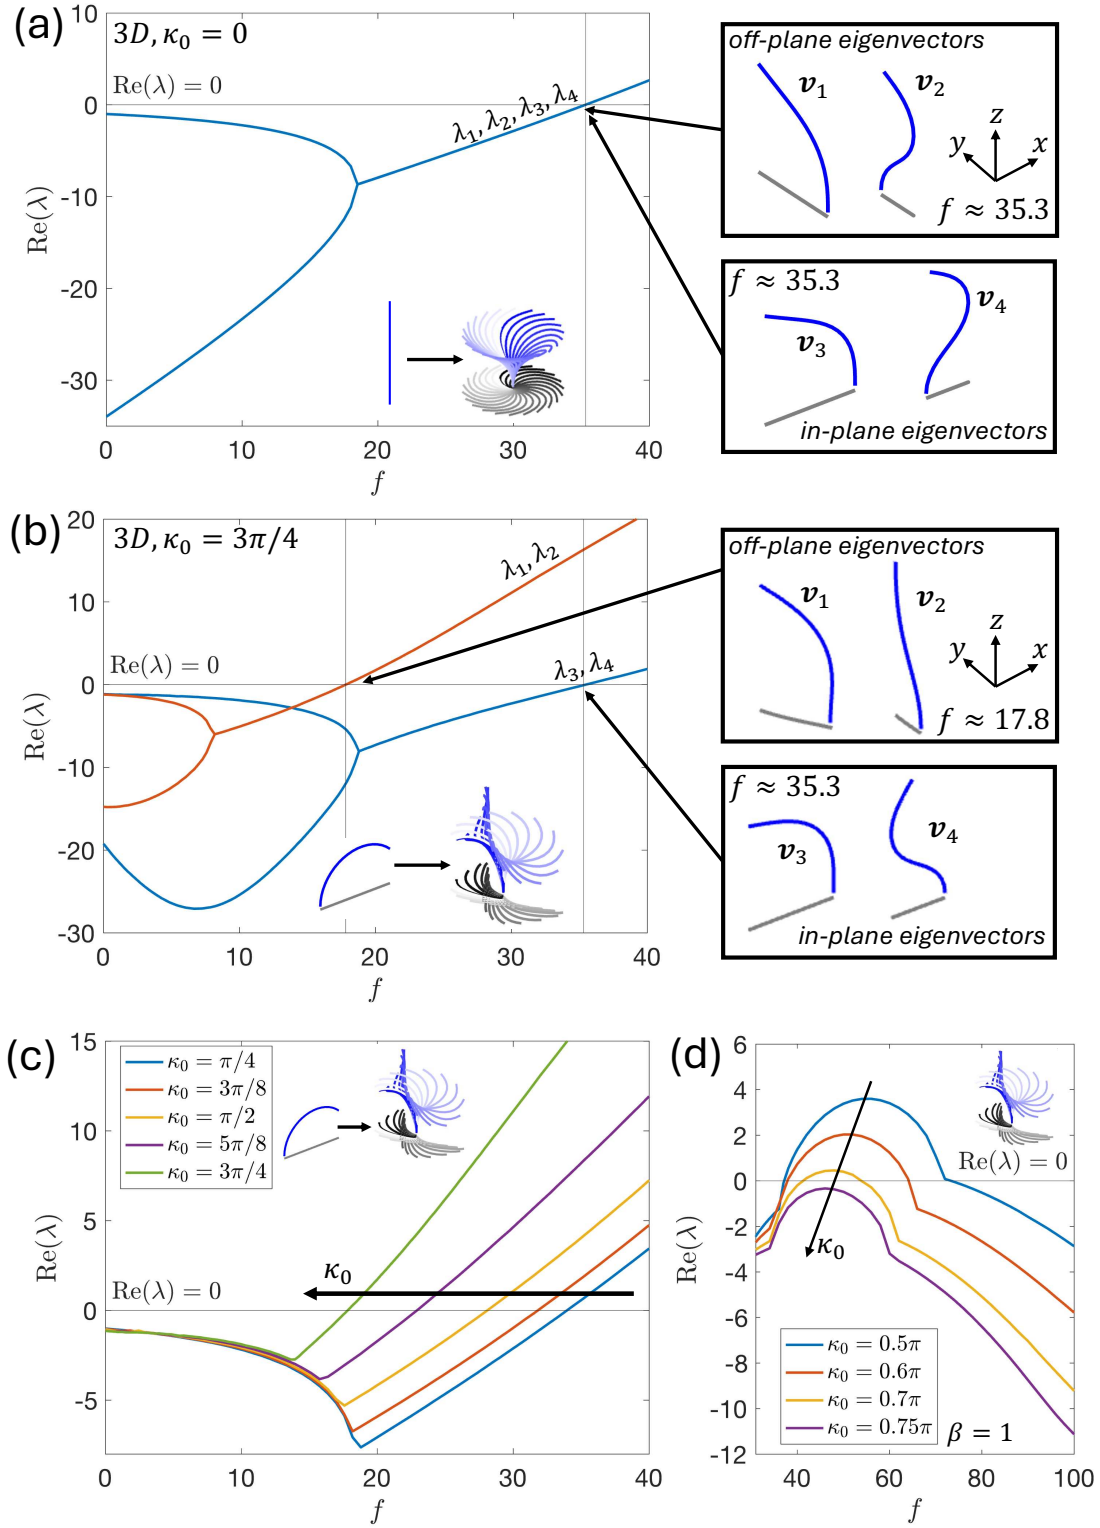


Figure 7.5: The real part of the four largest eigenvalues, λ_i for $i = 1, ..., 4$, associated to the steady state for $\beta = 1$, (a) $\kappa_0 = 0$, and (b) $\kappa_0 = 3\pi/4$. Analysing the corresponding eigenvectors at each bifurcation, obtained as described in the main text and displayed on the right, allows us to identify the branches corresponding to off-plane and in-plane beating. (c) The real part of the largest eigenvalue associated to the steady state for various values of κ_0. We see that the eigenvalue becomes unstable for smaller f as we increase κ_0. (d) The real part of the largest eigenvalue associated to the P1 state for various values of κ_0. For $\kappa_0 > \kappa_0^*$, we see the eigenvalue remains stable.

Event	κ_0		
	$\pi/4$	$\pi/2$	$3\pi/4$
Steady state $\rightarrow$ P1	$f \approx 33.9$	$f \approx 28.1$	$f \approx 17.8$
P1 $\rightarrow$ Whirling	$f \approx 36.3$	$f \approx 37.5$	N/A
Whirling $\rightarrow$ QP	$f \approx 126.4$	$f \approx 100.3$	N/A
QP $\rightarrow$ P1	$f \approx 122.7$	$f \approx 73.3$	N/A

Table 7.1: The critical forcing at which each bifurcation occurs in Figure 7.4 for three values of intrinsic curvature, κ_0.

eigenmode. As the magnitude of the state variable relates to angles of rotation of the local basis (see Section 4.3.1), taking larger values of C results in more exaggerated filament deformation. We show these eigenvectors in Figure 7.5 for $C = 5$. We observe that the first complex conjugate pair is associated to beating in the $(x, z)-$plane, and the second is associated to beating in the perpendicular plane, as shown by the eigenvectors in Figure 7.5a (right). It is these modes that combine to give rise to the stable whirling solution observed after the bifurcation, as well as the unstable planar beating branch. For larger f, whirling becomes unstable and gives rise to QP1 through a supercritical Hopf bifurcation. The QP1 branch exists for a small region of f, and vanishes through a pitchfork bifurcation when planar beating stabilises for larger f (see Figure 7.4a, inset).

Asymmetric planar solutions unstable

Regardless of κ_0, the planar solution exists and can be tracked using JFNK. In all cases, it can be tracked back to $f \approx f^*$. As discussed above, this planar solution is stable for a certain region of f when $\kappa_0 = 0$. However for $\kappa_0 > 0$, this branch remains unstable for all values of the forcing (see Figure 7.4), matching observations from initial value problems.

Initial buckling event

As we increase κ_0, we see a change in the initial buckling event; as shown in Figure 7.4a-d left inset, the initial buckling happens at smaller values of f as we increase κ_0. The critical forcing values for different values of κ_0 are shown in Table 7.1. We also now observe two solution

branches bifurcating from the steady state at different values of the forcing, and note that whirling no longer bifurcates from the steady state for $\kappa_0 > 0$. Instead, we find the P1 branch bifurcates from the steady state initially at a critical forcing $f < f^*$, followed by the unstable planar beating branch at $f = f^*$. To explore these changes in the bifurcation diagrams, we analyse the eigenmodes associated to the steady state.

We plot the real parts of the four dominant eigenmodes of the steady state for $\kappa_0 = 3\pi/4$ in Figure 7.5b, including visualisations of the eigenvectors obtained as discussed previously. We see that the two modes corresponding to buckling in either plane separate; the mode corresponding to in-plane, asymmetric beating still becomes unstable at $f = f^*$, whereas the second mode, corresponding to off-plane beating, becomes unstable earlier. We note that the off-plane eigenmodes do not both lie in the $(y, z)-$plane. This is because the intrinsic curvature breaks the symmetry about the (y, z)-plane, and so we expect one of these eigenmodes to reflect this broken symmetry. This can be seen in the off-plane eigenvectors (see Figure 7.5b, top right). Hence for $\kappa_0 > 0$, at the initial bifurcation, it is only the off-plane modes that are becoming unstable which lead to P1 through a Hopf bifurcation. In Figure 7.5c, we plot the real part of the dominant eigenmode for various values of κ_0. We see that the off-plane modes become unstable earlier as we increase the intrinsic curvature, meaning the steady state buckles at smaller forcing values, as we observed earlier through initial value problems (see Figure 7.3a).

The off-plane eigenmodes destabilising for lower f as we increase κ_0 appears to be a result of the choice of frame in which the curvature is defined to act. We can choose the direction in which the curvature exists, i.e. in the material frame, so that the moment is given by Eq (7.2), or in the lab frame, so that the moment is defined by

$$\boldsymbol{M}(s,t) = K_B \left(\hat{\boldsymbol{t}} \cdot \frac{\partial \hat{\boldsymbol{\nu}}}{\partial s} \right) \hat{\boldsymbol{\mu}} + \beta K_B \left(\hat{\boldsymbol{\mu}} \cdot \frac{\partial \hat{\boldsymbol{t}}}{\partial s} \right) \hat{\boldsymbol{\nu}} + K_T \left(\hat{\boldsymbol{\nu}} \cdot \frac{\partial \hat{\boldsymbol{\mu}}}{\partial s} \right) \hat{\boldsymbol{t}} - K_B L \kappa_0 \hat{\boldsymbol{y}}. \tag{7.4}$$

We show the stability analysis for both cases for $\kappa_0 = 3\pi/4$ in Figure 7.6a. Here, we see that in both cases, the in-plane buckling occurs at $f \approx 35.5$. When the intrinsic curvature is incorporated in the material frame, the off-plane buckling occurs much earlier at $f \approx 17.8$, contrasting $f \approx 36.7$ for the lab frame case. This suggests that the earlier off-plane buckling

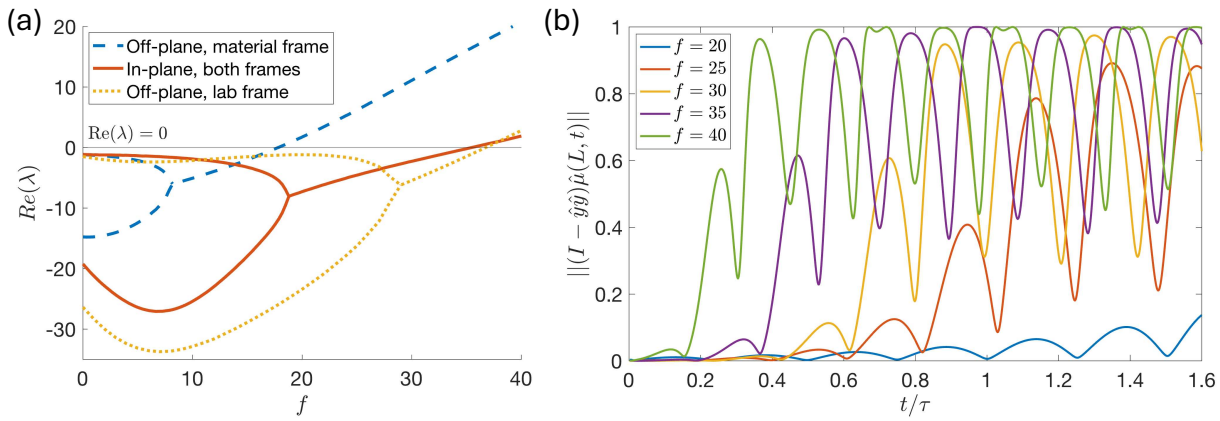


Figure 7.6: (a) The real part of the dominant four eigenvalues from a linear stability analysis of the steady state for $\kappa_0 = 3\pi/4$, $\beta = 1$ associated to the off-plane buckling (blue dashed/yellow dotted) or in-plane buckling (red) when the curvature is introduced in either frame. (b) The off-plane component of the direction of the intrinsic curvature for various values of the forcing. These are obtained from initial value problems, as described in the text. We note that if the curvature is fixed in the lab frame, this would be zero for all time.

arises due to the way in which the intrinsic curvature is tied to the filament frame; small perturbations to the filament generate changes to the filament's local frame, and thus to the direction of intrinsic curvature through $\hat{\boldsymbol{\mu}}$. When the intrinsic curvature is defined in the material frame, a small off-plane change in $\hat{\boldsymbol{\mu}}$ leads to off-plane torque,

$$\tau_\perp = -K_B \kappa_0 ||(\boldsymbol{I} - \hat{\boldsymbol{y}}\hat{\boldsymbol{y}})\hat{\boldsymbol{\mu}}||. \tag{7.5}$$

In essence, the filament now has a now non-planar intrinsic curvature, causing off-plane torques which yield an earlier buckling. The effect of these torques is amplified as we increase κ_0, explaining why the filament buckles earlier as we increase the intrinsic curvature.

In Figure 7.6b, we consider these off-plane effects by calculating the in-plane contributions of $\hat{\boldsymbol{\mu}}$, and hence of the intrinsic curvature. This data is taken from initial value problems in the case where the intrinsic curvature is defined in the material frame, i.e. using Eq (7.2). The initial filament configuration is given by the steady state, as found by JFNK, with an additional small perturbation to the Lie algebra of $\sim O(10^{-3})$. We let this simulation run until periodic oscillations are achieved. As predicted by the stability analysis, the perturbations grow faster as f increases, leading to quicker saturation into the periodic beating we call P1. We can contrast this with the lab-frame intrinsic curvature, where the internal moment is defined by Eq (7.4),

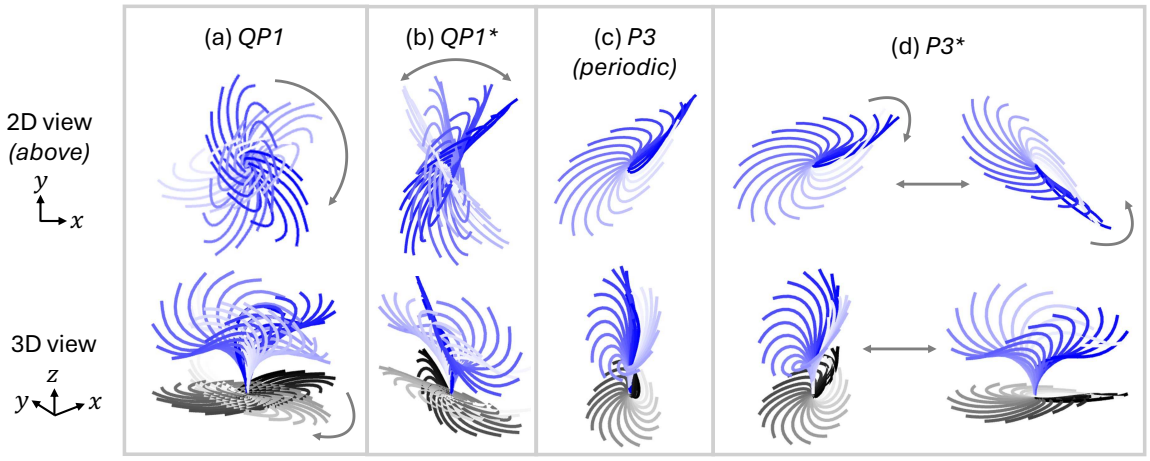


Figure 7.7: Plots of filament dynamics for solutions obtained via bisection from two views: a 2D view from above (top) and a 3D view (bottom). The passage of time is indicated by an increase in line boldness. Plots for (a) QP1, $(\kappa_0, \beta, f) = (0, 1, 138)$, (b) QP1*, $(\kappa_0, \beta, f) = (\pi/4, 1, 123)$, (c) P3, $(\kappa_0, \beta, f) = (\pi/2, 1, 80)$, (d) P3*, $(\kappa_0, \beta, f) = (\pi/2, 1, 90)$

in which the off-plane contribution of the torque is always $\tau_\perp = -K_B\kappa_0||(\boldsymbol{I} - \hat{\boldsymbol{y}}\hat{\boldsymbol{y}})\hat{\boldsymbol{y}}|| \equiv 0$.

Immediately following the initial buckling bifurcation for $\kappa_0 > 0$, we see from Figures 7.4b-d that the P1 solution is stable. For moderate κ_0, the P1 solution becomes unstable for larger f, giving way to the stable whirling solution (see Figures 7.4b,c, left inset). This can again be linked to the initial buckling event. As the whirling solution is known to arise from a combination of both the in-plane and off-plane buckling eigenmodes (see Section 5.5), it follows that the whirling solution is always born after the second buckling mode becomes unstable, branching from the P1 solution branch.

Bistability

Recall that for $\kappa_0 = 0$, we observe a stable branch (QP1) connecting the whirling and beating solutions. As we increase κ_0, we see that the stable branch connecting P1 and whirling exists for smaller ranges of f, until it is no longer stable. In this case, i.e. for moderate κ_0 values, we observe bistability between the P1 and whirling solutions (see Figures 7.4b,c). As such, we use bisection to unearth the unstable solution branches which exist in the region.

For $\kappa_0 = \pi/4$, the bistable region is between $122.7 \lesssim f \lesssim 126.4$. Bisection in this region unveils two distinct solutions (see Figure 7.4b, right inset); QP1, as in the isotropic case, and a new

solution, which we call QP1*. We show these dynamics in Figure 7.7a,b. In this new solution, the filament has an almost-elliptical beat like QP1, although now the overall beat rotates from side-to-side, rather than unidirectionally like QP1.

For $\kappa_0 = \pi/2$, we see a larger region of bistability between whirling and P1, for $73.3 \lesssim f \lesssim 100.3$. Bisection in this region reveals three solution branches (see Figure 7.4c, right inset). The first branch is a periodic solution close to the P1 bifurcation which we call P3; here the tip traces a tear-drop shape in the $(x, y)-$plane, as shown in Figure 7.7c. For larger f, bisection instead reveals the second branch, separated by the black line in Figure 7.4c inset; a solution which we call P3*. Here, the filament beats like P3, but now the tear-drop shape traced by the filament tip moves with time. The beat repeatedly changes direction, each time beginning to align the major axis of the shape traced by the filament tip with the $\hat{\boldsymbol{x}}-$axis before changing direction again, as illustrated in Figure 7.7d. The final branch, close to the whirling bifurcation, is QP1 as in the isotropic case, again shown in Figure 7.7a.

In both cases, the sharp turns in the bifurcation diagrams indicate further bifurcations may be occurring in this region, but due to the quasiperiodic nature of these solutions, we are unable to identify these additional bifurcations through JFNK or bisection alone.

Loss of whirling solution

In the previous sections we identify three key bifurcations; P1 becoming unstable, P1 subsequently restabilising, and whirling becoming unstable. As we increase κ_0, these three bifurcations occur within a smaller range of f, until the intrinsic curvature reaches a critical value, and the bifurcations coalesce. For larger κ_0, for instance $\kappa_0 = 3\pi/4$ as shown in Figure 7.4d, the whirling solution vanishes completely, and we only observe the stable P1 solution branch after the initial buckling bifurcation.

We can explore the disappearance of the whirling solution by performing a Floquet analysis on the P1 solution for various values of κ_0. The real part of the dominant eigenvalue for each κ_0 is shown in Figure 7.5d. We see that the two bifurcations, corresponding to P1 becoming

unstable and then restabilising, get closer together as we increase κ_0, until they merge at a critical value, $\kappa_0 = \kappa_0^*$. At this value, the whirling solution collapses onto the P1 branch, also terminating QP1 and QP1* in the process. Beyond this κ_0, only P1 is stable in this region of f. We can identify the loss of the whirling solution by again revisiting the initial buckling event. For $\kappa_0 > \kappa_0^*$, it appears that the growth rate associated to the off-plane modes is so large compared to the growth rate of the in-plane modes that the whirling solution never exists. Instead, we only observe P1, matching observations from initial value problems.

7.4 Aniostropic filaments with no intrinsic curvature

7.4.1 Initial value problems

In the previous section we investigated the effect of intrinsic curvature defined in the plane defined by $\hat{\boldsymbol{\nu}}$ and $\hat{\boldsymbol{t}}$. Now, we set $\kappa_0 = 0$ and focus on the effect of introducing anisotropy through the bending modulus by introducing the parameter β, which describes the ratio of the bending moduli about the $\hat{\boldsymbol{\nu}}-$ and $\hat{\boldsymbol{\mu}}-$axes. Throughout this paper we focus on $\beta > 1$, but the $\beta < 1$ case follows immediately by considering a rotation of the axes, and redefining the non-dimensional follower force as $\hat{f} = \beta f$ (c.f. Eq (4.5)). In Figure 7.8a, we show the results of initial value problems for $\beta \in [1, 1.5]$ and $f \in [0, 300]$. Regardless of β and f, the steady state is the vertically upright filament. For $\beta = 1$, this buckles into whirling, while for $\beta > 1$ the steady state buckles into planar beating in the $(x, z)-$plane. For moderate values of β, beating becomes unstable as we increase the forcing, with whirling appearing in its place. Close to the bifurcation in which planar beating appears to become unstable, whirling is very similar to in-plane beating, with a small off-plane component. This causes the tip to trace an ellipse with a major axis that aligns with the $\hat{\boldsymbol{x}}-$axis, as shown in Figure 7.8b, middle. The large amplitude in the $\hat{\boldsymbol{x}}-$direction close to the bifurcation suggests that we can expect the whirling solution to branch from the planar beating solution branch, instead of the steady state. Further from the bifurcation, however, the off-plane component grows and whirling resembles the isotropic case more closely (see Figure 7.8b, right). For larger forcing values, whirling becomes unstable and

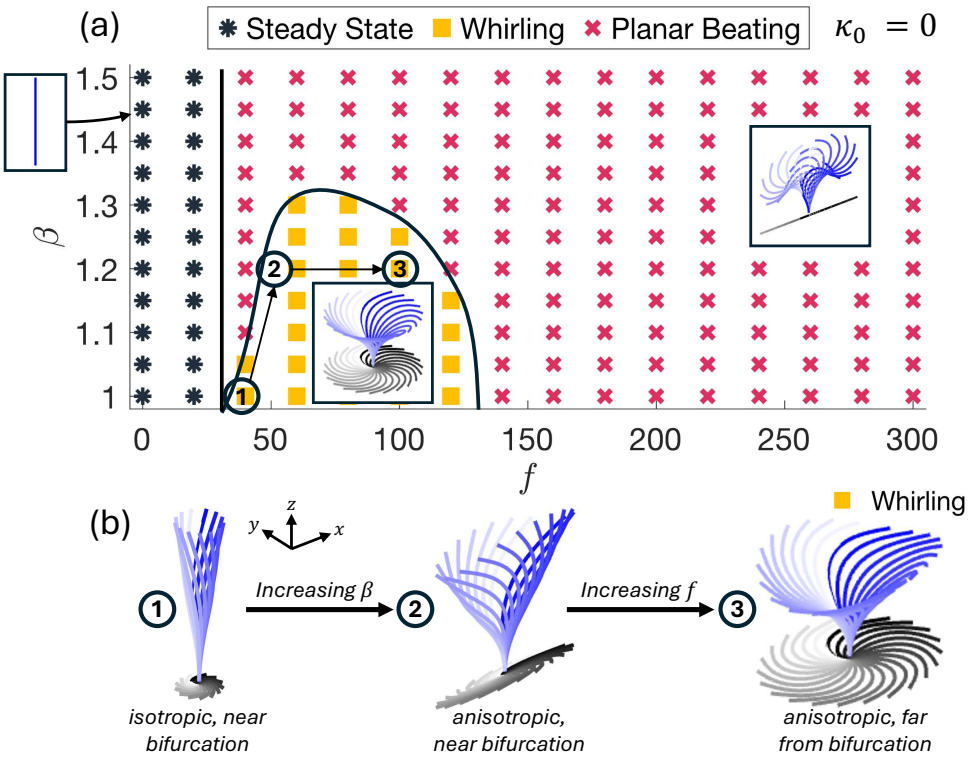


Figure 7.8: (a) Results of numerical simulations of an anisotropic filament with no intrinsic curvature ($\kappa_0 = 0, \beta \geq 1$). The black lines are drawn to indicate the approximate boundaries of the solution regions. (b) Snapshots of filaments over one period undergoing whirling dynamics in the isotropic case, close to the bifurcation ($\beta = 1, f = 36$), an anisotropic case near the bifurcation ($\beta = 1.2, f = 46$) and an anisotropic case further from the bifurcation ($\beta = 1.2, f = 100$). The passage of time is indicated by the increasing boldness of the lines. The three dynamics are indicated on the phase diagram in (a) using numbered circles.

planar beating in the $(x, z)-$plane reappears. As we increase β, the region in which whirling is stable appears to decrease, until $\beta = 1.35$, where we no longer observe whirling through initial value problems. For these β values, we only observe planar beating in the $(x, z)-$plane after buckling.

7.4.2 Bifurcation and stability analyses

Using JFNK to track the steady and periodic solutions and then performing a stability analysis on the subsequent branches allows us to generate the bifurcation diagrams shown in Figure 7.9. We provide the corresponding bifurcation values in Table 7.2, and outline the key features of the bifurcation diagrams below.

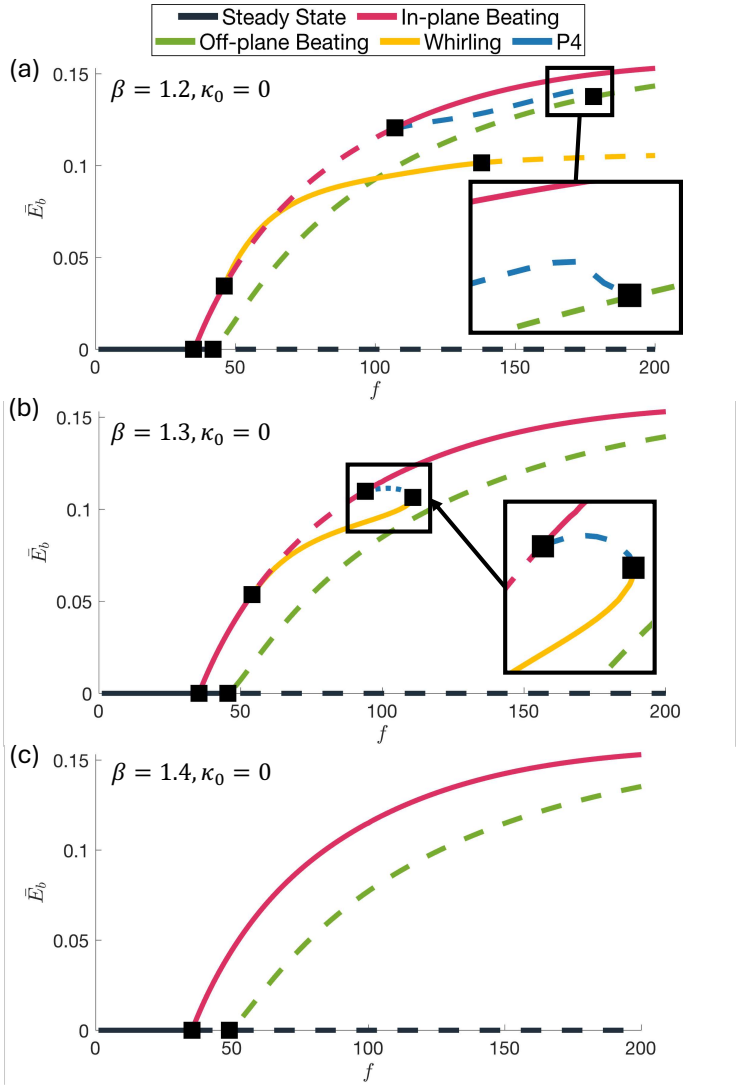


Figure 7.9: (a) Bifurcation diagrams showing the different solutions as we vary f, and the corresponding mean bending energy of these solutions, $\bar{E}_b$, for an anisotropic filament with no intrinsic curvature ($\kappa_0 = 0$) for (a) $\beta = 1.2$, (b) $\beta = 1.3$ and (c) $\beta = 1.4$. Full (dashed) lines refer to stable (unstable) solutions. Square markers are used to indicate the location of bifurcations as we vary f. We see the separation of the in-plane and off-plane solution branches as β increases, and observe bistability between in-plane beating and whirling (inset). For larger values of β, i.e. (c), the whirling solution branch vanishes completely.

