

Hydrodynamics of transiently cross-linked actin networks: theory, numerics, and emergent behaviors

by

Ondrej Maxian

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Mathematics
New York University
May, 2023

Professor Aleksandar Donev

© Ondrej Maxian

All Rights Reserved, 2023

*“For the foolishness of God is wiser than human wisdom,
and the weakness of God is stronger than human strength”*

1 Corinthians 1:25

Acknowledgements

This dissertation is the culmination of five years of work at NYU’s Courant Institute. It would not have been possible without my advisors, Aleks Donev and Alex Mogilner. I want to especially thank Aleks for teaching me almost everything I know about numerical analysis, and for providing the guidance necessary to make the vision for this project a reality. I also want to thank Alex Mogilner for teaching me how to think about mathematical biology, and for taking the lead on the more biology-focused parts of the project. I am indebted to Brennan Sprinkle and Raul Perez Peláez, who both provided support with the numerical and coding parts of this project. Finally, I wish to acknowledge everyone who offered me advice these last few years. The list is simply too long to remember without leaving someone out, but includes those who helped with certain parts of the project, those who gave professional advice, and those who taught me how to conduct research and manage my time. Yes, that includes inanimate objects like the swimming pool!

A note on publications

Most of the material in this thesis has already appeared in the following peer-reviewed publications:

- [A] Ondrej Maxian, Alex Mogilner, and Aleksandar Donev. Integral-based spectral method for inextensible slender fibers in Stokes flow. *Physical Review Fluids*, 6(1):014102, 2021.
- [B] Ondrej Maxian, Raul P Peláez, Alex Mogilner, and Aleksandar Donev. Simulations of dynamically cross-linked actin networks: morphology, rheology, and hydrodynamic interactions. *PLOS Computational Biology*, 17(12):e1009240, 2021.
- [C] Ondrej Maxian, Aleksandar Donev, and Alex Mogilner. Interplay between brownian motion and cross-linking controls bundling dynamics in actin networks. *Biophysical Journal*, 121(7):1230–1245, 2022.

- [D] Ondrej Maxian, Brennan Sprinkle, Charles S Peskin, and Aleksandar Donev. Hydrodynamics of a twisting, bending, inextensible fiber in Stokes flow. *Physical Review Fluids*, 7(7):074101, 2022.
- [E] Ondrej Maxian and Aleksandar Donev. Slender body theories for rotating filaments. *Journal of Fluid Mechanics*, 952:A5, 2022.
- [F] Ondrej Maxian, Brennan Sprinkle, and Aleksandar Donev. Bending fluctuations in semiflexible, inextensible, slender filaments in Stokes flow: towards a spectral discretization. *Journal of Chemical Physics*, 158:154114, 2023.

In particular, the material of Part I comes from [E] (Chapter 2) and the first part of [D] (Chapter 3). In Part II, we present the second part of [D] in Chapter 6. Then in Chapter 7, we present material from [A], but updated with the new discretization we developed in [F] (the numerical tests in this chapter are therefore new material). Part II concludes in Chapter 8 with [F], augmented with new material relating to hydrodynamic interactions among multiple filaments. In Part III, we present biological results from [A] (this forms a small part of Chapter 13) and [B] and [C].

In Chapter 7 and Section 9.2, we present new material on evaluating nonlocal (hydrodynamic and steric) interactions between fibers. This material will appear in a forthcoming publication

- [G] Ondrej Maxian and Aleksandar Donev. Large-scale simulations of suspensions of semiflexible fibers with Brownian fluctuations, hydrodynamic interactions, and steric repulsion.

Abstract

The cell cytoskeleton is at once complex and simple. On the one hand, it is made of small semiflexible filaments, which are tied together by cross linkers and motors that follow simple physical laws of motion. On the other hand, extracting any two of these components gives an in vitro system which can display a complex array of shape and structural dynamics, and a diverse range of viscoelastic behavior. The picture in vivo is even more interesting, as the filaments can form protrusions and rings that give the cell its structure and enable it to move and divide. Because cytoskeletal systems are at their core made of simple constituents, our approach to understanding them is to break the system into each individual component, model that component with fidelity to the true dynamics, and observe what happens in simulations.

Three questions immediately arise when constructing a model of such filaments. First, what equations do we use to describe the evolution of the fibers in a viscous fluid? Second, how do we discretize these equations in space and time to enable simulation of large-scale cytoskeletal systems? And third, what do the simulations tell us about how the microscopic dynamics enables large-scale pattern formation? This dissertation is divided into three parts, corresponding to each of these three questions.

In part one, we discuss the equations we use to simulate filament hydrodynamics. We consider inextensible filaments which resist both twist and bend while immersed in a viscous fluid. We use a model of the fiber as a series of infinitely many regularized singularities, which we show is asymptotically equivalent to the more-commonly used slender body theory as the fiber slenderness $\epsilon \rightarrow 0$. After formulating the filament hydrodynamics, we discuss the model equations for fiber mechanics, choosing to track twist via a scalar angle along the fiber centerline, and not an orthonormal triad.

In part two, we give our numerical discretization of these equations. Our main advance here is to use *spectral* discretizations, which in deterministic contexts are known to give high

accuracy. We begin by considering the full equations for twist and bend, and show that twist elasticity is negligible for free fibers. After this, we give a more specialized discretization for free bent filaments, which is well-suited for nonsmooth (fluctuating) problems. We show that the spectral method, though designed for deterministic problems (and indeed quite successful in that context), can also reduce the number of degrees of freedom required to simulate *fluctuating* fibers. The increase in lengthscale leads to increases in the time step size of two to three orders of magnitude, which enables large-scale simulations of suspensions of slender fluctuating fibers.

In part three, we apply our numerical methods to simulate the rheology and morphology of cross-linked actin networks. We model cross linkers as transient springs which form between the filaments, and simulate their attachment and detachment using the kinetic Monte Carlo method and the next reaction method. We show that, unabated by another process, thermal ratcheting by the cross linkers gives bundles of actin filaments, whose formation is accelerated by Brownian fluctuations (and semiflexible fluctuations when the fiber persistence length is small enough). We conclude by studying how the morphology of the fiber suspension affects its viscoelastic behavior, finding three fundamental timescales in a system of deterministic filaments. We show that hydrodynamic interactions reduce stress in systems of bundled filaments, but have little impact otherwise.

Contents

Acknowledgements	iv
Abstract	vi
List of Figures	xv
List of Tables	xxi
Introduction and program	1
I Hydrodynamics of fibers in flow: theory	3
1 Introduction	4
1.1 The mobility equations	5
1.1.1 Singularity representations	7
1.1.2 Boundary integral representations	10
1.1.3 Regularized classical SBT for translating filaments	14
1.2 Coupling fiber mechanics and hydrodynamics	16
1.2.1 Inextensibility	16
1.2.2 Coupling to twist elasticity	17
2 Rotne-Prager-Yamakawa (RPY) Mobility	20
2.1 Slender body theory from RPY singularities	22
2.1.1 Translation from force	28

2.1.2	Translation from torque	32
2.1.3	Rotation from force	35
2.1.4	Rotation from torque	36
2.1.5	Comparison to Keller-Rubinow-Johnson	38
2.1.6	Comparing to three-dimensional boundary integral	41
2.2	Extension to multiple filaments	45
3	Fiber mechanics	48
3.1	Inextensibility	49
3.1.1	Tension-based formulation	50
3.1.2	Kinematic formulation	51
3.1.3	Principle of virtual work	54
3.2	Bend and twist elasticity	56
3.2.1	Fiber geometry: the material frame and twist angle	57
3.2.2	The Kirchhoff model	60
3.2.3	The Euler model	64
3.2.4	Comparing the Kirchhoff and Euler formulations	68
4	Conclusions and governing equations	79
4.1	Summary of equations	82
4.2	Future directions	84
II	Spectral methods for fiber dynamics	86
5	Introduction	87
5.1	Single fiber hydrodynamics	88
5.2	Nonlocal hydrodynamics	92
5.3	Thermal fluctuations	94

6	Numerics for a single filament	99
6.1	Spectral spatial discretization	99
6.1.1	Inextensibility constraint	102
6.1.2	Spectral fibers from blob-link fibers	103
6.1.3	Discretizing the internal fiber force and torque (density)	111
6.1.4	Discretization of kinematics	117
6.2	Discretization of mobility for a single fiber	119
6.2.1	Translational mobility	120
6.2.2	Rot-trans mobility	126
6.2.3	Trans-rot mobility	127
6.2.4	Rot-rot mobility	128
6.2.5	Convergence of quadratures	129
6.3	Well-posedness/convergence of the static problem	132
6.3.1	Eigenvalues of \mathbf{M}_{tt}	133
6.3.2	Strong convergence of velocity	135
6.3.3	Weak convergence of $\boldsymbol{\lambda}$	140
6.4	Temporal discretization	141
6.4.1	Updating $\boldsymbol{\tau}$ and \mathbf{X}	143
6.4.2	Updating the Bishop frame	143
6.5	Results	145
6.5.1	Role of twist in the relaxation of a bent filament	145
6.5.2	Twirling to whirling instability	147
7	Free fibers with nonlocal hydrodynamics	156
7.1	Reformulation in terms of force	157
7.1.1	Saddle point system as gradient descent dynamics	158
7.2	Hydrodynamic interactions in linear time	160

7.2.1	Treating the self term differently	162
7.2.2	A hybrid approach: near fiber quadrature	164
7.2.3	Fast direct summation	171
7.2.4	The choice of mobility	177
7.3	Temporal integration	178
7.3.1	Second-order discretization	180
7.3.2	First-order discretization	182
7.3.3	Preconditioning	183
7.4	Suspensions of bent filaments	184
7.4.1	Spatio-temporal accuracy: comparing to strong formulation	184
7.4.2	Spatio-temporal convergence: sheared fibers	187
7.4.3	Stability	192
7.4.4	Bundled systems	199
8	Spectral discretization for Brownian fluctuations	203
8.1	Langevin equation for semiflexible filaments	204
8.1.1	Gibbs-Boltzmann distribution	204
8.1.2	Overdamped Langevin dynamics	206
8.2	Temporal integration	209
8.2.1	Implicit methods	211
8.2.2	RFD scheme	212
8.2.3	Midpoint scheme	215
8.3	Small equilibrium fluctuations of a curved fiber	223
8.3.1	Set-up for small fluctuations	223
8.3.2	MCMC estimation of reference covariance	227
8.3.3	Dynamics	231
8.3.4	Temporal accuracy	234

8.3.5	Importance of drift terms	237
8.4	Equilibrium statistical mechanics of free fibers	239
8.4.1	Quantifying the Gibbs-Boltzmann distribution with MCMC	241
8.4.2	Sampling with the midpoint temporal integrator	242
8.5	Dynamics of relaxation to equilibrium	243
8.5.1	Comparing special quadrature to other mobilities	250
8.6	Nonlocal hydrodynamics	255
8.6.1	Equilibrium fluctuations	256
8.6.2	Out of equilibrium dynamics	259
9	Other nonlocal interactions: cross linking and sterics	262
9.1	Cross linking model	262
9.1.1	Markov chain for transient linkers	263
9.1.2	Keeping the CL dynamics in detailed balance	266
9.1.3	Computing forces from fiber connectivity	267
9.1.4	Fiber turnover	270
9.1.5	Temporal integration	271
9.2	Steric interactions	271
9.2.1	Global uniform point resampling	273
9.2.2	Segment-based algorithm	273
10	Conclusions	285
10.1	Future directions	289
III	Macroscopic behavior of cross-linked actin gels	294
11	Introduction	295
11.1	Rheology	296

11.2	Morphology	298
11.3	Other modeling studies	299
11.4	Program	302
11.4.1	Network statistics	302
11.4.2	Simulation parameters	304
11.4.3	Calculating stress and SAOS rheology	305
12	Bundling dynamics	308
12.1	Non-Brownian filaments	309
12.2	Thermal fluctuations speed up bundling	316
12.3	Dependence on actin and CL concentration	320
12.4	Faster bundling for shorter filaments	324
12.5	Semiflexible fibers	326
12.5.1	Visualizing the bundling process	327
12.5.2	Quantifying the role of semiflexible bending fluctuations	329
12.6	Influence of steric interactions	332
12.7	Bundling, turnover times, and the steady state	336
13	Rheology of transiently cross-linked actin networks	339
13.1	Permanent cross linkers	340
13.2	Steady states with transient CLs	347
13.2.1	Homogeneous meshwork morphology	348
13.2.2	Bundles embedded in meshwork morphology (B-In-M)	349
13.2.3	Other systems (parameter variations)	351
13.2.4	The linear response regime (LRR)	352
13.2.5	Timescales of stress relaxation	354
13.3	Viscoelastic moduli	358

13.3.1	Short timescales: elastic links and viscous fibers	359
13.3.2	Viscosity on long timescales	361
13.3.3	Intermediate timescales	364
13.4	Role of nonlocal hydrodynamic interactions	366
13.4.1	Rescaling of time cannot explain long-timescale moduli	368
13.4.2	Bundled fibers: when flow reduces stress	370
13.5	A generalized Maxwell continuum framework	372
14	Discussion and conclusions	377
14.1	Future directions	383
	Bibliography	387

List of Figures

1.1	Regularizing the local drag coefficient in SBT	15
2.1	Accurate estimation of the SBT rot-trans coupling coefficient from boundary integral calculations	45
3.1	Comparing the Kirchhoff and Euler models for fiber-fluid coupling in the absence of twisting	75
3.2	Dependence of perpendicular torques in the Kirchhoff model on fiber aspect ratio, with decay to 0 as $\epsilon \rightarrow 0$	76
3.3	Comparing the Kirchhoff and Euler models for fiber-fluid coupling in the absence of bending	77
5.1	Conceptual picture of mixed discrete-continuous filament discretization. . .	96
6.1	Transferring the fiber discretization for the blob-link model to a spectral grid	101
6.2	Comparing the strict minimizer and its approximation for the spectral coarse- grained representation of a blob-link fiber	105
6.3	Smoothing boundary value problem for fiber coarse-graining, $\ell_p/L = 1$. . .	107
6.4	Smoothing boundary value problem for fiber coarse-graining, $\ell_p/L = 1/2$. .	108
6.5	Optimal smoothing parameter in coarse-graining BVP	109
6.6	Using a least-squares fit for coarse graining	110
6.7	Convergence of the nearly singular quadrature schemes for RPY integrals .	130

6.8	Fiber shapes for static convergence tests	133
6.9	Eigenvalues of discrete translational mobility matrix	134
6.10	Weak convergence of constraint forces $\boldsymbol{\lambda}$	137
6.11	Convergence of translational and parallel rotational velocities for the Euler model	138
6.12	Numerically solving the ODE for the Bishop frame	145
6.13	Error in neglecting twist in bent fibers.	146
6.14	Finding the stability boundary for the twirling-overwhirling instability . . .	150
6.15	Snapshots of nonlinear overwhirling dynamics	152
6.16	Steady state overwhirling frequencies, amplitudes, and shapes	154
7.1	Summary of how we use the fiber configuration to compute the forces \mathbf{F} and mobility $\widetilde{\mathbf{M}}$	158
7.2	Accuracy of direct quadrature on nonlocal RPY integral	166
7.3	Accuracy in computing near-fiber quadrature	171
7.4	The sheared/parallelepiped periodic simulation cell	173
7.5	The rectangular/parallelogram lattice for the sheared unit cell test	177
7.6	Results for four fibers in gravity, comparing to tension-based formulation . .	185
7.7	Starting and ending configuration for three fibers in shear flow	188
7.8	Spatio-temporal convergence for three sheared fibers with RPY special quadra- ture	189
7.9	Error for three sheared fibers with RPY special quadrature vs. oversampled	190
7.10	Convergence of GMRES for various local mobilities $\widetilde{\mathbf{M}}_L$ in dense suspensions	194
7.11	Diagnosing convergence problems with the special quadrature mobility . . .	195
7.12	Eigenvalues of the full nonlocal mobility for various levels of special quadra- ture truncation	196
7.13	Stability and accuracy of nonlocal hydrodynamics in bundled systems . . .	200

7.14	GMRES convergence at various times in deterministic bundled systems . . .	202
8.1	Penalty-bound blob-link chains	226
8.2	Variance of linearized modes in blob-link vs. spectral discretization for penalty-bound fiber	229
8.3	Timescales in the linearized (penalty-bound) fiber problem	233
8.4	Convergence of equilibrium covariance in our midpoint temporal integrator, with Δt expressed in terms of relaxation timescales	236
8.5	Convergence for changing bending stiffness	237
8.6	Variance for linearized fluctuations with/without stochastic drift terms . . .	238
8.7	Comparing spectral and blob-link discretizations of free fibers via MCMC .	240
8.8	Tangent vector correlations in blob-link and spectral free fibers	241
8.9	Convergence of midpoint temporal integrator to MCMC results for end-to-end distance as time step size decreases	244
8.10	Snapshots of fibers relaxing to their equilibrium fluctuations	246
8.11	Comparing the spectral and blob-link trajectories for relaxation of a straight filament to its equilibrium fluctuations	247
8.12	Fiber relaxation when $\hat{\epsilon} = 10^{-3}$ (only possible with spectral method)	249
8.13	End-to-end distance trajectory comparing special quadrature with oversampled quadrature	251
8.14	End-to-end distance trajectory comparing special quadrature with direct quadrature	252
8.15	End-to-end distance trajectory comparing special quadrature with local drag	254
8.16	Studying temporal convergence of special quadrature and local drag	255
8.17	Convergence of end-to-end distance distribution in midpoint integrator with nonlocal hydrodynamics	257

8.18	Experimenting with GMRES tolerance for saddle point system with nonlocal hydrodynamics	258
8.19	How the number of GMRES iterations depends on the concentration	259
8.20	How the number of GMRES iterations depends on the morphology	260
9.1	Cross linking model	263
9.2	Schematic of segment-based steric interaction algorithm, which has cost independent of ϵ	274
9.3	Error in estimating deviation from straight segments using segment midpoint	280
9.4	Error in using segment-based algorithm to compute steric forces in bundled configurations	282
9.5	An example where a large number of segments is needed	283
9.6	An example where segment-based quadrature can be inaccurate	284
12.1	Snapshots of dynamics on small and large scales	310
12.2	Bundling statistics over time in differently sized domains	311
12.3	Bundling statistics with varied bending stiffness (looking for the rigid limit), non-Brownian fibers	313
12.4	Histograms of bundle size over time for various bending stiffness	315
12.5	Statistics for the bundling process with and without thermal fluctuations . .	318
12.6	Snapshots of network at critical bundling time with and without Brownian motion	320
12.7	Changing the statistic used to define a bundle	321
12.8	Bundling time scales for range of initial mesh sizes and CL binding rates . .	322
12.9	Snapshots of network at critical bundling time with various values of CL binding rate k_{on}	323
12.10	Effect of changing filament length for rigid fibers with and without Brownian motion	325

12.11	Snapshots of bundling in semiflexible filament networks	328
12.12	Determining the role of bending fluctuations in bundled networks through statistics	331
12.13	Snapshots of bundling with steric interactions included	333
12.14	Bundling statistics with steric interactions, and changes from the no sterics case	335
12.15	Number of contacts in simulations with steric interactions	336
12.16	Adding turnover gives steady-state morphology controlled by ratio of bundling to turnover time	337
12.17	Snapshots of dynamic steady states for various turnover times	338
13.1	Estimating relaxation timescale in permanent networks	342
13.2	Steady state fiber configurations for permanently cross-linked network of 700 fibers and 8400 CLs	343
13.3	Elastic and viscous moduli for permanently cross-linked network of 700 fibers and 8400 CLs.	343
13.4	How hydrodynamic interactions affect the constraint forces	346
13.5	The steady state configurations for our morphology studies	348
13.6	Bundling statistics in dynamic meshwork steady state, repeated across dif- ferent domain lengths	349
13.7	Bundling statistics in dynamic bundle-in-mesh steady state, repeated across different domain lengths	350
13.8	Side and top view of homogeneous meshwork with smaller bending stiffness	352
13.9	Spectrum of the stress in bundle-in-mesh geometry	354
13.10	Snapshots from the stress relaxation test in a homogeneous meshwork and B-In-M geometry	356
13.11	Normalized stress profiles over time in the stress relaxation test	357

13.12	Elastic and viscous modulus in steady state transiently cross-linked networks	360
13.13	Elastic and viscous modulus due to the links alone, scaling on short and long times	362
13.14	Trying to find a scaling for the elastic modulus at medium frequencies	365
13.15	Quantifying changes in viscoelastic moduli due to hydrodynamics	368
13.16	Simplified bundle model explains reduction in stress with hydrodynamics	370
13.17	Continuum model that could enable larger scale simulations	371
13.18	Fitting the continuum model	373
13.19	Fitting the elastic and viscous modulus with varied number of generalized Maxwell elements	375

List of Tables

7.1	Number of hydrodynamic evaluations needed for stability in deterministic fiber suspensions	198
11.1	Simulation parameters for transiently cross-linked actin networks.	304
13.1	Description of dynamic steady states for rheology	347
13.2	Network structure for the bundle-in-mesh network with $\omega = 1$ Hz	353
13.3	Network statistics for $\omega = 0.5$ Hz and varying turnover times and mobility models (looking in particular at hydrodynamic interactions)	369
13.4	Parameters in the three-timescale generalized Maxwell model for our steady state systems	376

Introduction and program

Interactions of long, thin, inextensible filaments with a viscous fluid abound in biology, engineering, physics, and medicine. In biology, the swimming mechanisms of flagellated organisms have been of interest for decades, with an initial cluster of studies on how force and torque balances lead to swimming [1–4], and a more recent focus on flagellar bundling and propulsion [5–7]. In physics and engineering, suspensions of high-aspect-ratio fibers have been observed to display non-Newtonian, viscoelastic behavior both experimentally [8] and computationally [9, 10].