Focusing first on the initial bifurcation, we see that the critical forcing value at which the steady state buckles remains at $f = f^*$ regardless of β. At this point, the planar beating branch in the $(x, z)-$plane, which we will henceforth call 'in-plane' beating (opposed to 'off-plane' beating, in the $(y, z)-$plane), is born through a Hopf bifurcation. We can also identify the unstable off-plane beating branch for each β, which also appears to connect to the trivial steady state, bifurcating at $f = \beta f^*$. These solutions are present regardless of β, as shown in Figure 7.9. These branches are linked to the initial buckling event; introducing $\beta \neq 1$ has the effect of

Event	β		
	1.2	1.3	1.4
Steady state $\rightarrow$ in-plane beating	$f \approx f^*$	$f \approx f^*$	$f \approx f^*$
In-plane beating $\rightarrow$ Whirling	$f \approx 46.0$	$f \approx 54.0$	N/A
Whirling unstable	$f \approx 138.4$	$f \approx 110.8$	N/A
In-plane beating stable	$f \approx 106.2$	$f \approx 93.9$	N/A

Table 7.2: The critical forcing at which each bifurcation occurs in Figure 7.9 for three values of the bending moduli ratio, β.

separating the critical values at which the beating in either plane becomes unstable, similar to the effect of incorporating $\kappa_0 > 0$. This is a consequence of the non-dimensionalisation of the follower force, as given in Eq (4.5). Accordingly, we expect eigenmodes to become unstable at $f = f^*$ (corresponding to buckling in the $(x, z)-$plane) and $f = \beta f^*$ (corresponding to buckling in the $(y, z)-$plane). Thus, the magnitude of β dictates which plane we initially buckle into, and furthermore when successive buckling modes become activated. To verify the effect of β on the initial buckling bifurcation, we plot the dominant eigenvalues for $\beta = 1.5$ in Figure 7.10a. As expected, we find that the first bifurcation occurs at $f = f^*$, with corresponding planar eigenmodes lying in the $(x, z)-$plane. Therefore, the filament is expected to undergo planar deformations in this plane, as is confirmed by the numerical simulations. The second mode, corresponding to buckling in the perpendicular plane, is found to occur at $f = \beta f^* \approx 53.0$.

For moderate β (see Figures 7.9a, b) we find that, for a larger value of the forcing, in-plane beating becomes unstable and the whirling solution is born. Between the bifurcations in which in-plane beating becomes stable again and whirling becomes unstable, there is a region of bistability. Bisection in this region yields a single solution branch, which we call P4. Here, the filament undergoes elliptical whirling oscillations, where the major axis of the ellipse has components in both the $\hat{\boldsymbol{x}}$ and $\hat{\boldsymbol{y}}$ planes. We plot P4 over one period in Figure 7.10b. As P4 is periodic, we can track the behaviour using JFNK. In doing so, we find that P4 connects to the off-plane beating branch, rather than the whirling branch, as shown in Figure 7.9a (inset). Floquet analysis of the off-plane beating branch confirms this, showing that the second largest eigenvalue, λ_2, becomes stable when P4 collapses onto the branch. It is likely that other unstable solutions exist in this region, but we are unable to find these using bisection or JFNK alone. By performing a Floquet analysis on the off-plane beating branch for various values

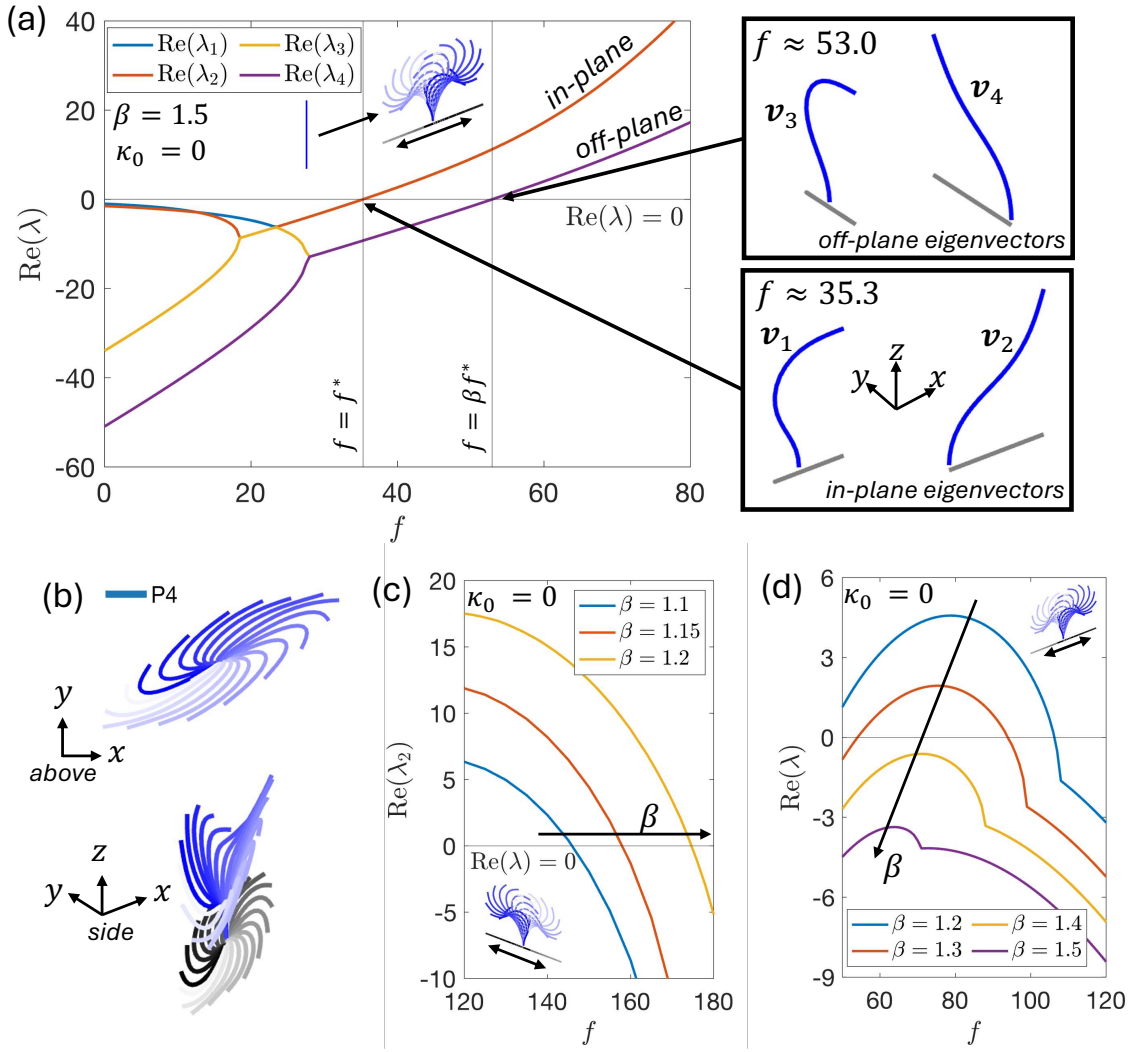


Figure 7.10: (a) The real part of the four largest eigenvalues associated to the steady state for $\beta = 1.5$. Analysing the corresponding eigenvectors, obtained as described in the text and visualised on the right, allows us to identify the branches corresponding to in-plane beating (bottom right), and off-plane beating (top right). (b) Plots of P4 dynamics for $(\kappa_0, \beta, f) = (0, 1.2, 120)$ from above and a side view. (c) The real part of the second largest eigenvalue associated to off-plane beating for various values of β. We associate this eigenvalue becoming stable with P4 collapsing onto the off-plane beating branch. (d) The real part of the largest eigenvalue associated to the in-plane beating state for various values of β. For $\beta > \beta^*$, we see the eigenvalue remains stable.

of β, as shown in Figure 7.10c, we see that λ_2 becomes stable at higher f as we increase β, indicating that P4 joins the off-plane beating branch at higher values of f as we increase β. This trend continues until a critical value of β, beyond which P4 now connects to the whirling branch, after which both branches terminate. Consequently, for these values of β, the whirling branch only exists when it is stable (see Figure 7.9b).

The three bifurcation points corresponding to in-plane beating becoming unstable, becoming stable again, and the whirling branch disappearing all occur within smaller ranges of f as we increase β, until the three bifurcations collide and we see only one stable branch - in-plane beating, as shown in Figure 7.9c for $\beta = 1.4$. Indeed, a stability analysis of the in-plane beating solution, shown in Figure 7.10d, confirms that the range of f for which beating is unstable decreases with increasing β, up to a critical value of $\beta = \beta^*$ where the two pitchfork bifurcations collide, which likely also coalesces with the P4-whirling bifurcation. Beyond this, beating is stable for all values of f in our tested range and, as in the intrinsic curvature case, whirling ceases to exist entirely. Both the Floquet analysis and initial value problems suggest that the whirling branch collapses onto the beating branch at $\beta = \beta^*$, resulting in the loss of P4 and whirling as distinct solutions. The disappearance of whirling can again be explained through the large separation in critical force values between in-plane and off-plane beating solutions bifurcating from the trivial state.

7.5 Using bending modulus anisotropy to stabilise asymmetric beating

Following the previous section, we see that increasing $\beta > 1$ has the effect of stabilising planar solutions. Thus we now combine the effects of $\kappa_0 \neq 0$ and $\beta \neq 1$, with the goal of recovering the planar asymmetric beating observed in 2D simulations of filaments with intrinsic curvature, which were found to be unstable to 3D perturbations in Section 7.3 for $\beta = 1$. Both the P1 and whirling states must be eradicated for a collapse onto planar dynamics. Following the results of the previous sections, we anticipate that the stability of the P1 branch and existence of whirling solution are related to the eigenmodes of the initial buckling event. As increasing κ_0 appears to cause an earlier buckling of the off-plane eigenmodes, while increasing β delays the onset of this instability, we expect that increasing the intrinsic curvature will require a larger amount of anisotropy to stabilise the planar modes. As such, we focus on $\kappa_0 = 3\pi/4$, anticipating this to provide an upper bound for the value of β which causes a collapse to 2D dynamics, and

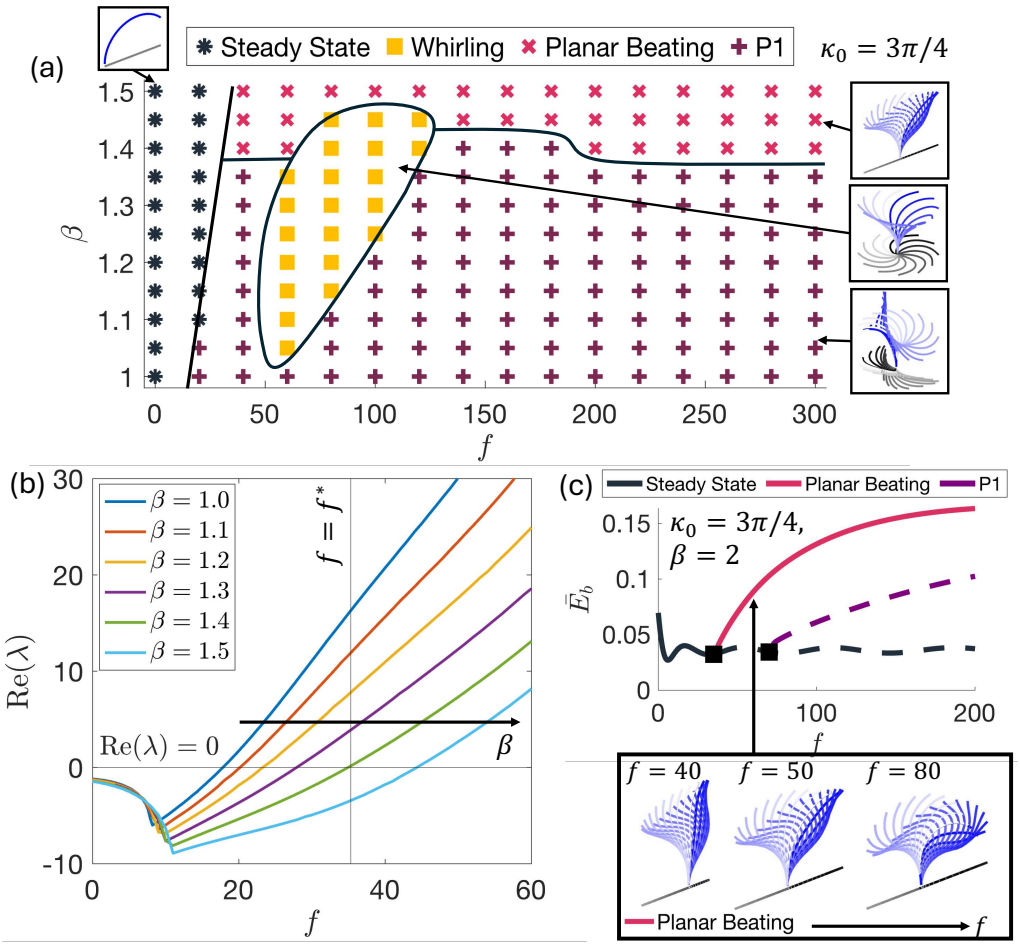


Figure 7.11: (a) Results of numerical simulations of an anisotropic filament with intrinsic curvature ($\kappa_0 = 3\pi/4, \beta \geq 1$). The black lines are drawn to indicate the approximate boundaries of the solution regions. (b) The eigenvalue associated to off-plane beating for several values of β. For $\beta \lesssim 1.4$, this eigenmode becomes unstable before the in-plane eigenmode (i.e. before $f = f^*$). For $\beta \gtrsim 1.4$, the in-plane eigenmode becomes unstable first. (c) The bifurcation diagram showing the different solutions as we vary f, and the corresponding mean bending energy of these solutions, $\bar{E}_b$, for $\kappa_0 = 3\pi/4$, $\beta = 2$. We see that only asymmetric planar beating (bottom) is stable.

increase β from $\beta = 1$ to $\beta = 1.5$. The results from initial value problems over this parameter range are shown in Figure 7.11a.

As discussed in previous sections, κ_0 and β have opposite effects on the initial buckling modes; increasing the intrinsic curvature (i.e. increasing κ_0) causes the off-plane modes to buckle earlier, whereas increasing the anisotropy through the bending moduli (i.e. increasing β) causes later buckling of these modes. In both cases, in-plane buckling happens at similar critical forcing values, regardless of κ_0 or β. Therefore, fixing κ_0 and increasing β shifts the off-plane beating modes to buckle at higher values of the forcing. This is shown in Figure 7.11b, which displays

the eigenvalues corresponding to off-plane buckling for various values of β. As we can see, at $\beta \approx 1.4$ the off-plane buckling occurs at $f \approx f^*$, the critical forcing at which in-plane buckling occurs. Beyond this β, the off-plane buckling occurs for $f > f^*$, indicating that the initial buckling of the filament will be to in-plane beating, and will occur at $f \approx f^*$. Accordingly, as we increase β from 1, the initial value problems in Figure 7.11a show that the initial buckling into P1 happens at a critical value of the forcing which increases with β until it reaches $f \approx f^*$ at $\beta \approx 1.4$, beyond which the filament buckles into in-plane asymmetric beating at $f \approx f^*$.

Furthermore, as increasing β increases the critical forcing value at which off-plane beating bifurcates from the steady state, increasing the bending modulus anisotropy results in the birth of whirling again (as illustrated by the initial value problems in Figure 7.11a). For these intermediate values of β, Floquet analysis indicates bistability between P1, whirling and in-plane beating at several points, likely linked to the existence of distinct unstable quasiperiodic and periodic solutions joining the three branches. However as β is increased further, and the off-plane buckling modes become unstable for larger f, and we reach a scenario where the whirling branch dies again above a critical value of β. This occurs through the same mechanism as the $\kappa_0 = 0$ and $\beta = 1$ cases.

As expected, increasing the bending modulus anisotropy further leads to a complete collapse onto planar dynamics. While initial value problems from the upright-state indicate that this collapse has occurred for $\beta = 1.5$ (see Figure 7.11a), Floquet analysis shows that there are still regions where P1 is stable for this choice of parameters, alongside in-plane beating. By $\beta = 2$, however only in-plane beating is stable, showcasing the same asymmetric beating dynamics as the planar case for $\kappa_0 = 3\pi/4$, with particular asymmetry close to the bifurcation. We show the bifurcation diagram for $\beta = 2$, as well as the planar beating dynamics, in Figure 7.11c.

7.6 Discussion

In this chapter, we investigated the effects of asymmetry and structural anisotropy in a follower-force driven filament, through the introduction of both intrinsic curvature and bending modulus

anisotropy. In both cases, we observed that the initial buckling changed from the double Hopf bifurcation observed in the symmetric case, to a single Hopf, suggesting that the double Hopf is a special case of these systems and is related to symmetry. When varying only the intrinsic curvature, in 2D simulations we obtain the same asymmetric beating reported in the literature [Wang et al., 2023]. However in 3D, we see that this beating is inherently unstable, with the filament instead buckling into a fully 3D beating occurring primarily perpendicular to the plane the filament would deform in the absence of the follower force (P1). Increasing the forcing, we observe a tilted whirling and discuss how this changes with both f and κ_0. For moderate curvatures, we observe that P1 and whirling become bistable, making the stable QP1 branch connecting whirling to beating in the zero-curvature case now unstable, and introducing further unstable solutions which we identify through bisection. Similar results hold when varying the bending modulus anisotropy, replacing P1 with planar beating in the plane of preferential bending. In both cases, we observe that as we increase the amount of asymmetry, whirling does not only become unstable, but rather is completely eliminated as a solution. Instead, we observe only P1 for the intrinsic curvature case, or in-plane beating for the anisotropic bending modulus case. Finally, we show how introducing anisotropy through the bending modulus on a filament with intrinsic curvature can re-introduce whirling as a solution for moderate β values, and ultimately stabilise the planar asymmetric beating observed in 2D for isotropic filaments with intrinsic curvature for larger β values. We find that a difference by a factor of 2 in the bending moduli is sufficient for this collapse into planar beating for the highest value of κ_0 tested for the follower force model; a remarkably reasonable ratio.

Changing the intrinsic curvature has the effect of adding asymmetry into the beating pattern, motivating its inclusion in recent numerical studies [Sartori et al., 2016, Chakrabarti and Saintillan, 2019a,b, Chakrabarti et al., 2022, Cass and Bloomfield-Gadêlha, 2023, Wang et al., 2023]. These 2D works use values $\kappa_0 \in [2, 2.71]$ to match the beating pattern to experiments, in particular to wild-*Chlamydomonas* flagella [Sartori et al., 2016, Chakrabarti and Saintillan, 2019b]. However this is not the only potential cause of asymmetry in ciliary beating; some works suggest that the bending modulus in each plane may not be constant, instead utilising a curvature dependent bending modulus [Han and Peskin, 2018, Moreau et al., 2024]. In such

a model, the in-plane bending modulus can take one of two values depending on the sign of the local curvature. Other works suggest that the asymmetry is a result of bias in the rates at which the dynein attach and detach from the microtubules, depending on their placement in the axoneme [Chakrabarti and Saintillan, 2019a,b, Chakrabarti et al., 2022]. In each of these models, numerical studies result in asymmetric dynamics exhibiting distinct recovery and effective strokes. Our results indicate that these asymmetric beating solutions could be unstable in 3D, but that they could be stabilised by incorporating anisotropy to the model, for instance structural anisotropy through the bending modulus.

Incorporating anisotropy through the bending modulus in either direction is motivated by the internal structure of the cilium. Cilia have a preferential bending direction both due to the orientation of the central pair of microtubules [Gibbons, 1981, Marshall and Nonaka, 2006, Gilpin et al., 2020] and the stiffness of the 5-6 bridge [Afzelius, 1959, Gibbons, 1981] which limits sliding between microtubule doublets 5 and 6 in the axoneme. The effect of the bridge is captured in our work by allowing the perpendicular stiffness to differ from the in-plane stiffness, where 'in-plane' is aligned with the preferential beating plane of the cilium. Structural anisotropy was utilised in [Rallabandi et al., 2022], where it was shown that the transition from planar to whirling dynamics is delayed as we increase the anisotropy. Based on our results, we suspect that increasing the anisotropy further in this study would eradicate the 3D whirling dynamics completely, and only planar solutions would remain. It is expected that larger values of β may be required for different forcing mechanisms. For example, in Rallabandi et al. [2022], values up to $\beta = 10$ were considered and 3D dynamics were still observed, while we found values as small as $\beta = 2$ were sufficient to suppress 3D dynamics entirely for the follower force model. In Ishijima and Hiramoto [1994], experimental evidence on sand dollar sperm flagella reports a value of $\beta = 13.8$. This is much higher than the threshold for entirely planar dynamics that we find, and so could explain the planar beating observed by these cells [Ishijima and Hiramoto, 1994]. In Lindemann and Lesich [2016], a wooden axonemal model supplemented with a stronger 5-6 bridge (by using a silicon adhesive between doublets 5 and 6) and a central pair was used to generate an estimate of $\beta = 2.6$ for the axoneme.

Although the follower force model differs from the internal shear-driven sliding mechanism

which generates ciliary beating, the dynamics observed in this paper are reminiscent of those observed biologically. For instance the tilted whirling dynamics of nodal cilia found in the embryo [Nonaka et al., 1998, Smith et al., 2008a, 2011] is similar to the asymmetric whirling observed for moderate κ_0 and β. Respiratory cilia on the other hand, undergo asymmetric near-planar beating with distinct power and recovery strokes [Chilvers and O'Callaghan, 2000]. While we do not observe these two components of the beat with this simple model, for large enough β and moderate κ_0, we do observe asymmetric planar beating. Structurally, these cilia are different; nodal cilia lack the central pair of MTs in the axoneme, which respiratory cilia do have. The existence of a central pair would increase the structural anisotropy, and hence increase β. This is consistent with our findings that filaments with lower β values can access tilted whirling as a state, whereas higher β-filaments can only achieve planar beating.

Throughout this thesis we have utilised the follower force model, due to its simplicity and the fact that it yields cilia-like dynamics through instabilities, which provides an environment to explore the effect of anisotropies on these states. Some other forcing mechanisms more accurately describe the shear-driven activation of cilia, instead considering the forcing with a multi-level approach; considering how microscale attachment and detachment of dynein motors in the axoneme can lead to large scale ciliary beating. In such models, anisotropy can be included through the bending modulus [Rallabandi et al., 2022], or by including stiff springs to represent the 5-6 bridge in the axoneme [Han and Peskin, 2018]. Then asymmetric beating can be achieved through the intrinsic curvature [Sartori et al., 2016, Cass and Bloomfield-Gadêlha, 2023], or through biased feedback mechanisms which control the forcing output of the motors [Chakrabarti and Saintillan, 2019a,b, Chakrabarti et al., 2022]. Exploring how anisotropies through the curvature and stiffness affect dynamics in models with more accurate forcing mechanisms would give more accurate estimates for the κ_0 and β values necessary to recover beats closer to those observed in the literature, and also perhaps give more knowledge of the core ingredients needed to obtain such beats.

The anisotropies studied in this work are assumed to be constant along the filament length, but in reality this may not be the case. Towards the bottom of each cilium, connecting the axoneme to the basal body, is a region called the transition zone, in which several structural changes from

the '9+2' axoneme can be found [Mercey et al., 2024]. Incorporating the structural properties of this region is of interest, as defective proteins in the transition zone are known to cause a range of human diseases, called ciliopathies [Gonçalves and Pelletier, 2017]. Investigating the role of this coupling at the base, either through different boundary conditions at the wall or through spatially-varying anisotropies, will provide an insight to the impact of these localised anisotropies on ciliary dynamics. Another future avenue of interest is considering the collective motion of asymmetrically beating filaments, investigating if hydrodynamic interactions between cilia can lead to synchronisation. We begin to explore such systems in the next chapter.

Chapter 8

Collective Dynamics of Instability-Driven Filaments

8.1 Introduction

Thus far, we have examined the behaviour of individual follower force–driven filaments, generating robust asymmetric planar beating patterns through the introduction of intrinsic curvature and anisotropy in the bending modulus. We now shift our focus to the collective dynamics that emerge when many such filaments interact, for instance representing an array of cilia. Cilia exist collectively in many contexts; numerous microswimmers, including *Paramecium* [Tamm, 1972] and *Volvox carteri* [Brumley et al., 2012], are covered with dense bands or sheets of cilia whose coordinated beating propels the swimmer through flow. On the surfaces of these cells, neighbouring cilia maintain a constant phase difference in their motion, producing a travelling wave of phase along the domain. This coordinated motion is known as a metachronal wave (MCW). MCWs have been shown to be much more efficient at fluid propulsion than synchronised beating states [Osterman and Vilfan, 2011, Ding et al., 2014], allowing the cilia to collectively pump fluid at a much larger scale. MCWs are not confined to cellular locomotion: they also occur within the body, where cilia lining epithelial cells undergo MCWs to perform essential transport functions. In the respiratory system, for example, this beating drives mu-

cociliary clearance — a process in which mucus is transported to protect the lungs and clean the airways [Sleigh et al., 1988, Smith et al., 2008b, Lai et al., 2009]. In the brain, cilia collectively beat to regulate the flow of cerebrospinal fluid through the ventricles [Worthington and Cathcart, 1966, Faubel et al., 2016], facilitating the delivery of essential nutrients [Sawamoto et al., 2006, Guldbrandsen et al., 2014]. MCWs have also been observed in artificial carpets of magnetically activated cilia [Gu et al., 2020, Milana et al., 2020].

Although MCWs have been observed in a variety of biological contexts, the mechanism causing their development is not fully understood. It is commonly thought that MCWs are a consequence of the hydrodynamic coupling between neighbouring cilia [Gueron et al., 1997, Niedermayer et al., 2008, Brumley et al., 2015]. However this may not be the only driving force; *Chlamydomonas reinhardtii* are known to control their two flagella through basal coupling [Quaranta et al., 2015, Wan and Goldstein, 2016], and it is unknown whether this, or other factors, have a role in MCW generation more generally.

Due to the prevalence of MCWs throughout these biological systems, there has been much interest in capturing these dynamics numerically. Numerical studies have been developed which span a range of detail; at the coarsest scale, the squirmer model [Lighthill, 1952, Blake, 1971b, Pedley et al., 2016, Ishikawa, 2024] aims to capture the essence of these collectively beating cilia. The squirmer model considers a single ciliated cell, and imposes a velocity on the cell surface which represents the flow generated by the beating cilia across the surface. Analytic solutions for the corresponding fluid velocity can be found [Lauga, 2020]. Therefore this simple and analytically tractable model provides a pathway to investigate the overall dynamics of a ciliated swimmer. However to understand the emerging dynamics and synchronisation of collections of cilia, the individual filament phase must be allowed to evolve, and so more detail is required.

The next step up in ciliary modelling comes in the form of rotor models [Uchida and Golestanian, 2010, Ghorbani and Najafi, 2017, Meng et al., 2021, Kanale et al., 2022, Bennett, 2025], in which an individual cilium is represented by a single particle moving on a prescribed path. Such models have been used to show that hydrodynamic interactions are sufficient to recover

MCWs. While these types of modelling techniques allow us more insight into hydrodynamic coordination in a broader sense, by permitting MCWs to form in response to the other cilia rather than imposing them directly, they neglect the individual dynamics of each cilium and, in turn, do not allow for changes in the beat pattern of an individual cilium. To shed light on what occurs at the individual level, and how collective dynamics can vary the beat pattern to achieve the final state, we must resolve along the filament length.

An issue with modelling the cilium as an elastic filament is that it is not immediately clear how best to drive the filament to generate cilia-like dynamics on an individual level. Some studies circumvent this by utilising data from real life cilia, obtained using high-speed imaging techniques, for instance on cilia in the lung [Fulford and Blake, 1986, Ding et al., 2014], on the surface of *Paramecium* [Solovev and Friedrich, 2022] or on the surface of coral larvae [Poon et al., 2023]. These models are used to either entirely prescribe MCWs, and then to calculate the corresponding flow fields to explore mixing and fluid transport [Ding et al., 2014], or are constructed in such a way that the filament phases are allowed to evolve [Solovev and Friedrich, 2022]. In the latter case, the filament configuration will be defined by a phase which is allowed to vary in response to external forces (such as the hydrodynamic effects of neighbouring cilia). This approach has been used to show emerging MCWs [Solovev and Friedrich, 2022]. While these models allow us to explore the onset of MCWs and the resulting flow fields, confining the configuration of the filaments to a pre-defined set of states is restrictive, and neglects the scenario in which the individual dynamics of the cilium may change in response to its neighbours, which may occur to facilitate more effective fluid propulsion.

Therefore, an interesting focus in this field centres on collective dynamics in active filament models, where filaments are free to deform in response to both internal and external forces. The symmetric follower force model considered in Chapter 5 has been briefly used to investigate coordination in arrays of filaments in unbounded periodic domains [Su and Keaveny, 2024], as well as planar beating filaments in open domains with a no-slip planar wall, and on spherical surfaces [Westwood and Keaveny, 2021]. However MCWs were not observed in either case and, regardless, the reciprocal motion of the filament in these cases leads to no net fluid flow. In [De Canio, 2019], collections of follower force-driven filaments driven with backflow inside a

sphere were found to lead to MCWs if the forcing was distributed along the filament length and incorporated backflow, for particular values of the forcing and density of filaments. More complex models which account for the internal shear-driven actuation of cilia with enforced asymmetry through a curvature-dependent bending modulus [Han and Peskin, 2018], or biased motor protein kinetics [Chakrabarti et al., 2022] have been used to explore MCW formation with these non-reciprocal dynamics. In [Han and Peskin, 2018], it was found through 3D numerical simulations that spontaneous symmetry breaking from the synchronous beating can occur, while in [Chakrabarti et al., 2022], symplectic waves with multiple wavelengths were found to exist through 2D initial value problems. It was also found that different wavelength solutions could be obtained through a so-called 'cut-and-stitch' method, where an integer property called the charge (related to the wavelength of the solution) would remain approximately conserved throughout the simulation [Chakrabarti et al., 2022].

The fully three dimensional collective dynamics of arrays of active filaments, particularly in the context of follower force-driven filaments, is largely unexplored. The follower force model neglects the internal molecular motor-level forcing which drives cilia, providing a simple forcing mechanism that still gives rise to self-sustained oscillations. Using the previous sections we can incorporate anisotropy and asymmetry, motivated by structural properties of the axoneme, to generate non-reciprocal cilia-like beating dynamics. These dynamics lead to net fluid flow. The follower force model is a computationally scaleable set-up, meaning we are able to consider large arrays of follower force-driven filaments and explore their coordination. In this section we investigate the collective dynamics of an array of hydrodynamically coupled follower force-driven filaments. This will allow us to explore the existence and formation of MCWs, and analyse the fluid pumping abilities of different wavelength solutions. To avoid edge effects, we will use periodic boundary conditions. In Section 8.2 we discuss how this is implemented, and in particular, how we incorporate a no-slip boundary. We also outline how the non-zero net flow we anticipate the filaments will generate modifies the theory, and discuss the required extensions to the filament model when simulating collections of filaments. Next, in Section 8.3, we verify the accuracy of the hydrodynamic wall. Finally, in Section 8.4, we show preliminary results on arrays of asymmetrically beating filaments, crucially showing bistability between

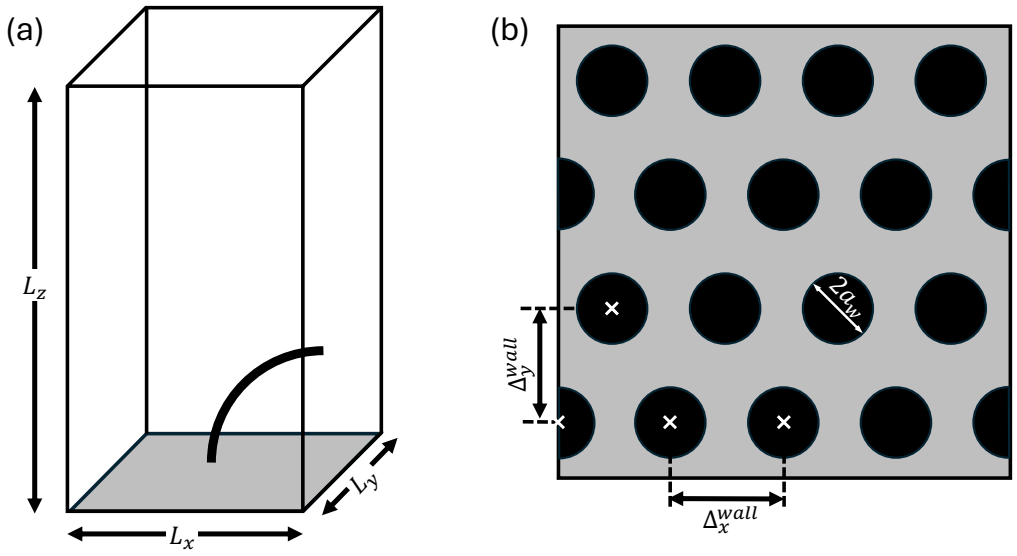


Figure 8.1: Schematics of (a) the computational domain of size $L_x \times L_y \times L_z$ with $N_f = 1$ filament, and (b) the particles making up the wall at $z = 0$. The particles have radius a_w, and are seeded hexagonally with separations Δ_x^{wall} and Δ_y^{wall} in the $\hat{\boldsymbol{x}}-$ and $\hat{\boldsymbol{y}}-$directions respectively.

different wavelength solutions if the number of filaments within the domain exceeds a critical value, and investigate how to robustly construct MCWs.

8.2 Numerical methods

In this chapter we consider N_f active filaments submerged in fluid in a domain, V, of size $L_x \times L_y \times L_z$, as illustrated in Figure 8.1a for $N_f = 1$. Each filament is driven at its tip by a single follower force, and has asymmetric and anisotropic properties following the methods outlined in Chapter 7. Otherwise, the intrinsic properties of each filament is identical to the previous chapters. To avoid edge effects in an array of filaments, we enforce periodic boundary conditions in the $\hat{\boldsymbol{x}}$ and $\hat{\boldsymbol{y}}$ directions. In the $\hat{\boldsymbol{z}}$ direction, we incorporate no-slip planar walls at $z = 0$ and $z = L_z$.

In the subsequent sections, we first outline the implementation of the hydrodynamic walls at $z = 0$ and $z = L$, before discussing how the filament methodology described in Section 4.2 is modified to incorporate the new boundary conditions.

8.2.1 Boundary implementation

Due to its computational efficiency, the end goal is to use Fast FCM to calculate the hydrodynamic coupling between the filament segments, as outlined in Section 2.2.3. Consequently, the domain we consider will have triply periodic boundary conditions, and so the notion of a wall is different to the previous open domain case, where the no-slip wall could be incorporated through the addition of image terms in the RPY formulation (see Section 2.2.1). One method to incorporate the wall is through the hydrodynamic solver; considering the steps outlined in Section 2.2.2, the inclusion of a no-slip wall in FCM would adapt the flow-solving step. The Fourier transform would be taken in the two periodic directions, instead of all three as in Section 2.2.2, and the resulting second order ODE can then be solved for the fluid velocity, as in Hashemi et al. [2023]. The resulting domain would have doubly periodic boundary conditions, say in the $\hat{\boldsymbol{x}}-$ and $\hat{\boldsymbol{y}}-$directions, with a no-slip condition satisfied exactly, say at $z = 0$. However this does not interface with Fast FCM, meaning our resulting simulations would be limited by the computational limitations associated with regular FCM.

Otherwise, in the triply periodic set-up, incorporating a wall at $z = 0$ naturally introduces a second wall at the top of the domain. Thus, incorporating a wall now corresponds to channel flow. Such walls can be constructed in the Stokes solver [Lomholt and Maxey, 2003], or by discretisation [Nott and Brady, 1994, Balboa Usabiaga et al., 2017]. We follow the latter method due to its flexibility in creating boundaries, and take an immersed boundary method approach to implementing a wall. We choose the height of the channel to be large to minimise the effect of the upper boundary. We now discuss how this works in practice.

We wish to implement a no-slip wall along $z = 0$, corresponding to the boundary condition:

$$\boldsymbol{u}(z = 0) = \mathbf{0}. \tag{8.1}$$

Following the rigid multiblob method [Balboa Usabiaga et al., 2017], we discretise the wall into N_w particles of hydrodynamic radius a_w. We take the particles to be seeded in a hexagonal lattice, with adjacent particles separated by distances Δ_x^{wall} and Δ_y^{wall} in the $\hat{\boldsymbol{x}}-$ and $\hat{\boldsymbol{y}}-$directions

respectively, as sketched in Figure 8.1b. We define the position of the centre of each particle to be $\boldsymbol{Y}_i^w(t)$, for $i = 1, ..., N_w$, and the corresponding velocity of the particle as $\boldsymbol{V}_i^w(t)$. In order to enforce Eq (8.1), we may either enforce the velocity of each particle to be zero, i.e. $\boldsymbol{V}_i^w(t) = \boldsymbol{0}$ for all t, or, equivalently, fix the position of each particle. We proceed with the latter, which corresponds to the constraint

$$\boldsymbol{C}_i = \boldsymbol{Y}_i^w(t) - \boldsymbol{Y}_i^w(0) = \boldsymbol{0}, \quad i = 1, ..., N_w, \tag{8.2}$$

for all t. Each constraint then gives rise to a constraint force,

$$\boldsymbol{F}_i^{C,wall} = \frac{\partial}{\partial \boldsymbol{Y}_i^w} \left(\boldsymbol{\lambda}_i^w \cdot \boldsymbol{C}_i \right) = \boldsymbol{\lambda}_i^w, \tag{8.3}$$

where $\boldsymbol{\lambda}_i^w$ is the vector Lagrange multiplier associated to the constraint $\boldsymbol{C}_i = \boldsymbol{0}$. An alternative view of these vectors is that they define the force required on the particle to ensure it has zero velocity.

Therefore we can consider the no-slip wall as a discrete collection of particles. Taking a force balance for particle $i = 1, \cdots, N_w$ yields:

$$\boldsymbol{F}_i^{C,wall} + \boldsymbol{F}_i^{H,wall} = \boldsymbol{0}, \tag{8.4}$$

where $\boldsymbol{F}_i^{H,wall}$ is the hydrodynamic force exerted by the fluid onto particle i. By Newton's third law, this implies that the force exerted by the particle on the fluid is given by

$$-\boldsymbol{F}_i^{H,wall} = \boldsymbol{\lambda}_i^w. \tag{8.5}$$

We now discuss how the above formulation interfaces with the active filament model outlined in Section 4.2.

8.2.2 Modified filament model

We consider N_f filaments, each comprised of N_p segments, clamped to a wall discretised into N_w particles. The wall particles are seeded as described in the previous section. Filament segments and the wall particles are chosen to have the same hydrodynamic radius, i.e. $a = a_w$.

Force and torque balances for each filament segment are given by Eqs (4.6) and (4.7), where the bending moment which gives rise to the elastic torque is defined by Eq (7.2). A similar force balance on each wall particle is given in Eq (8.5), which shows that the hydrodynamic force exerted by the particle on the fluid balances the Lagrange multiplier. We then augment this onto our mobility problem, Eq (4.11), such that this now reads:

$$\begin{bmatrix} \boldsymbol{V}^{fil} \\ \boldsymbol{V}^{wall} \\ \boldsymbol{\Omega}^{fil} \\ \boldsymbol{\Omega}^{wall} \end{bmatrix} = \mathcal{M} \begin{bmatrix} -\boldsymbol{F}^{H} \\ \boldsymbol{\lambda}^{w} \\ -\boldsymbol{T}^{H} \\ \boldsymbol{0}^{w} \end{bmatrix} \tag{8.6}$$

where $\boldsymbol{V}^{fil}$ and $\boldsymbol{\Omega}^{fil}$ are $3N_pN_f \times 1$ vectors containing the velocities and angular velocities of the filament segments, $\boldsymbol{V}^{wall}$ and $\boldsymbol{\Omega}^{wall}$ are $3N_w \times 1$ vectors containing the velocities and angular velocities of the wall particles, $-\boldsymbol{F}^H$ and $-\boldsymbol{T}^H$ are $3N_pN_f \times 1$ vectors containing the hydrodynamic forces and torques exerted by the filament segments onto the surrounding fluid, $\boldsymbol{\lambda}_w$ is the $3N_w \times 1$ vector of hydrodynamic forces exerted by the wall particles on the fluid, and $\boldsymbol{0}^w$ is a $3N_w \times 1$ vector of zeroes, stating that the wall particles do not contribute a torque to the flow. The mobility matrix, $\mathcal{M}$, is a $6(N_pN_f + N_w) \times 6(N_pN_f + N_w)$ matrix, coupling all of the filament segments and wall particles through the flow. To enforce the periodicity of the domain, we perform this mobility relation using FCM or Fast FCM, as described in Sections 2.2.2 and 2.2.3, depending on the system size. Any simulations presented in Sections 8.2 and 8.3 are generated using regular FCM. The collective dynamics results in Section 8.4 use Fast FCM.

With the velocities and angular velocities of each particle in hand, we now need to integrate forwards in time. Proceeding as described in Section 4.2, Eqs (4.21), (4.22) and (4.23) hold

for each filament segment, which can be discretised to provide $6N_p$ equations that describe the filament configuration; Eqs (4.24), (4.25) and (4.26). The wall particles each satisfy the continuous equation

$$\frac{d\boldsymbol{Y}_i^w}{dt} = \boldsymbol{V}_i^w \tag{8.7}$$

for $i = 1, \cdots, N_w$. Discretising in time, again using a second-order backwards diffential formula for consistency with the filament model, and using our constraint gives, for each particle,

$$\boldsymbol{Y}_i^{w,j+1} - \frac{4}{3}\boldsymbol{Y}_i^{w,j} + \frac{1}{3}\boldsymbol{Y}_1^{w,j-1} = \boldsymbol{0}. \tag{8.8}$$

The nonlinear system given by Eqs (4.24), (4.25), (4.26) and (8.8) can be solved using Broyden's method. At each timestep, Broyden's method calculates the Lie algebra elements, internal forces and clamping and tethering lagrange multipliers for each filament, as well as the constraint forces for each wall particle, reaching a total system size of $6N_pN_f + 3N_w$.

8.2.3 Non-zero mean flow

There is one subtle point missed by the methods outlined thus far. In periodic domains, Stokes equations are defined up to an additive constant - a spatially-invariant translational velocity. Specifically, if $\boldsymbol{u}(\boldsymbol{x}, t)$ is a solution of the Stokes equations in a periodic domain, then $\tilde{\boldsymbol{u}}(\boldsymbol{x}, t) = \boldsymbol{u}(\boldsymbol{x}, t) + \boldsymbol{V}^\infty(t)$ is also. When the hydrodynamics are solved using FCM or Fast FCM, this ambiguity is resolved by imposing the additional constraint that there is no mean flow out of the domain,

$$\int_{\delta V} \tilde{\boldsymbol{u}}(\boldsymbol{x}, t) \cdot \hat{\boldsymbol{n}}\, dS = 0, \tag{8.9}$$

where δV is the boundary of the domain, V, and $\hat{\boldsymbol{n}}$ is the outward facing normal. This constraint fixes $\boldsymbol{V}^\infty$, and therefore allows for a unique solution to be sought. We call this the no-flux, zero net flow, or zero-wavenumber condition, as it corresponds to the constraint that the Fourier transform of the fluid velocity vanishes when the wavenumber is zero.

However we anticipate this constraint will not be satisfied in our case; at the very least, we

expect non-reciprocal motion to generate net fluid flow, so over one period we expect the total flow out of the domain to be non-zero. To see where this issue arises, we consider a single filament in a triply periodic domain, with a wall discretised into N_w particles. Integrating the Stokes equations over the domain gives rise to the following constraint

$$\mathbf{0} = \int_V (\nabla \cdot \boldsymbol{\sigma} + \boldsymbol{f})\, d^3\boldsymbol{x} \tag{8.10}$$

$$= \sum_{i=1}^{N_p} \boldsymbol{F}_i + \sum_{i=1}^{N_w} \boldsymbol{\lambda}_i^w \tag{8.11}$$

which does not necessarily hold. In this case, enforcing the zero-wavenumber condition leads to unphysical flow fields. To illustrate this, we plot the flow fields corresponding to a single filament at two points within its beat in Figures 8.2a and c. We observe flow in the upper region of the domain, moving in the opposite direction to the filament beat. This backflow compensates for the net fluid that the filament is 'pushing' out of the domain, to enforce the zero-wavenumber condition. The speed of the backflow is time-dependent, oscillating in time as the total integral of the force within the domain oscillates. The speed of the backflow reaches its maximum at the same time the net filament force peaks.

This backflow is not an artefact of the physical system, and steps must be taken to circumvent this issue. In the literature, some works address this by including the inertia term ($\rho\frac{\partial \boldsymbol{u}}{\partial t}$) in the Navier-Stokes equation and using a method like the immersed boundary method to solve for the flow field [Han and Peskin, 2018]. Another method finds a reference velocity, $\boldsymbol{V}^\infty$, which enforces Eq (8.11) [Teran and Peskin, 2009]. The idea behind this is that the component of the force which doesn't integrate to zero generates an overall flow which we can solve for, and use to enforce Eq (8.11). We follow the idea of the latter, and seek to solve for the mean flow, $\boldsymbol{V}^\infty$ such that the forcing integrates to zero.

In particular, at each timestep, we calculate the net flow, $\boldsymbol{V}^\infty$, which enforces Eq (8.11). We then use FCM or Fast FCM to solve for the flow field in the frame where the system is force-free, $\boldsymbol{u}(\boldsymbol{x})$, and exploit the translational invariance of Stokes equations to calculate the fluid velocity

in the stationary frame,

$$\tilde{\boldsymbol{u}}(\boldsymbol{x}) = \boldsymbol{u}(\boldsymbol{x}) + \boldsymbol{V}^{\infty}. \tag{8.12}$$

An alternative physical viewpoint is that the wall at $z = 0$ is a rigid body, with velocity $\boldsymbol{V}^{\infty}$ and angular velocity $\boldsymbol{\Omega}^{\infty}$, such that each particle making up the wall satisfies

$$\boldsymbol{V}_i^w = \boldsymbol{V}^{\infty} + \boldsymbol{\Omega}^{\infty} \times \boldsymbol{Y}_i^w, \tag{8.13}$$

and equally the base segment of each filament similarly satisfies

$$\boldsymbol{V}_{1+N_p(k-1)} = \boldsymbol{V}^{\infty} + \boldsymbol{\Omega}^{\infty} \times \boldsymbol{Y}_{1+N_p(k-1)} \tag{8.14}$$

for $k = 1, ..., N_f$. The remaining segment positions take the rigid body motion into account through Eq (4.23). We can then calculate these rigid body velocities using different constraints, but in our case these relevant constraints will be that the total force integrates to zero, i.e. Eq (8.11), and that the body does not rotate, i.e. $\boldsymbol{\Omega}^{\infty} = \boldsymbol{0}$.

One method to incorporate this mean flow calculation is as a linear system which would be solved at each timestep, for instance using GMRES. This would occur within the mobility calculation, with the particle velocities adapted to include this mean flow. However this is computationally expensive; applying the mobility relation is the most expensive component of the numerical filament model, and using GMRES on such a linear system will involve evaluating this relationship for several GMRES iterations, within each Broyden's iteration. This slows computation massively, particularly for larger systems where we have on the order of thousands of particles.

Instead, we can offset the entire calculation to Broyden's method for limited additional cost. We can augment the nonlinear system given by Eqs (4.24), (4.25), (4.26) and (8.8), modified to include the rigid body motion, with the additional constraints for the mean flow, Eq (8.11),

to give the complete set of $6N_pN_f + 3N_w + 3$ equations:

$$\boldsymbol{Y}_1^{j+1} - \frac{4}{3}\boldsymbol{Y}_1^{j} + \frac{1}{3}\boldsymbol{Y}_1^{j-1} - \frac{2}{3}\Delta t\,(\boldsymbol{V}^{\infty})^{j+1} = \boldsymbol{0}, \tag{8.15}$$

$$\boldsymbol{u}_i^{j+1} - \frac{1}{3}\boldsymbol{u}_i^{j} - \frac{2}{3}\Delta t\,\exp^{-1}_{\boldsymbol{u}_i^{j+1}}\left(\boldsymbol{\Omega}_i^{j+1}\right) = \boldsymbol{0}, \tag{8.16}$$

$$\boldsymbol{Y}_1^{j+1} + \frac{\Delta L}{2}\sum_{m=2}^{i}\left(\hat{\boldsymbol{t}}_{m-1}^{j+1} + \hat{\boldsymbol{t}}_{m}^{j+1}\right) - \frac{4}{3}\boldsymbol{Y}_i^{j} + \frac{1}{3}\boldsymbol{Y}_i^{j-1} - \frac{2}{3}\Delta t\boldsymbol{V}_i^{j+1} = \boldsymbol{0}, \tag{8.17}$$

for each filament, where Eq (8.16) holds for $i = 1, ..., N_p$ and Eq (8.17) holds for $i = 2, ..., N_p$,

$$\boldsymbol{Y}_1^{w,j+1} - \frac{4}{3}\boldsymbol{Y}_1^{w,j} + \frac{1}{3}\boldsymbol{Y}_1^{w,j-1} - \frac{2}{3}\Delta t\,(\boldsymbol{V}^{\infty})^{j+1} = \boldsymbol{0}, \tag{8.18}$$

for $i = 1, ..., N_w$, and

$$\sum_{j=1}^{N_f}\sum_{i=1}^{N_p}\boldsymbol{F}_{fil\ j,\ seg\ i} + \sum_{i=1}^{N_w}\boldsymbol{\lambda}_i^{w} = \boldsymbol{0}. \tag{8.19}$$

This nonlinear system is solved to calculate the $6N_pN_f$ unknowns associated to the filaments (the $N_p - 1$ segment orientations and $N_p + 1$ Lagrange multipliers for each filament), the $3N_w$ unknowns corresponding to the wall particles, and 3 unknowns in the mean flow $\boldsymbol{V}^{\infty}$ at each timestep. We could also include the net angular velocity in the system, for nominal extra cost. This is solved via Broyden's method, which requires an approximate Jacobian based on the system. For this purpose we again utilise the diagonal mobility approximation leading to a sparse approximate Jacobian, and so use Schur complements to apply the inverse of this matrix efficiently. Further discussion about the approximate Jacobian can be found in Appendix C. At the end of each timestep, the net velocity is subtracted from the fluid velocity and the particle velocities, exploiting the translational invariance of Stokes flow, to return to the frame where the filament is clamped in space.

To emphasise the importance of incorporating the mean flow, we run simulations with and without the mean flow amendment. In these simulations, we model a single filament with aspect ratio $\alpha = 44$ and length $L = \alpha a$, with parameters $(f, \kappa_0, \beta) = (40, 3\pi/4, 2)$. The filament is in a domain of size $L \times L \times 2L$, with a wall at $z = 0$ discretised into $N_w = 100$ equispaced particles, seeded hexagonally as shown in Figure 8.1. The hydrodynamics are calculated via

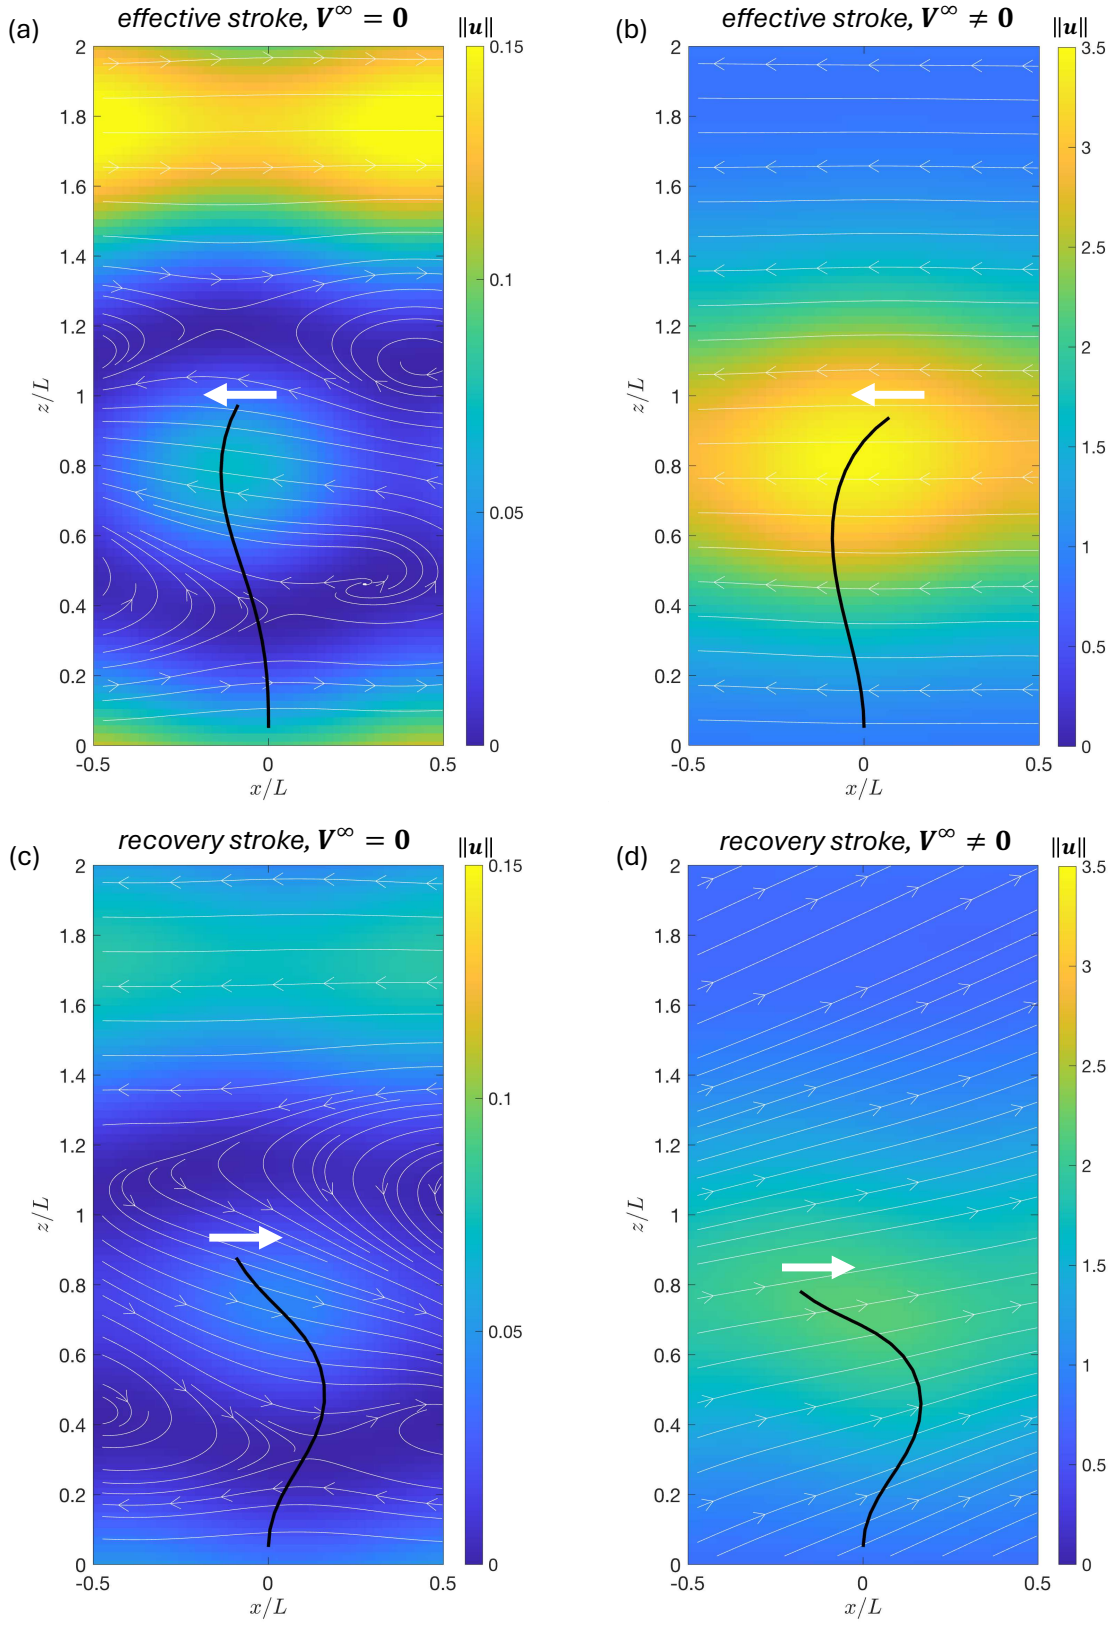


Figure 8.2: Visualisations of the flow field of a filament during part of the effective (recovery) stroke when the net flow is (a) ((c)) neglected or (b) ((d)) included. The thin white lines visualise the streamlines of the flow and the thick white arrow refers to the direction of beating. The filament is undergoing beating oscillations obtained using the parameters $(f, \kappa_0, \beta) = 40, 3\pi/4, 2)$, and is in a domain of size $L \times L \times 2L$ with a wall at $z = 0$ constructed of $N_w = 100$ particles.

regular FCM. We plot the flow field of the filaments at two points in their beat in Figure 8.2. In Figure 8.2a and b, the filament is undergoing part of the effective stroke in the direction indicated by the white arrow, and in Figure 8.2c and d, the filament is mid-recovery stroke. In both Figures 8.2a and c, the net velocity is neglected. As can be seen, this leads to much smaller fluid velocities, as well as a region of backflow towards the top of the domain. The magnitude of the backflow is related to the magnitude of the total force on all the particles, and acts in the opposite direction to the filament motion to compensate for the zero-wavenumber condition. Figures 8.2b and d show the flow fields obtained when the net velocity is included. The backflow no longer exists, and the fluid velocity is localised around the filament tip, as expected. The fluid velocity is also much larger than in the case where the net velocity is neglected (approx 40x the magnitude).