Our particular area of interest is the simulation of semi-flexible filaments that make up the cell cytoskeleton. These inextensible filaments, which include microtubules and actin filaments, maintain the cell’s structure, control the mechanics of the cell division process, and have aspect ratios from 10^2 to 10^4 [11]. In vivo, actin filaments are generally bound together into networks by cross-linking proteins, the properties of which determine the viscoelastic behavior of the cytoskeleton [12–14]. While there has been much work recently on microtubule systems [15–17], there has yet to be, to our knowledge, a systematic study of the influence of hydrodynamic interactions on the mechanics and rheology of cross-linked actin networks.

This dissertation is the result of a systematic program to perform such a study using a new computational framework. It is divided into three parts, corresponding to (I) theoretical advances in our understanding of slender filament hydrodynamics, (II) numerical method development, and (III) practical application to the simulation of cross-linked actin networks. We separate the first two parts into three sections, corresponding to the internal mechanics of actin filaments, as follows:

1. The resistance of the filaments to bending and straining. Because biological filaments are nearly inextensible, to avoid numerical stiffness it is best to treat inextensibility as a *constraint*.

2. The resistance of the filaments to twisting. While negligible in some applications, the twist elasticity of individual actin filaments may be important in developing macroscopic chirality in cells.
3. Thermal fluctuations of the filaments. This includes both rigid body diffusion and semiflexible *bending* fluctuations.

Our goal in particular is to develop spectral numerical methods that successfully resolve the hydrodynamic interactions in all three of these cases, while reducing the number of degrees of freedom required for simulations.

Part I

Hydrodynamics of fibers in flow: theory

Chapter 1

Introduction

This part of the thesis is concerned with our theoretical understanding of how slender filaments behave when immersed in a viscous fluid. The first issue we treat is the hydrodynamic mobility: given a force and torque acting on the filament surface, how does the filament move in a Stokes fluid? We then move on to discussing fiber mechanics, in which we address the tendency of the filament to resist bending, stretching, and twisting. Because biological filaments are nearly inextensible [18], we treat inextensibility as a constraint, thereby formulating a set of constrained PDEs which govern the dynamics of the filament centerlines.

In this introduction, we review the literature on these topics. We begin by reviewing slender body theory (SBT) as formulated in the 1970s [19–21] to study the *resistance* problem of computing the forces and torques on a slender fiber bending and twisting in a viscous fluid. In our review, we highlight the issues involved when trying to invert these SBTs and apply them to the *mobility* problem, which computes the translational and angular velocities given a set of forces and torques. This sets us up for Chapter 2, in which we propose a new SBT based on positioning regularized singularities along the fiber centerline. The second part of the introduction sets us up for Chapter 3, where we study the mechanics of filaments immersed in fluid. Our main contribution here is to formulate inextensibility without the direct use of a line tension function [22], and to harmonize two competing approaches in

the literature for how twist elasticity is handled. We discuss these two approaches in this introduction.

1.1 The mobility equations

Given the vital role of slender filaments in biology, it is not surprising that their interaction with their surrounding fluid medium has been a subject of intense study in physics and applied mathematics. Because most of the applications of interest involve small lengthscales, the relevant fluid equations are the Stokes equations, which are linear and therefore possible to solve analytically through the use of the Stokes Green’s function. However, because the body geometry is slender, it is expensive and often intractable to solve the Stokes equations on the surface of the slender filament. This intractability led Batchelor to develop the first so-called “slender body theory” (SBT) [23], in which the body is approximated via a one-dimensional line of Stokeslet singularities. Subsequent work by Keller and Rubinow and Johnson in the 1970s [19–21] placed Batchelor’s work on more rigorous asymptotic footing. A recent resurgence in theoretical analysis has placed these singularity methods back in the context of three-dimensional well-posed partial differential equations (PDEs) [24–27]. In particular, it has been shown that the SBT solution is an $\mathcal{O}(\epsilon)$ approximation to a Stokes PDE with mixed Dirichlet-Neumann boundary data and a boundary condition that the fiber maintain the integrity of its cross section [24–26]. It has also been shown [27, 28] that SBT can be derived from a three-dimensional boundary integral equation. We note a second class of recent work on SBTs for different geometries, including ribbons [29] and bodies with non-circular cross section [30], but such works are beyond the scope of this thesis.

Our particular focus in this introduction is solving the mobility problem: given a force and torque on the fiber cross section, what is the resulting velocity? To answer this question, we will survey a few SBTs, which give a well-defined solution for the velocity when the fiber translates, but suffer from asymptotic inconsistencies for fiber rotation. This section is not

meant to be a comprehensive list of previous SBTs (see [31] for a such a review), but rather a focus on two different techniques and the challenges involved in applying them to a rotating filament.

To establish some notation, we let $\mathbb{X}(s)$ be the fiber centerline, where $s \in [0, L]$ is an arclength parameterization of the curve, and thus $\boldsymbol{\tau}(s) = \partial_s \mathbb{X}(s)$ is the unit-length tangent vector. The fiber surface $\widehat{\mathbb{X}}$ is then defined as

$$\widehat{\mathbb{X}}(s, \theta) = \mathbb{X}(s) + a\rho(s) (\mathbf{e}_n(s) \cos \theta + \mathbf{e}_b(s) \sin \theta) := \mathbb{X}(s) + a\rho(s) \widehat{\mathbf{r}}(s, \theta), \quad (1.1)$$

where $a\rho(s)$, $0 \leq \rho(s) \leq 1$, is the radius of the (circular) cross section at s , $\mathbf{e}_n(s) = \partial_s^2 \mathbb{X}(s)/v(s)$ is the fiber normal, $v(s)$ is the centerline curvature at s , and $\mathbf{e}_b = \boldsymbol{\tau} \times \mathbf{e}_n$ is the binormal. The maximum fiber radius is a , which gives an aspect ratio $\epsilon := a/L$. Unless otherwise specified, we will assume that we are working in a region of the fiber where $\rho(s) \equiv 1$; thus $a\rho = a$.

We will assume that the fiber centerline is translating with velocity $\mathbf{U}(s)$ while its cross section rigidly rotates with angular velocity $\Psi^\parallel(s)$ around the tangent vector $\boldsymbol{\tau}(s)$, so that the velocity on the surface of the fiber is given by

$$\widehat{\mathbf{U}}(s, \theta) = \mathbf{U}(s) + a\Psi^\parallel(s) (\boldsymbol{\tau}(s) \times \widehat{\mathbf{r}}(s, \theta)) \quad (1.2)$$

The goal of our analysis is to find an equation relating this motion to the total force density $\mathbf{f}(s)$ and parallel torque density $n^\parallel(s)$ on the fiber cross section. We assume here that the no-slip condition applies to the fluid velocity $\mathbf{u}(\mathbf{x})$, so that

$$\mathbf{u}(\widehat{\mathbb{X}}(s, \theta)) = \widehat{\mathbf{U}}(s, \theta). \quad (1.3)$$

In the case of a resistance problem, \mathbf{U} and Ψ^\parallel are known quantities, and the goal is to solve

for \mathbf{f} and n^\parallel asymptotically in ϵ . In simulations, however, we typically have a model of the force and parallel torque from the internal fiber mechanics (see Chapter 3). Therefore, we think of \mathbf{f} and n^\parallel as fixed, and the goal is to compute a consistent translational and rotational velocity at point s . We will see that this is quite difficult to do because of the coupling between rotation and translation for curved fibers.

1.1.1 Singularity representations

We first try to apply an older SBT approach to the spinning filament problem: the singularity approach of Johnson [20], and later Götze [32].¹ The approach of Johnson is to look for a series of singularities along the fiber centerline to represent the flow field globally (everywhere outside the fiber). The singularities are chosen so that the boundary condition (1.3) is satisfied with an error $\mathcal{O}(\epsilon)$.

Using superposition, we can break the boundary velocity (1.2) into a velocity with $\mathbf{U} = \mathbf{0}$ and another with $\Psi^\parallel = 0$. For $\Psi^\parallel = 0$, Johnson shows that the boundary condition (1.3) can be satisfied to $\mathcal{O}(\epsilon)$ by setting the global flow outside of the fiber to be an integral of Stokeslets and doublets,

$$\mathbf{u}^{(u)}(\mathbf{x}) = \frac{1}{8\pi\mu} \int_0^L \left(\mathbb{S}(\mathbf{x}, \mathbb{X}(s)) + \frac{(a\rho(s))^2}{2} \mathbb{D}(\mathbf{x}, \mathbb{X}(s)) \right) \mathbf{f}(s) ds \quad (1.4)$$

$$\text{where} \quad \mathbb{S}(\mathbf{x}, \mathbb{X}(s)) = \left(\frac{\mathbf{I}}{r} + \frac{\mathbf{r}\mathbf{r}^T}{r^3} \right), \quad \mathbb{D}(\mathbf{x}, \mathbb{X}(s)) = \left(\frac{\mathbf{I}}{r^3} - 3\frac{\mathbf{r}\mathbf{r}^T}{r^5} \right),$$

μ is the background fluid viscosity, and $\mathbf{r} = \mathbf{x} - \mathbb{X}(s)$ with $r = \|\mathbf{r}\|$. The doublet strength is necessary to cancel the $\mathcal{O}(1)$ angular-dependent flow induced by the Stokeslet on a filament cross section. Since the doublet has strength $\mathcal{O}(\epsilon^2)$, it makes a negligible contribution to the outer expansion, but corrects the Stokeslet flow in the inner expansion. The resulting

¹The approach of Keller and Rubinow [19] is technically distinct from that of Johnson, since the former only use the singularity representation in the outer expansion, but the end result is the same, and so we discuss both in this section.

asymptotic evaluation of $\mathbf{u}^{(u)}$ on the filament cross section $\widehat{\mathbb{X}}(s, \theta)$ is independent of θ to $\mathcal{O}(\epsilon)$, and is given by

$$8\pi\mu\mathbf{u}^{(u)}\left(\widehat{\mathbb{X}}(s, \theta)\right) = \mathbf{U}_f^{(\text{LD})}(s) + \mathbf{U}_f^{(\text{FP})} + \mathcal{O}(\epsilon), \quad (1.5)$$

$$\mathbf{U}_f^{(\text{LD})}(s) = \ln\left(\frac{4s(L-s)}{a^2}\right) (\mathbf{I} + \boldsymbol{\tau}(s)\boldsymbol{\tau}(s)) \mathbf{f}(s) + (\mathbf{I} - 3\boldsymbol{\tau}(s)\boldsymbol{\tau}(s)) \mathbf{f}(s) \quad (1.6)$$

$$\mathbf{U}_f^{(\text{FP})}(s) = \int_0^L \left(\mathbb{S}(\mathbb{X}(s), \mathbb{X}(t)) \mathbf{f}(t) - \left(\frac{\mathbf{I} + \boldsymbol{\tau}(s)\boldsymbol{\tau}(s)}{|t-s|} \right) \mathbf{f}(s) \right) dt \quad (1.7)$$

where the doublet contributes the $(\mathbf{I} - 3\boldsymbol{\tau}\boldsymbol{\tau})$ term [20, Eqs. (10–12)], and the term $\mathbf{U}_f^{(\text{FP})}$ is a finite part integral which is only well-defined as a difference of two terms (the outer and inner expansions). The expression (1.5) has an error on the cross section of order ϵ , which means there will be a nonzero $\mathcal{O}(1)$ rotational velocity of the cross section. To zero out the rotational velocity, Johnson shows that it is necessary to introduce a rotlet, a source, two stresslets, and two quadrupoles, with strengths given as a function of the $\mathcal{O}(\epsilon)$ velocity generated by the surface flow (1.5). While these singularities zero out the $\mathcal{O}(\epsilon)$ term, they *do not* impart any additional translational velocity on the cross section, meaning that (1.5) gives the *translational* velocity on the cross section to order $\epsilon^2 \log \epsilon$.

Let us now consider how the fluid flow outside the fiber relates to its angular velocity. Johnson first sets

$$n^\parallel = 4\pi\mu a^2 \Psi^\parallel, \quad (1.8)$$

which is the result for an infinite, straight cylinder [33, Eq. (62)]. The singularity representation Johnson uses to match the angular velocity everywhere is then given by a line integral of rotlets,

$$8\pi\mu\mathbf{u}^{(\Psi)}(\mathbf{x}) = \int_0^L \mathbb{R}(\mathbf{x}, \mathbb{X}(s)) n^\parallel(s) ds, \quad \mathbb{R}(\mathbf{x}, \mathbb{X}(s)) = \frac{\boldsymbol{\tau}(s) \times \mathbf{r}}{r^3} \quad (1.9)$$

Asymptotically, this velocity is equal to

$$8\pi\mu\mathbf{u}^{(\Psi)}\left(\widehat{\mathbb{X}}(s,\theta)\right)\approx\frac{2n^{\parallel}}{a}\left(\boldsymbol{\tau}\times\hat{\mathbf{r}}\right)+\frac{n^{\parallel}}{2}\left(\ln\left(\frac{4s(L-s)}{a^2}\right)-2\right)\left(\boldsymbol{\tau}\times\partial_s\boldsymbol{\tau}\right)+\mathbf{U}_n^{(\text{FP})}\quad(1.10)$$

$$+n^{\parallel}v\cos\theta\left(\boldsymbol{\tau}\times\hat{\mathbf{r}}\right)+\mathcal{O}\left(\epsilon\log\epsilon\right),$$

$$\text{where}\quad\mathbf{U}_n^{(\text{FP})}(s)=\int_0^L\left(\mathbb{R}\left(\mathbb{X}(s),\mathbb{X}(t)\right)n^{\parallel}(t)-\frac{1}{2}\left(\frac{\boldsymbol{\tau}(s)\times\partial_s\boldsymbol{\tau}(s)}{|t-s|}\right)n^{\parallel}(s)\right)\quad(1.11)$$

on the fiber surface. Substituting (1.8) into (1.10), we see that the first term is exactly the angular velocity of (1.2), while the second and third terms on the first line of (1.10) give an additional constant velocity on the cross section (the rot-trans coupling term). However, we still have an angular-dependent term on the second line of (1.10): since $\boldsymbol{\tau}\times\hat{\mathbf{r}}=\mathbf{e}_b\cos\theta-\mathbf{e}_n\sin\theta$, the term in the second line has the angle dependence

$$n^{\parallel}v\cos\theta\left(\boldsymbol{\tau}\times\hat{\mathbf{r}}\right)=\frac{n^{\parallel}v}{2}\left(\mathbf{e}_b+\mathbf{e}_b\cos 2\theta-\mathbf{e}_n\sin 2\theta\right).\quad(1.12)$$

In Johnson's analysis, which is based on the resistance problem, the angular velocity Ψ^{\parallel} and translational velocity \mathbf{U} are independent of ϵ , and so the contribution of all subleading terms in (1.10) to translational velocity on the cross section is still ϵ^2 relative to the contribution of forcing. However, if force and torque are of the same order in ϵ (as in the mobility problem when a fiber is driven by a torque n^{\parallel} independent of ϵ), the second line will make an $\mathcal{O}(1)$ contribution to the cross-sectional velocity, violating the boundary condition (1.3). Thus, there is a $\mathcal{O}(1)$ angle dependence coming from the rotlet term which is unresolved. This dependence is the reason why the final integral equation of Keller and Rubinow [19, Eq. (28)] for twisting filaments contains terms which depend on the angle θ at which the inner expansion (of a rotating, translating cylinder) is matched with the outer expansion (of Stokeslets and rotlets), meaning there is no general solution for the Stokeslet and rotlet strength which is independent of the matching angle.

In addition, there is also a lack of symmetry in the final result of Johnson, since torque makes an $\mathcal{O}(\log \epsilon)$ contribution to \mathbf{U} in (1.10), but force does not contribute to Ψ^\parallel in (1.8). If the fiber is spun by a motor driven by a constant torque n^\parallel , the angular velocity $\Psi^\parallel \sim n^\parallel/\epsilon^2$. Symmetry tells us that the contribution of force to Ψ^\parallel is $\mathcal{O}(\log \epsilon)$, which means that the rotational velocity from force is $\Psi_f^\parallel \sim f \log \epsilon$, which is much smaller than the rotational velocity from torque $\Psi_n^\parallel \sim n/\epsilon^2$, if f and n scale similarly with ϵ . In the case when Ψ^\parallel does not scale with ϵ , however, the torque n^\parallel is order ϵ^2 and therefore makes a negligible contribution to \mathbf{U} , but the force might make a nontrivial contribution to Ψ^\parallel . Thus there are two main issues when applying singularity representations to twisting filaments: first, the representation of the rotation in terms of rotlets is insufficient to satisfy the boundary condition (1.3) to $\mathcal{O}(\epsilon)$, and, second, there is a lack of symmetry in the rot-trans coupling terms.

1.1.2 Boundary integral representations

Koens and Lauga [27] showed that the translational SBT of Johnson and Keller and Rubinow can be obtained from asymptotic expansion of the single layer² boundary integral equation [34, Sec. 4.1]

$$8\pi\mu\widehat{\mathbf{U}}(s, \theta) = \int_0^{2\pi} \int_0^L \mathbb{S}(\widehat{\mathbf{X}}(s, \theta), \widehat{\mathbf{X}}(s', \theta')) \widehat{\mathbf{f}}(s', \theta') ds' d\theta'. \quad (1.13)$$

Here we assume the fiber is a thin shell filled with fluid, so that $\widehat{\mathbf{f}}(s, \theta)$, which is the jump in surface traction times the surface area element [34, p. 105], is also equal to the force density on the fiber surface. Koens and Lauga show that the constant cross-sectional $\mathbf{U}(s)$ in (1.2)

²Koens and Lauga begin with the full boundary integral representation, but the single layer representation is sufficient as long as there is no change in the fiber volume. However, it is important to point out that $\widehat{\mathbf{f}}$ in (1.13) is not in general the force density on the fiber surface. Notable exceptions are a filament undergoing rigid body motion or a “filament” whose interior is a fluid of viscosity equal to that of the fiber exterior.

is related to the total force on a cross section,

$$\mathbf{f}(s) = \int_0^{2\pi} \widehat{\mathbf{f}}(s, \theta) d\theta \quad (1.14)$$

by the classical SBT (1.5), and that, for a *straight* cylinder, the torque-angular velocity relationship (1.8) is recovered [27, Eq. (4.11)] for the total parallel torque on a cross section,

$$n^{\parallel}(s) = \boldsymbol{\tau}(s) \cdot \int_0^{2\pi} a\rho(s) \left(\widehat{\mathbf{r}}(s, \theta) \times \widehat{\mathbf{f}}(s, \theta) \right) d\theta. \quad (1.15)$$

Here we briefly discuss how the approach of Koens and Lauga might be extended to account for rotation in curved filaments. The main issue is that, if the boundary integral equation (1.13) is only expanded to leading order (i.e., with an error of $\mathcal{O}(\epsilon)$), it is impossible to make observations about the rotational velocity, since the velocity on the cross section induced by an angular velocity Ψ^{\parallel} is $\sim \epsilon\Psi^{\parallel}$. Thus more terms are necessary in the asymptotics. Unfortunately, because the expansions must be conducted around both θ and s in (1.13), even expanding a single layer formulation to the next order is complex, as can be seen in the results of slender phoretic theory [35], which completely extends the approach of Koens and Lauga to next order for the case of a body which swims via a fluid slip along a self-generated surface chemical concentration gradient.