8.3 Discretisation verification - Poiseuille flow

Before proceeding with the results for the collective dynamics of filaments, we first verify the methodology of implementing the no-slip wall at $z = 0$, and investigate the effect of the wall discretisation using the example of Poiseuille flow.

For this example, we consider pressure-driven flow through a channel. The height of the channel is H, and the flow is driven by a pressure gradient in the $\hat{\boldsymbol{x}}-$direction,

$$\nabla p = -G\hat{\boldsymbol{x}}, \tag{8.20}$$

where G is the strength of the pressure gradient.

We consider a triply periodic domain of size $L \times L \times H$, and construct a wall from N_w particles as discussed in the previous section. We define the position of each particle to be $\boldsymbol{Y}_i^w$ for $i = 1, ..., N_w$. Each wall particle has an unknown Lagrange multiplier, $\boldsymbol{\lambda}_i^w$ for $i = 1, ..., N_w$, which acts as a force on the particle to enforce the prescribed velocity. In these simulations we use regular FCM as a hydrodynamic solver, due to its convenience in visualising the flow field.

In a similar manner to the previous section, we expect a non-zero net flow out of the domain. Therefore care needs to be taken to ensure the driving pressure term is handled correctly - in this case, if the zero-wavenumber condition is imposed, the resulting system would have zero flow. Alike our approach in the filament methodology, due to the translational invariance of Stokes flow, we split this problem into two; first, we solve for the flow in a frame such that the forces over the domain integrate to zero. This frame will be moving with velocity $\boldsymbol{V}^{\infty}(G)$ with respect to the rest frame; a velocity that will be determined as part of the problem. Consequently, the wall particles will have velocity $-\boldsymbol{V}^{\infty}(G)$ in our translated frame. Defining $\boldsymbol{\lambda} = (\boldsymbol{\lambda}_1^w, ..., \boldsymbol{\lambda}_{N_w}^w)$, this corresponds to the mobility relation

$$-\boldsymbol{K}\boldsymbol{V}^{\infty} = \mathcal{M}\boldsymbol{\lambda} \tag{8.21}$$

where

$$\boldsymbol{K} = \begin{pmatrix} \boldsymbol{I}_3 \\ \vdots \\ \boldsymbol{I}_3 \end{pmatrix} \in \mathbb{R}^{3N_w \times 3} \tag{8.22}$$

is a matrix of N_w stacked 3×3 identity matrices. In this case, $\mathcal{M}$ contains only the force-velocity entries of the mobility matrix, i.e. such that, for each particle i, the mobility relation reads

$$-\boldsymbol{V}^{\infty} = \sum_{j=1}^{N_w} \boldsymbol{M}_{ij}^{\boldsymbol{VF}} \boldsymbol{\lambda}_j^w. \tag{8.23}$$

By construction, in this frame we expect the total force (including the driving force, through the pressure gradient) to integrate to zero. Specifically, we expect

$$\begin{aligned} \boldsymbol{0} &= \int_V (\boldsymbol{f} - \boldsymbol{\nabla} p)\, d\boldsymbol{x} \\ &= \int_V \left(\sum_{i=1}^{N_b} \boldsymbol{\lambda}_{\boldsymbol{i}}^w \delta(\boldsymbol{x} - \boldsymbol{Y}_i^w) + G\hat{\boldsymbol{i}} \right) d\boldsymbol{x} \\ &= \boldsymbol{K}^T \boldsymbol{\lambda} + GL^2 H \hat{\boldsymbol{i}}. \end{aligned} \tag{8.24}$$

Equations (8.21) and (8.24) can be expressed as a saddle point system,

$$\begin{bmatrix} \mathcal{M} & \boldsymbol{K} \\ \boldsymbol{K}^T & \boldsymbol{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{\lambda} \\ \boldsymbol{V}^\infty \end{bmatrix} = \begin{bmatrix} \boldsymbol{0} \\ -GL_xL_yL_z\hat{\boldsymbol{i}} \end{bmatrix} \tag{8.25}$$

in which $\boldsymbol{\lambda}$ and $\boldsymbol{V}^\infty$ are our unknowns. This system is solved using GMRES. The corresponding fluid velocity in the translating frame, $\tilde{\boldsymbol{u}}(\boldsymbol{x}, t)$, can be obtained while applying FCM. Then the fluid velocity in the rest frame can be recovered through a translation, following the linearity of the Stokes equations:

$$\boldsymbol{u}(\boldsymbol{x}) = \tilde{\boldsymbol{u}}(x) + \boldsymbol{V}^\infty. \tag{8.26}$$

We note that the particle velocities in this frame are then given by

$$\boldsymbol{V}^w = \mathcal{M}\boldsymbol{\lambda} + \boldsymbol{K}\boldsymbol{V}^\infty = \boldsymbol{0} \tag{8.27}$$

by definition, consistent with our constraint.

8.3.1 Comparison to analytic results

To verify the wall methodology and implementation, we begin by conducting a numerical study in which the dimensions of the domain and the wall density are fixed. We use a domain of size $100a \times 100a \times 100a$ (i.e. $L = H = 100a$), and construct a wall at $z = 0$, discretised into $N_w = 625$ particles, with inter-particle separations of $\Delta_x^{wall} = \Delta_y^{wall} = 4$.

We can compare our results to the analytical solution, where no-slip conditions are applied to the top and bottom wall. In this case, Stokes equations can be solved to give the (dimensional) fluid velocity,

$$\boldsymbol{u}(z) = \frac{G}{2\eta} z(H - z)\hat{\boldsymbol{i}}. \tag{8.28}$$

In our comparison, we fix the pressure gradient to be such that the norm of the analytical

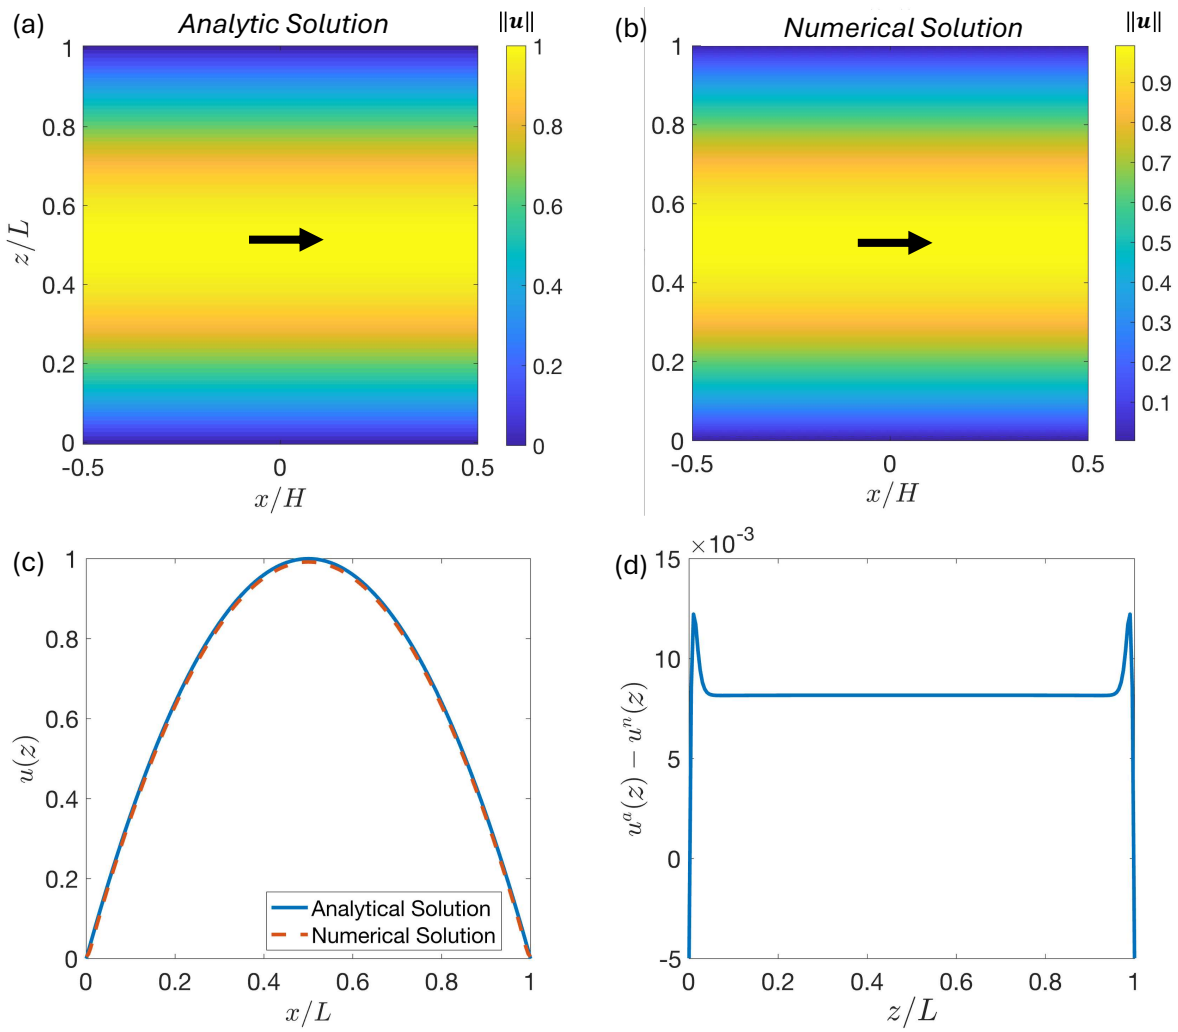


Figure 8.3: Heatmaps showing the (a) analytic solution and (b) numerical solution for the fluid velocity in Poiseuille flow with $G = 8\eta/H^2$. The domain is of size $100a \times 100a \times 100a$. The wall at $z = 0$ (and therefore, also at $z = H$ by periodicity) is discretised into 625 hexagonally seeded particles, separated by distances $\Delta_x^{wall} = \Delta_y^{wall} = 4a$. (c) The velocity profiles of the analytic and numerical solutions for this fluid velocity and (d) the error between the two solutions.

maximum velocity, $\max_{z\in[0,H]} \boldsymbol{u}(z) = (GH^2/8\eta)\hat{\boldsymbol{x}}$, is unity. This corresponds to the choice:

$$G = 8\eta/H^2. \tag{8.29}$$

We plot the resulting flow field for the analytic and numerical solutions in Figure 8.3a and b respectively. We also show the respective velocity profiles in Figure 8.3c, and show the difference between the analytic and numerical solutions in Figure 8.3d. Visually these are incredibly similar, and we find that the relative error in this case is ~ 0.011. Therefore it is clear that this methodology is able to replicate a no-slip wall with reasonable agreement. We

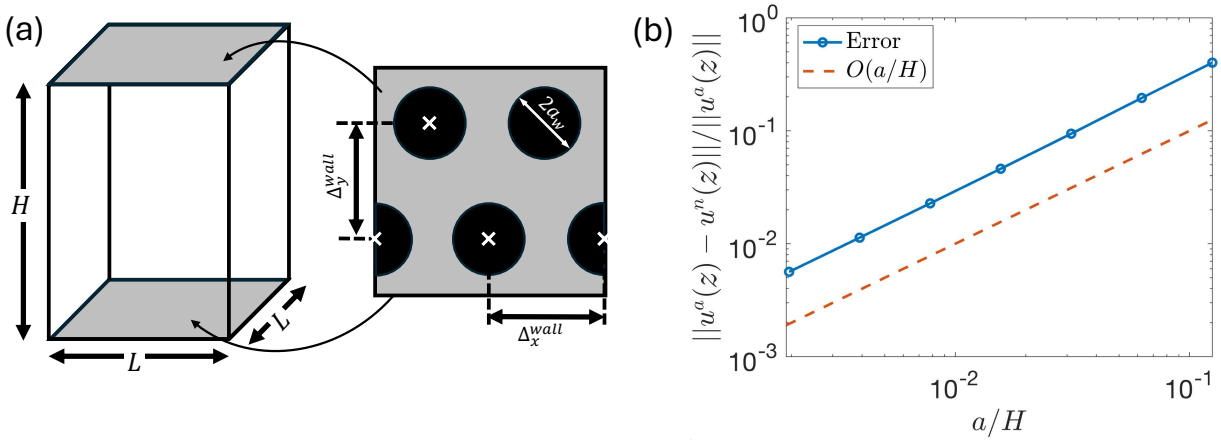


Figure 8.4: (a) A schematic showing the set-up for simulations where we vary the domain height. We fix the (x, y)-domain and the wall density such that $N_w = 4$, $\Delta_x^{wall} = \Delta_y^{wall} = 4a$, and $L = \Delta_x^{wall}\sqrt{N_w}$. We vary the height of the channel, H. (b) The relative pointwise error between the analytic velocity profile, $\boldsymbol{u}^a(z)$, and the numerical velocity profile, $\boldsymbol{u}^b(z)$, as we vary the height of the box. We see that the error scales like $\epsilon \sim O(a/H) \sim O(1/\sqrt{N_w})$.

now investigate how the effective particle radius, a/H, and the wall density can impact the numerical accuracy of this approximation.

8.3.2 Effect of the relative particle radius

We now focus on refining the wall to confirm convergence to the analytic solution. We do this by considering the case where $a/H \to 0$ while the density of particles is fixed, which is equivalent to taking $N_w \to \infty$. This is achieved by considering a domain of size $L \times L \times H$. To consider the smallest size problem, which also encapsulates the hexagonal packing of the wall particles, we consider $N_w = 4$ as shown in Figure 8.4a. The particles are separated by a distance $\Delta_x^{wall} = \Delta_y^{wall} = 4a$, and consequently the domain length is taken to be $L = 8a$.

We then increase the height of the domain, hence reducing the particle size relative to the channel height. We plot the relative error between the analytical solution, $\boldsymbol{u}^a(z)$, and the numerical solution, $\boldsymbol{u}^n(z)$, given by

$$\varepsilon = \frac{||\boldsymbol{u}^a(z) - \boldsymbol{u}^n(z)||}{||\boldsymbol{u}^a(z)||}, \tag{8.30}$$

in Figure 8.4b. We see this scales like $\varepsilon \sim a/H$, as predicted in the literature [Balboa Usabiaga

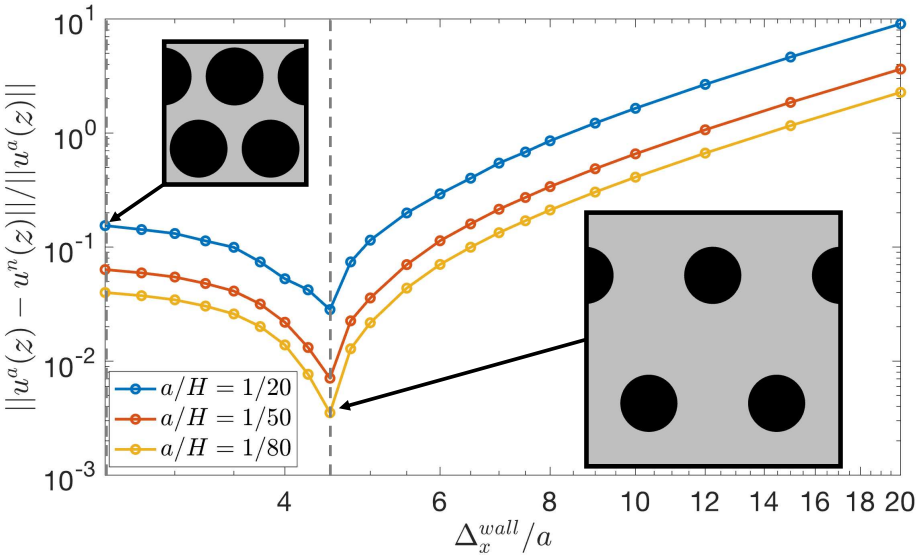


Figure 8.5: The relative pointwise error in the analytic velocity profile, $\boldsymbol{u}^a(z)$, and the numerical velocity profile, $\boldsymbol{u}^n(z)$, as we vary the spacing between the particles. We see that this reaches a minimum around $\Delta_x^{wall} \approx 4.5a$ regardless of the domain height. Inset: Plot of the $z = 0$ plane in the domain, showing the wall particles when the particle separation is $\Delta_x^{wall} = 2.5a$ (left) and $\Delta_x^{wall} = 4.5a$ (right).

et al., 2017].

We note that we can relate the relative radius of the particles to the number of particles via the relation $a/H \sim (N_w)^{-1/2}$. To see this, consider how the number of wall particles changes as we double the height of the domain, or equivalently halve the radius of each particle, keeping the relative distance between the particles fixed. Therefore, we can conclude that the error scales with the number of particles through the relation $\varepsilon \sim N_w^{-1/2}$.

8.3.3 Effect of the wall density

We now investigate the effect of refining the wall by instead varying the density of particles. We perform a study in which both the particle radius and the height of the domain are fixed, hence fixing a/H, and vary the distance between particles. Again, we consider a domain of size $L \times L \times H$, as illustrated in Figure 8.4a.

Once more, we consider the scenario where the wall is discretised into $N_w = 4$ particles within the domain. We take particle separation to be the same in either direction, i.e. $\Delta_x^{wall} = \Delta_y^{wall}$,

such that the length of the domain in either direction is then given by $L = 2\Delta_x^{wall}$. We vary Δ_x^{wall} and plot the relative pointwise error in Figure 8.5 for three domain heights. The error is largest for large distances between particles, and hits a local minimum at $\Delta_x^{wall} \approx 4.5a$. We see that the error reduces as we increase the domain height, as found in the previous section, reaching a minimum at the same value regardless of the domain height. As such, it's clear that the discretisation of the wall into particles is effective, particularly for particle separations of $\Delta_x^{wall} = \Delta_y^{wall} \approx 4.5a$, and when the particle height, relative to the domain height, is smaller.

8.4 Collective dynamics of filaments

8.4.1 Simulation parameters

In this section we present results for an array of asymmetric and anisotropic follower force-driven filaments in a triply periodic domain. To allow us to consider larger domains through the use of larger grid spacings, we use Fast FCM as the hydrodynamic solver. The use of a coarser grid speeds up computational time, and bypasses the memory concerns associated to using larger grids with regular FCM. The domain has dimensions $L_x \times L_y \times L_z$, as sketched in Figure 8.6a. We incorporate a hydrodynamic wall at $z = 0$, which we discretise into N_w hexagonally packed particles, as in Section 8.2.1. Following Figure 8.5, we choose the separation to be $\Delta_x^{wall} = \Delta_y^{wall} = 4.4a$ to be close to the minimiser of the pointwise error. To minimise the impact of the upper wall on dynamics, and to reduce the ratio a/H, we maintain a box height of $L_z = 10L$. This corresponds to choosing $a/H = 1/10\alpha \approx 0.002$.

We model N_f filaments, separated by a distance Δ_x^{fil} in the $\hat{\boldsymbol{x}}-$direction, and Δ_y^{fil} in the $\hat{\boldsymbol{y}}-$direction. We seed the filaments in a rectangular array of size $N_f^x \times N_f^y$, where N_f^x and N_f^y are the number of filaments in the $x-$ and $y-$directions respectively, so that the total number of filaments in $N_f = N_f^x N_f^y$. We illustrate this for $N_f^x = N_f^y = 3$ in Figure 8.6a. Correspondingly, the domain length and width are given by $L_x = N_f^x \Delta_x^{fil}$ and $L_y = N_f^y \Delta_y^{fil}$. There are two key types of MCWs which can occur in such a system: (i) waves parallel to the beat plane, called

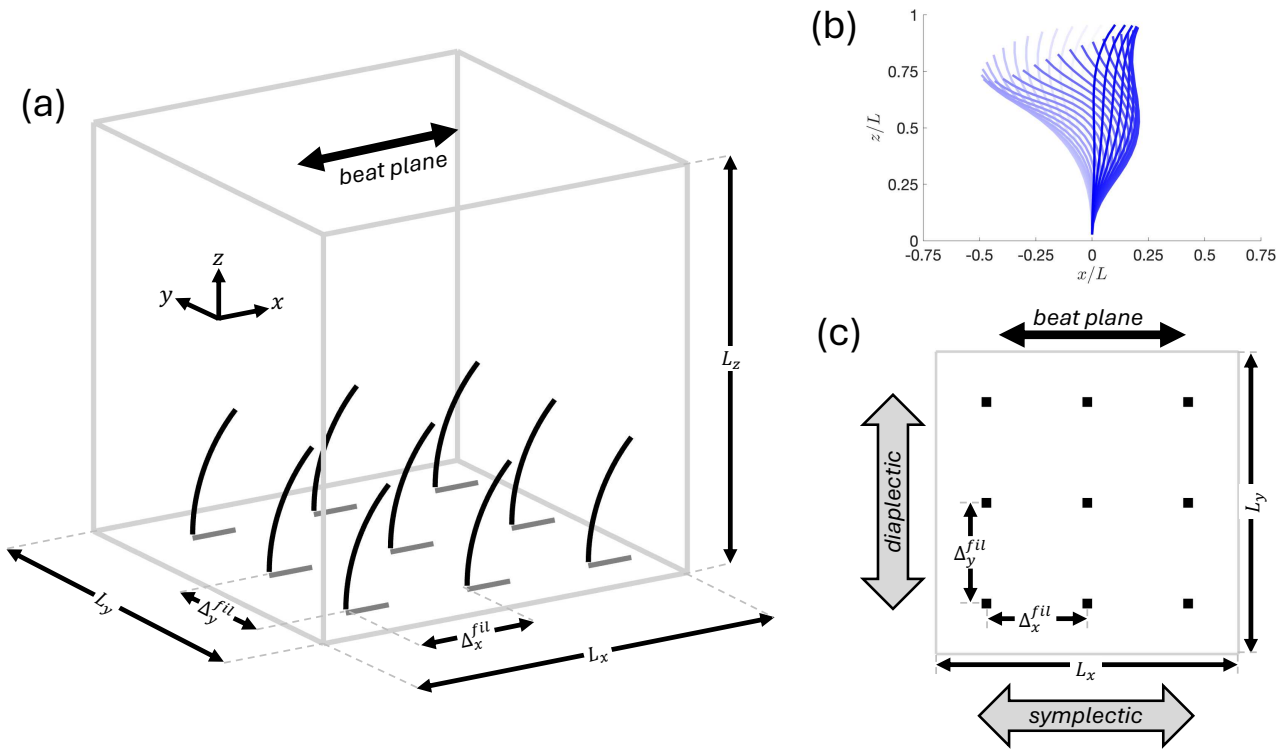


Figure 8.6: (a) A schematic of the domain, with dimensions $L_x \times L_y \times L_z$, and the 3×3 rectangular array of filaments. The beat plane of each filament, defined for $(f, \kappa_0, \beta) = (40, 3\pi/4, 2)$, is aligned with the $\hat{\boldsymbol{x}}-$axis. The distance between neighbouring filaments in the $\hat{\boldsymbol{x}}$ and $\hat{\boldsymbol{y}}$ directions is given by Δ_x^{fil} and Δ_y^{fil} respectively. (b) Schematic of the dynamics of a single filament in an open domain with a no-slip wall at $z = 0$, calculated using the RPY tensor with a no-slip wall and for the parameters $(f, \kappa_0, \beta) = (40, 3\pi/4, 2)$, as found in Chapter 7. The passage of time is indicated by the increasing boldness of the lines. (c) A view of the domain from above. Square markers represent the filament seeding. There are two main types of MCWs. If the MCW direction aligns with the beat plane, we define the wave as symplectic. If the wave direction is perpendicular to the beat plane, we observe a diaplectic wave.

symplectic waves, or (ii) waves beat perpendicular to the beat plane, called diaplectic waves. We label these in Figure 8.6c. These each have sub-types, depending on the direction of wave propagation relative to the effective beat direction. By confining $N_f^x = 1, N_f^y > 1$, we consider a diaplectic row of filaments in which the MCWs, if present, propagate perpendicular to the beat plane of the filament. Similarly, choosing $N_f^x > 1, N_f^y = 1$ corresponds to a symplectic row of filaments, in which the MCWs, if present, propagate parallel to the beat plane of the filament.

We fix the material parameters associated to each filament, but emphasise that the filament dynamics are able to evolve in response to the hydrodynamic effect of the neighbouring filaments. At the level of the individual filament, we take $K_B = K_T$, and set the filament aspect

ratio to be $\alpha = 44$. Informed by Chapter 7, we define the dynamics of each filament using the parameters $f = 40$, $\kappa_0 = 3\pi/4$ and $\beta = 2$, which gives rise to an asymmetric planar beat for a single filament in an infinite domain, bounded by a no-slip wall at $z = 0$. For completeness, we plot this dynamic over one period in Figure 8.6b.

In the following sections, we fix the individual filament dynamics as described above, and vary the domain length, filament separation and initial conditions to explore the stable solutions in the state space. We now outline the key results.

8.4.2 MCWs stable above a critical wavelength

To begin, we fix the filament spacing, $\Delta_x^{fil} = L$ and $\Delta_x^{fil} = L/2$, and consider one dimensional arrays of filaments, either of dimension $1 \times N_f$ (diaplectic) or $N_f \times 1$ (symplectic). We run initial value problems with different initial conditions, systematically increasing the number of filaments (and consequently, the domain length or width) from $N_f \in \{5, 10, ..., 40\}$.

We choose a variety of initial conditions; where all filaments are vertically upright with small perturbations to the forcing on each filament, where all filaments have a random phase (based on the dynamics of a single filament, see Figure 8.6b), where all filaments have the same phase, or where a MCW is prescribed (again, based on the dynamics of a single filament). The MCW initial condition is prescribed to have a wavelength equal to the length of the domain. Denoting the period of the solution in the single filament case as T, each initial value problem is run for $t = 250T$ to ensure convergence to the final state. In our experience, all simulations converged within the first 100 periods. In the first case, for initially vertically upright filaments, the filaments all first deform in response to the intrinsic curvature which dominates over the small perturbations, and undergo synchronous beating. Similarly, synchronous initial conditions lead to synchronous beating. Thus we can conclude that synchronous beating solution is a stable attractor in the state space.

However using initial conditions corresponding to MCWs, different dynamics are observed. We summarise the results in Figure 8.7. For small arrays of filaments, the dynamics quickly

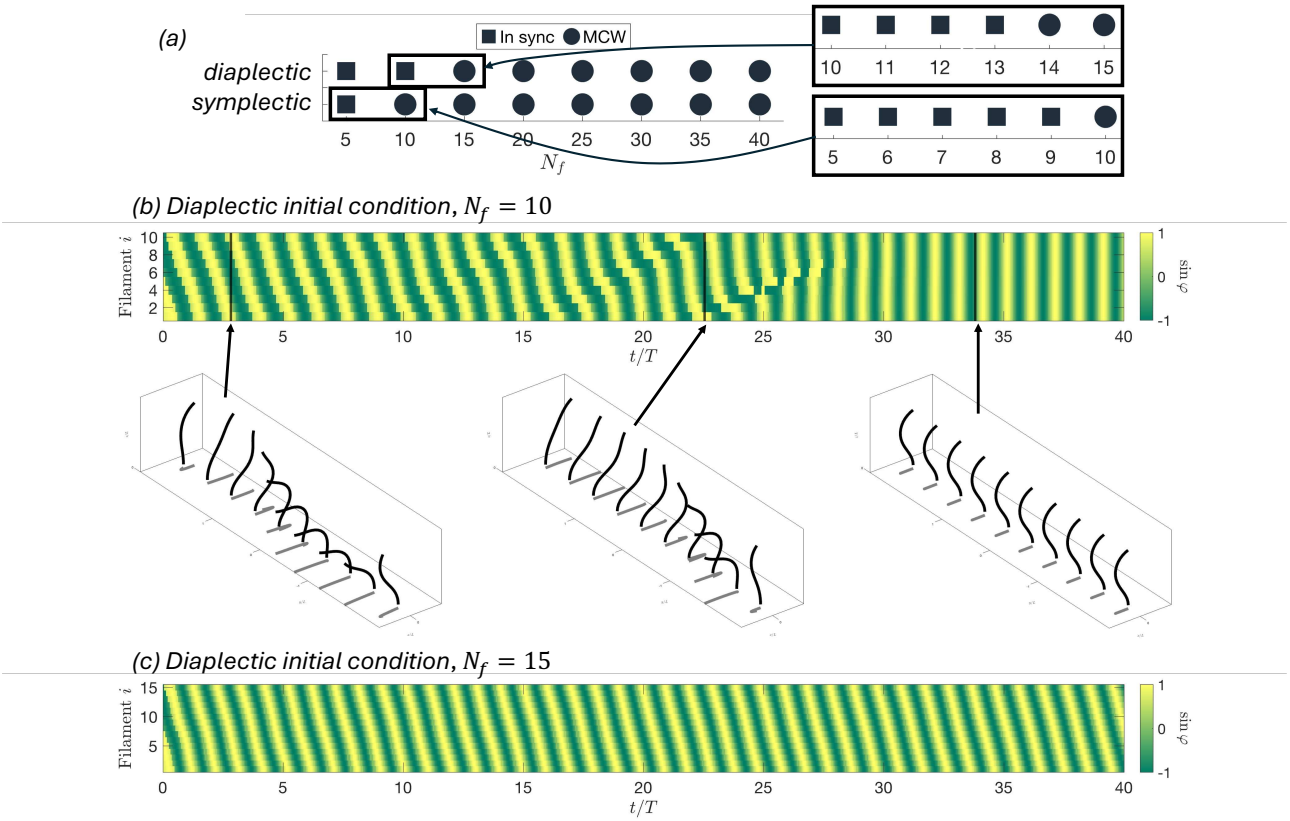


Figure 8.7: (a) Results of initial value problems of arrays of $1 \times N_f$ and $N_f \times 1$ filaments, corresponding to diaplectic (top) and symplectic (bottom) waves respectively. The initial condition for each simulation is a MCW, based on the dynamics of a single cilium in an infinite domain, bound by a wall at $z = 0$, as found in Chapter 7. The inter-filament spacing is fixed, so increasing N_f also corresponds to increasing the length of the domain. For diaplectic (symplectic) waves, metachronal initial conditions lead to synchronous beating for $N_f \leq 13$ ($N_f \leq 9$). For larger numbers of filaments, the array converges to a MCW. (b) Heat maps of filament phase over 40 periods of the simulation for a 1×10 array of filaments, and visualisations of the dynamics at three points in time. We see that the filaments synchronise their beat. (c) Heat maps of filament phase over 40 periods of the simulation for a 1×15 array of filaments. We see that the solution converges to a metachronal beat within the first 5 periods.

synchronise. In Figure 8.7b, we plot the heatmap associated to the phase of the solution in an array of 1×10 filaments. Here, the phase is defined by the beat of a single filament in the array once the solution has converged to its final state. We see that the filaments initially undergo a MCW, i.e. a travelling wave of phase along the domain length. At $t \approx 10T$, we see a disruption in the phase, which travels along the domain as time progresses. By $t \approx 30T$, all the filaments have synchronised. Increasing the number of filaments (and hence, the domain length), we see that synchronisation does not occur; in fact, the MCW remains stable for all time (see Figure 8.7c for the phase of an array of 1×15 filaments over time). For diaplectic arrays of filaments, MCWs appear to be stable for $N_f \geq 14$ filaments, whereas for symplectic arrays this value

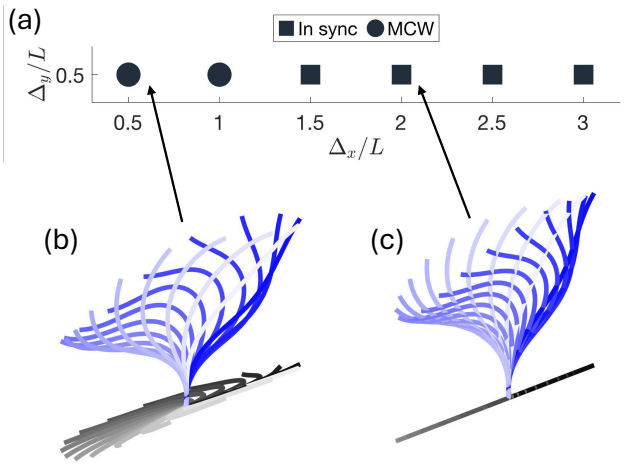


Figure 8.8: (a) The results of initial value problems with metachronal initial conditions on a 1×15 array of filaments. We fix the spacing of the filaments in the $\hat{\boldsymbol{y}}-$direction to be $\Delta_y^{fil} = L/2$, and vary the inter-filament spacing in the $\hat{\boldsymbol{x}}-$direction, Δ_x^{fil}. We observe MCWs below a critical spacing. Above this spacing, the filaments synchronise. The dynamics of one filament in the array over one period for (b) $\Delta_x^{fil} = L/2$ and (c) $\Delta_x^{fil} = 2L$. The passage of time is indicated by the increasing boldness of the lines. We note that decreasing the separation between filaments changes the dynamics of the filaments at the individual level, adding an aplanar component to the beat.

changes to $N_f \geq 10$. Therefore it appears that there is a minimum wavelength required for the MCW to be stable - below this, only synchronous states are observed.

Using initial conditions where all the filaments have a random phase usually leads to synchronous beating. However increasing the number of filaments, a subset of these random phase initial conditions lead to MCWs. This indicates that these solutions are truly bistable, and that the synchronous beating solution likely has a much larger basin of attraction than the other states.

8.4.3 Dependence on inter-filament spacing

With the knowledge that MCWs exist and are stable for particular choices of parameters, it is of interest how robust these solutions are to changes in the filament spacing. To investigate this, we consider a diaplectic row of $N_f = 15$ filaments, i.e. an array of 1×15 filaments, and vary the inter-filament spacings, Δ_x^{fil} and Δ_y^{fil}, one at a time. This also impacts the domain size, through the relations $L_x = \Delta_x^{fil}$, $L_y = 15\Delta_y^{fil}$. We run initial value problems with metachronal

initial conditions, with a wavelength equal to the domain length, $\lambda = L_y$.

Based on the previous section, MCW solutions exist and are stable for $\Delta_x^{fil} = L$, $\Delta_y^{fil} = L/2$. We observe that fixing $\Delta_x^{fil} = L$ and changing $\Delta_y^{fil} \in \{L/2, L, ..., 3L\}$ makes no quantitative change to the dynamics - MCWs still exist, and are stable, implying a weak dependence on the spacing perpendicular to the beat plane. However fixing $\Delta_y^{fil} = L/2$ and varying $\Delta_x^{fil} \in \{L/2, L, ..., 3L\}$, we note an exchange of stability. We plot the results of these initial value problems in Figure 8.8a. Above a critical separation, MCWs appear to be unstable, and we observe that the filaments quickly synchronise. This is likely a result of the weaker hydrodynamic interactions as the distance between filaments is increased. When the filaments are closer together, the hydrodynamic response from a filament on its neighbours is stronger and MCWs appear to be stable.

Despite using the same filament material parameters in each simulation, as we vary the inter-filament spacing in the $\hat{\boldsymbol{x}}-$direction, we note that the individual dynamics of each filament changes. For larger separations, we observe asymmetric planar beating similar (but not identical) to the individual filament case. We plot the beat of a single filament in the array over one period in Figure 8.8c. As we increase the separation, this beat becomes closer to the single filament beat shown in Figure 8.7b. However, decreasing the separation, we observe that the hydrodynamic effect of the neighbouring filaments changes the individual beat, drawing each filament out of its planar orbit, as shown for a single filament in the array in Figure 8.8b.

8.4.4 Multistability between different wavelength solutions

In the previous sections we saw that, if enough filaments are present and are sufficiently close, we observe both synchronous states and MCWs as stable solutions, depending on the prescribed initial condition. To test the robustness of the different wavelength solutions, we test a diaplectic array of 1×60 filaments with various initial conditions. We fix the spacing between filaments to be $\Delta_x^{fil} = L$, $\Delta_y^{fil} = L/2$, therefore considering domains of size $L_x \times L_y \times L_z = L \times 30L \times 10L$. We summarise some key findings from these initial value problems below.

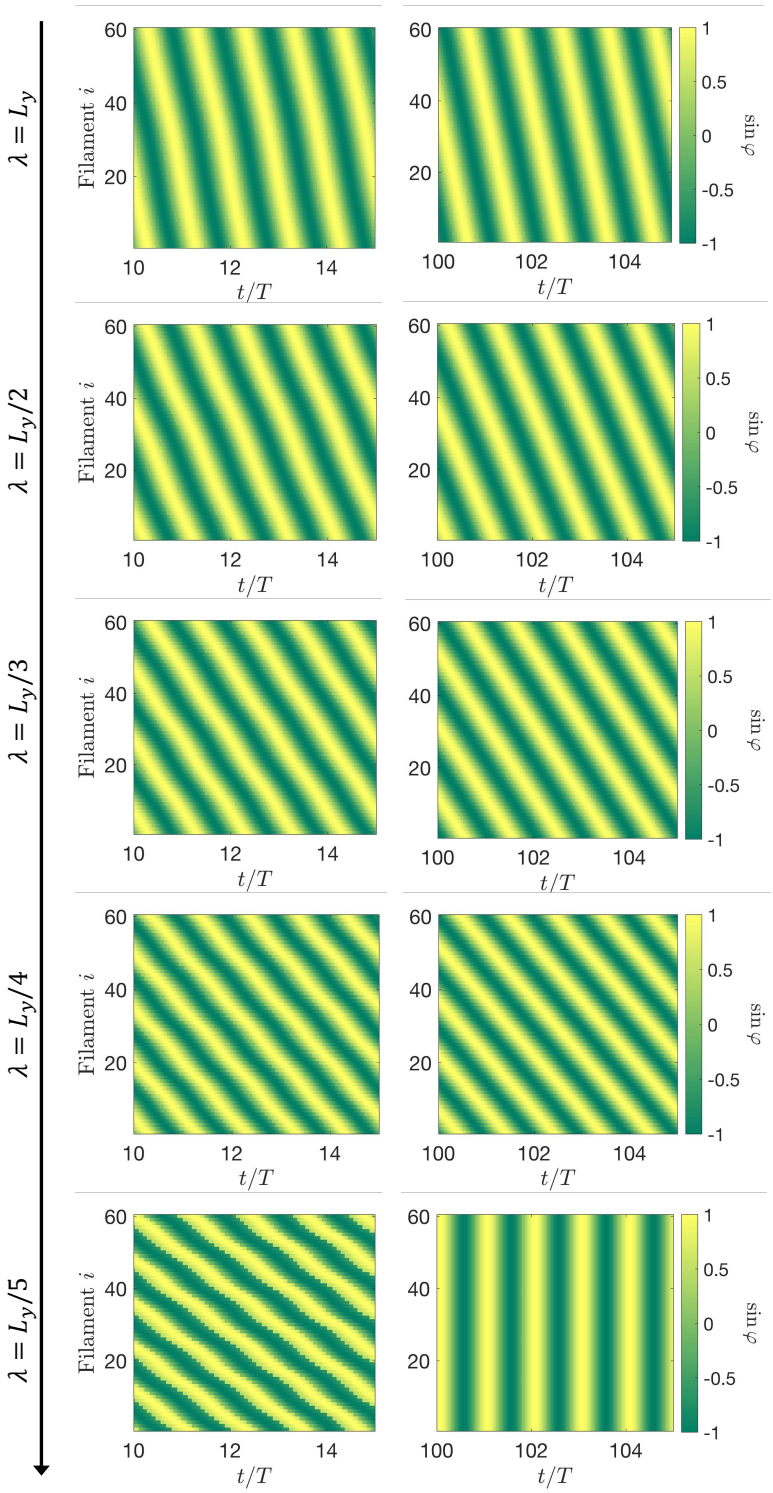


Figure 8.9: Heat maps of filament phase at two points in time for 1×60 arrays of filaments with various initial conditions. The left column shows times $t/T \in [10, 15]$, and the right column shows times $t/T \in [100, 105]$. The initial conditions are all MCWs. From top to bottom, the wavelength of the corresponding initial condition is given as $\lambda/L_y = \{1, 1/2, 1/3, 1/4, 1/5\}$. We see that if the wavelength of the solution exceeds the critical value, as found in Figure 8.7a, the MCW is stable. Below this, the dynamics converge to the synchronised solution.

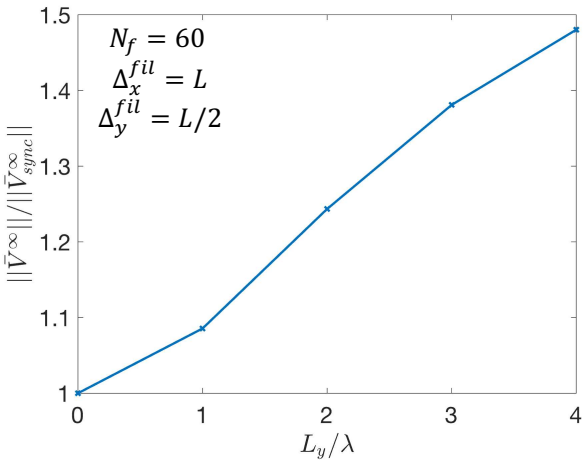


Figure 8.10: The norm of the mean net velocity, $||\bar{\boldsymbol{V}}^\infty||$, for the different stable wavelength solutions found in Figure 8.9. The net velocity is normalised by the norm of the mean for the synchronous beating case. Synchronous beating is categorised by the wavelength $L_y/\lambda = 0$. We observe that the MCW solutions pump more fluid than the synchronous beating solution. Smaller wavelength solutions are more effective at pumping fluid.

First, we investigate the robustness of the different metachronal wavelengths observed in Section 8.4.2. We do this by running initial value problems of MCWs with various wavelengths, again based on the dynamics of a single filament in an open domain with a no-slip wall at $z = 0$. We investigate the wavelengths $\lambda/L_y \in \{1/5, 1/4, 1/3, 1/2, 1\}$. We plot the heatmaps associated to the filament phase for a region of time close to the start of the simulation, and after 100 periods in Figure 8.9. We observe that MCWs remain stable in all tested cases except for $\lambda/L_y = 1/5$, implying that MCWs are stable if the wavelength is sufficiently long. As the domain length is linearly related to the number of filaments, we can relate this wavelength to the number of filaments per full cycle of phase, $\lambda/N_f \in \{12, 15, 20, 30, 60\}$. Since our interfilament spacings are consistent with the previous section, we can relate this result to our findings there. Our results are in agreement with Figure 8.7; wavelengths corresponding to $N_f \leq 13$ filaments (with our chosen spacing) are unstable, whereas wavelengths with $N_f \geq 14$ are stable. Crucially, this result implies multiple MCW solutions (as well as the synchronous beating state) are all stable at this particular choice of parameters.

We can evaluate the fluid-pumping ability of the different stable solutions through the parameter $\boldsymbol{V}^\infty$, which expresses the net fluid pumped out of the domain. For MCW solutions, this value is constant over time. When the filaments beat in sync, the net fluid velocity oscillates depending

on the phase of the beat. After the initial value problems in Figure 8.9 have converged, we can evaluate the net velocity over one period for each of the solutions. The final solutions correspond to MCWs with wavelength $\lambda/L_y \in \{\infty, 1/4, 1/3, 1/2, 1\}$ (the first corresponds to the synchronous beating state reached via the $\lambda/L_y = 1/5$ initial condition). We calculate the mean net velocity over one period, and calculate the norm. We plot these values, $||\bar{\boldsymbol{V}}^{\infty}||$, normalised by the value for the synchronous beating state in Figure 8.10. As reported in the literature, we find that MCW solutions produce stronger flow fields than synchronous beating. We find that the smaller wavelength solutions are able to pump more fluid, with the smallest wavelength able to pump almost 50% more fluid than the synchronous state over each period.

Finally, we explore how we can robustly construct initial conditions which lead to MCWs, and in turn, compare our results to those for two dimensional arrays of filaments undergoing symplectic waves in the literature. For this purpose, we follow the ideas presented in Chakrabarti et al. [2022] and introduce a quantity called the charge, which corresponds to the winding number of the phase along the domain length:

$$Q = \frac{1}{2\pi} \int_0^{L_y} \frac{\partial \phi(x)}{\partial x} \, dx, \tag{8.31}$$

where $\phi(x)$ denotes the phase of the filament at position $x \in [0, L_y)$. If the MCW spans the entire length of the domain, this is related to the solution wavelength via

$$Q = L_y/\lambda. \tag{8.32}$$

In Chakrabarti et al. [2022], it was found the charge is approximately conserved in these systems. As such, the authors introduce a method where MCW solutions with a particular wavelength can be obtained by superimposing solutions with smaller wavelengths, provided the initial charge is the same as the charge of the desired outcome. For instance, defining an initial condition where each half of the domain has a MCW of wavelength $\lambda = L_x/2$, but propagating

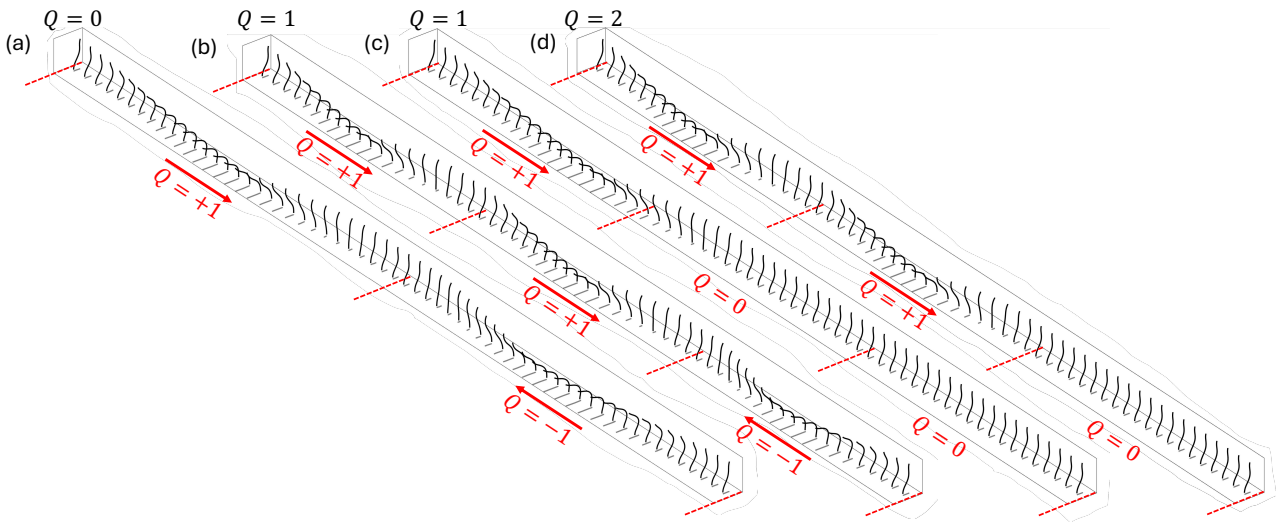


Figure 8.11: Initial conditions for four simulations following the cut-and-stitch approach for MCW generation outlined in Chakrabarti et al. [2022]. (a) Two MCWs spanning half the domain propagating in opposite directions, leading to a total charge of $Q = 0$. (b) Three MCWs, each spanning a third of the domain. Two are propagating in one direction, and the third in the opposite direction, leading to a total charge of $Q = 1$. (c) One MCW spanning a third of the domain. The rest of the filaments are in phase, leading to a total charge of $Q = 1$. (d) Two MCWs, each spanning a third of the domain. The rest of the filaments are in phase, leading to a total charge of $Q = 2$.

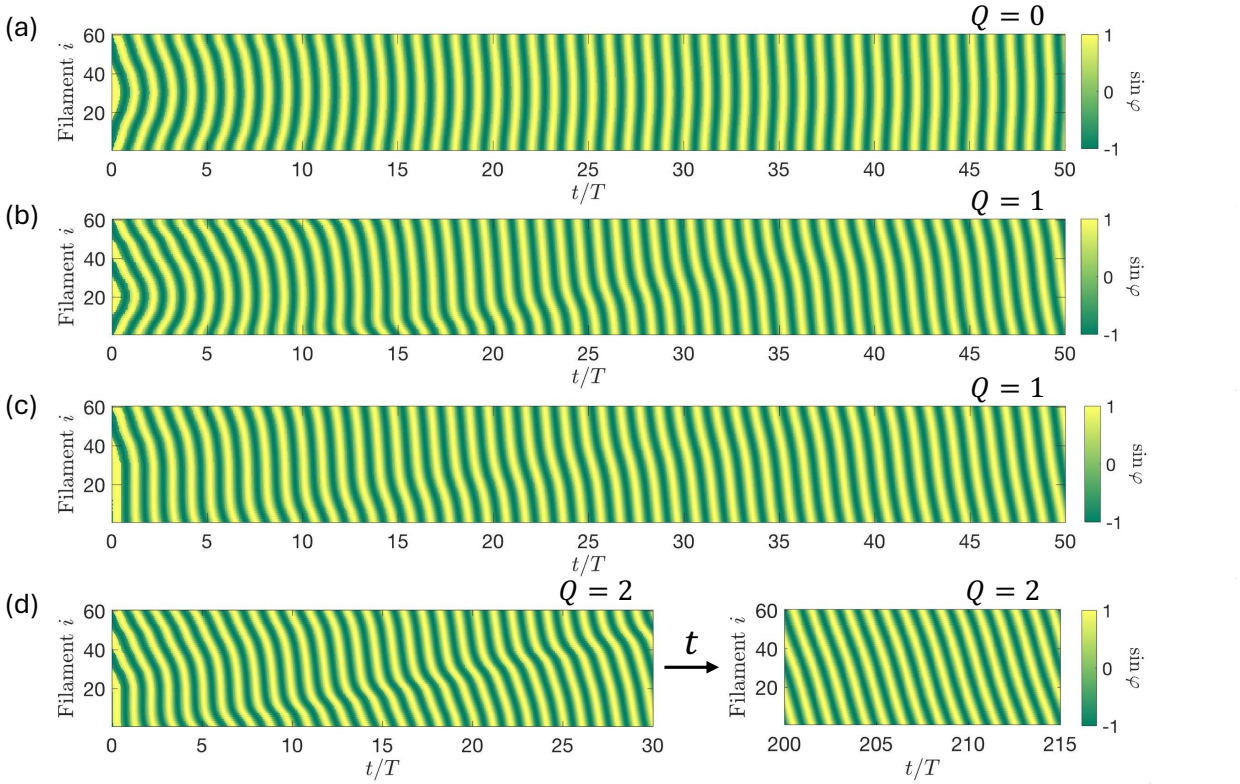


Figure 8.12: Heatmaps showing the phase evolution of each filament starting from the initial conditions outlined in Figure 8.11. We see the solutions converge to (a) a synchronous beating state with charge $Q = 0$, or MCWs with charge (b) $Q = 1$, (c) $Q = 1$ or (d) $Q = 2$. In all cases, we find that this charge is conserved over time.

in the opposite direction, as shown in Figure 8.11a, we have

$$\frac{1}{2\pi}\int_0^{L_y/2}\frac{\partial\phi(x)}{\partial x}\,dx = +1 \tag{8.33}$$

$$\frac{1}{2\pi}\int_{L_y/2}^{L_y}\frac{\partial\phi(x)}{\partial x}\,dx = -1 \tag{8.34}$$

which implies $Q = 0$. If the charge is conserved, we expect the final state to therefore be a synchronised beating state. We plot the heatmap of the phase over time for this solution in Figure 8.12a. We see that the 'bump' corresponding to these opposing MCWs smooths with time, resulting in a single synchronised beating state, agreeing with this hypothesis [Chakrabarti et al., 2022].

Similarly, we can construct different initial conditions which lead to the same state. In Figure 8.11b and c, we show two initial conditions which both correspond to a total charge of $Q = 1$. Both states split the domain into three. In the first initial condition, we concatenate three MCWs of wavelength $\lambda = L_y/3$; the first two propagating in one direction, and the final in the opposite direction, leading to a total charge of $Q = 1$. In the second, we use one MCW of wavelength $\lambda = L_y/3$, with the rest of the filaments in phase. In both cases we find that the solution converges to a MCW with wavelength $\lambda = L_y$ (and hence $Q = 1$), as shown in Figures 8.12b and c. As a final example, we also show an initial condition where we concatenate two MCWs of wavelength $\lambda = L_y/3$, with the final third of the domain containing entirely in phase filaments. This leads to a total initial charge of $Q = 2$. We plot the corresponding heatmap of phase over time in Figure 8.12d. We see that defects propagate along the domain and, after a long time, the solution converges to a MCW with wavelength $\lambda = L_y/2$, thus conserving the charge. Therefore, as reported in Chakrabarti et al. [2022], we find that MCW solutions of a particular wavelength can be robustly generated, at parameter values where such a solution exists and is stable, by concatenating other MCW solutions with a total charge equivalent to the desired final state.

8.5 Discussion

In this chapter, we have extended the previous chapters to consider arrays of asymmetrically beating filaments. To do this, we began by outlining how the methodology extends when considering triply periodic boundary conditions, opposed to an open domain where we would observe edge effects. In particular, we discussed how to bypass the zero-wavenumber condition applied in FCM and Fast FCM, as well as outlining how to include a boundary through an immersed boundary-type approach and verifying how the discretisation of the boundary affects the error. With this in hand, we used initial value problems to analyse the dynamics of arrays of filaments with various initial conditions.

We observed that synchronised beating states are always stable, regardless of the number of filaments and the filament spacing. MCWs were also found to be stable above a critical wavelength, leading to bistability between synchronous beating and MCWs for a range of parameters. Considering larger arrays, we saw further solutions are also stable, corresponding to smaller wavelength MCWs. This result is interesting, as it raises questions about the role of hydrodynamics in MCW formation. Hydrodynamic interactions clearly facilitate the existence and stability of these difference wavelength solutions, but our results suggest that there must be a further mechanism, such as basal coupling, which acts to select the correct state. Therefore, it appears likely that it is not hydrodynamics alone which are used to generate MCWs of various wavelength in these systems.

While each filament is free to evolve in response to its neighbours, we fixed the material properties of each individual filament so that the forcing was sufficient to induce oscillations, and the intrinsic curvature and bending modulus anisotropy were large enough to allow for asymmetric planar beating, ensuring net fluid flow over one period of oscillation. However it is also interesting to note how the state space changes as these parameters are changed. Preliminary results suggest that κ_0 has little effect on the formation of MCWs; these waves are still possible for specific parameter ranges, even for $\kappa_0 = 0$. Fixing the intrinsic curvature and bending modulus anisotropy but increasing the forcing, we still observe MCWs, but the individual dynamics appear to change and are no longer planar. Fixing f and κ_0 while varying β is much more com-

plex; reducing β means the whirling solutions are no longer suppressed, and so the individual dynamics of each filament changes. On a larger scale, the coordination of the filaments changes also; for smaller values of β, our results suggest that the filaments only beat in synchrony. Even in the symmetric, isotropic case (i.e. $\kappa_0 = 0, \beta = 1$), the dynamics change significantly; initial value problems indicate that the whirling states exist for wider ranges of the forcing, and that MCWs in the whirling solution are stable for particular parameter values. Further work is needed to evaluate how the collective dynamics change in response to the individual properties of each filament, and to understand exactly when these different solutions are stable.

The tracking and stability analysis tools developed and used previously in this thesis could also be used in the collective dynamics case presented here; in the same way that the internal forces in each filament are calculated as part of the numerical method, and so are not needed in the state variable defining the setup, the forces on the wall particles are also not necessary to include in the state variable. Therefore our state variable would still contain only the effective Lie algebra elements of the segments of each filament, again not including the clamped base segments, leading to a variable of unknowns of size $3N_f(N_p - 1)$ for steady states, or $3N_f(N_p - 1) + 1$ for time-periodic states. The computational expense in performing the analysis arises from the repeated application of the initial value problem with various timesteps. While this makes the analysis slower, it is certainly still possible, particularly when focusing on smaller arrays of $N_f \sim O(10)$ filaments, where MCWs have been observed. Therefore these tools could be use to investigate the onset of MCWs, in particular to evaluate the nature of the bifurcations which give rise to MCW solutions.