To get a flavor for the challenges involved, we consider an asymptotic expansion of the *isotropic part* of the single-layer potential (c.f. [27, Eq. (5.1)])

$$\mathbf{U}_{i1}(s, \theta) := \int_0^{2\pi} \int_0^L \frac{\mathbf{f}(s', \theta')}{R(\widehat{\mathbf{X}}(s, \theta), \widehat{\mathbf{X}}(s', \theta'))} d\theta' ds', \quad (1.16)$$

in an inner region where $s - s' = \mathcal{O}(a)$. To expand the isotropic single layer (1.1.2), we begin

with an inner expansion of the displacement $\mathbf{R}(s, s', \theta, \theta') = \widehat{\mathbb{X}}(s, \theta) - \widehat{\mathbb{X}}(s', \theta')$,

$$\mathbf{R}(s, s', \theta, \theta') = \widehat{\mathbb{X}}(s, \theta) - \widehat{\mathbb{X}}(s', \theta') = \mathbb{X}(s) - \mathbb{X}(s') + \hat{\mathbf{e}}_\rho(s, \theta) - \hat{\mathbf{e}}_\rho(s', \theta'), \quad (1.17)$$

where $\hat{\mathbf{e}}_\rho(s, \theta)$ is the unit vector going from $\mathbb{X}(s)$ to $\widehat{\mathbb{X}}(s, \theta)$. Since the geometry is more complicated in this true three-dimensional case, it is helpful to parameterize $\hat{\mathbf{e}}_\rho$ by an angle $\theta_i(s)$ satisfying $\partial_s \theta_i = \mathbf{e}_b \cdot \partial_s \mathbf{e}_n$ (the curve torsion) [27, Eq. (2.2)]. This gives [27, Sec. 2]

$$\begin{aligned} \hat{\mathbf{e}}_\rho(s, \theta) &= \mathbf{e}_n \cos(\theta - \theta_i(s)) + \mathbf{e}_b \sin(\theta - \theta_i(s)) \\ \partial_s \hat{\mathbf{e}}_\rho(s, \theta) &= -v(s) \cos(\theta - \theta_i) \boldsymbol{\tau}(s) \end{aligned} \quad (1.18)$$

Substituting into (1.17), we can now write the first two terms in the Taylor expansion of \mathbf{R} and $1/R$ as

$$\begin{aligned} \mathbf{R}(\widehat{\mathbb{X}}(s, \theta), \widehat{\mathbb{X}}(s', \theta')) &\approx a(-\xi \boldsymbol{\tau} + \Delta \hat{\mathbf{e}}_\rho(s, \theta, \theta')) + a^2 \left(\frac{\xi^2}{2} \partial_s \boldsymbol{\tau} - \xi v \cos(\theta' - \theta_i) \boldsymbol{\tau} \right) \\ \frac{1}{R(\widehat{\mathbb{X}}(s, \theta), \widehat{\mathbb{X}}(s', \theta'))} &= \frac{1}{a\sqrt{\xi^2 + \Delta \hat{\mathbf{e}}_\rho \cdot \Delta \hat{\mathbf{e}}_\rho}} + \frac{-\xi^2 v \cos(\theta' - \theta_i) + \xi^2/2 \boldsymbol{\tau}' \cdot \Delta \hat{\mathbf{e}}_\rho}{2\sqrt{\xi^2 + \Delta \hat{\mathbf{e}}_\rho \cdot \Delta \hat{\mathbf{e}}_\rho}^3} \end{aligned} \quad (1.19)$$

where $\Delta \hat{\mathbf{e}}_\rho(s, \theta, \theta') = \hat{\mathbf{e}}_\rho(s, \theta) - \hat{\mathbf{e}}_\rho(s, \theta')$ and $\xi = (s' - s)/a$ is order 1. Note the appearance of the centerline curvature (scalar v and vector $\partial_s \boldsymbol{\tau}$) at this order, which is vital to yield proper rot-trans coupling. Combining (1.19) with the Taylor series representation $\mathbf{f}(s', \theta') = \mathbf{f}(s, \theta') + a\xi \partial_s \mathbf{f}(s, \theta')$, we have the inner expansion of the integrand in (1.16) as

$$\begin{aligned} \frac{\mathbf{f}(s', \theta')}{R} &= \frac{\mathbf{f}(s, \theta')}{a\sqrt{\xi^2 + \Delta \hat{\mathbf{e}}_\rho \Delta \hat{\mathbf{e}}_\rho}} + \frac{\xi \partial_s \mathbf{f}(s, \theta')}{\sqrt{\xi^2 + \Delta \hat{\mathbf{e}}_\rho \Delta \hat{\mathbf{e}}_\rho}} \\ &+ \frac{-\xi^2 v \cos(\theta' - \theta_i) + \xi^2/2(\partial_s \boldsymbol{\tau} \cdot \Delta \hat{\mathbf{e}}_\rho)}{2\sqrt{\xi^2 + \Delta \hat{\mathbf{e}}_\rho \Delta \hat{\mathbf{e}}_\rho}^3} \mathbf{f}(s, \theta') + \mathcal{O}(\epsilon). \end{aligned} \quad (1.20)$$

The second line already makes it clear how rot-trans coupling is going to arise. If $\mathbf{f}(s, \theta') =$

$\mathbf{f}(s)$ (no angular dependence), then there will still be a term in the velocity proportional to $\cos\theta\mathbf{f}(s)$, which is rotational velocity coming from translational force, or rot-trans coupling.

To complete the expansion, we proceed with the integration of (1.20) from $\xi = -s/a$ to $\xi = (L - s)/a$. The integration can be done exactly, and then expanded in ϵ to yield simpler expressions [27, Table 1]. After using (1.18) to determine $\partial_s\boldsymbol{\tau} \cdot \Delta\hat{\mathbf{e}}_\rho = 0$, then integrating what remains, we have

$$\begin{aligned} \mathbf{U}_{i1}(s, \theta) = & \int_0^{2\pi} \left(\ln \left(\frac{4s(L-s)}{a^2(\Delta\hat{\mathbf{e}}_\rho \cdot \Delta\hat{\mathbf{e}}_\rho)} \right) \mathbf{f}(s, \theta') + (L-2s) \partial_s \mathbf{f}(s, \theta) \right. \\ & \left. - \frac{\epsilon v \cos(\theta' - \theta_i(s)) \mathbf{f}(s, \theta')}{2} \left(\ln \left(\frac{4s(L-s)}{a^2(\Delta\hat{\mathbf{e}}_\rho \cdot \Delta\hat{\mathbf{e}}_\rho)} \right) - 2 \right) \right) d\theta' + \mathcal{O}(\epsilon^2). \end{aligned} \quad (1.21)$$

We now let $\tilde{\theta} = \theta - \theta_i(s)$, and substitute the one-term Fourier series representation

$$\hat{\mathbf{f}}(s, \tilde{\theta}) = \frac{1}{2\pi} \mathbf{f}(s) + \frac{1}{2\pi} \left(\mathbf{f}_{c,1}(s) \cos \tilde{\theta} + \mathbf{f}_{s,1}(s) \sin \tilde{\theta} \right), \quad (1.22)$$

into (1.21). To integrate over θ' , we use $\Delta\hat{\mathbf{e}}_\rho \cdot \Delta\hat{\mathbf{e}}_\rho = 2(1 - \cos(\theta - \theta'))$. After integrating over θ' , the final inner expansion for the surface velocity is

$$\begin{aligned} \mathbf{U}_{i1}(s, \theta) = & \ln \left(\frac{4s(L-s)}{a^2} \right) \mathbf{f}(s) + (L-2s) \partial_s \mathbf{f}(s) + \mathbf{f}_{c,1}(s) \cos \tilde{\theta} - \frac{av}{2} \left(\mathbf{f}(s) \cos \tilde{\theta} \right. \\ & \left. + \left(\frac{1}{4} \cos 2\tilde{\theta} + \frac{1}{2} \left(\ln \left(\frac{4s(L-s)}{a^2} \right) - 2 \right) \right) \mathbf{f}_{c,1}(s) - \frac{1}{2} \mathbf{f}_{s,1}(s) \cos \tilde{\theta} \sin \tilde{\theta} \right) + \mathcal{O}(\epsilon^2). \end{aligned} \quad (1.23)$$

We can make three important observations from (1.1.2): first, the term $\mathbf{f}(s)v \cos\theta$ demonstrates rot-trans coupling (rotational motion from constant forcing) which results from including additional terms in the SBT expansion, or more specifically from accounting for the curvature of the fiber in the inner expansion. Likewise, the *torque* (1.15) generated from the traction scales as $a\mathbf{f}_{c/s,1}$, and thus (1.1.2) tells us that torque makes an $\mathcal{O}(\log \epsilon)$ contribution to $\mathbf{U}(s)$, as we expect from the result (1.10) of singularity methods.

While the result of Koens and Lauga is vital to understanding translational SBT, there are clearly limitations to extending their approach to rotation. Even for the simplest possible term in the boundary integral formulation (isotropic part of the single layer), there are many terms that have to be tracked if we expand to $\mathcal{O}(\epsilon)$, i.e., to an error of $\mathcal{O}(\epsilon^2)$. More importantly, it is not even clear if we can solve (1.1.2), since the Fourier modes no longer decouple as they do for translation.

1.1.3 Regularized classical SBT for translating filaments

In addition to not having a consistent treatment of rotation-translation coupling, there is a second problem with SBT for translation only. In this case, the results of Keller and Rubinow [19], Johnson [20, 21], and Koens and Lauga [27] all agree on the equations (1.5)–(1.7). Looking at (1.6) in particular, we see that the local drag coefficient

$$c(s) := \ln \left(\frac{4s(L-s)}{a(s)^2} \right) \quad (1.24)$$

is singular without proper decay of $a(s)$ at $s = 0$ and $s = L$. Clearly, if $a(s)$ decays near the fiber endpoints as $2\epsilon\sqrt{s(L-s)}$, then the leading order coefficient (1.24) is finite at the fiber endpoints [20]. This fact is the basis for a general assumption across the SBT literature that the filaments have ellipsoidal shape, so that in most studies $c(s) = -\ln(\epsilon^2)$ is constant for all s [15, 22, 36].

Actin filaments are best modeled as cylinders with constant radius, so that $a(s) = r = \epsilon L$ on $s \in [0, L]$. In this case, the coefficient (1.24) becomes singular at the filament ends, and some resolution is needed to make the problem well-posed and numerically tractable. In this thesis, we present two possible alternatives to (1.24). The first, which also gives an SBT with consistent rot-trans coupling, is based on regularized singularities and is deferred to the next chapter.




Figure 1.1: Regularized drag coefficients and effective radius functions. (Left:) The regularized local drag coefficients $\bar{c}(s; \delta)$ for various δ . (Right:) The effective fiber radius $\bar{a}(s; \delta)$ as defined in (1.27). We show only one side of the fiber, with the other being the former's mirror image about $s/L = 0.5$.

The second approach is to simply modify the local drag coefficient by effectively tapering the fiber radius over a distance $\sim \delta L$ near the endpoints. Specifically, to regularize $c(s)$, we set $\eta = 2s/L - 1$, so that $\eta \in [-1, 1]$, and compute a weight function

$$w(s; \delta) = \tanh\left(\frac{\eta(s) + 1}{\delta}\right) - \tanh\left(\frac{\eta(s) - 1}{\delta}\right) - 1, \quad (1.25)$$

which is 1 near the fiber center ($\eta = 0$) and zero at the fiber ends ($\eta = \pm 1$). We then assign to each s a regularized fiber centerline coordinate by

$$\bar{s}(s; \delta) = w(s; \delta)s + (1 - w(s; \delta))^2 \frac{\delta L}{2} \quad (1.26)$$

on $0 \leq s \leq L/2$, with the corresponding reflection for $s > L/2$. The regularized coefficient for a given δ is then given by

$$\bar{c}(s; \delta) := c(\bar{s}(s; \delta)) = \ln\left(\frac{4\bar{s}(L - \bar{s})}{(\epsilon L)^2}\right) := \ln\left(\frac{4s(L - s)}{\bar{a}(s; \delta)^2}\right). \quad (1.27)$$

Figure 1.1 shows how the choice of δ impacts the local drag coefficient $\bar{c}(s)$ and effective radius function $\bar{a}(s)$. We see that the fibers are cylindrical ($\bar{a}(s) = \epsilon L$) for $s/L \gtrsim \delta$, while for $s/L \lesssim \delta$ the effective radius smoothly decays to zero, as it would for an ellipsoidal fiber. Larger values of δ yield smoother radius functions and smoother local drag coefficients. When we use this approach, we will use $\delta = 0.1$, unless otherwise stated.

1.2 Coupling fiber mechanics and hydrodynamics

The second broad topic we treat in this part of the thesis is how to couple the internal fiber mechanics to the hydrodynamics. Actin filaments are nearly inextensible (stretching stiffness on the order 10^5 pN/ μ m [37]), and also resist bending and twisting elasticity. As such, they are best modeled as inextensible Kirchhoff rods [38, c. 3]. While we will also take this as our base model, we make two important advances in coupling the elasticity to the fluid mechanics: first, we offer a novel theoretical treatment of inextensibility which is based on the principle of virtual work. Second, we reconcile two approaches in the literature for the coupling of twist elasticity and rotation to the surrounding fluid.

1.2.1 Inextensibility

In the traditional approach to inextensibility, which was introduced by Tornberg and Shelley [22] and has since been used by a number of others [15, 39, 40], an auxiliary (integro-) differential equation is derived for the line tension in the filament, which acts as a Lagrange multiplier to preserve inextensibility. This equation, which involves multiplications of high-order (as high as four) derivatives of the fiber position function \mathbb{X} , is ill-suited for numerical discretization, especially with spectral methods, where aliasing can become a problem [15]. But more importantly, it is not clear how to adapt the continuum line tension equation to the discrete nature of the simulation, and so Tornberg and Shelley [22] insert an additional

penalty term into the line tension equation to reduce the discrete extensibility. For fibers tugged by cross-linkers or strong extensional flows, this penalty parameter will be large, introducing artificial stiffness into the problem. And finally, the coupling of the trajectory of \mathbb{X} with an auxiliary equation makes it difficult to write a closed-form equation for the overall evolution of \mathbb{X} . Such an equation is necessary when we try to incorporate thermal fluctuations (see Chapter 8).

There is therefore a need for a new formulation for inextensible filaments. In our approach, we make the observation that the constraint $\boldsymbol{\tau}(s, t) \cdot \boldsymbol{\tau}(s, t) = 1$ implies that the velocity of the fiber tangent vector $\partial_t \boldsymbol{\tau}$, is orthogonal to itself ($\partial_t \boldsymbol{\tau} \cdot \boldsymbol{\tau} = 0$), which implies that

$$\partial_t \boldsymbol{\tau} = \boldsymbol{\Omega} \times \boldsymbol{\tau}, \quad (1.28)$$

for some rotational velocity $\boldsymbol{\Omega}(s)$. In discrete time, the tangent vectors will be updated via an actual rotation by $\boldsymbol{\Omega}(s)$, which keeps them exactly on the constraint. The fiber positions will then be obtained by integration, which is an approach similar to that of [41], but unique because we consider the fiber as a continuum, rather than a collection of discrete line segments. To close our formulation, we treat the force due to tension as a Lagrange multiplier and enforce the principle that the constraint forces do no work [42]. By imposing this principle, the constraint forces can be eliminated, and we can obtain a closed form equation for the evolution of \mathbb{X} which is sorely needed when we consider thermal fluctuations.

1.2.2 Coupling to twist elasticity

In addition to being inextensible, biological filaments also resist bend and twist elasticity. Most studies that couple the fiber to the surrounding fluid use some form of the Kirchhoff rod model [43] to treat the internal mechanics, in which the inextensible and unsharable rod resists both bending (along two principal axes) and twisting (along the third axis parallel to the fiber centerline). As we discuss, there is an ambiguity in the equilibrium equations

of the Kirchhoff rod in which the torque density perpendicular to the fiber centerline can be recast into a force density; many studies refer to this parallel-torque-only version as “the Kirchhoff rod model” [44, Sec. 2.3]. In either case, the Kirchhoff model is a special case of the Cosserat model [28, 45], which also allows for extensibility and shear of the rod, and has recently been used to model cytoskeletal fibers without dynamics [46].

A theoretical ambiguity that we consider is how the fiber elasticity couples to the fluid mechanics when the fiber resists twist. When we consider bending forces only, then the formulation of slender body theory ensures that the rod moves with the local fluid velocity. But what about when we need to track twist elasticity – does the cross section also rotate with the local angular fluid velocity? If it does, then for the rod to be unsharable, the fluid velocity has to be constrained so that the translational velocity $\mathbf{U}(s)$ and angular velocity $\mathbf{\Psi}(s)$ are consistent with each other. Specifically, there are two possible velocities of the tangent vector: $\partial_t \boldsymbol{\tau} = \partial_s \mathbf{U}$, and $\partial_t \boldsymbol{\tau} = \mathbf{\Psi} \times \boldsymbol{\tau}$. If we require that the two match, i.e., if we impose the constraint $\partial_s \mathbf{U} = \mathbf{\Psi} \times \boldsymbol{\tau}$, we get what we refer to as the *Kirchhoff model* for historical reasons [41, 47, 48]. To avoid constraints, Lim et al. [49] introduced the *generalized* Kirchhoff rod model, in which the cross section is shearable and the fiber is extensible, both with an associated penalty. This model has been used frequently in the literature on fluid-fiber coupling [50, 51], but its consistency with the constrained Kirchhoff formulation as all penalty parameters go to infinity has not yet been established. Of course, imposing the constraint $\partial_s \mathbf{U} = \mathbf{\Psi} \times \boldsymbol{\tau}$ requires knowledge of the full rotational and translational fluid velocity, which so far has not been possible using SBT, since SBT cannot properly account for rot-trans coupling.

In the second class of fluid coupling methods, the redundancy in the possible motions is treated not by enforcing a constraint, but by ignoring the rotational fluid velocity in perpendicular directions. The rod then has only four possible motions: translational velocity $\mathbf{U}(s)$ and rotational velocity $\Omega^\parallel(s)$ about its axis. In this model, the tangent vector $\boldsymbol{\tau}$ is

updated by the translational velocity, and then a local triad or some other measure of twist is rotated around $\boldsymbol{\tau}$ by the velocity Ω^\parallel . This decouples solving for the translational and rotational velocity, simplifying the problem and reducing the size of the linear system to be solved. This model, which we shall refer to as the *Euler model*, is presented in the review article by Powers [33], and has been used in the simulation of both discrete [52–56] and continuous [57] elastic rods. All of these methods use an SBT-style mobility, in which nonlocal hydrodynamics is included in the translation-translation (trans-trans) coupling (to give the translational velocity \mathbf{U}), but only local drag is included for rotation-rotation (rot-rot, to give the parallel rotational velocity Ω^\parallel). Indeed, a perpendicular angular velocity is not available because no slender body theory exists for translation-rotation (trans-rot) or rotation-translation (rot-trans) coupling, so only parallel rotational velocity can be used.

To compare the two models of fluid coupling, we use both asymptotic arguments and numerical tests to show that the Kirchhoff and Euler models give the same translational velocity \mathbf{U} a distance $\mathcal{O}(1)$ away from the fiber endpoints as the slenderness $\epsilon \rightarrow 0$. The key idea is that the perpendicular torque in the Kirchhoff model, which enforces the constraint $\partial_s \mathbf{U} = \boldsymbol{\Psi} \times \boldsymbol{\tau}$, can be as small as $\mathcal{O}(\epsilon^2)$ and still generate a total angular velocity $\boldsymbol{\Psi}$ that is consistent with \mathbf{U} . As a result, these Lagrange multipliers have strength of order ϵ^2 , and therefore have a minimal impact on $\partial_s \mathbf{U}$, so they can be dropped in the Euler model without changing the resulting centerline velocity \mathbf{U} to $\mathcal{O}(\epsilon^2)$. Thus this part of the dissertation establishes the asymptotic equivalence of the Euler and Kirchhoff models as $\epsilon \rightarrow 0$. This sets us up for Part II, where we will design numerical methods around the simpler Euler model.

Chapter 2

Rotne-Prager-Yamakawa (RPY)

Mobility

The previous chapter showed that there has yet to be a consistent way to reduce the hydrodynamics of a spinning, bending filament from three dimensions to one in the case of the *mobility* problem, where the force and torque on the fiber centerline are prescribed and we seek the translational and rotational velocities. Since one-dimensional SBT fails, it seems logical to fall back to the three-dimensional boundary integral calculation in Section 1.1.2.

One potential approach for solving the three-dimensional boundary integral equation is the so-called “tubular body theory” (TBT) of Koens [58], which is designed to *exactly* solve for the traction on the fiber surface through a series of Fredholm integral equations of the second kind. These integral equations are similar to that of SBT, but are instead forced by a function of the previous order traction. The approach, however, offers little insight into the local resistance behavior under high rotation, and it is unclear how many iterations would be required to accurately capture the nonlocal translation-rotation coupling that is not correctly treated by SBT. For this reason, we seek an alternative approximation which takes advantage of the fiber slenderness to minimize the number of nonlocal integral evaluations and allows for dynamic simulations with many fibers. Our idea in particular is to

use regularized singularities to represent the mobility relationship, then compare the results we obtain to full three-dimensional boundary integral calculations.

A key contribution of this dissertation is to model the fiber using a line integral of regularized singularities, in particular Rotne-Prager-Yamakawa (RPY) singularities. Mathematically, the approach of the original slender body pioneers [19, 20, 23] was to distribute singularities along the centerline, then evaluate the flow from these singularities on a fiber cross section. The flow is nonsingular since it is *evaluated* a fiber radius away from the location of the singularities. In regularized singularities, we instead distribute the force and torque over a “cloud” or “blob” region of characteristic size $\hat{a} \sim a$, which is defined by some regularized delta function $\delta_{\hat{a}}$. Throughout this thesis, $\delta_{\hat{a}}$ will be a surface δ function on a sphere of radius \hat{a} , but other approaches are also possible. In this case, the velocity field is nonsingular because the force and torque are *distributed* over a region of finite size approximately equal to the fiber radius. We can then convolve the velocity field (or half the vorticity field) with the regularized delta function again to obtain the velocity (or angular velocity) of the fiber centerline. In this way, we derive fundamental regularized singularities for all components of the mobility operator, then perform asymptotics on line integrals of these regularized singularities. The sum of these components gives a *model* of the filament dynamics with the correct physical behavior that can be used in practice without approximating precisely a specific geometry; after all, the geometry of an actin filament [11] or flagella [59, Fig. 5.3] is not perfectly spheroidal or even tubular, and so one could argue that the regularized singularity model is just as much of an approximation as SBT.

Representing filaments using regularized singularities is not a novel idea, as it was first employed for numerical purposes in the immersed boundary (IB) method of the early 1970s to model the flow patterns inside the heart [60]. Later variations, including the force coupling method [61], method of regularized Stokeslets [62, 63], and Rotne-Prager-Yamakawa (RPY) tensor [64, 65], focused on choosing the regularized delta function such that the Stokes

equations are solvable analytically. While not a unique or necessarily superior choice, we prefer to use the RPY kernel because of its simple far-field analytical form (as a Stokeslet plus a source doublet), its symmetric positive definite property, and because of its long history of use in polymer suspensions (e.g., [47, 66–68]). The analysis here could be conducted for any regularization function where the mobility between two points can be determined analytically.

Recently, the accuracy of using regularized singularities to approximate slender *translating* filaments has come under scrutiny both from a numerical [69] and analytical [70–72] perspective. Generally speaking, the recent analysis has affirmed [70, 71] that regularization methods can effectively match the slender body theory *centerline and far-field fluid velocity* with a judicious choice of regularization radius \hat{a} , while others [71] have shown that the global flow field *near* the slender body cannot be matched unless very specific conditions on the regularization function are met [72, Sec. 4]. While the latter restriction is somewhat disappointing, it has little effect in practice: since regularized singularities can match the SBT filament centerline velocity, the only impact of the flow near the filament being incorrect is in computing disturbance flows on other nearly touching filaments. These disturbance flows change sharply over distances of $\mathcal{O}(\epsilon)$, which makes a one-dimensional SBT approach for nearly touching fibers quite challenging if not impractical. One of the goals of this paper is therefore to determine how the *centerline* motion induced by RPY singularities compares to that of SBT when rotation is considered. For instance, do they reproduce the relationship between torque and rotational velocity for a straight cylinder? And what about the feedback of rotational velocity on translational motion for curved filaments?