More generally, an interesting avenue of research would be to perform an in-depth exploration into the impact of filament spacing and heterogeneity of the filament patches, or the effect of dysfunctional filaments within the array, on the formation of MCWs. In many biological scenarios, cilia exist in dense patches; when cilia are dysfunctional, they can lead to a wide range of diseases called ciliopathies [Tobin and Beales, 2009]. Understanding the impact of immotile or incorrectly functioning cilia will enhance our understanding on how the symptoms of ciliopathies arise, and give insights into the impacts of dysfunctional cilia that are difficult to observe experimentally due to the small lengthscales involved in these problems.

Another research avenue of interest is to incorporate more biological features into the modelling choice to generate models more relevant to specific biological contexts. For instance ctenophores are marine invertebrates which use plates of synchronously beating cilia for locomotion [Tamm, 2014]. The plates of cilia coordinate such that a constant phase difference is observed between neighbouring plates, causing a MCW along the ctenophore length. Each cilium on the comb plates has a '9+3' axoneme [Afzelius, 1961], suggesting anisotropy at the individual level is important to ctenophore motility. However it is not immediately clear how the structure of these comb plates, in particular how each cilium is bound to the plate, will affect both the individual and collective dynamics. Furthermore, the lengthscales and velocity scales associated to ctenophore motion imply that the low Reynolds approximation we make throughout this thesis may no longer hold. As such, symmetry breaking is no longer a necessity for generating fluid flow. Therefore in order to model the locomotion of ctenophores, considerations would have to be taken at multiple levels - from the individual anisotropies of each filament, to the seeding and boundary conditions of each filament, to the properties of the surrounding fluid. The study of collective dynamics in the context of microswimmers, particularly those with rich and diverse ciliary superstructures such as ctenophores, would be an interesting future avenue of research.

Chapter 9

Conclusions and Outlook

In this thesis, we have explored the dynamics of an active filament model through the use of initial value problems, and the application of a combination of tools from computational dynamical systems. These tools have allowed us to track steady and time-periodic solution branches into regions where these solutions are unstable, and to evaluate their stability. In doing so, we are able to go beyond initial value problems to analyse secondary and tertiary bifurcations, which provide stronger insights into the state space. The follower force model is useful in exploring filament deformation due to its simplicity, and the self-sustained oscillations it induces. As such, the model is the ideal conduit for exploring the impact of factors such as the domain, external forcings or mechanical properties of the filament.

The initial chapters focused on the individual dynamics of follower-force driven filaments. In Chapter 5, the groundwork was set by introducing the state space in the simplest case, and highlighting how the tracking tools we use can help develop a more complete understanding of the state space. In particular, this allowed us to classify the bifurcations between existing solutions, identify new solutions and reconcile the differences between 2D and 3D dynamics. In Chapter 6, we presented the results of various modelling choices associated to differences in the material properties of the filament, the domain, or in the details of the forcing mechanism. We showed that the relative value of the twist modulus had only small effects on the filament dynamics. For low to moderate forcing values, we found that the dynamics in the absence of

the hydrodynamic wall, or with the follower forces distributed along the filament length, were qualitatively the same as the single follower force case, motivating further studies into the model presented in Chapter 5. In Chapter 7, we introduced asymmetry through an imposed intrinsic curvature, and anisotropy through direction-dependent bending moduli. Through this study, we showed that whirling solutions can be eliminated if the anisotropy or intrinsic curvature is too large. Asymmetric planar beating patterns arise and are stable in 2D through the inclusion of intrinsic curvature, but are unstable in 3D unless structural anisotropy is included. Finally, in Chapter 8 we focus on collective dynamics by modelling arrays of asymmetrically beating filaments, presenting preliminary results on bistability between different coordinated states.

While the main biological motivation for these asymmetrically beating filaments is ciliary beating, with arrays of these filaments representing carpets of cilia alike those in the respiratory system [Smith et al., 2008b], as previously mentioned the follower force model does not accurately represent the forcing mechanism driving cilia dynamics. Recall that to achieve such complex ciliary dynamics, dynein motors across the axoneme cooperate in some way to generate the overall bending (and subsequently, beating) of the cilium. There are several theories for how this coordination takes place, where dynein motors are regulated by a key property of the axoneme such as the local curvature or the relative sliding between doublets. As discussed in Chapter 1, while many numerical studies incorporate some manifestation of a regulatory feedback loop [Lindemann, 1994a,b, Dillon and Fauci, 2000, Lindemann, 2002, Dillon et al., 2003, 2007, Riedel-Kruse et al., 2007, Müller et al., 2008, Woolley, 2010, Sartori et al., 2016, Chakrabarti and Saintillan, 2019a,b, Chakrabarti et al., 2022, Gallagher et al., 2023], it is not clear exactly what the forcing mechanism or regulation should be. A natural next step for this work would be to map out the state space of a filament with a modified forcing mechanism, which more accurately encapsulate the dynein motor forcing, and perform a thorough analysis of the changes to the dynamics upon the introduction of different regulation techniques for dynein motility. Such a study would provide a comprehensive overview of the precise impact of the different regulatory feedback loops, and allow for a more conclusive answer surrounding the nature of cilia actuation.

Regardless of its direct relevance to the shear-driven actuation of cilia, the follower force model

provides an excellent set-up to explore the impact of the material properties of the filament due to its simplicity, which allows for a clearer and more direct exploration into the effect of factors such as the material properties of the filament. A natural extension of the results presented in this thesis on the level of individual filaments would be to explore different types of asymmetries and anisotropies on a single filament and how they impact the resulting dynamics. For instance, incorporating curvature dependence to material properties of the filament, such as the bending modulus [Moreau et al., 2024], or allowing for spatial variations to properties such as the bending modulus or intrinsic curvature to represent distinct parts of the axoneme, such as the transition zone in the cilium [Mercey et al., 2024].

While our discussions thus far have focused primarily on a biological context, a second application of the results presented in this thesis for individual filaments comes from biomimetic devices based on synthetic, filament-like colloidal structures that can be actuated by magnetic fields [Elgeti et al., 2015] or diffusiophoresis [Michelin, 2023]. These artificial microswimmers emulate the swimming of ciliated cells, with potential applications in biomedical sciences, for instance in targeted drug delivery [Bunea and Taboryski, 2020]. The follower force model is a simple forcing mechanism that can be used to provide design insights into recovering self-sustained buckling dynamics. Based on values in the literature for phoretic particles [Nishiguchi et al., 2018], we can extract the particle radius and velocity to be $a = 3.17\mu$m and $V = 10\mu\text{ms}^{-1}$, respectively. Using the viscosity of water, $\eta = 1$cP, we can estimate the corresponding follower force by calculating the magnitude of the force required to keep a self-propelled colloid fixed, $F = 6\pi\eta aV \approx 5.98 \times 10^{-13}$ N. Using the bending modulus reported for an artificial microswimmer in Dreyfus et al. [2005], $K_B \approx 3.3 \times 10^{-22}\text{Jm}^{-1}$, we can extract the corresponding non-dimensional follower force for a chain of length $L = 50a$, i.e. a chain constructed of 25 particles, to be $f \approx 45.5$. This value is above the critical threshold, and so it is plausible that time-dependent dynamics can be induced in colloidal chains using available mechanisms. Furthermore, in the case where all the particles in the chain are active, the more relevant model would be a distribution of follower forces along the filament length which, following the same non-dimensionalisation as above, has a critical buckling value of $f \approx 4$. In this case, we anticipate the chain would be able to access different states associated with higher forcing

values. Thus, we suspect that one could, for instance, introduce asymmetries and anisotropies in the ways outlined in Chapter 7 to uncover similar dynamics similar to those reported in this chapter, or use the collective dynamics results presented in Chapter 8 to generate arrays of synthetic cilia which beat in metachronal formations to carry large cargo, or pump fluid.

On the collective level, the results presented in Chapter 8 raise questions about the mechanisms controlling MCW formation and stability. Many studies suggest that MCWs are possibly a result of the hydrodynamic interactions between neighbouring filaments [Gueron et al., 1997, Brumley et al., 2015, Meng et al., 2021]. Our results agree, showing that MCWs are possible for particular ranges of the parameters. However the bistability we observe, particularly in larger arrays of filaments where MCWs of several wavelengths are all stable at the same parameter values, indicates that while hydrodynamics certainly facilitate these solutions, there must be an additional mechanism which works to select the appropriate solution. *Chlamydomonas* are known to control its two flagella through basal coupling [Quaranta et al., 2015, Wan and Goldstein, 2016], and it is plausible that a similar tool is present in other ciliated systems. Further studies will be necessary to investigate the impact of these mechanisms on MCW formation and regulation.

There are many natural extensions of the results presented in Chapter 8. For instance, further initial value problems could be run which investigate the impact of inter-filament spacing, the heterogeneity of the filament seeding, or to consider larger arrays of filaments. As an example, it would be interesting to explore the stability and basins of attraction of different metachronal solutions in an $N \times N$ array of filaments, and to explore whether further solutions exist which are hidden in the $1 \times N$ and $N \times 1$ arrays considered in this work. The application of the tracking and stability analysis tools utilised in previous chapters would also allow us to draw conclusions about the onset of these different solutions, and provide a clearer view of the solution state space in collections of filaments.

The impact of individual filament dynamics on the collective scale are also of interest; not only in terms of the three parameters which define the individual beat, f, κ_0 and β, but also in terms of parameters which were fixed in our case, such as the preferred beating direction.

How would the dynamics change if the beat direction no longer aligned with the $\hat{\boldsymbol{x}}-$axis, but rather at an angle in the $(x, y)-$plane? What if the filaments have different preferred bending directions? Exploring how synchronisation changes in response to these factors will provide a stronger foundation to our knowledge of collective coordination, and more broadly, to the role of individuality on a larger scale.

More generally, there are several extensions to this work that would allow us to align with different biological contexts more closely. One simple example is the geometry of the surface - many surfaces on which cilia exist are not entirely planar. For instance, a spherical surface might be desirable to explore how cilia coordinate outside [Westwood and Keaveny, 2021] or inside [De Canio, 2019] curved surfaces. To incorporate a different surface, the methodology described in Section 8.2.1 mostly holds. The surface could be discretised into several particles, and the filaments would now be seeded above the new surface. Studies would need to be performed to ensure the size and density of the particles is sufficient for the discretisation error to be small. Alternatively, including passive filaments above an array of active filaments could represent mucin fibres in mucus, and therefore be used to model processes such as mucociliary clearance [Smith et al., 2008a]. Further studies into the fluid pumping abilities of these arrays of cilia would also be an interesting avenue, particularly in the presence of immotile filaments as a way to understand the impact of defective cilia and how these can give rise to ciliopathies [Tobin and Beales, 2009].

To summarise, the results presented in this thesis provide the groundwork underpinning instability-driven filaments in the particular context of the follower force model. We have thoroughly mapped out the state space of the simplest incarnation of this model and some of its extensions, crucially providing a simple mechanism to generate asymmetric beating patterns in instability-driven filaments, and presented preliminary results on the collective dynamics of these filaments. Overall there are several interesting, and broad, potential avenues of research that build on this work. For instance, using these computational dynamical systems tools on different forcing mechanisms, to generate a stronger understanding of the exact impact of these different modelling choices. Another example is to incorporate more biological effects, such as additional passive fibres, or immotile cilia, to investigate how these defects can affect bodily

processes. Alternatively, the key insights from these model could be applied to inform designs in biomimetics, with applications in artificial microswimmers, for instance appropriate for drug delivery. Considering the extensions outlined above will not only help provide a stronger fundamental understanding of cilia actuation, but also of the role of filament dynamics in fluid propulsion and life processes more generally.

Bibliography

B. Afzelius. Electron microscopy of the sperm tail; results obtained with a new fixative. *The Journal of biophysical and biochemical cytology*, 5(2), 1959. doi: 10.1083/jcb.5.2.269.

B. Afzelius. The fine structure of the cilia from ctenophore swimming-plates. *The Journal of biophysical and biochemical cytology*, 9, 1961. doi: 10.1083/jcb.9.2.383.

B. Afzelius. A human syndrome caused by immotile cilia. *Science*, 193(4250), 1976. doi: 10.1126/science.1084576.

C. Ainsworth. Cilia: Tails of the unexpected. *Nature*, 448(7154), 2007. doi: 10.1038/448638a.

R. D. Allen, D. G. Weiss, J. H. Hayden, D. T. Brown, H. Fujiwake, and M. Simpson. Gliding movement of and bidirectional transport along single native microtubules from squid axoplasm: Evidence for an active role of microtubules in cytoplasmic transport. *Journal of Cell Biology*, 100(5), 1985. doi: 10.1083/jcb.100.5.1736.

F. Balboa Usabiaga, B. Kallemov, B. Delmotte, A. Bhalla, B. Griffith, and A. Donev. Hydrodynamics of suspensions of passive and active rigid particles: a rigid multiblob approach. *Communications in Applied Mathematics and Computational Science*, 11(2):217–296, 2017. doi: 10.2140/camcos.2016.11.217.

P. V. Bayly and S. K. Dutcher. Steady dynein forces induce flutter instability and propagating waves in mathematical models of flagella. *Journal of the Royal Society, Interface*, 13(123), 2016. doi: 10.1098/rsif.2016.0523.

R. R. Bennett. Direction selection of metachronal waves in hydrodynamic coordination of cilia. *Physical Review E*, 111(3):034402, 2025. doi: 10.1103/PhysRevE.111.034402.

R. R. Bennett and R. Golestanian. A steering mechanism for phototaxis in Chlamydomonas. *Journal of the Royal Society Interface*, 12(104), 2015. doi: 10.1098/rsif.2014.1164.

J. R. Blake. A note on the image system for a stokeslet in a no-slip boundary. *Mathematical Proceedings of the Cambridge Philosophical Society*, 70(2), 1971a. doi: 10.1017/S0305004100049902.

J. R. Blake. A spherical envelope approach to ciliary propulsion. *Journal of Fluid Mechanics*, 46(1), 1971b. doi: 10.1017/S002211207100048X.

L. Bourdieu, T. Duke, M. B. Elowitz, D. A. Winkelmann, S. Leibler, and A. Libchaber. Spiral defects in motility assays: A measure of motor protein force. *Physical Review Letters*, 75(1), 1995. doi: 10.1103/PhysRevLett.75.176.

D. Bray. *Cell Movements : From Molecules to Motility*. Garland Science, 2nd edition, 2000.

C. Brennen and H. Winet. Fluid Mechanics of Propulsion by Cilia and Flagella. *Annual Review of Fluid Mechanics*, 9(1):339–398, 1977. doi: 10.1146/ANNUREV.FL.09.010177.002011.

C. J. Brokaw. Stochastic simulation of processive and oscillatory sliding using a two-headed model for axonemal dynein. *Cell Motility and the Cytoskeleton*, 47(2), 2000. doi: 10.1002/1097-0169(200010)47:2¡108::aid-cm2¿3.3.co;2-z.

C. G. Broyden. A Class of Methods for Solving Nonlinear Simultaneous Equations. *Mathematics of Computation*, 19(92), 1965. doi: 10.2307/2003941.

D. R. Brumley, M. Polin, T. J. Pedley, and R. E. Goldstein. Hydrodynamic Synchronization and Metachronal Waves on the Surface of the Colonial Alga Volvox carteri. *Physical Review Letters*, 109(26):268102, 2012. doi: 10.1103/PhysRevLett.109.268102.

D. R. Brumley, M. Polin, T. J. Pedley, and R. E. Goldstein. Metachronal waves in the flagellar beating of Volvox and their hydrodynamic origin. *Journal of the Royal Society Interface*, 12 (108), 2015. doi: 10.1098/rsif.2014.1358.

A. I. Bunea and R. Taboryski. Recent advances in microswimmers for biomedical applications. *Micromachines*, 11(12), 2020. doi: 10.3390/mi11121048.

J. F. Cass and H. Bloomfield-Gadêlha. The reaction-diffusion basis of animated patterns in eukaryotic flagella. *Nature Communications*, 14(1), 2023. doi: 10.1038/s41467-023-40338-2.

B. Chakrabarti and D. Saintillan. Hydrodynamic Synchronization of Spontaneously Beating Filaments. *Physical Review Letters*, 123(20), 2019a. doi: 10.1103/PhysRevLett.123.208101.

B. Chakrabarti and D. Saintillan. Spontaneous oscillations, beating patterns, and hydrodynamics of active microfilaments. *Physical Review Fluids*, 4(4), 2019b. doi: 10.1103/PhysRevFluids.4.043102.

B. Chakrabarti, S. Fürthauer, and M. J. Shelley. A multiscale biophysical model gives quantized metachronal waves in a lattice of beating cilia. *Proceedings of the National Academy of Sciences*, 119(4), 2022. doi: 10.1073/pnas.2113539119.

P. Chełminiak, J. M. Dixon, and J. A. Tuszyński. Torsional elastic deformations of microtubules within continuous sheet model. *The European Physical Journal E 2010 31:2*, 31(2):215–227, 2010. doi: 10.1140/EPJE/I2010-10562-X.

M. A. Chilvers and C. O'Callaghan. Analysis of ciliary beat pattern and beat frequency using digital high speed imaging: comparison with the photomultiplier and photodiode methods. *Thorax*, 55(4):314–317, 2000. doi: 10.1136/THORAX.55.4.314.

C. Cho and R. D. Vale. The mechanism of dynein motility: Insight from crystal structures of the motor domain. *Biochimica et Biophysica Acta - Molecular Cell Research*, 1823(1), 2012. doi: 10.1016/j.bbamcr.2011.10.009.

B. Clarke, Y. Hwang, and E. E. Keaveny. Bifurcations and nonlinear dynamics of the follower force model for active filaments. *Physical Review Fluids*, 9(7):073101, 2024. doi: 10.1103/PHYSREVFLUIDS.9.073101/FIGURES/11/MEDIUM.

B. Clarke, Y. Hwang, and E. E. Keaveny. Structural Anisotropy Stabilises Asymmetric Beating in Instability Driven Filaments. *Journal of Fluid Mechanics*, 1023, 2025. doi: 10.1017/jfm.2025.10787.

G. M. Cooper. *The Cell: A Molecular Approach*. Sinauer Associates, 2nd edition, 2000.

R. Cortez. The method of regularized stokeslets. *SIAM Journal on Scientific Computing*, 23 (4), 2001. doi: 10.1137/S106482750038146X.

R. Cortez, L. Fauci, and A. Medovikov. The method of regularized Stokeslets in three dimensions: Analysis, validation, and application to helical swimming. *Physics of Fluids*, 17(3), 2005. doi: 10.1063/1.1830486.

Dartmouth Electron Microscope Facility, Dartmouth College. URL `www.dartmouth.edu/emlab/`.

G. De Canio. *Motion of filaments induced by molecular motors: from individual to collective dynamics*. PhD thesis, University of Cambridge, 2019.

G. De Canio, E. Lauga, and R. E. Goldstein. Spontaneous oscillations of elastic filaments induced by molecular motors. *Journal of the Royal Society Interface*, 14(136), 2017. doi: 10.1098/rsif.2017.0491.

J. E. Dennis and R. B. Schnabel. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. SIAM, 1996. doi: 10.1137/1.9781611971200.

R. H. Dillon and L. J. Fauci. An integrative model of internal axoneme mechanics and external fluid dynamics in ciliary beating. *Journal of Theoretical Biology*, 207(3), 2000. doi: 10.1006/jtbi.2000.2182.

R. H. Dillon, L. J. Fauci, and C. Omoto. Mathematical modeling of axoneme mechanics and fluid dynamics in ciliary and sperm motility. *Dynamics of Continuous, Discrete and Impulsive Systems Series A: Mathematical Analysis*, 10(5), 2003.

R. H. Dillon, L. J. Fauci, C. Omoto, and X. Yang. Fluid dynamic models of flagellar and ciliary beating. *Annals of the New York Academy of Sciences*, 1101, 2007. doi: 10.1196/annals.1389.016.

Y. Ding, J. C. Nawroth, M. J. McFall-Ngai, and E. Kanso. Mixing and transport by ciliary carpets: A numerical study. *Journal of Fluid Mechanics*, 743, 2014. doi: 10.1017/jfm.2014.36.

K. Drescher, R. E. Goldstein, and I. Tuval. Fidelity of adaptive phototaxis. *Proceedings of the National Academy of Sciences of the United States of America*, 107(25), 2010. doi: 10.1073/pnas.1000901107.

R. Dreyfus, J. Baudry, M. L. Roper, M. Fermigier, H. A. Stone, and J. Bibette. Microscopic artificial swimmers. *Nature*, 437(7060), 2005. doi: 10.1038/nature04090.

J. Elgeti, R. G. Winkler, and G. Gompper. Physics of microswimmers - Single particle motion and collective behavior: A review. *Reports on Progress in Physics*, 78(5), 2015. doi: 10.1088/0034-4885/78/5/056601.

S. Fatehiboroujeni, A. Gopinath, and S. Goyal. Nonlinear oscillations induced by follower forces in prestressed clamped rods subjected to drag. *Journal of Computational and Nonlinear Dynamics*, 13(12), 2018. doi: 10.1115/1.4041681.

S. Fatehiboroujeni, A. Gopinath, and S. Goyal. Three-dimensional nonlinear dynamics of prestressed active filaments: Flapping, swirling, and flipping. *Physical Review E*, 103(1), 2021. doi: 10.1103/PhysRevE.103.013005.

R. Faubel, C. Westendorf, E. Bodenschatz, and G. Eichele. Cilia-based flow network in the brain ventricles. *Science*, 353(6295), 2016. doi: 10.1126/science.aae0450.

A. N. Fehr, C. L. Asbury, and S. M. Block. Kinesin steps do not alternate in size. *Biophysical Journal*, 94(3), 2008. doi: 10.1529/biophysj.107.126839.

H. Felgner, R. Frank, and M. Schliwa. Flexural rigidity of microtubules measured with the use of optical tweezers. *Journal of Cell Science*, 109(2):509–516, 1996. doi: 10.1242/JCS.109.2.509.

Y. Feng and S. Mitran. Data-driven reduced-order model of microtubule mechanics. *Cytoskeleton*, 75(2), 2018. doi: 10.1002/cm.21419.

Y. Fily, P. Subramanian, T. M. Schneider, R. Chelakkot, and A. Gopinath. Buckling instabilities and spatio-temporal dynamics of active elastic filaments. *Journal of the Royal Society Interface*, 17(165), 2020. doi: 10.1098/rsif.2019.0794.

S. Forth and T. M. Kapoor. The mechanics of microtubule networks in cell division. *Journal of Cell Biology*, 216(6), 2017. doi: 10.1083/jcb.201612064.

L. Fourriere, A. J. Jimenez, F. Perez, and G. Boncompain. The role of microtubules in secretory protein transport. *Journal of Cell Science*, 133(2), 2020. doi: 10.1242/jcs.237016.

L. A. Fox and W. S. Sale. Direction of force generated by the inner row of dynein arms on flagellar microtubules. *The Journal of cell biology*, 105(4), 1987. doi: 10.1083/jcb.105.4.1781.

B. M. Friedrich and F. Jülicher. Flagellar synchronization independent of hydrodynamic interactions. *Physical Review Letters*, 109(13), 2012. doi: 10.1103/PhysRevLett.109.138102.

G. R. Fulford and J. R. Blake. Muco-ciliary transport in the lung. *Journal of Theoretical Biology*, 121(4), 1986. doi: 10.1016/S0022-5193(86)80098-4.

E. A. Gaffney, H. Gadêlha, D. J. Smith, J. R. Blake, and J. C. Kirkman-Brown. Mammalian sperm motility: Observation and theory. *Annual Review of Fluid Mechanics*, 43, 2011. doi: 10.1146/annurev-fluid-121108-145442.

M. T. Gallagher, J. C. Kirkman-Brown, and D. J. Smith. Axonemal regulation by curvature explains sperm flagellar waveform modulation. *PNAS Nexus*, 2(3), 2023. doi: 10.1093/pnasnexus/pgad072.

P. Garbacz. Introduction to Molecular Chirality. In *Physical Principles of Chirality in NMR*, volume 34, chapter 1. Royal Society of Chemistry, 2024. ISBN 978-1-83916-548-1. doi: 10.1039/BK9781839169588-00001.

A. Gennerich and R. D. Vale. Walking the walk: how kinesin and dynein coordinate their steps. *Current Opinion in Cell Biology*, 21(1), 2009. doi: 10.1016/j.ceb.2008.12.002.

A. Ghorbani and A. Najafi. Symplectic and antiplectic waves in an array of beating cilia attached to a closed body. *Physical review. E*, 95(5-1), 2017. doi: 10.1103/PhysRevE.95.052412.

I. R. Gibbons. Cilia and flagella of eukaryotes. *Journal of Cell Biology*, 91(3 II), 1981. doi: 10.1083/jcb.91.3.107s.

I. R. Gibbons and A. J. Rowe. Dynein: A protein with adenosine triphosphatase activity from cilia. *Science*, 149(3682), 1965. doi: 10.1126/science.149.3682.424.

W. Gilpin, V. N. Prakash, and M. Prakash. Vortex arrays and ciliary tangles underlie the feeding-swimming trade-off in starfish larvae. *Nature Physics*, 13(4), 2017. doi: 10.1038/nphys3981.

W. Gilpin, M. S. Bull, and M. Prakash. The multiscale physics of cilia and flagella. *Nature Reviews Physics*, 2(2), 2020. doi: 10.1038/s42254-019-0129-0.

F. Gittes, B. Mickey, J. Nettleton, and J. Howard. Flexural rigidity of microtubules and actin filaments measured from thermal fluctuations in shape. *The Journal of cell biology*, 120(4): 923–934, 1993. doi: 10.1083/JCB.120.4.923.

F. Gittes, E. Meyhöfer, S. Baek, and J. Howard. Directional loading of the kinesin motor molecule as it buckles a microtubule. *Biophysical Journal*, 70(1), 1996. doi: 10.1016/S0006-3495(96)79585-1.

D. M. Goedecke and T. C. Elston. A model for the oscillatory motion of single dynein molecules. *Journal of Theoretical Biology*, 232(1), 2005. doi: 10.1016/j.jtbi.2004.07.017.

R. E. Goldstein and J. W. van de Meent. A physical perspective on cytoplasmic streaming. *Interface Focus*, 5(4), 2015. doi: 10.1098/rsfs.2015.0030.

M. Golubitsky and I. Stewart. Hopf Bifurcation in the presence of symmetry. *Archive for Rational Mechanics and Analysis*, 87(2), 1985. doi: 10.1007/BF00280698.

J. Gonçalves and L. Pelletier. The ciliary transition zone: Finding the pieces and assembling the gate. *Molecules and Cells*, 40(4), 2017. doi: 10.14348/molcells.2017.0054.

G. Goshima and R. D. Vale. The roles of microtubule-based motor proteins in mitosis: Comprehensive RNAi analysis in the Drosophila S2 cell line. *The Journal of Cell Biology*, 162(6), 2003. doi: 10.1083/jcb.200303022.

J. Gray and G. J. Hancock. The Propulsion of Sea-Urchin Spermatozoa. *Journal of Experimental Biology*, 32(4), 1955. doi: 10.1242/jeb.32.4.802.

S. P. Gross, M. Vershinin, and G. T. T. Shubeita. Cargo Transport: Two Motors Are Sometimes Better Than One. *Current Biology*, 17(12):R478–R486, 2007. doi: 10.1016/j.cub.2007.04.025.

H. Gu, Q. Boehler, H. Cui, E. Secchi, G. Savorana, C. De Marco, S. Gervasoni, Q. Peyron, T. Y. Huang, S. Pane, A. M. Hirt, D. Ahmed, and B. J. Nelson. Magnetic cilia carpets with programmable metachronal waves. *Nature Communications*, 11(1), 2020. doi: 10.1038/s41467-020-16458-4.

E. Guazzelli and J. F. Morris. *A physical introduction to suspension dynamics*. Cambridge University Press, 2011. doi: 10.1017/CBO9780511894671.

P. Guedes-Dias and E. L. Holzbaur. Axonal transport: Driving synaptic function. *Science*, 366 (6462), 2019. doi: 10.1126/science.aaw9997.

S. Gueron, K. Levit-Gurevich, N. Liron, and J. J. Blum. Cilia internal mechanism and metachronal coordination as the result of hydrodynamical coupling. *Proceedings of the National Academy of Sciences of the United States of America*, 94(12), 1997. doi: 10.1073/pnas.94.12.6001.

A. Guldbrandsen, H. Vethe, Y. Farag, E. Oveland, H. Garberg, M. Berle, K.-M. Myhr, J. A. Opsahl, H. Barsnes, and F. S. Berven. In-depth characterization of the cerebrospinal fluid (CSF) proteome displayed through the CSF proteome resource (CSF-PR). *Molecular & cellular proteomics : MCP*, 13(11):3152–63, 2014. doi: 10.1074/mcp.M114.038554.

H. O. Gutzeit and R. Koppa. Time-lapse film analysis of cytoplasmic streaming during late oogenesis of Drosophila. *Development*, 67(1):101–111, 1982. doi: 10.1242/DEV.67.1.101.

J. Han and C. S. Peskin. Spontaneous oscillation and fluid–structure interaction of cilia. *Proceedings of the National Academy of Sciences of the United States of America*, 115(17), 2018. doi: 10.1073/pnas.1712042115.

A. J. Hanson. *Visualizing Quaternions*. Morgan Kaufmann, 2006.

M. A. Hartman and J. A. Spudich. The myosin superfamily at a glance. *Journal of Cell Science*, 125:1627–1632, 2012. doi: 10.1242/jcs.094300.

A. Hashemi, R. P. Peláez, S. Natesh, B. Sprinkle, O. Maxian, Z. Gan, and A. Donev. Computing hydrodynamic interactions in confined doubly periodic geometries in linear time. *Journal of Chemical Physics*, 158(15), 2023. doi: 10.1063/5.0141371.

G. Herrmann and R. W. Bungay. On the Stability of Elastic Systems Subjected to Nonconservative Forces. *Journal of Applied Mechanics*, 31(3):435–440, 1964. doi: 10.1115/1.3629660.

N. Hirokawa, Y. Tanaka, and Y. Okada. Left-right determination: involvement of molecular motor KIF3, cilia, and nodal flow. *Cold Spring Harbor Perspectives in Biology*, 1(1), 2009. doi: 10.1101/cshperspect.a000802.

N. Hirokawa, S. Niwa, and Y. Tanaka. Molecular motors in neurons: Transport mechanisms and roles in brain function, development, and disease. *Neuron*, 68(4), 2010. doi: 10.1016/j.neuron.2010.09.039.

E. L. Holzbaur and S. S. Scherer. Microtubules, Axonal Transport, and Neuropathy. *New England Journal of Medicine*, 365(24), 2011. doi: 10.1056/nejmcibr1112481.

S. H. Hwang, M. Litt, and W. C. Forsman. Rheological properties of mucus. *Rheologica Acta*, 8(4), 1969. doi: 10.1007/BF01976227.

R. E. Isele-Holder, J. Elgeti, and G. Gompper. Self-propelled worm-like filaments: spontaneous spiral formation, structure, and dynamics. *Soft Matter*, 11(36), 2015. doi: 10.1039/c5sm01683e.

S. Ishijima and Y. Hiramoto. Flexural Rigidity of Echinoderm Sperm Flagella. *Cell Structure and Function*, 19(6), 1994. doi: 10.1247/csf.19.349.

T. Ishikawa. Axoneme structure from motile cilia. *Cold Spring Harbor Perspectives in Biology*, 9(1), 2017. doi: 10.1101/cshperspect.a028076.

T. Ishikawa. Fluid Dynamics of Squirmers and Ciliated Microorganisms. *Annual Review of Fluid Mechanics*, 56, 2024. doi: 10.1146/annurev-fluid-121021-042929.

T. Ishikawa, H. Ueno, T. Omori, and K. Kikuchi. Cilia and centrosomes: Ultrastructural and mechanical perspectives. *Seminars in Cell and Developmental Biology*, 110, 2021. doi: 10.1016/j.semcdb.2020.03.007.

R. E. Johnson. An improved slender-body theory for Stokes flow. *Journal of Fluid Mechanics*, 99(2), 1980. doi: 10.1017/S0022112080000687.

M. Jorissen and M. Jaspers. Cilia, ciliary movement, and mucociliary transport. In *Nasal Physiology and Pathophysiology of Nasal Disorders*. Springer, 2023. doi: 10.1007/978-3-031-12386-3_3.

A. V. Kanale, F. Ling, H. Guo, S. Fürthauer, and E. Kanso. Spontaneous phase coordination and fluid pumping in model ciliary carpets. *Proceedings of the National Academy of Sciences of the United States of America*, 119(45), 2022. doi: 10.1073/pnas.2214413119.

R. Kawamura, A. Kakugo, Y. Osada, and J. P. Gong. Microtubule bundle formation driven by ATP: The effect of concentrations of kinesin, streptavidin and microtubules. *Nanotechnology*, 21(14), 2010. doi: 10.1088/0957-4484/21/14/145603.

J. B. Keller and S. I. Rubinow. Slender-body theory for slow viscous flow. *Journal of Fluid Mechanics*, 75(4), 1976. doi: 10.1017/S0022112076000475.

S. Kim and S. J. Karrila. *Microhydrodynamics*. Elsevier, 1991. ISBN 9780750691734. doi: 10.1016/C2013-0-04644-0.

A. Kis, S. Kasas, B. Babić, A. J. Kulik, W. Benoît, G. A. Briggs, C. Schönenberger, S. Catsicas, and L. Forró. Nanomechanics of Microtubules. *Physical Review Letters*, 89(24), 2002. doi: 10.1103/PhysRevLett.89.248101.

N. Klughammer, J. Bischof, N. D. Schnellbächer, A. Callegari, P. Lénárt, and U. S. Schwarz. Cytoplasmic flows in starfish oocytes are fully determined by cortical contractions. *PLoS Computational Biology*, 14(11), 2018. doi: 10.1371/journal.pcbi.1006588.

S. Klumpp and R. Lipowsky. Cooperative cargo transport by several molecular motors. *Proceedings of the National Academy of Sciences of the United States of America*, 102(48), 2005. doi: 10.1073/pnas.0507363102.

D. E. Kretschmann. Mechanical Properties of Wood. In *Wood Handbook - Wood as an engineering material*, chapter 5. U.S. Department of Agriculture, 2010.

W. Kühne. *Untersuchungen über das Protoplasma und die Contractilität*. Leipzig: W. Engelmann, 1864.

S. Kunishima and H. Saito. Advances in the understanding of MYH9 disorders. *Current Opinion in Hematology*, 17(5), 2010. doi: 10.1097/MOH.0b013e32833c069c.

C. Kural, H. Kim, S. Syed, G. Goshima, V. I. Gelfand, and P. R. Selvin. Cell Biology: Kinesin and dynein move a peroxisome in vivo: A tug-of-war or coordinated movement? *Science*, 308(5727), 2005. doi: 10.1126/science.1108408.

Y. Kuramoto. *Chemical Oscillations, Waves and Turbulence.* Springer Berlin, Heidelberg, 1984. doi: 10.1007/978-3-642-69689-3.

Y. A. Kuznetsov. Elements of Applied Bifurcation Theory, Second Edition. *Kuznetsov, Y.A. 1998. Elements of Applied Bifurcation Theory*, 112, 1998.

S. K. Lai, Y.-Y. Wang, D. Wirtz, and J. Hanes. Micro- and macrorheology of mucus. *Advanced Drug Delivery Reviews*, 61(2), 2009. doi: 10.1016/j.addr.2008.09.012.

A. T. Lam, C. Curschellas, D. Krovvidi, and H. Hess. Controlling self-assembly of microtubule spools via kinesin motor density. *Soft Matter*, 10(43), 2014. doi: 10.1039/c4sm01518e.

L. D. Landau and E. M. Lifshitz. *Theory of Elasticity*, volume 7. Elsevier, 3rd edition, 1986. ISBN 9780080570693. doi: 10.1016/C2009-0-25521-8.

A. Laskar and R. Adhikari. Filament actuation by an active colloid at low Reynolds number. *New Journal of Physics*, 19(3), 2017. doi: 10.1088/1367-2630/aa5f80.

E. Lauga. *The Fluid Dynamics of Cell Motility*. Cambridge University Press, 2020. doi: 10.1017/9781316796047.

E. Lauga and T. R. Powers. The hydrodynamics of swimming microorganisms. *Reports on Progress in Physics*, 72(9), 2009. doi: 10.1088/0034-4885/72/9/096601.

M. C. Ledbetter and K. R. Porter. A "microtubule" in plant cell fine structure. *The Journal of Cell Biology*, 19(1):239, 1963. doi: 10.1083/JCB.19.1.239.

C. Leduc, O. Campàs, K. B. Zeldovich, A. Roux, P. Jolimaitre, L. Bourel-Bonnet, B. Goud, J. F. Joanny, P. Bassereau, and J. Prost. Cooperative extraction of membrane nanotubes by molecular motors. *Proceedings of the National Academy of Sciences of the United States of America*, 101(49), 2004. doi: 10.1073/pnas.0406598101.

J. Leite, D. S. Osorio, A. F. Sobral, A. M. Silva, and A. X. Carvalho. Network contractility during cytokinesis—from molecular to global views. *Biomolecules*, 9(5), 2019. doi: 10.3390/biom9050194.

M. J. Lighthill. On the squirming motion of nearly spherical deformable bodies through liquids at very small reynolds numbers. *Communications on Pure and Applied Mathematics*, 5(2), 1952. doi: 10.1002/cpa.3160050201.

J. Lin and D. Nicastro. Asymmetric distribution and spatial switching of dynein activity generates ciliary motility. *Science*, 360(6387), 2018. doi: 10.1126/science.aar1968.

C. B. Lindemann. A model of flagellar and ciliary functioning which uses the forces transverse to the axoneme as the regulator of dynein activation. *Cell Motility and the Cytoskeleton*, 29 (2), 1994a. doi: 10.1002/cm.970290206.

C. B. Lindemann. A 'geometric clutch' hypothesis to explain oscillations of the axoneme of cilia and flagella. *Journal of Theoretical Biology*, 168(2), 1994b. doi: 10.1006/jtbi.1994.1097.

C. B. Lindemann. Geometric clutch model version 3: The role of the inner and outer arm dyneins in the ciliary beat. *Cell Motility and the Cytoskeleton*, 52(4), 2002. doi: 10.1002/cm.10049.

C. B. Lindemann and K. A. Lesich. Functional anatomy of the mammalian sperm flagellum. *Cytoskeleton*, 73(11), 2016. doi: 10.1002/cm.21338.

F. Ling, H. Guo, and E. Kanso. Instability-driven oscillations of elastic microfilaments. *Journal of the Royal Society Interface*, 15(149), 2018. doi: 10.1098/rsif.2018.0594.

K. G. Link, R. D. Guy, B. Thomases, and P. E. Arratia. Effect of fluid elasticity on the emergence of oscillations in an active elastic filament. *Journal of The Royal Society Interface*, 21(214), 2024. doi: 10.1098/rsif.2024.0046.

S. Lomholt and M. R. Maxey. Force-coupling method for particulate two-phase flow: Stokes flow. *Journal of Computational Physics*, 184(2), 2003. doi: 10.1016/S0021-9991(02)00021-9.

R. Mallik, B. C. Carter, S. A. Lex, S. J. King, and S. P. Gross. Cytoplasmic dynein functions as a gear in response to load. *Nature*, 427(6975), 2004. doi: 10.1038/nature02293.

J. L. Maravillas-Montero and L. Santos-Argumedo. The myosin family: unconventional roles of actin-dependent molecular motors in immune cells. *Journal of Leukocyte Biology*, 91(1), 2012. doi: 10.1189/jlb.0711335.

Marcos, H. C. Fu, T. R. Powers, and R. Stocker. Bacterial rheotaxis. *Proceedings of the National Academy of Sciences of the United States of America*, 109(13), 2012. doi: 10.1073/pnas.1120955109.

W. F. Marshall and S. Nonaka. Cilia: Tuning in to the Cell's Antenna. *Current Biology*, 16 (15), 2006. doi: 10.1016/j.cub.2006.07.012.

P. K. Mattila and P. Lappalainen. Filopodia: Molecular architecture and cellular functions. *Nature Reviews Molecular Cell Biology*, 9(6), 2008. doi: 10.1038/nrm2406.

M. R. Maxey and B. K. Patel. Localized force representations for particles sedimenting in stokes flow. *International Journal of Multiphase Flow*, 27(9), 2001. doi: 10.1016/S0301-9322(01)00014-3.

F. Meng, R. R. Bennett, N. Uchida, and R. Golestanian. Conditions for metachronal coordination in arrays of model cilia. *Proceedings of the National Academy of Sciences of the United States of America*, 118(32), 2021. doi: 10.1073/pnas.2102828118.

O. Mercey, S. Mukherjee, P. Guichard, and V. Hamel. The molecular architecture of the ciliary transition zones. *Current Opinion in Cell Biology*, 88:102361, 2024. doi: 10.1016/j.ceb.2024.102361.

S. Michelin. Self-Propulsion of Chemically Active Droplets. *Annual Review of Fluid Mechanics*, 55, 2023. doi: 10.1146/annurev-fluid-120720-012204.

E. Milana, R. Zhang, M. R. Vetrano, S. Peerlinck, M. de Volder, P. R. Onck, D. Reynaerts, and B. Gorissen. Metachronal patterns in artificial cilia for low Reynolds number fluid propulsion. *Science Advances*, 6(49), 2020. doi: 10.1126/sciadv.abd2508.

C. Moreau, B. J. Walker, R. N. Poon, D. Soto, D. I. Goldman, E. A. Gaffney, and K. Y. Wan. Minimal design of a synthetic cilium. *Physical Review Research*, 6(4):L042061, 2024. doi: 10.1103/PhysRevResearch.6.L042061.

M. J. Müller, S. Klumpp, and R. Lipowsky. Tug-of-war as a cooperative mechanism for bidirectional cargo transport by molecular motors. *Proceedings of the National Academy of Sciences of the United States of America*, 105(12), 2008. doi: 10.1073/pnas.0706825105.

D. Needleman and M. Shelley. The stormy fluid dynamics of the living cell. *Physics Today*, 72 (9):32–38, 2019. doi: 10.1063/PT.3.4292.

T. Niedermayer, B. Eckhardt, and P. Lenz. Synchronization, phase locking, and metachronal wave formation in ciliary chains. *Chaos*, 18(3), 2008. doi: 10.1063/1.2956984.

D. Nishiguchi, J. Iwasawa, H. R. Jiang, and M. Sano. Flagellar dynamics of chains of active Janus particles fueled by an AC electric field. *New Journal of Physics*, 20(1), 2018. doi: 10.1088/1367-2630/aa9b48.

R. Niwayama, K. Shinohara, and A. Kimura. Hydrodynamic property of the cytoplasm is sufficient to mediate cytoplasmic streaming in the Caenorhabiditis elegans embryo. *Proceedings*

of the National Academy of Sciences of the United States of America, 108(29), 2011. doi: 10.1073/pnas.1101853108.

S. Nonaka, Y. Tanaka, Y. Okada, S. Takeda, A. Harada, Y. Kanai, M. Kido, and N. Hirokawa. Randomization of Left–Right Asymmetry due to Loss of Nodal Cilia Generating Leftward Flow of Extraembryonic Fluid in Mice Lacking KIF3B Motor Protein. *Cell*, 95(6):829–837, 1998. doi: 10.1016/S0092-8674(00)81705-5.

P. R. Nott and J. F. Brady. Pressure-driven flow of suspensions: Simulation and theory. *Journal of Fluid Mechanics*, 275(6), 1994. doi: 10.1017/S0022112094002326.

N. Osterman and A. Vilfan. Finding the ciliary beating pattern with optimal efficiency. *Proceedings of the National Academy of Sciences of the United States of America*, 108(38), 2011. doi: 10.1073/pnas.1107889108.

I. R. Ozdemir, B. Clarke, Y. Hwang, and E. E. Keaveny. A Comparison of Instabilities and Dynamic States in Active Filament Models. *arXiv preprint*, 2025. doi: 10.48550/arXiv.2508.12831.

A. R. Palmer. Animal asymmetry. *Current Biology*, 19(12), 2009. doi: 10.1016/j.cub.2009.04.006.

F. Pampaloni, G. Lattanzi, A. Jonáš, T. Surrey, E. Frey, and E. L. Florin. Thermal fluctuations of grafted microtubules provide evidence of a length-dependent persistence length. *Proceedings of the National Academy of Sciences of the United States of America*, 103(27), 2006. doi: 10.1073/pnas.0603931103.

T. J. Pedley, D. R. Brumley, and R. E. Goldstein. Squirmers with swirl: A model for Volvox swimming. *Journal of Fluid Mechanics*, 798, 2016. doi: 10.1017/jfm.2016.306.

C. S. Peskin. Flow patterns around heart valves: a digital computer method for solving the equations of motion. *IEEE Transactions on Biomedical Engineering*, BME-20(4), 1973.

C. S. Peskin. The immersed boundary method. *Acta Numerica*, 11, 2002. doi: 10.1017/S0962492902000077.

R. N. Poon, T. A. Westwood, H. Laeverenz-Schlogelhofer, E. Brodrick, J. Craggs, E. E. Keaveny, G. Jékely, and K. Y. Wan. Ciliary propulsion and metachronal coordination in reef coral larvae. *Physical Review Research*, 5(4), 2023. doi: 10.1103/PhysRevResearch.5.L042037.

C. Pozrikidis. *Boundary Integral and Singularity Methods for Linearized Viscous Flow*. Cambridge University Press, 1992. doi: 10.1017/CBO9780511624124.

E. M. Purcell. Life at low Reynolds number. *American Journal of Physics*, 45(1), 1977. doi: 10.1119/1.10903.

G. Quaranta, M. E. Aubin-Tam, and D. Tam. Hydrodynamics Versus Intracellular Coupling in the Synchronization of Eukaryotic Flagella. *Physical Review Letters*, 115(23), 2015. doi: 10.1103/PhysRevLett.115.238101.

M. E. Quinlan. Cytoplasmic Streaming in the Drosophila Oocyte. *Annual Review of Cell and Developmental Biology*, 32, 2016. doi: 10.1146/annurev-cellbio-111315-125416.

M. S. Quraishi, N. S. Jones, and J. Mason. The rheology of nasal mucus: a review. *Clinical Otolaryngology & Allied Sciences*, 23(5):403–413, 1998. doi: 10.1046/J.1365-2273.1998.00172.X.

B. Rallabandi, Q. Wang, and M. Potomkin. Self-sustained three-dimensional beating of a model eukaryotic flagellum. *Soft Matter*, 18(28), 2022. doi: 10.1039/d2sm00514j.

I. H. Riedel-Kruse, A. Hilfinger, J. Howard, and F. Jülicher. How molecular motors shape the flagellar beat. *HFSP Journal*, 1(3):192, 2007. doi: 10.2976/1.2773861.

J. Rotne and S. Prager. Variational Treatment of Hydrodynamic Interaction in Polymers. *The Journal of Chemical Physics*, 50(11), 1969. doi: 10.1063/1.1670977.

C. R. Safinya, P. J. Chung, C. Song, Y. Li, H. P. Miller, M. C. Choi, U. Raviv, K. K. Ewert, L. Wilson, and S. C. Feinstein. Minireview - Microtubules and Tubulin Oligomers: Shape Transitions and Assembly by Intrinsically Disordered Protein Tau and Cationic Biomolecules. *Langmuir*, 35(48):15970–15978, 2019. doi: 10.1021/ACS.LANGMUIR.9B02208/ASSET/IMAGES/LARGE/LA9B02208_0006.JPEG.

W. S. Sale and P. Satir. Direction of active sliding of microtubules in Tetrahymena cilia. *Proceedings of the National Academy of Sciences of the United States of America*, 74(5), 1977. doi: 10.1073/pnas.74.5.2045.

P. Sartori, V. F. Geyer, A. Scholich, F. Jülicher, and J. Howard. Dynamic curvature regulation accounts for the symmetric and asymmetric beats of Chlamydomonas flagella. *eLife*, 5 (MAY2016), 2016. doi: 10.7554/eLife.13258.

P. Satir. Studies on cilia II. examination of the distal region of the ciliary shaft and the role of the filaments in motility. *The Journal of Cell Biology*, 26(3), 1965. doi: 10.1083/jcb.26.3.805.

P. Satir. Studies on cilia III. Further Studies on the Cilium Tip and a Sliding Filament Model of Ciliary Motility. *The Journal of Cell Biology*, 39(1):77–94, 1968. doi: 10.1083/JCB.39.1.77.

P. Satir and T. Matsuoka. Splitting the ciliary axoneme: Implications for a 'switch-point' model of dynein arm activity in ciliary motion. *Cell Motility and the Cytoskeleton*, 14(3), 1989. doi: 10.1002/cm.970140305.

P. Satir, T. Heuser, and W. S. Sale. A structural basis for how motile cilia beat. *Bioscience*, 64(12), 2014. doi: 10.1093/biosci/biu180.

K. Sawamoto, H. Wichterle, O. Gonzalez-Perez, J. A. Cholfin, M. Yamada, N. Spassky, N. S. Murcia, J. M. Garcia-Verdugo, O. Marin, J. L. Rubenstein, M. Tessier-Lavigne, H. Okano, and A. Alvarez-Buylla. New neurons follow the flow of cerebrospinal fluid in the adult brain. *Science*, 311(5761), 2006. doi: 10.1126/science.1119133.

M. Schliwa and G. Woehlke. Molecular motors. *Nature*, 422(6933):759–765, 2003. doi: 10.1038/nature01601.

P. J. Schmid and D. S. Henningson. *Stability and Transition in Shear Flows*, volume 142. Springer New York, 2001. ISBN 978-1-4612-6564-1. doi: 10.1007/978-1-4613-0185-1.

S. Schmitz, G. Munira, S. Howell, L. J. Calder, M. Gaeb, J. C. Pinder, A. A. Holder, and C. Veigel. Malaria Parasite Actin Filaments are Very Short. *Journal of Molecular Biology*, 349(1), 2005. doi: 10.1016/j.jmb.2005.03.056.

O. Schnitzer. Onset of spontaneous beating and whirling in the follower force model of an active filament. *Journal of Fluid Mechanics*, 1007:A65, 2025. doi: 10.1017/jfm.2025.128.

S. F. Schöller. *Simulations of passive and active filament suspensions*. PhD thesis, Imperial College London, 2019.

S. F. Schöller, A. K. Townsend, T. A. Westwood, and E. E. Keaveny. Methods for suspensions of passive and active filaments. *Journal of Computational Physics*, 424, 2021. doi: 10.1016/j.jcp.2020.109846.

K. Sekimoto, N. Mori, K. Tawada, and Y. Y. Toyoshima. Symmetry breaking instabilities of an in vitro biological system. *Physical Review Letters*, 75(1), 1995. doi: 10.1103/PhysRevLett.75.172.

D. Sept, J. Xu, T. D. Pollard, and J. A. McCammon. Annealing accounts for the length of actin filaments formed by spontaneous polymerization. *Biophysical Journal*, 77(6), 1999. doi: 10.1242/jcs.051326.

M. J. Shelley. The Dynamics of Microtubule Motor-Protein Assemblies in Biology and Physics. *Annual Review of Fluid Mechanics*, 48, 2016. doi: 10.1146/annurev-fluid-010814-013639.

C. Shingyoji, H. Higuchi, M. Yoshimura, E. Katayama, and T. Yanagida. Dynein arms are oscillating force generators. *Nature*, 393(6686), 1998. doi: 10.1038/31520.

M. P. Singh, R. Mallik, S. P. Gross, and C. C. Yu. Monte Carlo modeling of single-molecule cytoplasmic dynein. *Proceedings of the National Academy of Sciences of the United States of America*, 102(34), 2005. doi: 10.1073/pnas.0501570102.

Y. M. Sirenko, M. A. Stroscio, and K. . . Kim. Elastic vibrations of microtubules in a fiuid. *Physical Review E*, 53(1), 1996.

J. D. Skufca, J. A. Yorke, and B. Eckhardt. Edge of chaos in a parallel shear flow. *Physical Review Letters*, 96(17), 2006. doi: 10.1103/PhysRevLett.96.174101.

M. A. Sleigh, J. R. Blake, and N. Liron. The propulsion of mucus by cilia. *American Review of Respiratory Disease*, 137(3), 1988. doi: 10.1164/ajrccm/137.3.726.

D. J. Smith, J. R. Blake, and E. A. Gaffney. Fluid mechanics of nodal flow due to embryonic primary cilia. *Journal of the Royal Society Interface*, 5(22), 2008a. doi: 10.1098/rsif.2007.1306.

D. J. Smith, E. A. Gaffney, and J. R. Blake. Modelling mucociliary clearance. *Respiratory Physiology and Neurobiology*, 163(1-3), 2008b. doi: 10.1016/j.resp.2008.03.006.

D. J. Smith, A. A. Smith, and J. R. Blake. Mathematical embryology: the fluid mechanics of nodal cilia. *Journal of Engineering Mathematics*, 70:255–279, 2011. doi: 10.1007/s10665-010-9383-y.

D. J. Smith, T. D. Montenegro-Johnson, and S. S. Lopes. Symmetry-Breaking Cilia-Driven Flow in Embryogenesis. *Annual Review of Fluid Mechanics*, 51:105–128, 2019. doi: 10.1146/ANNUREV-FLUID-010518-040231.

A. Solovev and B. M. Friedrich. Synchronization in cilia carpets: Multiple metachronal waves are stable, but one wave dominates. *New Journal of Physics*, 24(1), 2022. doi: 10.1088/1367-2630/ac2ae4.

J. M. Squire. Muscle contraction: Sliding filament history, sarcomere dynamics and the two Huxleys. *Global Cardiology Science and Practice*, 2016(2), 2016. doi: 10.21542/gcsp.2016.11.

D. B. Stein, G. De Canio, E. Lauga, M. J. Shelley, and R. E. Goldstein. Swirling Instability of the Microtubule Cytoskeleton. *Physical Review Letters*, 126(2), 2021. doi: 10.1103/PhysRevLett.126.028103.

R. Stocker, J. R. Seymour, A. Samadani, D. E. Hunt, and M. F. Polz. Rapid chemotactic response enables marine bacteria to exploit ephemeral microscale nutrient patches. *Proceedings of the National Academy of Sciences of the United States of America*, 105(11), 2008. doi: 10.1073/pnas.0709765105.

J. T. Stuart. On the non-linear mechanics of wave disturbances in stable and unstable parallel flows Part 1. The basic behaviour in plane Poiseuille flow. *Journal of Fluid Mechanics*, 9(3), 1960. doi: 10.1017/S002211206000116X.

H. Su and E. E. Keaveny. Accelerating the force-coupling method for hydrodynamic interactions in periodic domains. *Journal of Computational Physics*, 510:113060, 2024. doi: 10.1016/J.JCP.2024.113060.

S. S. Suarez and A. A. Pacey. Sperm transport in the female reproductive tract. *Human Reproduction Update*, 12(1), 2006. doi: 10.1093/humupd/dmi047.

K. Svoboda and S. M. Block. Force and velocity measured for single kinesin molecules. *Cell*, 77(5), 1994. doi: 10.1016/0092-8674(94)90060-4.

J. W. Swan and J. F. Brady. Simulation of hydrodynamically interacting particles near a no-slip boundary. *Physics of Fluids*, 19(11), 2007. doi: 10.1063/1.2803837.

S. L. Tamm. Ciliary motion in Paramecium: A Scanning Electron Microscope Study. *The Journal of Cell Biology*, 55(1):250–255, 1972. doi: 10.1083/jcb.55.1.250.

S. L. Tamm. Cilia and the life of ctenophores. *Invertebrate Biology*, 133(1), 2014. doi: 10.1111/ivb.12042.

J. M. Teran and C. S. Peskin. Tether force constraints in Stokes flow by the immersed boundary method on a periodic domain. *SIAM Journal on Scientific Computing*, 31(5), 2009. doi: 10.1137/080720217.

J. L. Tobin and P. L. Beales. The nonmotile ciliopathies. *Genetics in Medicine*, 11(6), 2009. doi: 10.1097/GIM.0b013e3181a02882.

L. N. Trefethen and D. Bau. *Numerical Linear Algebra*. SIAM, 1997. doi: 10.1137/1.9780898719574.

N. Uchida and R. Golestanian. Synchronization and collective dynamics in a carpet of microfluidic rotors. *Physical Review Letters*, 104(17), 2010. doi: 10.1103/PhysRevLett.104.178103.

A. Vilfan and F. Jülicher. Hydrodynamic flow patterns and synchronization of beating cilia. *Physical Review Letters*, 96(5), 2006. doi: 10.1103/PhysRevLett.96.058102.