2.1 Slender body theory from RPY singularities

The idea of the regularized singularity approach is to model the force and torque on the fiber centerline with a regularized delta function $\delta_{\hat{a}}$, where \hat{a} is the “regularization radius” that is

chosen to model the three-dimensional problem. Our choice of kernel is the Rotne-Prager-Yamakawa (RPY) tensor [64, 65], for which $\delta_{\hat{a}}(\mathbf{r}) = \delta(r - \hat{a})/(4\pi\hat{a}^2)$, i.e., the regularized delta function is a surface delta function on a sphere of radius \hat{a} . This choice is more faithful to the original cylindrical geometry than delta functions with infinite support [62, 63, 73] or those whose support is tied to a numerical fluid grid [74]. In particular, since each regularized singularity is a sphere, it is not hard to imagine this resulting in a cylinder of constant radius \hat{a} along a length L , with hemispherical caps at each end. Still, since the regularized delta function is the same in the axial and radial directions, the spheres are not rings as we would need to make a proper cylinder. Also, the RPY tensor itself is an approximation to the dynamics of spheres, since it does not include stresslet terms required to keep the spheres rigid in flow [75], even though stresslets are of the same multipole order as torques, which are included. Because of these approximations, the intuitive picture of a cylinder composed of surface spheres is not quite the correct one, and we will require $\hat{a} > a$ for our RPY formulas to match those of SBT.

To obtain the centerline velocity of a filament made of regularized singularities, we first solve for the fluid velocity $\mathbf{u}(x)$ and corresponding pressure $\pi(\mathbf{x})$ by¹ solving the Stokes equations

$$\nabla\pi(\mathbf{x}) = \mu\nabla^2\mathbf{u}(\mathbf{x}) + \int_0^L \left(\left(\mathbf{f}(s) + \frac{\nabla}{2} \times n^{\parallel}(s)\boldsymbol{\tau}(s) \right) \delta_{\hat{a}}(\mathbf{x} - \mathbb{X}(s)) \right) ds, \quad (2.1)$$

and $\nabla \cdot \mathbf{u} = 0$. The translational and parallel rotational velocity of the fiber centerline are then given by convolving the pointwise velocity $\mathbf{u}(\mathbf{x})$ and half vorticity once more with the

¹It is important to emphasize again that $\mathbf{u}(\mathbf{x})$ and $\pi(\mathbf{x})$ are expected to be good approximations of the true fluid velocity and pressure only sufficiently far from the fiber centerline, but not $\mathcal{O}(\epsilon)$ away from it.

regularized delta function,

$$\mathbf{U}(s) = \int \mathbf{u}(\mathbf{x}) \delta_{\hat{a}}(\mathbb{X}(s) - \mathbf{x}) d\mathbf{x} \quad (2.2)$$

$$\Psi^{\parallel}(s) = \boldsymbol{\tau}(s) \cdot \int \delta_{\hat{a}}(\mathbb{X}(s) - \mathbf{x}) \frac{\nabla}{2} \times \mathbf{u}(\mathbf{x}) d\mathbf{x}. \quad (2.3)$$

This yields a mobility that is guaranteed to be symmetric positive definite [76, Footnote 3]. Here we only consider the parallel component of rotational velocity, since, for unsharable filaments, the other components can be derived from the translational velocity $\mathbf{U}(s)$ (see Chapter 3).

Using superposition, the translational and rotational velocity of the fiber centerline can be defined as integrals of the fundamental *regularized* singularities

$$\begin{aligned} \mathbf{U}(s) &= \int_0^L (\mathbb{M}_{\text{tt}}(\mathbb{X}(s), \mathbb{X}(s')) \mathbf{f}(s') + \mathbb{M}_{\text{tr}}(\mathbb{X}(s), \mathbb{X}(s')) n^{\parallel}(s') \boldsymbol{\tau}(s')) ds' \quad (2.4) \\ &= (\mathcal{M}_{\text{tt}} \mathbf{f})(s) + (\mathcal{M}_{\text{tr}} n^{\parallel} \boldsymbol{\tau})(s) \\ \Psi^{\parallel}(s) &= \boldsymbol{\tau}(s) \cdot \int_0^L (\mathbb{M}_{\text{rt}}(\mathbb{X}(s), \mathbb{X}(s')) \mathbf{f}(s') + \mathbb{M}_{\text{rr}}(\mathbb{X}(s), \mathbb{X}(s')) n^{\parallel}(s') \boldsymbol{\tau}(s')) ds' \\ &= \boldsymbol{\tau}(s) \cdot ((\mathcal{M}_{\text{rt}} \mathbf{f})(s) + (\mathcal{M}_{\text{rr}} n^{\parallel} \boldsymbol{\tau})(s)) \end{aligned}$$

where the fundamental regularized singularities $\mathbb{M}_{\text{tt}}, \mathbb{M}_{\text{tr}}, \mathbb{M}_{\text{rt}},$ and \mathbb{M}_{rr} are obtained by convolving the Stokeslet with $\delta_{\hat{a}}$, with a half curl taken each time rotation comes in,

$$\begin{aligned} \mathbb{M}_{\text{tt}}(\mathbf{x}, \mathbf{y}) &= \int \delta_{\hat{a}}(\mathbf{x} - \mathbf{z}) \int \mathbb{S}(\mathbf{z}, \mathbf{w}) \delta_{\hat{a}}(\mathbf{w} - \mathbf{y}) d\mathbf{w} d\mathbf{z}, \quad (2.5) \\ \mathbb{M}_{\text{tr}}(\mathbf{x}, \mathbf{y}) &= \int \delta_{\hat{a}}(\mathbf{x} - \mathbf{z}) \int \mathbb{S}(\mathbf{z}, \mathbf{w}) \frac{\nabla \mathbf{w}}{2} \times \delta_{\hat{a}}(\mathbf{w} - \mathbf{y}) d\mathbf{w} d\mathbf{z}, \\ \mathbb{M}_{\text{rt}}(\mathbf{x}, \mathbf{y}) &= \int \delta_{\hat{a}}(\mathbf{x} - \mathbf{z}) \frac{\nabla \mathbf{z}}{2} \times \int \mathbb{S}(\mathbf{z}, \mathbf{w}) \delta_{\hat{a}}(\mathbf{w} - \mathbf{y}) d\mathbf{w} d\mathbf{z} = \mathbb{M}_{\text{tr}}^T(\mathbf{y}, \mathbf{x}), \\ \mathbb{M}_{\text{rr}}(\mathbf{x}, \mathbf{y}) &= \int \delta_{\hat{a}}(\mathbf{x} - \mathbf{z}) \frac{\nabla \mathbf{z}}{2} \times \int \mathbb{S}(\mathbf{z}, \mathbf{w}) \frac{\nabla \mathbf{w}}{2} \times \delta_{\hat{a}}(\mathbf{w} - \mathbf{y}) d\mathbf{w} d\mathbf{z}. \end{aligned}$$

Using this approach, the four components of the mobility operator relating force and torque to angular and rotational velocity can be written *separately*, giving four different integrals for the four different pieces of the velocity in (2.4). For the RPY tensor, the velocity at point \mathbf{x} due to a force \mathbf{F} or torque \mathbf{N} at point \mathbf{y} is given by

$$(8\pi\mu) \mathbb{M}_{\text{tt}}(\mathbf{x}, \mathbf{y}) \mathbf{F} = \begin{cases} \left(\frac{\mathbf{I} + \widehat{\mathbf{R}}\widehat{\mathbf{R}}}{R} + \frac{2\hat{a}^2}{3} \frac{\mathbf{I} - 3\widehat{\mathbf{R}}\widehat{\mathbf{R}}}{R^3} \right) \mathbf{F} & R > 2\hat{a} \\ \left(\left(\frac{4}{3\hat{a}} - \frac{3R}{8\hat{a}^2} \right) \mathbf{I} + \frac{R}{8\hat{a}^2} \widehat{\mathbf{R}}\widehat{\mathbf{R}} \right) \mathbf{F} & R \leq 2\hat{a} \end{cases} \quad (2.6)$$

$$(8\pi\mu) \mathbb{M}_{\text{tr}}(\mathbf{x}, \mathbf{y}) \mathbf{N} = \begin{cases} \frac{\mathbf{N} \times \widehat{\mathbf{R}}}{R^2} & R > 2\hat{a} \\ \frac{1}{2\hat{a}^2} \left(\frac{R}{\hat{a}} - \frac{3R^2}{8\hat{a}^2} \right) (\mathbf{N} \times \widehat{\mathbf{R}}) & R \leq 2\hat{a} \end{cases} \quad (2.7)$$

$$(8\pi\mu) \mathbb{M}_{\text{rt}}(\mathbf{x}, \mathbf{y}) \mathbf{F} = \begin{cases} \frac{\mathbf{F} \times \widehat{\mathbf{R}}}{R^2} & R > 2\hat{a} \\ \frac{1}{2\hat{a}^2} \left(\frac{R}{\hat{a}} - \frac{3R^2}{8\hat{a}^2} \right) (\mathbf{F} \times \widehat{\mathbf{R}}) & R \leq 2\hat{a} \end{cases} \quad (2.8)$$

$$(8\pi\mu) \mathbb{M}_{\text{rr}}(\mathbf{x}, \mathbf{y}) \mathbf{N} = \begin{cases} -\frac{1}{2} \left(\frac{\mathbf{I} - 3\widehat{\mathbf{R}}\widehat{\mathbf{R}}}{R^3} \right) \mathbf{N} & R > 2\hat{a} \\ \frac{1}{\hat{a}^3} \left(\left(1 - \frac{27R}{32\hat{a}} + \frac{5R^3}{64\hat{a}^3} \right) \mathbf{I} + \left(\frac{9R}{32\hat{a}} - \frac{3R^3}{64\hat{a}^3} \right) \widehat{\mathbf{R}}\widehat{\mathbf{R}} \right) \mathbf{N} & R \leq 2\hat{a}, \end{cases} \quad (2.9)$$

where $\mathbf{R} = \mathbf{x} - \mathbf{y}$, $R = \|\mathbf{R}\|$, $\widehat{\mathbf{R}} = \mathbf{R}/\|\mathbf{R}\|$, and $\widehat{\mathbf{R}}\widehat{\mathbf{R}}$ denotes the outer product of $\widehat{\mathbf{R}}$ with itself. Note that the RPY tensor changes form for the case of overlapping spheres ($R \leq 2\hat{a}$), making them nonsingular and once continuously differentiable at $R = 2\hat{a}$. Because the kernels are still nearly singular, in the limit $\hat{a}/L := \hat{\epsilon} \ll 1$ it is more efficient to evaluate the integrals (2.4) asymptotically, since some of the singularities can safely be assumed to only contribute to the “inner” expansion. Indeed, by doing asymptotics on each of the four

integrals in (2.5) separately, we get an approximation to each term of the mobility to some order in ϵ . This is in contrast to the SBTs of Chapter 1, where the translational velocity on the cross section was some combination of contributions from force and torque, and we were left with trying to determine the order of each term from the desired order of \mathbf{U} and Ψ^{\parallel} on the cross section.

In this chapter, we therefore perform asymptotic expansions on the integrals (2.4) to obtain an SBT for a twisting filament. Our strategy is standard matched asymptotics and similar to the approaches of Johnson [20] and Gotz [32] for SBT. Because we use regularized singularities, however, all of the integrals take place on the fiber centerline, and so the asymptotics are in one dimension, rather than on the fiber cross section (two dimensions). It therefore becomes much simpler to obtain expressions for the velocity at the endpoints, where we only have to redefine the domain of integration in the inner expansion. Thus a physical endpoint regularization, which is an ingredient of any SBT-based numerical method (see Section 1.1.3, for example), comes with our choice to use regularized singularities as an approximation to the three-dimensional fiber geometry.

For each integral in (2.4), we first compute an outer expansion by considering the region where $|s - s'|$ is $\mathcal{O}(1)$. Then in the inner expansion, we consider the part of the integrals where $|s - s'|$ is $\mathcal{O}(\hat{a})$. In this case, we follow [32] and introduce the rescaled variable

$$\xi = \frac{s' - s}{\hat{a}}, \tag{2.10}$$

so that ξ is $\mathcal{O}(1)$ in the inner expansion, and the domain of ξ is $[-s/\hat{a}, (L - s)/\hat{a}]$. With this definition, it is straightforward to integrate over ξ and obtain an inner expansion, which is done in two pieces because the RPY tensor changes at $s = 2\hat{a}$ (the case of overlapping

spheres).² Using the rescaled variable ξ defined in (2.10), we write the series expansions

$$\begin{aligned}
\mathbf{R}(s') &= \mathbb{X}(s) - \mathbb{X}(s') = -\xi \hat{a} \boldsymbol{\tau}(s) - \frac{\xi^2 \hat{a}^2}{2} \partial_s \boldsymbol{\tau}(s) + \mathcal{O}(\hat{a}^3) \\
\mathbf{R}\mathbf{R} &= \xi^2 \hat{a}^2 \boldsymbol{\tau}(s) \boldsymbol{\tau}(s) + \mathcal{O}(\hat{a}^3), \\
R^2 &= \mathbf{R} \cdot \mathbf{R} = \xi^2 \hat{a}^2 + \mathcal{O}(\hat{a}^3), \quad R = |\xi| \hat{a} + \mathcal{O}(\hat{a}^2), \\
R^{-1} &= \frac{1}{|\xi| \hat{a}} + \mathcal{O}(1), \quad R^{-3} = \frac{1}{|\xi|^3 \hat{a}^3} + \mathcal{O}(\hat{a}^{-2}), \quad R^{-5} = \frac{1}{|\xi|^5 \hat{a}^5} + \mathcal{O}(\hat{a}^{-4}). \\
\mathbf{f}(s') &= \mathbf{f}(s) + \xi \hat{a} \partial_s \mathbf{f}(s) + \mathcal{O}(\hat{a}^2) \quad n^{\parallel}(s') = n^{\parallel}(s) + \xi \hat{a} \partial_s n^{\parallel}(s) + \mathcal{O}(\hat{a}^2),
\end{aligned} \tag{2.11}$$

which we substitute into the corresponding RPY velocity formulas (2.14), (2.26), (2.37), and (2.43). What results is two separate integrals over ξ : one for $|\xi| > 2$ (when the RPY spheres do not overlap), and another for $|\xi| \leq 2$ (when they do). The bounds on the integrals are

$$\begin{aligned}
\int_{|\xi| > 2} f(\xi) d\xi &= \begin{cases} \int_{-s/\hat{a}}^{-2} f(\xi) d\xi + \int_2^{(L-s)/\hat{a}} f(\xi) d\xi & 2\hat{a} < s < L - 2\hat{a} \\ \int_2^{(L-s)/\hat{a}} f(\xi) d\xi & s \leq 2\hat{a} \\ \int_{-s/\hat{a}}^{-2} f(\xi) d\xi & s \geq L - 2\hat{a} \end{cases} \\
\int_{|\xi| \leq 2} f(\xi) d\xi &= \begin{cases} \int_{-2}^{-2} f(\xi) d\xi & 2\hat{a} < s < L - 2\hat{a} \\ \int_{-s/\hat{a}}^2 f(\xi) d\xi & s \leq 2\hat{a} \\ \int_{-2}^{(L-s)/\hat{a}} f(\xi) d\xi & s \geq L - 2\hat{a} \end{cases}
\end{aligned} \tag{2.12}$$

With these integration bounds, the inner velocity can be computed for any s on the filament.

Adding the inner and outer solutions together and subtracting the common part then gives

²The break up of the integral into two pieces is reminiscent of the classic approach of Lighthill [77] (see also [4, Fig. 6]), who represented the flow due to a translating filament as the sum of an integral of Stokeslets (at distances greater than some intermediate length $q \gg a$) and another integral of Stokeslets and doublets (at distances smaller than q). The RPY approach is fundamentally different, however, because here the RPY regularization imposes the splitting of the integral at the distance $2\hat{a}$. For translation-translation, the doublet term can still make a contribution to the flow field on this lengthscale; its precise $\mathcal{O}(1)$ contribution is computed in (2.17).

the final matched asymptotic solution. Unless otherwise specified, all expansions will be carried out to $\mathcal{O}(\hat{\epsilon})$, i.e., in this section \approx denotes equality up to $\mathcal{O}(\hat{\epsilon})$, where the regularized aspect ratio $\hat{\epsilon} = \hat{a}/L$ is distinguished from the actual physical aspect ratio $\epsilon = a/L$.

2.1.1 Translation from force

We begin by expanding the translational velocity from force, which we define as

$$\mathbf{U}_f = \int_0^L \mathbb{M}_{\text{tt}}(\mathbb{X}(s), \mathbb{X}(s')) \mathbf{f}(s') ds'. \quad (2.13)$$

When we substitute the definition of \mathbb{M}_{tt} for RPY singularities from (2.6) into (2.13), we have the integral

$$\begin{aligned} 8\pi\mu\mathbf{U}_f(s) = & \int_{R>2\hat{a}} \left(\mathbb{S}(\mathbb{X}(s), \mathbb{X}(s')) + \frac{2\hat{a}^2}{3}\mathbb{D}(\mathbb{X}(s), \mathbb{X}(s')) \right) \mathbf{f}(s') ds' \\ & + \int_{R\leq 2\hat{a}} \left(\left(\frac{4}{3\hat{a}} - \frac{3R(s')}{8\hat{a}^2} \right) \mathbf{I} + \frac{1}{8\hat{a}^2 R(s')} (\mathbf{R}\mathbf{R})(s') \right) \mathbf{f}(s') ds', \end{aligned} \quad (2.14)$$

where $\mathbf{R}(s') = \mathbb{X}(s) - \mathbb{X}(s')$, and $R = \|\mathbf{R}\|$. The separation of the integrals captures the change in the RPY tensor when $R \leq 2\hat{a}$. We will assume that s is approximately an arclength parameterization, so that we can replace the bound $R > 2\hat{a}$ with $|s - s'| > 2\hat{a}$ going forward. This incurs a relative error on the order $\hat{\epsilon}$, which is the same order to which we carry the asymptotic expansions.

2.1.1.1 Outer expansion

In the outer expansion, we consider the part of the integral (2.14) where $|s - s'|$ is $\mathcal{O}(1)$. In this case, the doublet term in (2.14) is insignificant and we obtain the outer velocity by

integrating the Stokeslet over the fiber centerline,

$$8\pi\mu\mathbf{U}_f^{(\text{outer})}(s) = \int_{|s-s'|>2\hat{a}} \mathbb{S}(\mathbb{X}(s), \mathbb{X}(s')) \mathbf{f}(s') ds'. \quad (2.15)$$

The part of the kernel (2.14) for $R \leq 2\hat{a}$ makes no contribution to the outer expansion since $|s - s'|$ is $\mathcal{O}(\hat{a})$ there.