D. Viswanath. Recurrent motions within plane Couette turbulence. *Journal of Fluid Mechanics*, 580, 2007. doi: 10.1017/S0022112007005459.

D. Viswanath. The critical layer in pipe flow at high Reynolds number. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1888): 561–576, 2009. doi: 10.1098/RSTA.2008.0225.

E. Wajnryb, K. A. Mizerski, P. J. Zuk, and P. Szymczak. Generalization of the Rotne-Prager-Yamakawa mobility and shear disturbance tensors. *Journal of Fluid Mechanics*, 731, 2013. doi: 10.1017/jfm.2013.402.

K. Y. Wan and R. E. Goldstein. Coordinated beating of algal flagella is mediated by basal coupling. *Proceedings of the National Academy of Sciences of the United States of America*, 113(20), 2016. doi: 10.1073/pnas.1518527113.

K. Y. Wan and R. E. Goldstein. Time Irreversibility and Criticality in the Motility of a Flagellate Microorganism. *Physical Review Letters*, 2018. doi: 10.1103/PhysRevLett.121.058103.

C. Wang, S. Gsell, U. D'Ortona, and J. Favier. Generalized-Newtonian fluid transport by an instability-driven filament. *Journal of Fluid Mechanics*, 965:A6, 2023. doi: 10.1017/JFM.2023.381.

T. A. Westwood. *Computational methods for flexible filaments and Brownian suspensions*. PhD thesis, Imperial College London, 2020.

T. A. Westwood and E. E. Keaveny. Coordinated motion of active filaments on spherical surfaces. *Physical Review Fluids*, 6(12), 2021. doi: 10.1103/PhysRevFluids.6.L121101.

C. H. Wiggins and R. E. Goldstein. Flexive and Propulsive Dynamics of Elastica at Low Reynolds Number. *Physical Review Letters*, 1998.

J. H. Wilkinson. *The Algebraic Eigenvalue Problem*. Clarendon Press, Oxford, 1965.

A. P. Willis. The Openpipeflow Navier–Stokes solver. *SoftwareX*, 6:124–127, 2017. doi: 10.1016/j.softx.2017.05.003.

A. P. Willis. Equilibria, periodic orbits and computing them. *arXiv preprint*, 2019. doi: 10.48550/arxiv.1908.06730.

E. Wojcik, R. Basto, M. Serr, F. Scaírou, R. Karess, and T. Hays. Kinetochore dynein: Its dynamics and role in the transport of the Rough deal checkpoint protein. *Nature Cell Biology*, 3(11), 2001. doi: 10.1038/ncb1101-1001.

L. G. Woodhams, Y. Shen, and P. V. Bayly. Generation of ciliary beating by steady dynein activity: the effects of inter-filament coupling in multi-filament models. *Journal of the Royal Society Interface*, 19(192), 2022. doi: 10.1098/RSIF.2022.0264.

D. M. Woolley. Flagellar oscillation: A commentary on proposed mechanisms. *Biological Reviews*, 85(3), 2010. doi: 10.1111/j.1469-185X.2009.00110.x.

W. C. Worthington and R. S. Cathcart. Ciliary currents on ependymal surfaces. *Annals of the New York Academy of Sciences*, 130(3), 1966. doi: 10.1111/j.1749-6632.1966.tb12638.x.

H. Yamakawa. Transport properties of polymer chains in dilute solution: hydrodynamic interaction. *The Journal of Chemical Physics*, 53(1), 1970. doi: 10.1063/1.1673799.

Y. N. Young. Dynamics of a semiflexible polar filament in Stokes flow. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, 82(1), 2010. doi: 10.1103/PhysRevE.82.016309.

Y. Zhang, Z. Guo, L. Zou, Y. Yang, L. Zhang, N. Ji, C. Shao, W. Sun, and Y. Wang. A comprehensive map and functional annotation of the normal human cerebrospinal fluid proteome. *Journal of Proteomics*, 119, 2015. doi: 10.1016/j.jprot.2015.01.017.

P. J. Zuk, E. Wajnryb, K. A. Mizerski, and P. Szymczak. Rotne-Prager-Yamakawa approximation for different-sized particles in application to macromolecular bead models. *Journal of Fluid Mechanics*, 741, 2014. doi: 10.1017/jfm.2013.668.

Appendix A

Derivatives of Oseen's Tensor

The solution to Stokes equations for a single point force located at the origin can be expressed in terms of a Green's function, called Oseen's tensor:

$$\mathcal{G}(\boldsymbol{x}) = \frac{1}{8\pi\eta}\left(\frac{1}{r}\boldsymbol{I} + \frac{1}{r^3}\boldsymbol{x}\boldsymbol{x}^T\right) \tag{A.1}$$

where $r = |\boldsymbol{x}|$.

For completeness, we show the following derivatives which may be helpful in the text:

$$\mathcal{G}_{ij} = \frac{1}{8\pi\eta}\left(\frac{1}{r}\delta_{ij} + \frac{1}{r^3}x_i x_j\right) \tag{A.2}$$

$$\frac{\partial}{\partial x_k}\mathcal{G}_{ij}(\boldsymbol{x}) = \frac{1}{8\pi\eta}\left(\frac{1}{r^3}(-\delta_{ij}x_k + \delta_{ik}x_j + \delta_{jk}x_i) - \frac{3}{r^5}x_i x_j x_k\right) \tag{A.3}$$

$$\nabla^2\mathcal{G}_{ij}(\boldsymbol{x}) = \frac{\partial^2}{\partial x_k^2}\mathcal{G}_{ij}(\boldsymbol{x}) = \frac{1}{4\pi\eta}\left(\frac{1}{r^3}\delta_{ij} - \frac{3}{r^5}x_i x_j\right) \tag{A.4}$$

$$\frac{\partial}{\partial x_k}\nabla^2\mathcal{G}_{ij}(\boldsymbol{x}) = \frac{3}{4\pi\eta}\left(-\frac{1}{r^5}(\delta_{ij}x_k + \delta_{ik}x_j + \delta_{jk}x_i) + \frac{5}{r^7}x_i x_j x_k\right) \tag{A.5}$$

$$\nabla^4\mathcal{G}_{ij}(\boldsymbol{x}) = \frac{\partial^2}{\partial x_k^2}\nabla^2\mathcal{G}_{ij}(\boldsymbol{x}) = 0 \tag{A.6}$$

We note that Eq (A.6) implies that subsequent derivatives of the Green's function are also zero.

Appendix B

Half-Space RPY Tensor

Due to the linearity of the Stokes equations, the mobility tensor in Eq (2.42) can be decomposed into two parts. The first corresponds to the mobility tensor for the collection of particles in an unbounded domain, which is then superposed with the corrections due to the presence of the wall [Swan and Brady, 2007]. Following Section 2.2.1, we can express the mobility for particle i as

$$\begin{bmatrix} \boldsymbol{V}_i \\ \boldsymbol{\Omega}_i \end{bmatrix} = \sum_{j=1}^{N} \begin{bmatrix} \boldsymbol{M}_{ij}^{\boldsymbol{VF}} & \boldsymbol{M}_{ij}^{\boldsymbol{VT}} \\ \boldsymbol{M}_{ij}^{\boldsymbol{\Omega F}} & \boldsymbol{M}_{ij}^{\boldsymbol{\Omega T}} \end{bmatrix} \begin{bmatrix} \boldsymbol{F}_j \\ \boldsymbol{T}_j \end{bmatrix}. \tag{B.1}$$

The $(n,m)^{th}$ entry of the matrices associated to the particle's self-mobility (i.e. for $i = j$), including the image terms, are given by [Swan and Brady, 2007]:

$$\left(\boldsymbol{M}_{ii}^{\boldsymbol{VF}}\right)_{nm} = \frac{1}{6\pi\eta a}\left(\delta_{nm} - \frac{1}{16}\left(\frac{9}{h_i} - \frac{2}{h_i^3} + \frac{1}{h_i^5}\right)(\delta_{nm} - \delta_{n3}\delta_{m3}) - \frac{1}{8}\left(\frac{9}{h_i} - \frac{4}{h_i^3} + \frac{1}{h_i^5}\right)\delta_{n3}\delta_{m3}\right) \tag{B.2}$$

$$\left(\boldsymbol{M}_{ii}^{\boldsymbol{\Omega F}}\right)_{nm} = \frac{1}{64\pi\eta a h_i^4}\epsilon_{3nm} \tag{B.3}$$

$$\left(\boldsymbol{M}_{ii}^{\boldsymbol{\Omega T}}\right)_{nm} = \frac{1}{8\pi\eta a^3}\delta_{nm} - \frac{1}{64\pi\eta a h_i^3}\left(\frac{5}{2}(\delta_{nm} - \delta_{n3}\delta_{m3}) + \delta_{n3}\delta_{m3}\right) \tag{B.4}$$

where $h_i = \boldsymbol{Y}_i \cdot \hat{\boldsymbol{e}}_z/a$ is the height of the particle above the wall, normalised by the particle's radius. The first terms for $\boldsymbol{M}_{ii}^{\boldsymbol{VF}}$ and $\boldsymbol{M}_{ii}^{\boldsymbol{\Omega T}}$ correspond to the contributions considering the particles in unbounded flow, and the subsequent terms arise from considerations of the addi-

tional image terms. The final sub-block, $\boldsymbol{M}_{ii}^{\boldsymbol{VT}}$, can be obtained by using the symmetry of the mobility matrix.

To express the tensors corresponding to particle-particle interactions ($i \neq j$), we first define terms related to the difference in displacements of particles i and j, $\boldsymbol{Y}_{ij} = \boldsymbol{Y}_i - \boldsymbol{Y}_j$, $r_{ij} = ||\boldsymbol{Y}_{ij}||$ and $\hat{\boldsymbol{Y}}_{ij} = (\hat{Y}_{ij}^1, \hat{Y}_{ij}^2, \hat{Y}_{ij}^3) = \boldsymbol{Y}_{ij}/r_{ij}$, as well as terms related to the difference in displacements of particle i and the reflected image of particle j, $\boldsymbol{R} = (\boldsymbol{Y}_i - \boldsymbol{Y}_j + 2h_j\delta_3)/a$, $R = ||\boldsymbol{R}||$ and $\boldsymbol{e} = \boldsymbol{R}/R$. Also defining $\hat{h} = h_j/(a\boldsymbol{R}\cdot\hat{\boldsymbol{e}}_z)$, these interaction terms are given by [Swan and Brady, 2007]:

$$
\begin{aligned}
\left(\boldsymbol{M}_{ij}^{\boldsymbol{VF}}\right)_{nm} = {} & \frac{1}{8\pi\eta r_{ij}}\left(\left(1+\frac{2a^2}{3r_{ij}^2}\right)\delta_{nm} + \left(1-\frac{2a^2}{r_{ij}^2}\hat{Y}_{ij}^i\hat{Y}_{ij}^j\right)\right) \\
& + \frac{1}{12\pi\eta a}\left[-\frac{1}{2}\left(\frac{3\left(1+2\hat{h}(1-\hat{h})e_3^2\right)}{R} + \frac{2\left(1-3e_3^2\right)}{R^3} - \frac{2\left(1-5e_3^2\right)}{R^5}\right)\delta_{nm}\right. \\
& -\frac{1}{2}\left(\frac{3\left(1-6\hat{h}(1-\hat{h})e_3^2\right)}{R} - \frac{6\left(1-5e_3^2\right)}{R^3} + \frac{10\left(1-7e_3^2\right)}{R^5}\right)e_n e_m \\
& +e_3\left(\frac{3\hat{h}\left(1-6(1-\hat{h})e_3^2\right)}{R} - \frac{6\left(1-5e_3^2\right)}{R^3} + \frac{10\left(2-7e_3^2\right)}{R^5}\right)e_n\delta_{m3} \\
& \left.+e_3\left(\frac{3\hat{h}}{R} - \frac{10}{R^5}\right)\delta_{n3}e_m - 2\left(\frac{3\hat{h}^2e_3^2}{R} + \frac{3e_3^2}{R^3} + \frac{(2-15e_3^2)}{R^5}\right)\delta_{n3}\delta_{m3}\right],
\end{aligned}
\tag{B.5}
$$

$$
\begin{aligned}
\left(\boldsymbol{M}_{ij}^{\boldsymbol{\Omega F}}\right)_{nm} = {} & \frac{1}{8\pi\eta r_{ij}^2}\epsilon_{nmk}\hat{Y}_k - \frac{1}{4\pi\eta a}\left(\frac{1}{2R^2}\epsilon_{nmk}e_k + \left(\frac{6\hat{h}e_3^2}{R^2} + \frac{(1-10e_3^2)}{R^4}\right)\epsilon_{3kn}e_k\delta_{m3}\right. \\
& \left.-e_3\left(\frac{3\hat{h}}{R^2} - \frac{5}{R^4}\right)\epsilon_{3kn}e_k e_m - e_3\left(\frac{\hat{h}}{R^2} - \frac{1}{R^4}\right)\epsilon_{3nm}\right),
\end{aligned}
\tag{B.6}
$$

$$
\begin{aligned}
\left(\boldsymbol{M}_{ij}^{\boldsymbol{\Omega T}}\right)_{nm} = {} & \frac{1}{16\pi\eta r_{ij}^3}\left(3\hat{Y}_{ij}^i\hat{Y}_{ij}^j - \delta_{nm}\right) + \frac{3}{8\pi\eta a}\left(\frac{(1-6e_3^2)}{6R^3}\delta_{nm} - \frac{1}{2R^3}e_n e_m\right. \\
& \left.+\frac{e_3}{R^3}\delta_{n3}e_m + \frac{1}{R^3}\epsilon_{3kn}\epsilon_{3lm}e_k e_l\right).
\end{aligned}
\tag{B.7}
$$

The first term in each expression originates from considerations of the hydrodynamics in an unbounded domain, while the following terms correspond to the presence of the no-slip wall, as derived in Swan and Brady [2007]. Again, the terms from the final sub-block, $\boldsymbol{M}_{nm}^{\boldsymbol{VT}}$, can be obtained using symmetry.

Appendix C

Approximate Jacobian

In order to use Broyden's method, we require an approximate form of the Jacobian corresponding to our nonlinear system. For the cases considered in Chapters 5, 6 and 7, where we have a single filament comprised of N_p particles clamped to a no-slip planar wall, we derive an approximate form of the Jacobian by using a diagonal approximation for the mobility,

$$\boldsymbol{V}_i = \frac{1}{6\pi\eta a}\boldsymbol{F}_i \tag{C.1}$$

$$\boldsymbol{\Omega}_i = \frac{1}{8\pi\eta a^3}\boldsymbol{T}_i \tag{C.2}$$

i.e. using Stokes drag. This simplifies the nonlinear system dramatically, and we can proceed to calculate the Jacobian of this approximate system following the normal procedure. For more details and for the explicit form of the approximate Jacobian, we direct the reader to Appendix B in Schöller et al. [2021]. The matrix inverse is applied using an LU factorisation.

In the case of Chapter 8, where we have N_f filaments and the additional equations corresponding to the wall particles and the rigid body velocities, the approximate Jacobian for Eqs (8.15)-(8.19) is constructed as follows. First, we assume a diagonal approximation for the the mobility of each filament, identically to the individual filament case, to generate N_f approximate Jacobian matrices of size $6N_p \times 6N_p$, $\boldsymbol{J}_i$ for $i = 1, ..., N_f$. The full Jacobian can be

expressed in the block form:

$$\tilde{\boldsymbol{J}} = \left(\begin{array}{c|c|c} \tilde{\boldsymbol{J}}^{fil} & \boldsymbol{0} & \boldsymbol{B}^{fil} \\ \hline \boldsymbol{0} & \tilde{\boldsymbol{J}}^{wall} & \boldsymbol{B}^{wall} \\ \hline \boldsymbol{D}^{fil} & \boldsymbol{D}^{wall} & \boldsymbol{C} \end{array}\right). \tag{C.3}$$

Here,

$$\tilde{\boldsymbol{J}}^{fil} = \begin{pmatrix} \boldsymbol{J}_1 & & \\ & \ddots & \\ & & \boldsymbol{J}_{N_f} \end{pmatrix} \in \mathbb{R}^{6N_pNf \times 6N_pN_f} \tag{C.4}$$

is a block diagonal matrix comprised of the dependence of the filament equations on the filament variables,

$$\tilde{\boldsymbol{J}}^{wall} = -\frac{1}{6\pi\eta a}\boldsymbol{I}_{3N_w \times 3N_w} \in \mathbb{R}^{3N_w \times 3N_w} \tag{C.5}$$

is a diagonal matrix which describes the dependence of the wall particle equations on the wall particle constraint forces, and

$$\boldsymbol{C} = \boldsymbol{0}_3 \in \mathbb{R}^{3\times 3} \tag{C.6}$$

relates the force-free equation to the rigid body velocity.

The off-diagonal blocks can be further decomposed into

$$\boldsymbol{B}^{fil} = \begin{pmatrix} \boldsymbol{B}_1^{fil} \\ \vdots \\ \boldsymbol{B}_{N_f}^{fil} \end{pmatrix} \in \mathbb{R}^{6N_pN_f \times 3}, \quad \boldsymbol{B}^{wall} = \begin{pmatrix} \boldsymbol{B}_1^{wall} \\ \vdots \\ \boldsymbol{B}_{N_w}^{wall} \end{pmatrix} \in \mathbb{R}^{3N_w \times 3} \tag{C.7}$$

where each block describes the dependence of the filament and wall particle equations, respectively, on the rigid body velocities. These are defined as:

$$\boldsymbol{B}_i^{fil} = \begin{pmatrix} \boldsymbol{I}_3 & \cdots & \boldsymbol{I}_3 & \boldsymbol{0}_3 & \cdots & \boldsymbol{0}_3 \end{pmatrix}^T \in \mathbb{R}^{6N_p \times 3}, \tag{C.8}$$

for $i = 1, ..., N_f$, and

$$\boldsymbol{B}_i^{wall} = \frac{3}{2\Delta t}\boldsymbol{I}_3 \in \mathbb{R}^{3\times 3}. \tag{C.9}$$

for $i = 1, ..., N_w$. The other off-diagonal components can be expressed:

$$\boldsymbol{D}^{fil} = \begin{pmatrix} \boldsymbol{D}_1^{fil} & \cdots & \boldsymbol{D}_{N_f}^{fil} \end{pmatrix} \in \mathbb{R}^{3\times 6N_pN_f}, \tag{C.10}$$

$$\boldsymbol{D}^{wall} = \begin{pmatrix} \boldsymbol{D}_1^{wall} & \cdots & \boldsymbol{D}_{N_w}^{wall} \end{pmatrix} \in \mathbb{R}^{3\times 3N_w} \tag{C.11}$$

describing the dependence of the force-free equation on the filament unknowns and the wall particle Lagrange multipliers respectively. The sub-matrices are defined as

$$\boldsymbol{D}_i^{fil} = \begin{pmatrix} \boldsymbol{I}_3 & \boldsymbol{0}_3 & \cdots & \boldsymbol{0}_3 & \left[\times(-f\hat{\boldsymbol{t}}_{fil\ i,\ seg\ N_p})\right] \end{pmatrix} \in \mathbb{R}^{3\times 6N_p} \tag{C.12}$$

for $i = 1, ..., N_f$, and

$$\boldsymbol{D}_i^{wall} = \boldsymbol{I}_3 \in \mathbb{R}^{3\times 3} \tag{C.13}$$

for $i = 1, ..., N_w$.

To find and apply the inverse of the approximate Jacobian, we use Schur complements. Specifically, defining

$$\boldsymbol{A} = \left(\begin{array}{c|c} \tilde{\boldsymbol{J}}^{fil} & \boldsymbol{0} \\ \hline \boldsymbol{0} & \tilde{\boldsymbol{J}}^{wall} \end{array}\right), \quad \boldsymbol{B} = \left(\begin{array}{c} \boldsymbol{B}^{fil} \\ \hline \boldsymbol{B}^{wall} \end{array}\right), \quad \boldsymbol{D} = \left(\begin{array}{c|c} \boldsymbol{D}^{fil} & \boldsymbol{D}^{wall} \end{array}\right), \tag{C.14}$$

then we can calculate the Schur complement to be

$$\boldsymbol{S} = \boldsymbol{C} - \boldsymbol{D}\boldsymbol{A}^{-1}\boldsymbol{B} \tag{C.15}$$

$$= -\sum_{i=1}^{N_f} \boldsymbol{D}_i^{fil}\boldsymbol{J}_i^{-1}\boldsymbol{B}_i^{fil} + 6\pi\eta a \sum_{i=1}^{N_w} \boldsymbol{D}_i^{wall}\boldsymbol{B}_i^{wall} \in \mathbb{R}^{3\times 3} \tag{C.16}$$

which can be calculated at the start of each timestep, and inverted easily due to its size. Then, at each Broyden's iteration, we want to solve:

$$\begin{pmatrix} \boldsymbol{x} \\ \boldsymbol{y} \end{pmatrix} = \tilde{\boldsymbol{J}}^{-1} \begin{pmatrix} \boldsymbol{u} \\ \boldsymbol{v} \end{pmatrix} \tag{C.17}$$

where $\boldsymbol{x}, \boldsymbol{u}$ are $(6N_pN_f + 3N_b) \times 1$ vectors, and $\boldsymbol{y}, \boldsymbol{v}$ are 3×1 vectors. Rearranging gives the expressions

$$\boldsymbol{y} = \boldsymbol{S}^{-1}(\boldsymbol{v} - \boldsymbol{D}\boldsymbol{A}^{-1}\boldsymbol{u}) \tag{C.18}$$

$$\boldsymbol{x} = \boldsymbol{A}^{-1}(\boldsymbol{u} - \boldsymbol{B}\boldsymbol{y}), \tag{C.19}$$

which can be calculated efficiently, as this is the result of multiplying known matrices and vectors.

Appendix D

Lie Algebra Relation to Local Basis

In Section 4.3.1, we defined a new variable which represents the filament configuration as a rotation of a reference state. For this reference state, we chose the vertically upright filament, described by the quaternion

$$\boldsymbol{q}_{VE} = \left(\frac{1}{\sqrt{2}}, 0, -\frac{1}{\sqrt{2}}, 0 \right) \tag{D.1}$$

which corresponds to the basis vectors, $\hat{\boldsymbol{t}} = \boldsymbol{e}_z$, $\hat{\boldsymbol{\mu}} = \boldsymbol{e}_x$, and $\hat{\boldsymbol{\nu}} = -\boldsymbol{e}_y$.

The variable we define is a Lie algebra element which encapsulates a rotation of this local basis, and hence (using the kinematic constraint), the full filament configuration. If this variable has the form

$$\boldsymbol{v} = v(v_1, v_2, v_3) \tag{D.2}$$

such that $v_1^2 + v_2^2 + v_3^2 = 1$, and so $v = |\boldsymbol{v}|$, then the corresponding quaternion associated to this Lie algebra element can be expressed

$$\boldsymbol{q}_{\boldsymbol{v}} = \left(\cos\left(\frac{v}{2}\right), \sin\left(\frac{v}{2}\right) v_1, \sin\left(\frac{v}{2}\right) v_2, \sin\left(\frac{v}{2}\right) v_3 \right). \tag{D.3}$$

Using the shorthand $C = \cos(v/2), S = \sin(v/2)$ for brevity, the quaternion with reference to

the reference base state is then given by

$$\boldsymbol{q} = \boldsymbol{q_v} \bullet \boldsymbol{q}_{VE} \tag{D.4}$$

$$= \frac{1}{\sqrt{2}} \left(C + Sv_2, S(v_1 + v_3), -C + Sv_2, S(v_3 - v_1)\right). \tag{D.5}$$

Using the rotation matrix given in Eq (4.13), we can then derive the corresponding local basis:

$$\hat{\boldsymbol{t}} = \begin{bmatrix} v_2 \sin(v) + (1 - \cos(v)) v_1 v_3 \\ -v_1 \sin(v) + (1 - \cos(v)) v_2 v_3 \\ v_3^2 + (1 - v_3^2) \cos(v) \end{bmatrix}, \tag{D.6}$$

$$\hat{\boldsymbol{\mu}} = \begin{bmatrix} -v_3 \sin(v) + (1 - \cos(v)) v_1 v_2 \\ v_2^2 + (1 - v_2^2) \cos(v) \\ v_1 \sin(v) + (1 - \cos(v)) v_2 v_3 \end{bmatrix}, \tag{D.7}$$

$$\hat{\boldsymbol{\nu}} = \begin{bmatrix} -v_1^2 - (1 - v_1^2) \cos(v) \\ -v_3 \sin(v) - (1 - \cos(v)) v_1 v_2 \\ v_2 \sin(v) - (1 - \cos(v)) v_1 v_3 \end{bmatrix}. \tag{D.8}$$

In particular, we derive the local basis in the following special cases: for $v > 0$, if $\boldsymbol{v} = (v, 0, 0)$ we find:

$$\hat{\boldsymbol{t}} = \begin{bmatrix} 0 \\ -\sin(v) \\ \cos(v) \end{bmatrix}, \quad \hat{\boldsymbol{\mu}} = \begin{bmatrix} 0 \\ \cos(v) \\ \sin(v) \end{bmatrix}, \quad \hat{\boldsymbol{\nu}} = \begin{bmatrix} -1 \\ 0 \\ 0 \end{bmatrix}. \tag{D.9}$$

If instead $\boldsymbol{v} = (0, v, 0)$, then the local basis is given by:

$$\hat{\boldsymbol{t}} = \begin{bmatrix} \sin(v) \\ 0 \\ \cos(v) \end{bmatrix}, \quad \hat{\boldsymbol{\mu}} = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \quad \hat{\boldsymbol{\nu}} = \begin{bmatrix} -\cos(v) \\ 0 \\ \sin(v) \end{bmatrix}, \tag{D.10}$$

or

$$\hat{\boldsymbol{t}} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \quad \hat{\boldsymbol{\mu}} = \begin{bmatrix} -\sin(v) \\ \cos(v) \\ 0 \end{bmatrix}, \quad \hat{\boldsymbol{\nu}} = \begin{bmatrix} -\cos(v) \\ -\sin(v) \\ 0 \end{bmatrix}, \tag{D.11}$$

in the case where $\boldsymbol{v} = (0, 0, v)$. This agrees with the intuition that the Lie algebra element, $\boldsymbol{v}$, corresponds to a rotation of the angle $v = |\boldsymbol{v}|$ about the vector $(v_1, v_2, v_3) = \boldsymbol{v}/v$. Then, in the case $v = 0$, it would be expected that there be no rotation. Substituting $\boldsymbol{v} = (0, 0, 0)$ into Eq (D.6), we indeed recover the usual local basis given by $\boldsymbol{q}_{VE}$:

$$\hat{\boldsymbol{t}} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \quad \hat{\boldsymbol{\mu}} = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \quad \hat{\boldsymbol{\nu}} = \begin{bmatrix} -1 \\ 0 \\ 0 \end{bmatrix}. \tag{D.12}$$

Appendix E

Permissions to Re-use Copyrighted Work

The material in Chapters 5 and 6 is partially based on published materials in Clarke et al. [2024]. This paper is available under the terms of the Creative Commons Attribution 4.0 International license, and so permission is not required to reuse the figures in this work.

01/09/2025, 14:01 RightsLink Printable License

ELSEVIER LICENSE
TERMS AND CONDITIONS

Sep 01, 2025

This Agreement between Bethany Clarke ("You") and Elsevier ("Elsevier") consists of your license details and the terms and conditions provided by Elsevier and Copyright Clearance Center.

License Number	6100210650407
License date	Sep 01, 2025
Licensed Content Publisher	Elsevier
Licensed Content Publication	Current Biology
Licensed Content Title	Cargo Transport: Two Motors Are Sometimes Better Than One
Licensed Content Author	Steven P. Gross,Michael Vershinin,George T. Shubeita
Licensed Content Date	Jun 19, 2007
Licensed Content Volume	17
Licensed Content Issue	12
Licensed Content Pages	9
Start Page	R478
End Page	R486
Type of Use	reuse in a thesis/dissertation

https://s100.copyright.com/AppDispatchServlet 1/7

01/09/2025, 14:01 RightsLink Printable License

Portion	figures/tables/illustrations
Number of figures/tables/illustrations	1
Format	both print and electronic
Are you the author of this Elsevier article?	No
Will you be translating?	No
Title of new work	Nonlinear dynamics, bifurcations and collective motion of active filaments
Institution name	Imperial College London
Expected presentation date	Sep 2025
Portions	Figure 1
The Requesting Person / Organization to Appear on the License	Bethany Clarke
	Miss. Bethany Clarke
Requestor Location	
Publisher Tax ID	GB 494 6272 12
Total	0.00 GBP

Figure E.1: Proof of permission to reuse Figure 1.1 from Gross et al. [2007].