2.1.1.2 Inner expansion

Beginning with the trans-trans mobility (2.14), we start with the part of the integral in the region $R > 2\hat{a}$. For this we first need to integrate the inner expansion of the Stokeslet

$$\begin{aligned} \int_{R>2\hat{a}} \mathbb{S}(\mathbb{X}(s), \mathbb{X}(s')) \mathbf{f}(s') ds' &\approx \int_{|\xi|>2} \left(\frac{\mathbf{I} + \boldsymbol{\tau}(s)\boldsymbol{\tau}(s)}{|\xi|\hat{a}} \right) \mathbf{f}(s)\hat{a} d\xi \\ &= (\mathbf{I} + \boldsymbol{\tau}(s)\boldsymbol{\tau}(s)) \mathbf{f}(s) \int_{|\xi|>2} \frac{1}{|\xi|} d\xi. \end{aligned} \quad (2.16)$$

Unlike the Stokeslet, the doublet term is only included in the inner expansion. Its expansion is given by

$$\begin{aligned} \frac{2\hat{a}^2}{3} \int_{R>2\hat{a}} \mathbb{D}(\mathbb{X}(s), \mathbb{X}(s')) ds' &\approx \frac{2\hat{a}^2}{3} \int_{|\xi|>2} \left(\frac{\mathbf{I} - 3\boldsymbol{\tau}(s)\boldsymbol{\tau}(s)}{|\xi|^3\hat{a}^3} \right) \mathbf{f}(s)\hat{a} d\xi \\ &= \frac{2(\mathbf{I} - 3\boldsymbol{\tau}(s)\boldsymbol{\tau}(s)) \mathbf{f}(s)}{3} \int_{|\xi|>2} \frac{1}{|\xi|^3} d\xi. \end{aligned} \quad (2.17)$$

It still remains to include in the inner expansion the term for $R \leq 2\hat{a}$. For this we have the inner expansion of the first term

$$\int_{R\leq 2\hat{a}} \left(\frac{4}{3\hat{a}} - \frac{3R(s')}{8\hat{a}^2} \right) \mathbf{f}(s') ds' \approx \mathbf{f}(s) \int_{|\xi|\leq 2} \left(\frac{4}{3} - \frac{3|\xi|}{8} \right) d\xi. \quad (2.18)$$

Likewise the second term for $R \leq 2\hat{a}$ has the inner expansion

$$\int_{R \leq 2\hat{a}} \frac{1}{8\hat{a}^2 R(s')} (\mathbf{R}\mathbf{R})(s') \mathbf{f}(s') ds' \approx \boldsymbol{\tau}(s)\boldsymbol{\tau}(s) \mathbf{f}(s) \int_{|\xi| \leq 2} \frac{|\xi|}{8} d\xi. \quad (2.19)$$

Adding the terms (2.16), (2.17), (2.18), and (2.19), and computing integrals using the bounds (2.12), we obtain the complete inner expansion

$$\begin{aligned} 8\pi\mu \mathbf{U}_f^{(\text{inner})}(s) &= \left(a_S(s) + a_{\text{CLI}}(s) + \frac{2\hat{a}^2}{3} a_D(s) \right) \mathbf{f}(s) \\ &\quad + \left(a_S(s) + a_{\text{CLT}}(s) - 2\hat{a}^2 a_D(s) \right) (\boldsymbol{\tau}(s)\boldsymbol{\tau}(s)) \mathbf{f}(s) \\ a_S(s) &= \begin{cases} \ln\left(\frac{(L-s)s}{4\hat{a}^2}\right) & 2\hat{a} < s < L - 2\hat{a} \\ \ln\left(\frac{(L-s)}{2\hat{a}}\right) & s \leq 2\hat{a} \\ \ln\left(\frac{s}{2\hat{a}}\right) & s \geq L - 2\hat{a} \end{cases} \\ a_D(s) &= \begin{cases} \frac{1}{4\hat{a}^2} - \frac{1}{2s^2} - \frac{1}{2(L-s)^2} & 2\hat{a} < s < L - 2\hat{a} \\ \frac{1}{8\hat{a}^2} - \frac{1}{2(L-s)^2} & s \leq 2\hat{a} \\ \frac{1}{8\hat{a}^2} - \frac{1}{2s^2} & s \geq L - 2\hat{a} \end{cases} \\ a_{\text{CLI}}(s) &= \begin{cases} \frac{23}{6} & 2\hat{a} < s < L - 2\hat{a} \\ \frac{23}{12} + \frac{4s}{3\hat{a}} - \frac{3s^2}{16\hat{a}^2} & s \leq 2\hat{a} \\ \frac{23}{12} + \frac{4(L-s)}{3\hat{a}} - \frac{3(L-s)^2}{16\hat{a}^2} & s \geq L - 2\hat{a} \end{cases} \\ a_{\text{CLT}}(s) &= \begin{cases} \frac{1}{2} & 2\hat{a} < s < L - 2\hat{a} \\ \frac{1}{4} + \frac{s^2}{16\hat{a}^2} & s \leq 2\hat{a} \\ \frac{1}{4} + \frac{(L-s)^2}{16\hat{a}^2} & s \geq L - 2\hat{a} \end{cases} \end{aligned} \quad (2.20)$$

The velocity $\mathbf{U}_f^{(\text{inner})}$ is twice continuously differentiable, since it is the integral of a function which is once continuously differentiable.

2.1.1.3 Common part and total velocity

The common part is the outer velocity written in terms of the inner variables,

$$8\pi\mu\mathbf{U}_f^{(\text{common})}(s) = \int_{|s-s'|>2\hat{a}} \left(\frac{\mathbf{I} + \boldsymbol{\tau}(s)\boldsymbol{\tau}(s)}{|s-s'|} \right) \mathbf{f}(s) ds'. \quad (2.21)$$

The total velocity is then the sum of the inner and outer expansions, with the common part subtracted,

$$\mathbf{U}_f(s) = \mathbf{U}_f^{(\text{inner})}(s) + \mathbf{U}_f^{(\text{outer})}(s) - \mathbf{U}_f^{(\text{common})}(s). \quad (2.22)$$

This can be written as

$$\mathbf{U}_f(s) = \mathbf{U}_f^{(\text{inner})}(s) + \frac{1}{8\pi\mu} \int_{|s-s'|>2\hat{a}} \left(\mathbb{S}(\mathbb{X}(s), \mathbb{X}(s')) \mathbf{f}(s') - \left(\frac{\mathbf{I} + \boldsymbol{\tau}(s)\boldsymbol{\tau}(s)}{|s-s'|} \right) \mathbf{f}(s) \right) ds', \quad (2.23)$$

where $\mathbf{U}_f^{(\text{inner})}$ is defined in (2.20). The velocity (2.23) has a form similar to that of SBT, but with different bounds on the integral. Unlike SBT, the integral (2.23) is not a finite part integral, but a nearly singular integral that makes sense for each term separately. Numerical evidence suggests that, unlike traditional SBT, our asymptotic formula (2.23) keeps the eigenvalues of the mobility above zero (see Section 6.3).

To compare with SBT, however, it is advantageous to observe that the integrand in (2.23) is $\mathcal{O}(\hat{\epsilon})$ when $R \leq 2\hat{a}$, and so we can add the excluded part back in without changing the

asymptotic accuracy. This gives a velocity of the exact same form as SBT,

$$8\pi\mu\mathbf{U}_f(s) = (\mathbf{I} + \boldsymbol{\tau}(s)\boldsymbol{\tau}(s)) \mathbf{f}(s) \ln\left(\frac{(L-s)s}{4\hat{a}^2}\right) + 4\mathbf{f}(s) \quad (2.24)$$

$$+ \int_0^L \left(\mathbb{S}(\mathbb{X}(s), \mathbb{X}(s')) \mathbf{f}(s') - \left(\frac{\mathbf{I} + \boldsymbol{\tau}(s)\boldsymbol{\tau}(s)}{|s-s'|} \right) \mathbf{f}(s) \right) ds',$$

in the fiber interior, where we have dropped $\mathcal{O}(\hat{a}^2)$ terms from (2.20). The comparison of this formula to translational SBT will be given in Section 2.1.5.

2.1.2 Translation from torque

To complete our formulation for translational velocity, we next consider an asymptotic expansion of the velocity $\mathbf{U}_n(s)$ due to a torque density $n^\parallel(s)$ on the fiber centerline,

$$\mathbf{U}_n(s) = \int_0^L \mathbb{M}_{\text{tr}}(\mathbb{X}(s), \mathbb{X}(s')) n^\parallel(s') \boldsymbol{\tau}(s') ds'. \quad (2.25)$$

When we substitute the definition of \mathbb{M}_{tr} for RPY singularities from [65, Eq. (3.16)] into (2.25), we have the integral

$$8\pi\mu\mathbf{U}_n(s) = \int_{R>2\hat{a}} \frac{\boldsymbol{\tau}(s') \times \mathbf{R}(s')}{R(s')^3} n^\parallel(s') ds' \quad (2.26)$$

$$+ \frac{1}{2\hat{a}^2} \int_{R\leq 2\hat{a}} \left(\frac{1}{\hat{a}} - \frac{3R(s')}{8\hat{a}^2} \right) (\boldsymbol{\tau}(s') \times \mathbf{R}(s')) n^\parallel(s') ds',$$

where $\mathbf{R}(s') = \mathbb{X}(s) - \mathbb{X}(s')$. The first term in (2.26) is the Rotlet, while the second term is the RPY tensor when $R \leq 2\hat{a}$. As before, we will use $R \approx |s - s'|$ to modify the bounds on the integrals.

2.1.2.1 Outer expansion

In the outer expansion, we consider the part of the integral (2.26) where $|s' - s|$ is $\mathcal{O}(1)$. In this case the correction term is insignificant and we get the outer velocity

$$8\pi\mu\mathbf{U}_n^{(\text{outer})}(s) = \int_{|s-s'|>2\hat{a}} \frac{\boldsymbol{\tau}(s') \times \mathbf{R}(s')}{R(s')^3} n^\parallel(s') ds'. \quad (2.27)$$

2.1.2.2 Inner expansion

Beginning with the rotlet term for $R > 2\hat{a}$, we write the rotlet in terms of the inner variables as

$$\boldsymbol{\tau}(s') \times \mathbf{R}(s') \approx \frac{\xi^2 \hat{a}^2}{2} (\boldsymbol{\tau}(s) \times \partial_s \boldsymbol{\tau}(s)) \quad (2.28)$$

$$\int_{R>2\hat{a}} \frac{\boldsymbol{\tau}(s') \times \mathbf{R}(s')}{R(s')^3} n^\parallel(s') ds' \approx \int_{|\xi|>2} \frac{(\boldsymbol{\tau}(s) \times \partial_s \boldsymbol{\tau}(s))}{2|\xi|\hat{a}} n^\parallel(s) \hat{a} d\xi. \quad (2.29)$$

We next need to account for the term for $R \leq 2\hat{a}$, or the second term in (2.26). Using (2.28), we have that

$$\frac{1}{2\hat{a}^2} \int_{R<2\hat{a}} \left(\frac{1}{\hat{a}} - \frac{3R}{8\hat{a}^2} \right) (\boldsymbol{\tau}(s') \times \mathbf{R}(s')) ds' \approx \frac{(\boldsymbol{\tau}(s) \times \partial_s \boldsymbol{\tau}(s))}{4\hat{a}} \int_{|\xi|<2} \left(1 - \frac{3}{8}|\xi| \right) \xi^2 \hat{a} d\xi. \quad (2.30)$$

Adding (2.29) and (2.30) and using (2.12) to compute the inner expansion integrals, we obtain

$$8\pi\mu\mathbf{U}_n^{(\text{inner})}(s) = \frac{n^{\parallel}(s)}{2} (\boldsymbol{\tau}(s) \times \partial_s \boldsymbol{\tau}(s)) c_c(s) \quad (2.31)$$

$$c_c(s) = \begin{cases} \ln\left(\frac{(L-s)s}{4\hat{a}^2}\right) + \frac{7}{6} & 2\hat{a} < s < L - 2\hat{a} \\ \ln\left(\frac{(L-s)}{2\hat{a}}\right) + \frac{7}{12} + \frac{s^3}{6\hat{a}^3} - \frac{3s^4}{64\hat{a}^4} & s \leq 2\hat{a} \\ \ln\left(\frac{s}{2\hat{a}}\right) + \frac{7}{12} + \frac{(L-s)^3}{6\hat{a}^3} - \frac{3(L-s)^4}{64\hat{a}^4} & s \geq L - 2\hat{a} \end{cases}$$

which is a twice continuously differentiable function.

2.1.2.3 Common part and total velocity

The common part for trans-rot coupling is the inner expansion written in terms of the outer variables,

$$8\pi\mu\mathbf{U}_n^{(\text{common})}(s) = \int_{|s-s'|>2\hat{a}} \left(\frac{\boldsymbol{\tau}(s) \times \partial_s \boldsymbol{\tau}(s)}{2|s-s'|} \right) n^{\parallel}(s) ds'. \quad (2.32)$$

To $\mathcal{O}(\hat{\epsilon})$, the velocity due to torque $\mathbf{n}(s) = n^{\parallel}(s)\boldsymbol{\tau}(s)$ on the fiber centerline is given by

$$\mathbf{U}_n(s) = \mathbf{U}_n^{(\text{inner})}(s) + \mathbf{U}_n^{(\text{outer})}(s) - \mathbf{U}_n^{(\text{common})}(s) \quad (2.33)$$

$$\mathbf{U}_n(s) = \mathbf{U}_n^{(\text{inner})}(s) + \frac{1}{8\pi\mu} \int_{R>2\hat{a}} \left(\frac{\boldsymbol{\tau}(s') \times \mathbf{R}(s')}{R(s')^3} n^{\parallel}(s') - \frac{\boldsymbol{\tau}(s) \times \partial_s \boldsymbol{\tau}(s)}{2|s-s'|} n^{\parallel}(s) \right) ds', \quad (2.34)$$

We can make this velocity (2.34) more SBT-like by changing the bounds on the integral. Because the outer and common part match to $\mathcal{O}(\hat{\epsilon})$ when $R \leq 2\hat{a}$, we can change the limit

of integration in the integral to obtain a finite part integral, which gives the final result

$$8\pi\mu\mathbf{U}_n(s) = \frac{1}{2} (\boldsymbol{\tau}(s) \times \partial_s \boldsymbol{\tau}(s)) \left(\ln \left(\frac{s(L-s)}{4\hat{a}^2} \right) + \frac{7}{6} \right) \quad (2.35)$$

$$+ \int_0^L \left(\left(\frac{\boldsymbol{\tau}(s') \times \mathbf{R}(s')}{R(s')^3} \right) n^\parallel(s') - \left(\frac{\boldsymbol{\tau}(s) \times \partial_s \boldsymbol{\tau}(s)}{2|s-s'|} \right) n^\parallel(s) \right) ds',$$

for the velocity in the fiber interior. This is of exactly the same form as the first line of the singularity representation (1.10), but without the additional angle-dependent terms in the second line of (1.10).

2.1.3 Rotation from force

We next calculate the parallel rotational velocity $\Psi_f^\parallel(s)$ due to the force density $\mathbf{f}(s)$ on the fiber centerline,

$$\Psi_f^\parallel(s) = \boldsymbol{\tau}(s) \cdot \int_0^L \mathbb{M}_{\text{rt}}(\mathbf{X}(s), \mathbf{X}(s')) \mathbf{f}(s') ds'. \quad (2.36)$$

When we substitute the definition of \mathbb{M}_{tr} from [65, Eq. (3.16)] into (2.36), we have

$$8\pi\mu\Psi_f^\parallel(s) = \int_{R>2\hat{a}} \frac{(\mathbf{f}(s') \times \mathbf{R}(s'))}{R(s')^3} \cdot \boldsymbol{\tau}(s) ds' \quad (2.37)$$

$$+ \frac{1}{2\hat{a}^2} \int_{R<2\hat{a}} \left(\frac{1}{\hat{a}} - \frac{3R(s')}{8\hat{a}^2} \right) (\mathbf{f}(s') \times \mathbf{R}(s')) \cdot \boldsymbol{\tau}(s) ds'.$$

After using the triple cross product identity to rewrite $(\mathbf{f}(s') \times \mathbf{R}(s')) \cdot \boldsymbol{\tau}(s) = (\mathbf{R}(s') \times \boldsymbol{\tau}(s)) \cdot \mathbf{f}(s')$, the symmetry of the kernels \mathbb{M}_{rt} and \mathbb{M}_{tr} makes the asymptotics in the inner, outer,

and common expansions much the same as the previous section. The final result is that

$$8\pi\mu\Psi_{\mathbf{f}}^{\parallel}(s) = \Psi_{\mathbf{f}}^{\parallel,\text{inner}}(s) + \Psi_{\mathbf{f}}^{(\text{FP})}(s) \quad (2.38)$$

$$\Psi_{\mathbf{f}}^{\parallel,\text{inner}}(s) = \frac{1}{2}c_c(s)(\boldsymbol{\tau}(s) \times \partial_s\boldsymbol{\tau}(s)) \cdot \mathbf{f}(s) \quad (2.39)$$

$$\Psi_{\mathbf{f}}^{(\text{FP})}(s) = \int_0^L \left(\frac{\mathbf{R}(s') \times \boldsymbol{\tau}(s)}{R(s')^3} \cdot \mathbf{f}(s') - \frac{\boldsymbol{\tau}(s) \times \partial_s\boldsymbol{\tau}(s)}{2|s-s'|} \cdot \mathbf{f}(s) \right) ds'. \quad (2.40)$$

where $c_c(s)$ is defined in (2.31).

It is straightforward to show the symmetry (self-adjointness in L^2) of the mobility operator,

$$\int_0^L \Psi_{\mathbf{f}}^{\parallel}(s)n^{\parallel}(s) ds = \int_0^L \mathbf{U}_{\mathbf{n}}(s) \cdot \mathbf{f}(s) ds. \quad (2.41)$$

The only nontrivial part of this calculation is the rotlet term in the finite part integrals, in which the order of integration in s and s' must be swapped.

2.1.4 Rotation from torque

The final asymptotic expansion we need is the rotational velocity $\boldsymbol{\Psi}(s)$ due to the torque $n^{\parallel}(s)$. Because the doublet singularity has already been expanded in the translational case in (2.17), it will be convenient to work with a full vector rotational velocity and torque

$$\boldsymbol{\Psi}_{\mathbf{n}}(s) = \int_0^L \mathbb{M}_{\text{rr}}(\mathbb{X}(s), \mathbb{X}(s')) \mathbf{n}(s') ds', \quad (2.42)$$

and specialize to the case of parallel velocity and torque later. We substitute the definition of \mathbb{M}_{rr} from [65, Eq. (3.14)] into (2.42) to obtain the vector rotational velocity,

$$\begin{aligned} 8\pi\mu\boldsymbol{\Psi}_{\mathbf{n}}(s) = & -\frac{1}{2} \int_{R>2\hat{a}} \mathbb{D}(\mathbb{X}(s), \mathbb{X}(s')) \mathbf{n}(s') ds' \\ & + \frac{1}{\hat{a}^3} \int_{R\leq 2\hat{a}} \left(\left(1 - \frac{27R(s')}{32\hat{a}} + \frac{5R(s')^3}{64\hat{a}^3} \right) \mathbf{I} + \left(\frac{9}{32\hat{a}R(s')} - \frac{3R(s')}{64\hat{a}^3} \right) (\mathbf{R}\mathbf{R})(s') \right) \mathbf{n}(s') ds'. \end{aligned} \quad (2.43)$$

The first term in (2.43) is the doublet, while the second term is the RPY tensor for $R \leq 2\hat{a}$.

2.1.4.1 Outer expansion

In the outer expansion, we consider the part of the integral (2.43) where $|s' - s|$ is $\mathcal{O}(1)$,

$$8\pi\mu\Psi_{\mathbf{n}}^{(\text{outer})}(s) = -\frac{1}{2} \int_{|s-s'|>2\hat{a}} \mathbb{D}(\mathbb{X}(s), \mathbb{X}(s')) \mathbf{n}(s') ds'. \quad (2.44)$$

Importantly, the outer expansion (and the common expansion) is $\mathcal{O}(1)$. The inner expansion will be $\mathcal{O}(\hat{\epsilon}^{-2})$, so we will wind up dropping the entire outer expansion.

2.1.4.2 Inner expansion and total velocity

The inner expansion for (2.43) contains two parts: the doublet, and terms for $R \leq 2\hat{a}$. We have already seen the expansion of the doublet in (2.17). Multiplying the result in (2.17) by the appropriate prefactors, we obtain

$$-\frac{1}{2} \int_{R>2\hat{a}} \left(\frac{\mathbf{n}(s')}{R(s')^3} - 3 \frac{\mathbf{R}(s') (\mathbf{R}(s') \cdot \mathbf{n}(s'))}{R(s')^5} \right) ds' \approx -\frac{(\mathbf{I} - 3\boldsymbol{\tau}(s)\boldsymbol{\tau}(s)) \mathbf{f}(s)}{2\hat{a}^2} \int_{|\xi|>2} \frac{1}{|\xi|^3} d\xi, \quad (2.45)$$

which which is accurate to $\mathcal{O}(\hat{\epsilon}^{-1})$, although in the fiber interior that term actually cancels, leaving an $\mathcal{O}(\log \hat{\epsilon}^{-1})$ dependence.

In a similar way, the inner expansions of the $R \leq 2\hat{a}$ terms in the fiber interior are

$$\begin{aligned} \frac{1}{\hat{a}^3} \int_{R<2\hat{a}} \left(1 - \frac{27R(s')}{32\hat{a}} + \frac{5R(s')^3}{64\hat{a}^3} \right) \mathbf{n}(s') ds' &\approx \frac{\mathbf{n}(s)}{\hat{a}^2} \int_{|\xi|\leq 2} \left(1 - \frac{27}{32}|\xi| + \frac{5}{64}|\xi|^3 \right) d\xi \quad (2.46) \\ &\frac{1}{\hat{a}^3} \int_{R<2\hat{a}} \left(\left(\frac{9}{32aR(s')} - \frac{3R(s')}{64\hat{a}^3} \right) (\mathbf{R}\mathbf{R})(s') \right) \mathbf{n}(s') ds' \\ &\approx \frac{\boldsymbol{\tau}(s) (\boldsymbol{\tau}(s) \cdot \mathbf{n}(s))}{\hat{a}^2} \int_{|\xi|\leq 2} \left(\frac{18 - 3|\xi|^2}{64} \right) |\xi| d\xi \end{aligned}$$

with an error of $\mathcal{O}(\ln \hat{\epsilon})$ in the fiber interior and $\mathcal{O}(\hat{\epsilon}^{-1})$ at the endpoints.

Using the integral bounds (2.12) to integrate and sum (2.45) and (2.46), the complete inner expansion for $\Psi_{\mathbf{n}}$ is given by

$$8\pi\mu\Psi_{\mathbf{n}}^{(\text{inner})}(s) = (p_I(s)\mathbf{I} + p_{\tau}(s)\boldsymbol{\tau}(s)\boldsymbol{\tau}(s))\mathbf{n}(s), \quad \text{where} \quad (2.47)$$

$$p_I(s) = \frac{1}{\hat{a}^2} \begin{cases} \frac{9}{8} + \frac{\hat{a}^2}{4} \left(\frac{1}{s^2} + \frac{1}{(L-s)^2} \right) & 2\hat{a} < s < L - 2\hat{a} \\ \frac{9}{16} + \frac{s}{\hat{a}} - \frac{27s^2}{64\hat{a}^2} + \frac{5s^4}{256\hat{a}^4} + \frac{\hat{a}^2}{4(L-s)^2} & s \leq 2\hat{a} \\ \frac{9}{16} + \frac{L-s}{\hat{a}} - \frac{27(L-s)^2}{64\hat{a}^2} + \frac{5(L-s)^4}{256\hat{a}^4} + \frac{\hat{a}^2}{4s^2} & s \geq L - 2\hat{a} \end{cases}$$

$$p_{\tau}(s) = \frac{1}{\hat{a}^2} \begin{cases} \frac{9}{8} - \frac{3\hat{a}^2}{4} \left(\frac{1}{s^2} + \frac{1}{(L-s)^2} \right) & 2\hat{a} < s < L - 2\hat{a} \\ \frac{9}{16} + \frac{9s^2}{64\hat{a}^2} - \frac{3s^4}{256\hat{a}^4} - \frac{3\hat{a}^2}{4(L-s)^2} & s \leq 2\hat{a} \\ \frac{9}{16} + \frac{9(L-s)^2}{64\hat{a}^2} - \frac{3(L-s)^4}{256\hat{a}^4} - \frac{3\hat{a}^2}{4s^2} & s \geq L - 2\hat{a} \end{cases}$$

which is twice continuously differentiable.

Because the outer expansion is $\mathcal{O}(\hat{\epsilon}^2)$ smaller than the inner expansion, the total asymptotic velocity is just given by the inner expansion,

$$\Psi_{\mathbf{n}} = \Psi_{\mathbf{n}}^{(\text{inner})}. \quad (2.48)$$

Specializing now to the case of parallel angular velocity from parallel torque, $\mathbf{n} = n^{\parallel}\boldsymbol{\tau}$, and $\Psi_{\mathbf{n}}^{\parallel} = \Psi_{\mathbf{n}} \cdot \boldsymbol{\tau}$, and dropping the negligible terms in the fiber interior from (2.47), we obtain the simple formula

$$8\pi\mu\Psi_{\mathbf{n}}^{\parallel}(s) = \frac{9n^{\parallel}(s)}{4\hat{a}^2} \quad (2.49)$$

in the fiber interior.

2.1.5 Comparison to Keller-Rubinow-Johnson

We recall that the goal of our regularized singularity approach is to find \hat{a} such that the RPY-based filament centerline velocity matches that of SBT. With this in mind, let us compare our RPY-based slender body theory to the more rigorous three-dimensional SBT results we have from Section 1.1. We first compare the theories when we have a well-defined SBT, i.e., for translation-translation and rotation-rotation. Then, we use the form of the translation-rotation RPY coupling and our results in Section 1.1 to postulate an SBT equation for trans-rot coupling with a single unknown parameter. We will attempt to determine the unknown parameter in the next section.

2.1.5.1 Trans-trans and rot-rot

All three-dimensional slender body theories agree that the translational velocity in the case when $\Psi^\parallel = 0$ is given by

$$8\pi\mu\mathbf{U}_f = \ln\left(\frac{4s(L-s)}{a^2}\right)(\mathbf{I} + \boldsymbol{\tau}\boldsymbol{\tau})\mathbf{f} + (\mathbf{I} - 3\boldsymbol{\tau}\boldsymbol{\tau})\mathbf{f} + \mathbf{U}_f^{(\text{FP})} + \mathcal{O}(\epsilon), \quad (2.50)$$

while our RPY-based theory (2.24) gives

$$8\pi\mu\mathbf{U}_f = \ln\left(\frac{s(L-s)}{4\hat{a}^2}\right)(\mathbf{I} + \boldsymbol{\tau}\boldsymbol{\tau})\mathbf{f} + 4\mathbf{f} + \mathbf{U}_f^{(\text{FP})} + \mathcal{O}(\epsilon). \quad (2.51)$$

For the case of rotation from parallel torque, all theories of Section 1.1 agree that

$$8\pi\mu\Psi_n^\parallel = \frac{2n^\parallel}{a^2}, \quad (2.52)$$

while our RPY-based theory gives (2.49).

While we can set \hat{a} to match the two translational theories or the two rotational theories, there is no choice of \hat{a} that will match both. There are thus two possible candidates for \hat{a} .

If we take

$$\hat{a} = \hat{a}_{\text{tt}} = a \frac{e^{3/2}}{4} \approx 1.12a, \quad (2.53)$$

then the two translational theories match, but for rotation we obtain

$$8\pi\mu\Psi_{\text{n}}^{\parallel} = \frac{9n^{\parallel}(s)}{4\hat{a}_{\text{tt}}^2} = \frac{36n^{\parallel}(s)}{e^3a^2} \approx 1.79\frac{n^{\parallel}}{a^2}, \quad (2.54)$$

which is a 10% error from the known formula for a cylinder (2.48).

On the other hand, if we take

$$\hat{a} = \hat{a}_{\text{rr}} = \sqrt{9/8}a \approx 1.06a, \quad (2.55)$$

the rotational mobilities match, and the resulting RPY translational mobility becomes

$$\begin{aligned} 8\pi\mu\mathbf{U} &= \ln\left(\frac{s(L-s)}{4\hat{a}_{\text{rr}}^2}\right) (\mathbf{I} + \boldsymbol{\tau}\boldsymbol{\tau}) \mathbf{f} + 4\mathbf{f} + \mathbf{U}_{\text{f}}^{(\text{FP})} \\ &\approx \ln\left(\frac{4s(L-s)}{a^2}\right) (\mathbf{I} + \boldsymbol{\tau}\boldsymbol{\tau}) \mathbf{f} + (1.1\mathbf{I} - 2.9\boldsymbol{\tau}\boldsymbol{\tau}) \mathbf{f} + \mathbf{U}_{\text{f}}^{(\text{FP})}, \end{aligned}$$

which gives a small $\mathcal{O}(1)$ error relative to translational SBT (2.50). The need to choose a different regularized singularity radius for the two different mobility components shows that, despite the intuitive picture of making a cylinder out of spheres, our choice of spherically-symmetric regularized singularities is insufficient to exactly model a cylinder. It is clear, however, that the RPY approximation captures the correct form of all the relevant terms, but with $\mathcal{O}(1)$ error in the coefficients.

2.1.5.2 Trans-rot coupling

The more compelling question is whether we can use our findings to gain insights into the unknown $\mathcal{O}(1)$ coefficient in translation-rotation SBT. Our RPY-based result is (rearranging

(2.35) using properties of logs and rounding the $\mathcal{O}(1)$ coefficient to two decimal places)

$$8\pi\mu\mathbf{U}_n \approx \left(\ln \left(\frac{4s(L-s)}{\hat{a}^2} \right) - 1.61 \right) (\boldsymbol{\tau} \times \partial_s \boldsymbol{\tau}) \frac{n^\parallel}{2} + \mathbf{U}_f^{(\text{FP})}, \quad (2.56)$$

with the corresponding symmetric result for Ψ_f^\parallel in (2.38). Since Section 2.1.5.1 shows that the RPY models capture the essential physics of the interaction of the fiber with the fluid, but with different coefficients, we can postulate that the rotation-translation coupling in a true tubular geometry is also of the form (2.56), but with different coefficients. Furthermore, since there is an $\mathcal{O}(1)$ error in the boundary condition (1.3) in the result (1.10) of Johnson, we can expect that any modifications to (1.10) will be $\mathcal{O}(1)$ (not $\log \epsilon$). Combining these two assumptions, we postulate a rotation-translation coupling term of the form

$$8\pi\mu\mathbf{U}_n = \left(\ln \left(\frac{4s(L-s)}{a^2} \right) - k \right) (\boldsymbol{\tau} \times \partial_s \boldsymbol{\tau}) \frac{n^\parallel}{2} + \mathbf{U}_n^{(\text{FP})} + \mathcal{O}(\epsilon) \quad (2.57)$$

for some unknown constant k , with the corresponding symmetric form for Ψ_f^\parallel . In the next section, we fit the constant k to three dimensional boundary integral results. This gives us an empirical three-dimensional SBT for rot-trans coupling to which we can compare the RPY theory. It is important to note that our work does not rigorously justify (2.57) for a specific three-dimensional fiber geometry.

2.1.6 Comparing to three-dimensional boundary integral

Our RPY-based analysis leaves one uncertain coefficient, the $\mathcal{O}(1)$ coefficient k in the rot-trans coupling term (2.57), which we now attempt to extract using boundary integral simulations on slender bodies. Because there is only one coefficient to determine, we can fix the fiber geometry, vary the aspect ratio ϵ , and check that the error in SBT decreases with ϵ .

It will be convenient³ to fix $L = 2$ and work on $s \in [-1, 1]$ with an ellipsoidally-decaying radius function

$$\rho(s) = \sqrt{1 - s^2}. \quad (2.58)$$

The radius at each point s is $a\rho(s)$, where we will vary a to study the limit $\epsilon = a/L \rightarrow 0$.

As discussed in Section 1.1.2, under the assumption that the fiber is an infinitely thin membrane containing fluid of viscosity equal to that of the exterior fluid, the force density multiplied by the surface area element on the tubular boundary can be obtained by solving the integral equation

$$8\pi\mu\widehat{\mathbf{U}}(s, \theta) = \int_{-1}^1 \int_0^{2\pi} \mathbb{S} \left(\widehat{\mathbb{X}}(s, \theta), \widehat{\mathbb{X}}(s', \theta') \right) \widehat{\mathbf{f}}(s', \theta') d\theta' ds', \quad (2.59)$$

for $\widehat{\mathbf{f}}(s, \theta)$, where $\widehat{\mathbb{X}}$ and $\widehat{\mathbf{U}}$ are defined in (1.1) and (1.2).

In order to compare the BI output to SBT, we use the following sequence of steps:

1. Given the surface velocity $\widehat{\mathbf{U}}(s, \theta) = \mathbf{U}(s) + a\Psi^{\parallel}(s) (\boldsymbol{\tau}(s) \times \widehat{\mathbf{r}}(s, \theta))$, solve (2.59) using the numerical method of Malhotra and Barnett [79] to obtain the surface force density $\widehat{\mathbf{f}}(s, \theta)$.
2. Compute the centerline force $\mathbf{f}(s)$ and parallel torque $n^{\parallel}(s)$ from $\widehat{\mathbf{f}}(s, \theta)$ via (1.14) and (1.15).

³While the result (2.56) was derived by assuming a constant regularization radius \hat{a} , it is unchanged to leading order when considering a spatially-varying radius function $\hat{a}(s)$, for which the polydisperse RPY kernel must be used [78]. The choice of an ellipsoidally-decaying radius function removes any singular behavior of (2.57) near the endpoints and speeds up the convergence of the boundary integral method (relative to a cylinder with spherical caps, for example).

3. Compute the SBT velocity induced from the BI force as

$$8\pi\mu\mathbf{U}_{\text{SB}} = \ln(\epsilon^{-2}) (\mathbf{I} + \boldsymbol{\tau}\boldsymbol{\tau}) \mathbf{f} + (\mathbf{I} - 3\boldsymbol{\tau}\boldsymbol{\tau}) \mathbf{f} + \mathbf{U}_{\text{f}}^{(\text{FP})} \quad (2.60)$$

$$+ (\ln(\epsilon^{-2}) - k) (\boldsymbol{\tau} \times \partial_s \boldsymbol{\tau}) \frac{n^{\parallel}}{2} + \mathbf{U}_{\text{n}}^{(\text{FP})},$$

$$8\pi\mu\Psi_{\text{SB}}^{\parallel} = \frac{2n^{\parallel}}{a^2\rho^2} + \frac{1}{2} (\ln(\epsilon^{-2}) - k) (\boldsymbol{\tau} \times \partial_s \boldsymbol{\tau}) \cdot \mathbf{f} + \Psi_{\text{f}}^{(\text{FP})}, \quad (2.61)$$

where the finite part integrals are defined in (1.7), (1.11), and (2.40), and we have used the local drag coefficients for ellipsoidally-decaying radius (2.58).

4. Calculate the L^2 norm of the error $\mathbf{U}_{\text{SB}} - \mathbf{U}$ and $\Psi_{\text{SB}}^{\parallel} - \Psi^{\parallel}$, and plot it as a function of ϵ .

Note that we apply the forward SBT operator, which gives velocity from force/torque, in (2.60) and (2.61). This avoids inverting the slender body mobility, which is an ill-posed problem [22, 32]. To evaluate the finite part integrals, we use the singular quadrature scheme developed in [80] (see Section 6.2.1).

We have verified that the error in \mathbf{U}_{SB} and $\Psi_{\text{SB}}^{\parallel}$ decreases with ϵ in cases when rot-trans coupling is negligible (i.e., cases when $\Psi_{\text{SB}}^{\parallel} = 0$ or the fiber is straight). Having done this, our focus will be on the case of curved filaments with parallel rotational velocity $\Psi^{\parallel} \sim 1/\epsilon^2$, so that the magnitude of the parallel torque n^{\parallel} is unchanged to leading order as we change ϵ . In this case, which is the one that exposed the flaws in the SBTs of Section 1.1, the rot-trans coupling term in (2.60) is important, while the one in (2.61) is negligible. Our goal is to find a k in (2.60) such that the formulas are accurate to $\mathcal{O}(\epsilon)$; that is, the relative errors in both $\|\mathbf{U}_{\text{SB}} - \mathbf{U}\|$ and $\Psi_{\text{SB}}^{\parallel} - \Psi^{\parallel}$ are $\mathcal{O}(\epsilon)$.

To extract the unknown coefficient k in the trans-rot coupling term in (2.60), we consider a helix with $L = 1$ and an arbitrary number of turns

$$\mathbb{X}(s) = \frac{1}{\sqrt{2}} \left(\frac{1}{\pi n_t} \cos(\pi n_t s), \frac{1}{\pi n_t} \sin(\pi s n_t), s \right), \quad (2.62)$$

which spins with angular velocity $\Psi^\parallel \equiv 1/a^2$ and without moving, $\mathbf{U} = 0$ in (1.2). Because the fiber is not straight, the torque required to generate Ψ^\parallel will induce a translational velocity, and a nonzero average force \mathbf{f} must cancel this translational velocity to satisfy the boundary condition (1.3). This means that the translation-rotation coupling term in (2.60) will play a key role in the accuracy of our SBT. Because of the large torque, the rotational velocity in (2.61) is dominated by the rotation-rotation mobility, and so to study translation-rotation coupling we will focus only on the translational velocity (2.60).

In the months leading up to the publication of this dissertation, new progress has been made in developing boundary integral solvers for flexible slender bodies in Stokes flow. The work of Malhotra and Barnett [79] in particular, which is based on a combined single-double layer formulation, allows us to extract the force and torque on the centerline with arbitrary accuracy.⁴ Thus the results we obtain using the method of [79] can be trusted with higher confidence than those of our study [E], which applied a first-order numerical method based on a singularity subtraction technique previously introduced by Batchelor [82] and used by Koens [58]. We find in particular that both numerical methods agree on the force for translation only and the torque for rotation, but discrepancies are produced in the rot-trans coupling terms. The source of these discrepancies is unclear at this time.

Armed with the novel numerical method of [79], we extract the centerline force and torque on the spinning helical geometry (2.62) with $n_t = 1$ and $n_t = 2$. The results in Fig. 2.1 show fairly conclusively that $k = 1$ gives an error which decays with rate ϵ . These results are *not* well-aligned with what we found in our first order numerical method in [E], where we found $k \approx 2.85$ to give the $\mathcal{O}(\epsilon)$ decay. Thus more work is required to establish definitively that the “true” rot-trans coupling coefficient in (2.57) is $k = 1$. A suggested program would be to repeat this analysis on hundreds of random fiber shapes, using the novel boundary integral solver in [79]. If $k = 1$ is indeed the correct coefficient, then using an RPY radius of $\hat{a} \approx 0.73a$

⁴For rigid fibers, the method of Keaveny and Shelley [81], which is based on using a double layer representation plus a completion flow, is another option.




Figure 2.1: Error in the rot-trans SBT relative to the boundary integral calculation for various ϵ and k for the n_t turn helix (2.62). For each fiber, we prescribe a rotational velocity $\Psi^\parallel \equiv 1/a^2$ with translational velocity $\mathbf{U} = \mathbf{0}$. We then solve the boundary integral equation (2.59) for the surface traction (using the numerical method of [79]), use (1.14) and (1.15) to obtain the centerline force density and parallel torque density, and then (2.60) to obtain an SBT velocity \mathbf{U}_{SB} from these densities. We plot the L^2 norm of \mathbf{U}_{SB} (which should be zero) for several different values of k . For $k \neq 1$, $n_t = 1$, and the error is clearly $\mathcal{O}(1)$ with respect to ϵ . For $k = 1$, the error decreases with ϵ at a rate of roughly ϵ^1 (dashed black line), independent of the coefficient n_t in (2.62).

in (2.56) gives the formula (2.57) with $k = 1$. This value of the RPY radius is comparable, but fairly smaller than, what we obtained by matching the trans-trans ($\hat{a}_{\text{tt}} \approx 1.12a$) and rot-rot ($\hat{a}_{\text{rr}} \approx 1.06a$) mobilities in Section 2.1.5.

2.2 Extension to multiple filaments

It remains to include in the fiber centerline velocity (2.4) the perturbed flow due to other filaments, i.e., to account for hydrodynamic interactions between fibers. For translation, the simplest approach for including hydrodynamic interactions in SBT is the one taken by Tornberg and Shelley [22], in which the fluid velocity due to one fiber (1.4) is simply evaluated on the centerline of the other fibers. This approach presents a contradiction: the contribution of the doublet term is $\mathcal{O}(\epsilon^2)$ if the fibers are $\mathcal{O}(1)$ apart, but the same term

makes an $\mathcal{O}(1)$ if the fibers are $\mathcal{O}(\epsilon)$ apart. In the latter case, the flow field due to fiber i will not be constant on the cross section of fiber j , and it becomes totally unclear how to assign the centerline velocity on fiber j due to fiber i . Recognizing that the disturbance velocity field only makes sense $\mathcal{O}(1)$ away from the fiber centerline, Nazockdast et al. dropped the doublet term completely and included only the Stokeslet term [15].

An important advantage to our regularized singularity approach is that the velocity (2.4) is well-defined for both a single fiber and multiple fibers. Indexing the i th fiber by the superscript $\mathbf{X}^{(i)}$, the full velocity at coordinate s on fiber i can be written as

$$\begin{aligned} \mathbf{U}^{(i)}(s) = \mathbf{u}_0 \left(\mathbf{X}^{(i)}(s) \right) + \sum_{j=1}^F \int_0^L \left(\mathbb{M}_{\text{tt}} \left(\mathbb{X}^{(i)}(s), \mathbb{X}^{(j)}(s') \right) \mathbf{f}^{(j)}(s') \right. \\ \left. + \mathbb{M}_{\text{tr}} \left(\mathbb{X}^{(i)}(s), \mathbb{X}^{(j)}(s') \right) n_{\parallel}^{(j)}(s') \boldsymbol{\tau}^{(j)}(s') \right) ds' \end{aligned} \quad (2.63)$$

$$\begin{aligned} \Psi_{\parallel}^{(i)}(s) = \Psi_0^{\parallel} \left(\mathbb{X}^{(i)}(s) \right) + \boldsymbol{\tau}^{(i)}(s) \cdot \sum_{j=1}^F \int_0^L \left(\mathbb{M}_{\text{rt}} \left(\mathbb{X}^{(i)}(s), \mathbb{X}^{(j)}(s') \right) \mathbf{f}^{(j)}(s') \right. \\ \left. + \mathbb{M}_{\text{rr}} \left(\mathbb{X}^{(i)}(s), \mathbb{X}^{(j)}(s') \right) n_{\parallel}^{(j)}(s') \boldsymbol{\tau}^{(j)}(s') \right) ds', \end{aligned} \quad (2.64)$$

where F is the number of fibers, and \mathbf{u}_0 and Ψ_0^{\parallel} represent background flows and rotational velocities. Simply put (2.63) and (2.64) say that the fiber velocity of fiber i is the integral of the RPY kernel over itself plus the same integral over all the other fibers. When the fibers are well-separated, the self integral will be nearly-singular (special quadrature required), while the integrals over the other fibers will be relatively smooth (direct quadrature acceptable).

We will compactly write the velocities (2.63) and (2.64) for the entire collection of fibers as

$$\mathbf{U} - \mathbf{u}_0(\mathbf{X}) = \mathcal{M}_{\text{tt}}[\mathbf{X}] \mathbf{f} + \mathcal{M}_{\text{tr}}[\mathbf{X}] n^{\parallel} \quad (2.65)$$

$$\Psi^{\parallel} - \Psi_0^{\parallel} = \mathcal{M}_{\text{rt}}[\mathbf{X}] \mathbf{f} + \mathcal{M}_{\text{rr}}[\mathbf{X}] n^{\parallel} \quad (2.66)$$

This mobility equation can be closed by defining a constitutive equation for the fiber force densities \mathbf{f} and torque densities n^{\parallel} . This is covered in Chapter 3.

Chapter 3

Fiber mechanics

In the last chapter, we looked at how to model the motion of a slender fiber in a viscous fluid when the force \mathbf{f} and torque n^{\parallel} on the centerline are prescribed functions of arclength s . In practice, these quantities come from the internal fiber mechanics, which is the focus of this chapter. We will restrict our focus to a single fiber, as the generalization to multiple fibers is straightforward.

At every instant in time, the fiber resists bending with force density (per unit length) $\mathbf{f}^{(\kappa)}$ and twisting with torque density $n^{(\gamma)}$. When a fiber is not straight, there is also a force density associated with twist elasticity, which we refer to as $\mathbf{f}^{(\gamma)}$. Inextensibility can be enforced by introducing Lagrange multiplier force densities $\boldsymbol{\lambda}$. Thus the PDE that we need to solve on every fiber to obtain the translational velocity is given by (using the abstract notation of (2.65)),

$$\partial_t \mathbb{X} - \mathbf{u}_0(\mathbb{X}, t) = \mathcal{M}_{\text{tt}}[\mathbb{X}] \left(\mathbf{f}^{(\kappa)}[\mathbb{X}] + \mathbf{f}^{(\gamma)}[\mathbb{X}] + \boldsymbol{\lambda} \right) + \mathcal{M}_{\text{tr}}[\mathbb{X}] n^{(\gamma)}[\mathbb{X}] \quad (3.1)$$

where the mobility operator \mathcal{M} is defined in (2.65) and the background flow function \mathbf{u}_0 can in general vary in time. The fibers are also constrained to be inextensible, so that for

every fiber

$$\boldsymbol{\tau}(s, t) \cdot \boldsymbol{\tau}(s, t) = 1, \tag{3.2}$$

for all s and t . We still need to specify boundary conditions for the evolution equation (3.1) and additional conditions on $\boldsymbol{\lambda}$ to make the solution unique, as we explain in Section 3.1. In addition, there is a question of how to track and evolve the *twist* density of the fiber. Should we fix an orthonormal triad to the fiber centerline [47, 49, 50] to track twist? Intuitively, it seems that a simpler metric would be a scalar angle by which the fiber is twisted relative to a twist-free frame. In Section 3.2, we discuss the compatibility of these two choices and give the equations governing the latter.

3.1 Inextensibility

The focus of this section is on a new formulation for inextensible filaments. In our approach, the fibers are evolved via a rotation of the tangent vector on the unit sphere, and the fiber positions are then obtained by integration. This approach is similar to that of [41], but unique because we consider the fiber as a continuum, rather than a collection of discrete line segments. In this way, we maintain strict inextensibility of the fibers without introducing a penalty parameter. To close our formulation, we treat the force due to tension as a Lagrange multiplier and enforce the principle that the constraint forces do no work [42]. This section presents two possibilities for the virtual work constraint which are equivalent in continuum: one is a “strong” formulation based on enforcing the virtual work constraint pointwise, while another is a “weak” formulation which enforces the virtual work constraint via integrals against a set of basis functions. The former is better suited for numerical calculations, and it is what we use for our discretization in Part II.

3.1.1 Tension-based formulation

Prior to deriving our new formulation of inextensibility, we review the formulation traditionally used in the literature, around which many previous numerical methods [15, 22, 39, 40] are designed. In this “tension-based” formulation [22, 82, 83], the inextensibility constraint (3.2) is differentiated with respect to time. Then, s and t derivatives are interchanged to yield

$$\partial_s (\partial_t \mathbb{X}) \cdot \boldsymbol{\tau} = 0. \quad (3.3)$$

The system is closed by substituting the mobility equation (3.1) into the differentiated inextensibility constraint (3.3) and assuming that $\boldsymbol{\lambda} = \partial_s (T\boldsymbol{\tau})$, where $T(s, t)$ is an unknown scalar tension [22]. This results in the *line tension equation*,

$$\partial_s \left(\mathbf{u}_0(\mathbb{X}) + \mathcal{M}_{tt}[\mathbb{X}] \left(\mathbf{f}^{(\kappa)}[\mathbb{X}] + \mathbf{f}^{(\gamma)}[\mathbb{X}] + \partial_s (T\boldsymbol{\tau}) \right) + \mathcal{M}_{tr}[\mathbb{X}] n^{(\gamma)}[\mathbb{X}] \right) \cdot \boldsymbol{\tau} = 0. \quad (3.4)$$

which is a second-order linear BVP in T . Despite being linear in T , (3.4) is still highly nonlinear in \mathbb{X} , since the bending force density $\mathbf{f}^{(\kappa)}$ involves fourth derivatives of \mathbb{X} . Even in the absence of any nonlocal hydrodynamic interactions and zero background flow ($\mathbf{u}_0 = \mathbf{0}$), the line tension equation still has terms of the form $\partial_s^3 \mathbb{X} \cdot \partial_s^3 \mathbb{X}$ (see [22, Eq. (13)]), which lead to aliasing errors in spectral numerical methods. In addition, solving the line tension equation leaves unclear exactly how the mobility matrix \mathcal{M}_{tt} is modified when the inextensibility constraint is considered, which makes it difficult to design schemes for fluctuating filaments which preserve fluctuation-dissipation balance. Because of these complications, we seek an alternative route where we dispense with the scalar T and solve directly for $\boldsymbol{\lambda}$, while at the same time obtaining a closed-form form mobility $\partial_t \mathbb{X} = \mathcal{N}\mathbf{f}$ in the case of inextensible filaments.

3.1.2 Kinematic formulation

In our approach, we evolve the tangent vector $\boldsymbol{\tau}(s, t)$, rather than $\mathbb{X}(s, t)$, which can be obtained from $\boldsymbol{\tau}$ via integration (plus a constant which also needs to be tracked). Considering the evolution of $\boldsymbol{\tau}(s, t)$, the differentiated inextensibility constraint (3.3) implies that, for every fiber,

$$\partial_t \boldsymbol{\tau}(s, t) = \boldsymbol{\Omega}(s, t) \times \boldsymbol{\tau}(s, t), \quad (3.5)$$

i.e., that the fiber evolution can be thought of as rotations of $\boldsymbol{\tau}$ on the unit sphere. The constrained motion (3.5) implies a corresponding constrained motion for \mathbb{X} ,

$$\partial_t \mathbb{X}(s, t) = \mathbf{U}_{\text{MP}} + \int_{L/2}^s \partial_t \boldsymbol{\tau}(s', t) dt := (\boldsymbol{\mathcal{K}}\boldsymbol{\alpha})(s), \quad (3.6)$$

where \mathbf{U}_{MP} is the velocity of the midpoint, at which we have chosen to center the kinematic operator to preserve symmetry. In this section, we present two possible definitions of $\boldsymbol{\mathcal{K}}$ which differ in their dimension, depending on whether we remove the null space of (3.5) explicitly. The first of these is the most straightforward for numerical calculations, while the second is subject to aliasing in spectral methods. The second might still be useful from a theoretical point of view, so we present it here.

3.1.2.1 A rank-deficient kinematic operator

In the first approach, we define $\boldsymbol{\alpha} = \{\mathbf{U}_{\text{MP}}, \boldsymbol{\Omega}\}$ and set

$$(\boldsymbol{\mathcal{K}}\boldsymbol{\alpha})(s) = \mathbf{U}_{\text{MP}} + \int_{L/2}^s \boldsymbol{\Omega}(s') \times \boldsymbol{\tau}(s') ds', \quad (3.7)$$

where here we allow $\boldsymbol{\Omega}$ to have a parallel component. Thus the null-space of $\boldsymbol{\mathcal{K}}$ is non-empty.

3.1.2.2 Modifications for clamped ends

While our default is to consider filaments with two free ends, in some parts of this dissertation we will also consider filaments with a *clamped* end at $s = 0$ and a free end at $s = L$. For these filaments, the \mathcal{K} operator has to be modified to account for motion at $s = 0$ being prohibited. For a clamped end, it is more logical to center the dynamics around $s = 0$ instead of $s = L/2$, and therefore define the kinematic operator $\mathcal{K} \equiv \mathcal{K}_c$ by

$$\mathcal{K}_c \alpha := \mathcal{K}_c \Omega = \int_0^s \Omega(s') \times \tau(s') ds'. \quad (3.8)$$

While it is also possible to build the clamped condition $\tau(0) \equiv \tau_0$ into \mathcal{K}_c by restricting $\Omega^\perp(s = 0) = \mathbf{0}$, doing this causes tension to be ill-defined at $s = 0$. We choose not to enforce $\Omega^\perp(0) = \mathbf{0}$ kinematically, but rather mechanically through the bending force. The ambiguity in enforcing $\Omega^\perp(0) = \mathbf{0}$ kinematically and/or mechanically has a physical origin in the nonzero length of a clamp. If we consider $s = 0$ to be “inside” the clamp, then $\Omega^\perp(0) = \mathbf{0}$, but if $s = 0$ is “outside” the clamp, then $\Omega^\perp(0)$ is free.

3.1.2.3 Removing the null space

We now present a formulation which explicitly removes the null space from \mathcal{K} by parameterizing $\Omega(s)$ using an orthonormal triad $(\tau(s), \mathbf{n}_1(s), \mathbf{n}_2(s))$ at each point. To do so, we represent the unit tangent vector $\tau(s, t)$ using spherical angles $\theta(s, t)$ and $\phi(s, t)$ as

$$\tau(s, t) = \begin{pmatrix} \cos \theta \cos \phi \\ \sin \theta \cos \phi \\ \sin \phi \end{pmatrix}, \quad (3.9)$$

where we define θ to be single-valued at $\phi = \pi/2$ by setting $\theta(\phi = \pm\pi/2) = 0$. A choice of normal vectors that are always orthonormal to $\boldsymbol{\tau}$ on the unit sphere is

$$\mathbf{n}_1 = \begin{pmatrix} -\sin \theta \\ \cos \theta \\ 0 \end{pmatrix}, \quad \mathbf{n}_2 = \begin{pmatrix} -\cos \theta \sin \phi \\ -\sin \theta \sin \phi \\ \cos \phi \end{pmatrix}. \quad (3.10)$$

Because \mathbf{n}_1 and \mathbf{n}_2 can be determined uniquely from $\boldsymbol{\tau}$, we denote them henceforth with $\mathbf{n}_j(\boldsymbol{\tau}(s,t))$, for $j = 1, 2$. Since θ is single-valued at $\phi = \pi/2$, each component of the orthonormal coordinate system $(\boldsymbol{\tau}, \mathbf{n}_1, \mathbf{n}_2)$ is a smooth function of \mathbb{X} when $\boldsymbol{\tau}$ is smooth.

If we separate the rotation of the centerline about itself from inextensible translational motions, $\boldsymbol{\Omega}(s,t)$ can be restricted to linear combinations of \mathbf{n}_1 and \mathbf{n}_2 . We let

$$\boldsymbol{\Omega}(s,t) := \boldsymbol{\Omega}(\boldsymbol{\tau}(s,t), \mathbf{g}(s,t)) := g_1(s,t)\mathbf{n}_2(\boldsymbol{\tau}(s,t)) - g_2(s,t)\mathbf{n}_1(\boldsymbol{\tau}(s,t)), \quad (3.11)$$

where $g_1(s,t)$ and $g_2(s,t)$ are two specific unknown functions and $\mathbf{g} = \{g_1, g_2\}$. Equation (3.11) implies that, by the right-handedness of the coordinate system $(\boldsymbol{\tau}, \mathbf{n}_1, \mathbf{n}_2)$,

$$\partial_t \boldsymbol{\tau} = \boldsymbol{\Omega} \times \boldsymbol{\tau} = g_1 \mathbf{n}_1 + g_2 \mathbf{n}_2. \quad (3.12)$$

Any inextensible velocity of the fiber centerline can now be written in the form

$$\mathbf{U}(s,t) = \partial_t \mathbb{X}(s,t) = \mathbf{U}_{\text{MP}}(t) + \int_{L/2}^s \sum_{j=1}^2 g_j(s',t) \mathbf{n}_j(\boldsymbol{\tau}(s',t)) ds'. \quad (3.13)$$

To derive a “weak” formulation of inextensibility, we expand the functions $g_1(s)$ and $g_2(s)$ in some basis as

$$g_j(s) = \sum_k \alpha_{jk} T_k(s), \quad \text{for } j = 1, 2, \quad (3.14)$$

where $T_k(s)$ are sufficiently smooth scalar-valued basis functions for $L^2 : [0, L]$. Substituting the basis function expansion (3.14) into the inextensible velocity (3.13), we obtain

$$\mathbf{U}(s) = \partial_t \mathbb{X}(s) = \mathbf{U}_{\text{MP}} + \int_{L/2}^s \sum_{j=1}^2 \sum_k \alpha_{jk} T_k(s') \mathbf{n}_j(\boldsymbol{\tau}(s')) ds' := (\mathcal{K}\boldsymbol{\alpha})(s). \quad (3.15)$$

where we have used $\boldsymbol{\alpha} = (\alpha_{jk}, \bar{\mathbf{U}})$ to parameterize the space of inextensible fiber motions.

3.1.3 Principle of virtual work

The kinematic formulation of Section 3.1.2 can still be used to solve for the line tensions and fiber velocities. In particular, by substituting the inextensible velocity (3.6) into the left hand side of the evolution equation (3.1) and setting $\boldsymbol{\lambda} = (T\boldsymbol{\tau})_s$, a PDE results with unknowns $\boldsymbol{\alpha}$ and T . We choose to close our formulation differently, in the process eliminating the need to solve for tension explicitly.

On every fiber, the principle of virtual work states that the constraint forces $\boldsymbol{\lambda}$ do no work for any choice of $\boldsymbol{\alpha}$ [42]. To impose the principle of virtual work, we use the L^2 inner product to compute the total power dissipated in the fluid from $\boldsymbol{\lambda}$

$$\mathcal{P} = \langle \mathcal{K}\boldsymbol{\alpha}, \boldsymbol{\lambda} \rangle = \int_0^L (\mathcal{K}\boldsymbol{\alpha})(s') \cdot \boldsymbol{\lambda}(s') ds'.$$

3.1.3.1 Rank-deficient formulation

In the case when \mathcal{K} is given by (3.7), the calculation of the work done by $\boldsymbol{\lambda}$ proceeds as

$$\langle \mathcal{K}\boldsymbol{\alpha}, \boldsymbol{\lambda} \rangle = \int_0^L \left(\mathbf{U}_{\text{MP}} + \int_{L/2}^s (\boldsymbol{\Omega}(s') \times \boldsymbol{\tau}(s')) ds' \right) \cdot \boldsymbol{\lambda}(s) ds \quad (3.16)$$

$$\langle \mathcal{K}\boldsymbol{\alpha}, \boldsymbol{\lambda} \rangle = \int_0^L \left(\mathbf{U}(s=0) + \int_0^s (\boldsymbol{\Omega}(s') \times \boldsymbol{\tau}(s')) ds' \right) \cdot \boldsymbol{\lambda}(s) ds \quad (3.17)$$

$$= \mathbf{U}_0 \cdot \int_0^L \boldsymbol{\lambda}(s') ds' + \int_0^L \boldsymbol{\Omega}(s) \cdot \left(\boldsymbol{\tau}(s) \times \int_s^L \boldsymbol{\lambda}(s') ds' \right) ds \quad (3.18)$$

$$:= \langle \boldsymbol{\alpha}, \mathcal{K}^* \boldsymbol{\lambda} \rangle = 0,$$

which holds for all $\boldsymbol{\alpha}$ and defines \mathcal{K}^* as the L^2 adjoint of \mathcal{K} . The third equality, which comes from changing the order of integration in the first, immediately gives the constraints

$$\mathcal{K}^* \boldsymbol{\lambda} := \begin{pmatrix} \int_0^L \boldsymbol{\lambda}(s) ds \\ \boldsymbol{\tau}(s) \times \int_s^L \boldsymbol{\lambda}(s') ds' \end{pmatrix} = \begin{pmatrix} \mathbf{0} \\ \mathbf{0} \end{pmatrix}, \quad (3.19)$$

where the second line holds for all s . The second constraint in (3.19) implies $\boldsymbol{\lambda} = \partial_s(T\boldsymbol{\tau})$, with the first constraint giving the boundary conditions $T(s=0, L) = 0$ on tension.¹ For clamped fibers, the corresponding virtual work constraint $\mathcal{K}_c^* \boldsymbol{\lambda} = \mathbf{0}$ is obtained from (3.19) by dropping the first row, since the fiber is no longer force-free.

3.1.3.2 L^2 weak formulation

We can now use the same ideas on the normal vector formulation of \mathcal{K} in Section 3.1.2.3 to derive the principle of virtual work in that case. Substituting the representation (3.15) into

¹In the case when the fiber tangent vector is constant $\boldsymbol{\tau}(s) \equiv \boldsymbol{\tau}$, then the first equation in (3.19) gives only $T(0) = T(L)$. This is consistent with the corresponding line tension equation (3.4), where T only appears differentiated, since for constant $\boldsymbol{\tau}$ we have $\boldsymbol{\lambda} = \boldsymbol{\tau} \partial_s T$.

the general inner product (3.1.3), we arrive at the power equation

$$\mathcal{P} = \mathbf{U}(0) \cdot \int_0^L \boldsymbol{\lambda}(s) ds + \int_0^L \left(\int_0^s \sum_{j=1}^2 \sum_k \alpha_{jk} T_k(s') \mathbf{n}_j(\boldsymbol{\tau}(s')) ds' \right) \cdot \boldsymbol{\lambda}(s) ds = 0. \quad (3.20)$$

Since the power from the constraint forces must be zero for any inextensible motion (any $\boldsymbol{\alpha}$), each term of the constraint (3.20) must be zero. This gives the set of constraints on every fiber

$$\boldsymbol{\mathcal{K}}^*[\mathbb{X}]\boldsymbol{\lambda} := \begin{pmatrix} \int_0^L \left(\int_0^s T_k(s') \mathbf{n}_1(\boldsymbol{\tau}(s')) ds' \right) \cdot \boldsymbol{\lambda}(s) ds \\ \int_0^L \left(\int_0^s T_k(s') \mathbf{n}_2(\boldsymbol{\tau}(s')) ds' \right) \cdot \boldsymbol{\lambda}(s) ds \\ \int_0^L \boldsymbol{\lambda}(s) ds \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ \mathbf{0} \end{pmatrix}, \quad (3.21)$$

where the first two constraints hold for all k and the last constraint holds for each of the three Cartesian directions. Equation (3.21) has the flavor of a Galerkin method: for every T_k we ask that the forcing $\boldsymbol{\lambda}$ perform no work against the inextensible motion $\int T_k \mathbf{n}_j$. For this reason, we refer to (3.21) as a “weak” formulation of the constraint $\boldsymbol{\mathcal{K}}^* \boldsymbol{\lambda} = \mathbf{0}$, in contrast to the “strong” pointwise formulation of (3.19). Moving forward, we will use only the strong formulation, but the weak formulation might be an interesting object for future analysis.

3.2 Bend and twist elasticity

Having looked at inextensibility, we are now ready to examine the coupling of inextensibility to twist and bend elasticity, and the fluid that surrounds the elastic filament. Our presentation is much the same as in the previous section: we first present an approach which is common in the literature but which we *do not* use (which we call the Kirchhoff model), and then present our approach (which we call the Euler model). The distinction between the two, as discussed in Chapter 1, is whether the fiber cross section rotates with the local fluid angular velocity, or whether we need to track perpendicular components of the angular fluid velocity at all. As for inextensibility, we establish that the two approaches are equivalent,

but in this case the equivalence is only in the limit when the slenderness $\epsilon \rightarrow 0$.

3.2.1 Fiber geometry: the material frame and twist angle

Once we start to track the twist of the fiber centerline, we need to introduce additional degrees of freedom beyond just the positions $\mathbb{X}(s)$. In previous immersed-boundary formulations where twist elasticity is coupled with fluid mechanics [49, 50], the rod is defined by an orthonormal material frame $(\mathbf{D}^{(1)}(s), \mathbf{D}^{(2)}(s), \mathbf{D}^{(3)}(s))$. In the standard model for slender fibers, it is assumed that the fibers are unsharable, so that cross sections of the fiber maintain their shape (e.g., circular) and remain perpendicular to the tangent vector. This implies the constraint

$$\mathbf{D}^{(3)} = \boldsymbol{\tau}. \quad (3.22)$$

The fiber cross sections are allowed to twist relative to the tangent vector. The twist density $\psi(s)$ is defined as

$$\psi := \mathbf{D}^{(2)} \cdot \partial_s \mathbf{D}^{(1)} \quad (3.23)$$

with boundary conditions

$$\psi(0) = \psi(L) = 0 \quad (3.24)$$

for a free filament [33]. The use of the material frame is somewhat cumbersome, however, since it requires us to keep track of three separate axes throughout the course of a simulation.

We will instead use a formulation based on the twist-free Bishop frame [84], in which the state variables are the tangent vector $\boldsymbol{\tau}(s)$ and a degree of twist $\theta(s)$ relative to the Bishop frame. The Bishop frame is an orthonormal frame $(\mathbf{b}^{(1)}(s), \mathbf{b}^{(2)}(s), \boldsymbol{\tau}(s))$ which satisfies

$$\mathbf{b}^{(2)} \cdot \partial_s \mathbf{b}^{(1)} = -\mathbf{b}^{(1)} \cdot \partial_s \mathbf{b}^{(2)} = 0, \quad (3.25)$$

making it twist free. The Bishop frame along the entire fiber can be constructed by choosing

$\mathbf{b}^{(1)}(0)$ and $\mathbf{b}^{(2)}(0)$ such that $(\mathbf{b}^{(1)}(0), \mathbf{b}^{(2)}(0), \boldsymbol{\tau}(0))$ is an orthonormal frame, and then solving the ordinary differential equation (ODE) [84]

$$\partial_s \mathbf{b}^{(1)} = (\boldsymbol{\tau} \times \partial_s \boldsymbol{\tau}) \times \mathbf{b}^{(1)}, \quad (3.26)$$

for $\mathbf{b}^{(1)}(s)$, with $\mathbf{b}^{(2)} = \boldsymbol{\tau}(s) \times \mathbf{b}^{(1)}$. The material frame and Bishop frame are related by [84]

$$\mathbf{D}^{(1)} = \mathbf{b}^{(1)} \cos \theta + \mathbf{b}^{(2)} \sin \theta, \quad \mathbf{D}^{(2)} = -\mathbf{b}^{(1)} \sin \theta + \mathbf{b}^{(2)} \cos \theta, \quad \mathbf{D}^{(3)} = \boldsymbol{\tau}, \quad (3.27)$$

so that $\theta(s)$ is the angle of rotation of the material frame relative to the Bishop frame. Solving (3.26) is also called “parallel transporting” the vector $\mathbf{b}^{(1)}$, since its tangential component is kept at zero by performing rotations about the binormal [84]. Using (3.27), it can be shown that the twist (3.23) simplifies to [84]

$$\psi = \partial_s \theta. \quad (3.28)$$

If θ is a constant, the material frame is just a constant rotation of the twist-free Bishop frame, so the twist density $\psi = \partial_s \theta = 0$ as well.

3.2.1.1 Evolution of the material frame and twist angle

Let us now consider what the inextensibility and unshearability constraint imply for the motion of the material frame, and consequently the evolution of the twist density (3.28). Because the material frame is orthonormal and stays orthonormal for all time, its velocity is constrained to satisfy

$$\partial_t \mathbf{D}^{(i)} = \boldsymbol{\Omega} \times \mathbf{D}^{(i)}, \quad (3.29)$$

for some $\boldsymbol{\Omega}(s)$. In the particular case of $i = 3$, we have the evolution equation (3.5), which is the bedrock of our formulation of inextensibility.

Let us now derive an evolution equation for the twist ψ . While the result we give is well-known [33], this derivation uses straightforward calculus, rather than a variational approach, and it is hence worth repeating. Because the material frame $(\mathbf{D}^{(1)}(s), \mathbf{D}^{(2)}(s), \mathbf{D}^{(3)}(s))$ is orthonormal, it evolves by the parallel transport equation

$$\partial_s \mathbf{D}^{(i)} = \boldsymbol{\omega} \times \mathbf{D}^{(i)}, \quad (3.30)$$

along the arclength of the filament. Substituting the definition of twist from (3.23), we obtain

$$\begin{aligned} \psi &= \partial_s \mathbf{D}^{(1)} \cdot \mathbf{D}^{(2)} \\ &= (\boldsymbol{\omega} \times \mathbf{D}^{(1)}) \cdot \mathbf{D}^{(2)} = (\mathbf{D}^{(1)} \times \mathbf{D}^{(2)}) \cdot \boldsymbol{\omega} = \boldsymbol{\tau} \cdot \boldsymbol{\omega} \end{aligned} \quad (3.31)$$

So $\omega^{\parallel} = \psi$. Likewise, setting $i = 3$ in (3.30) and using the constraint $\boldsymbol{\tau} = \mathbf{D}^{(3)}$ from (3.22) gives $\partial_s \boldsymbol{\tau} = \boldsymbol{\omega} \times \boldsymbol{\tau}$. Crossing both sides with $\boldsymbol{\tau}$, we obtain the total representation of $\boldsymbol{\omega}$ as [84]

$$\boldsymbol{\omega} = \boldsymbol{\tau} \times \partial_s \boldsymbol{\tau} + \boldsymbol{\tau} \psi \quad (3.32)$$

We now equate the t derivative of (3.30) with the s derivative of (3.29),

$$\begin{aligned} \partial_t (\boldsymbol{\omega} \times \mathbf{D}^{(i)}) &= \partial_s (\boldsymbol{\Omega} \times \mathbf{D}^{(i)}) \\ \partial_t \boldsymbol{\omega} \times \mathbf{D}^{(i)} + \boldsymbol{\omega} \times \partial_t \mathbf{D}^{(i)} &= \partial_s \boldsymbol{\Omega} \times \mathbf{D}^{(i)} + \boldsymbol{\Omega} \times \partial_s \mathbf{D}^{(i)} \\ \partial_t \boldsymbol{\omega} \times \mathbf{D}^{(i)} &= \partial_s \boldsymbol{\Omega} \times \mathbf{D}^{(i)} + \boldsymbol{\Omega} \times (\boldsymbol{\omega} \times \mathbf{D}^{(i)}) - \boldsymbol{\omega} \times (\boldsymbol{\Omega} \times \mathbf{D}^{(i)}). \end{aligned} \quad (3.33)$$

The last equality used the evolution equations (3.30) and (3.29) for $\mathbf{D}^{(i)}$. Using the triple

cross product identity $\mathbf{a} \times (\mathbf{b} \times \mathbf{c}) + \mathbf{c} \times (\mathbf{a} \times \mathbf{b}) + \mathbf{b} \times (\mathbf{c} \times \mathbf{a}) = \mathbf{0}$, we can write (3.33) as

$$\partial_t \boldsymbol{\omega} \times \mathbf{D}^{(i)} = \partial_s \boldsymbol{\Omega} \times \mathbf{D}^{(i)} + (\boldsymbol{\Omega} \times \boldsymbol{\omega}) \times \mathbf{D}^{(i)}, \quad (3.34)$$

which must hold for all i . This gives [33, Eq. (54)]

$$\partial_t \boldsymbol{\omega} = \partial_s \boldsymbol{\Omega} + (\boldsymbol{\Omega} \times \boldsymbol{\omega}). \quad (3.35)$$

To derive the evolution equation for the twist angle, we take the $\boldsymbol{\tau}$ component of (3.35) and use (3.29)

$$\partial_t (\boldsymbol{\omega} \cdot \boldsymbol{\tau}) - \boldsymbol{\omega} \cdot \partial_t \boldsymbol{\tau} = \partial_s (\boldsymbol{\Omega} \cdot \boldsymbol{\tau}) - \boldsymbol{\Omega} \cdot \partial_s \boldsymbol{\tau} + (\boldsymbol{\Omega} \times \boldsymbol{\omega}) \cdot \boldsymbol{\tau} \quad (3.36)$$

$$\partial_t (\boldsymbol{\omega} \cdot \boldsymbol{\tau}) = \boldsymbol{\omega} \cdot (\boldsymbol{\Omega} \times \boldsymbol{\tau}) + \partial_s (\boldsymbol{\Omega} \cdot \boldsymbol{\tau}) - \boldsymbol{\Omega} \cdot \partial_s \boldsymbol{\tau} + (\boldsymbol{\Omega} \times \boldsymbol{\omega}) \cdot \boldsymbol{\tau}.$$

The triple products cancel, and using (3.31), we get the twist evolution equation

$$\partial_t \psi = \partial_s \Omega^\parallel - (\boldsymbol{\Omega} \cdot \partial_s \boldsymbol{\tau}). \quad (3.37)$$

3.2.2 The Kirchhoff model

The distinguishing feature of what we refer to as the *Kirchhoff model* is that the rotational velocity of the tangent vector must be exactly equal to the local angular fluid velocity (2.3),

$$\boldsymbol{\Omega} = \boldsymbol{\Psi} \quad \partial_t \boldsymbol{\tau} = \boldsymbol{\Psi} \times \boldsymbol{\tau}. \quad (3.38)$$

However, because $\partial_t \boldsymbol{\tau} = \partial_s \mathbf{U}$, the assumption (3.38) implies a local constraint that relates the translational and rotational velocity of the rod

$$\partial_s \mathbf{U} = \boldsymbol{\Psi} \times \boldsymbol{\tau}. \quad (3.39)$$

Thus the fluid velocity is (locally) constrained so that the integrity of the rod cross section is preserved, i.e., that the rod is unshearable [33], and inextensible, since (3.39) implies the inextensibility constraint (3.2).

In the Kirchhoff model as presented in [47, 49, 50], the total force density $\mathbf{f}(s)$ and torque density $\mathbf{n}(s)$ exerted by the fiber on the fluid are given by the force and torque balances [47]

$$\mathbf{f} = \partial_s \mathbf{F}, \quad \mathbf{n} = \partial_s \mathbf{N} + \boldsymbol{\tau} \times \mathbf{F}. \quad (3.40)$$

Here the torque density $\partial_s \mathbf{N}(s)$ comes from the internal moments on the rod $\mathbf{N}(s)$, which are given in terms of the bending modulus κ and twist modulus γ , both of which we assume to be constant along the cylindrical rod. The internal moment $\mathbf{N}(s)$ can be written in terms of the material frame as [47, 49, 50] (see also [38, c. 3] for a more mechanics-based discussion)

$$\begin{aligned} \mathbf{N} &= \kappa \left(-\mathbf{D}^{(2)} \cdot \partial_s \boldsymbol{\tau} \right) \mathbf{D}^{(1)} + \kappa \left(\partial_s \boldsymbol{\tau} \cdot \mathbf{D}^{(1)} \right) \mathbf{D}^{(2)} + \gamma \psi \boldsymbol{\tau} \\ &= \kappa \left(\left(\mathbf{D}^{(1)} \times \mathbf{D}^{(2)} \right) \times \partial_s \boldsymbol{\tau} \right) + \gamma \psi \boldsymbol{\tau} \\ &= \kappa \left(\boldsymbol{\tau} \times \partial_s \boldsymbol{\tau} \right) + \gamma \psi \boldsymbol{\tau}, \end{aligned} \quad (3.41)$$

subject to the free fiber boundary conditions [33]

$$\mathbf{N}(s=0, L) = \mathbf{0} \quad \partial_s \boldsymbol{\tau}(s=0, L) = \mathbf{0}, \quad (3.42)$$

and twist boundary conditions (3.24). The torque density $\partial_s \mathbf{N}(s)$ in (3.40) is

$$\partial_s \mathbf{N} = \kappa \left(\boldsymbol{\tau} \times \partial_s^2 \boldsymbol{\tau} \right) + \gamma \left(\boldsymbol{\tau} \partial_s \psi + \psi \partial_s \boldsymbol{\tau} \right), \quad (3.43)$$

with the remaining torque density coming from the force $\mathbf{F}(s)$, which is subject to the free

fiber boundary conditions

$$\mathbf{F}(s = 0, L) = \mathbf{0}. \quad (3.44)$$

As discussed in [47, p. 5], the force \mathbf{F} is a Lagrange multiplier to enforce the position-based constraint $\partial_s \mathbb{X} = \boldsymbol{\tau}$. It turns out, however, that the final solution for \mathbf{F} depends on the internal torque (3.43), since the perpendicular torque from \mathbf{F} must eliminate the perpendicular torque coming from $\partial_s \mathbf{N}$ for the constraint (3.39) to hold. This is because perpendicular angular velocity $\boldsymbol{\Psi}_n^\perp \sim \hat{\epsilon}^{-2} \mathbf{n}^\perp$, while translational velocity $\mathbf{U}_n \sim \mathbf{n}^\perp \log \hat{\epsilon}$ (see the local drag estimates in Section 2.1). Thus, it is impossible for (3.39) to hold if there is a $\mathcal{O}(1)$ perpendicular torque on the fluid, so \mathbf{F} must cancel the $\mathcal{O}(1)$ perpendicular torque in $\partial_s \mathbf{N}$.

We now write an explicit formula for \mathbf{F} that yields a torque density \mathbf{n} which is only in the parallel direction, plus a Lagrange multiplier torque in the perpendicular direction. Because this formula eliminates perpendicular torque from the constitutive model, it also eliminates the dependency of torque on the bending modulus κ , and leaves only parallel torque depending on the twist modulus γ . The main step is to introduce the forces from bending $\mathbf{F}^{(\kappa)}(s)$ and twisting $\mathbf{F}^{(\gamma)}(s)$, such that when we take their cross product with $\boldsymbol{\tau}$ and add it to $\partial_s \mathbf{N}$ in (3.48), the perpendicular components of torque will be eliminated

$$\mathbf{F}^{(\kappa)} = -\kappa \partial_s^2 \boldsymbol{\tau}, \quad \mathbf{F}^{(\gamma)} = \gamma \psi (\boldsymbol{\tau} \times \partial_s \boldsymbol{\tau}). \quad (3.45)$$

We then redefine the force \mathbf{F} in the Kirchhoff model as

$$\mathbf{F} = \mathbf{F}^{(\kappa)} + \mathbf{F}^{(\gamma)} + \boldsymbol{\Lambda}, \quad (3.46)$$

where now $\boldsymbol{\Lambda}(s)$ is the Lagrange multiplier that enforces the velocity-based constraint (3.39). For free fibers, the force \mathbf{F} must vanish at the endpoints (see (3.44)), which, combined with

(3.42), leads to the free fiber boundary conditions

$$\partial_s \boldsymbol{\tau}(s=0, L) = \mathbf{0}, \quad \partial_s^2 \boldsymbol{\tau}(s=0, L) = \mathbf{0}, \quad \boldsymbol{\Lambda}(s=0, L) = \mathbf{0}, \quad \text{and} \quad \psi(s=0, L) = 0. \quad (3.47)$$

Substituting the new form of \mathbf{F} in (3.46) into (3.40), we have the torque density applied to the fluid as

$$\mathbf{n} = \partial_s \mathbf{N} + \boldsymbol{\tau} \times \left(\mathbf{F}^{(\kappa)} + \mathbf{F}^{(\gamma)} + \boldsymbol{\Lambda} \right) \quad (3.48)$$

$$= (\gamma \partial_s \psi) \boldsymbol{\tau} + \boldsymbol{\tau} \times \boldsymbol{\Lambda} = n^{\parallel} \boldsymbol{\tau} + \boldsymbol{\tau} \times \boldsymbol{\Lambda}^{\perp}, \quad (3.49)$$

where we defined the parallel torque coming from twist as [33]

$$n^{\parallel} = \gamma \partial_s \psi. \quad (3.50)$$

Throughout this section, we have assumed that the filament is intrinsically straight and untwisted; expressions for n^{\parallel} and \mathbf{F} in the case of an intrinsically bent and twisted rod are given in Section 3.2.3.1.

Note that only the (two) perpendicular components of $\boldsymbol{\Lambda}$ enter in the torque (3.49). The third Lagrange multiplier required to enforce (3.39) is the parallel component of $\boldsymbol{\Lambda}$, which enters in the force density applied to the fluid as $\boldsymbol{\lambda}(s)$,

$$\begin{aligned} \mathbf{f} &= \partial_s \left(\mathbf{F}^{(\kappa)} + \mathbf{F}^{(\gamma)} + \boldsymbol{\Lambda} \right) = \mathbf{f}^{(\kappa)} + \mathbf{f}^{(\gamma)} + \boldsymbol{\lambda}, \quad \text{where} \quad (3.51) \\ \mathbf{f}^{(\kappa)} &= -\kappa \partial_s^3 \boldsymbol{\tau}, \quad \mathbf{f}^{(\gamma)} = \gamma \partial_s (\psi (\boldsymbol{\tau} \times \partial_s \boldsymbol{\tau})), \quad \text{and} \quad \boldsymbol{\lambda} = \partial_s \boldsymbol{\Lambda}. \end{aligned}$$

In our reformulation of the Kirchhoff model, the force density on the fluid is given by (3.51) and the torque is given by (3.49). This makes it obvious that the only perpendicular torque acting on the fluid comes from the Lagrange multipliers $\boldsymbol{\Lambda}^{\perp}$, while the parallel torque comes

from twist elasticity.

3.2.2.1 Linear system for Λ

To derive a linear system for Λ , we first write the constraint (3.39) in terms of the total force density $\mathbf{f}(s)$ and torque density $\mathbf{n}(s)$ applied by the rod to the fluid using the mobility operators defined in (2.4),

$$\partial_s (\mathcal{M}_{\text{tt}} \mathbf{f} + \mathcal{M}_{\text{tr}} \mathbf{n}) = (\mathcal{M}_{\text{rt}} \mathbf{f} + \mathcal{M}_{\text{rr}} \mathbf{n}) \times \boldsymbol{\tau}. \quad (3.52)$$

We then substitute the representations for \mathbf{f} and \mathbf{n} from (3.51) and (3.49), and define a linear operator \mathcal{C} such that $(\mathcal{C}\mathbf{x})(s) = \boldsymbol{\tau}(s) \times \mathbf{x}(s)$. This yields the second-order nonlocal boundary value problem (BVP)

$$\begin{aligned} & [\partial_s (\mathcal{M}_{\text{tt}} \partial_s + \mathcal{M}_{\text{tr}} \mathcal{C}) + \mathcal{C} (\mathcal{M}_{\text{rt}} \partial_s + \mathcal{M}_{\text{rr}} \mathcal{C})] \Lambda = \\ & -\partial_s \left(\mathcal{M}_{\text{tt}} \left(\mathbf{f}^{(\kappa)} + \mathbf{f}^{(\gamma)} \right) + \mathcal{M}_{\text{tr}} (n^{\parallel} \boldsymbol{\tau}) \right) - \mathcal{C} \left(\mathcal{M}_{\text{rt}} \left(\mathbf{f}^{(\kappa)} + \mathbf{f}^{(\gamma)} \right) + \mathcal{M}_{\text{rr}} (n^{\parallel} \boldsymbol{\tau}) \right) \end{aligned} \quad (3.53)$$

which can be solved for the Lagrange multipliers $\Lambda(s)$. Using the force \mathbf{f} in (3.51) and torque in (3.49) then gives the resulting velocities in the Kirchhoff model,

$$\mathbf{U} = \mathcal{M}_{\text{tt}} \left(\mathbf{f}^{(\kappa)} + \mathbf{f}^{(\gamma)} + \boldsymbol{\lambda} \right) + \mathcal{M}_{\text{tr}} (n^{\parallel} \boldsymbol{\tau} + \boldsymbol{\tau} \times \Lambda), \quad (3.54)$$

$$\boldsymbol{\Omega} = \boldsymbol{\Psi} = \mathcal{M}_{\text{rt}} \left(\mathbf{f}^{(\kappa)} + \mathbf{f}^{(\gamma)} + \boldsymbol{\lambda} \right) + \mathcal{M}_{\text{rr}} (n^{\parallel} \boldsymbol{\tau} + \boldsymbol{\tau} \times \Lambda). \quad (3.55)$$

3.2.3 The Euler model

The main difference between the Kirchhoff model and what we refer to as the *Euler model* is how the rotation rate of the cross section (equivalently, of the tangent vector) is obtained. In the Kirchhoff model, it does not matter whether we use $\partial_s \mathbf{U}$ or $\boldsymbol{\Psi} \times \boldsymbol{\tau}$ to update the tangent vector, since by (3.39) they must be the same. In the Euler model, by contrast, we

do not set $\boldsymbol{\Omega}$ to be equal to $\boldsymbol{\Psi}$ as in (3.38), and consequently do not enforce the constraint (3.39) that follows from (3.38). Instead, we relate $\boldsymbol{\Omega}^\perp(s)$ to the *translational* velocity via

$$\boldsymbol{\Omega}^\perp = \boldsymbol{\tau} \times \partial_s \mathbf{U}, \quad \text{and} \quad \partial_t \boldsymbol{\tau} = \partial_s \mathbf{U} = \boldsymbol{\Omega}^\perp \times \boldsymbol{\tau}, \quad (3.56)$$

where $\boldsymbol{\Omega}_E^\perp(s) \cdot \boldsymbol{\tau} = 0$, and, unlike the Kirchhoff model, $\boldsymbol{\Omega}^\perp$ is not necessarily equal to $\boldsymbol{\Psi}^\perp$. The evolution of the fiber can now be obtained by solving for $\mathbf{U}(s)$ and specifying $\Omega^\parallel(s) = \Psi^\parallel(s)$.

In the Euler model, there is only one constraint, the inextensibility constraint (3.2). This means that the Lagrange multiplier to enforce this constraint should have only one independent component, as opposed to the three independent components for the Lagrange multiplier in the Kirchhoff model. Specifically, the Lagrange multiplier $\boldsymbol{\Lambda}(s)$ reduces to a scalar line tension $T(s)$,

$$\boldsymbol{\Lambda} = T \boldsymbol{\tau}, \quad (3.57)$$

with the free-fiber boundary condition (3.47) implying that

$$T(s = 0, L) = 0. \quad (3.58)$$

Since $\boldsymbol{\Lambda}$ is tangential, the torque on the fluid in (3.49) reduces to $\boldsymbol{n} = n^\parallel \boldsymbol{\tau}$, where n^\parallel is defined in (3.50). The force is still given by (3.51), but now with $\boldsymbol{\Lambda}$ having only one independent component. Put another way, the equations now reduce to an inextensible rod plus an additional force and parallel torque due to twisting, so that we can justify the assumption in Chapter 2 that the centerline torque is only in the parallel direction.

To be more precise, the translational velocity $\mathbf{U}(s)$ in the Euler model is obtained by solving a constrained system which includes only the parallel torque density. This constrained system is derived by substituting (3.57) into the Kirchhoff velocity (3.54), and then using the parameterization of the inextensible motions (3.6) (and resulting virtual work constraint) to

obtain the system

$$\mathcal{K}\alpha = \mathcal{M}_{\text{tt}} \left(\mathbf{f}^{(\kappa)} + \mathbf{f}^{(\gamma)} + \boldsymbol{\lambda} \right) + \mathcal{M}_{\text{tr}} (n^{\parallel} \boldsymbol{\tau}) \quad (3.59)$$

$$\mathcal{K}^* \boldsymbol{\lambda} = \mathbf{0}.$$

The evolution of the twist angle in (3.37) requires the rotation rate $\Omega^{\parallel}(s)$, which is computed in a post-processing step as the parallel component of the angular velocity in the Kirchhoff model (3.55),

$$\Omega^{\parallel} = \left(\mathcal{M}_{\text{rt}} \left(\mathbf{f}^{(\kappa)} + \mathbf{f}^{(\gamma)} + \boldsymbol{\lambda} \right) + \mathcal{M}_{\text{tr}} (n^{\parallel} \boldsymbol{\tau}) \right) \cdot \boldsymbol{\tau}. \quad (3.60)$$

The Euler model *does not* enforce the constraint (3.39), since only the translational velocity \mathbf{U} and parallel rotational velocity Ω^{\parallel} are tracked. In the next section, we study the effect of this on the velocity of the fiber centerline.

We emphasize that our choice to name the model of this section the Euler model and that of the previous section the Kirchhoff model is historical and arbitrary. The reason we choose the naming we do is because the Kirchhoff model as formulated in Section 3.2.2 is what has been used before in the dynamics of rods with twist [47, 49, 50], while the Euler model is the one is closest to what has been done before without twist [15, 22], and is in our opinion easier to extend to include twist.

3.2.3.1 The Euler model with intrinsic curvatures

In (3.51) and (3.59), we derived the Euler model from the Kirchhoff model by first setting $\mathbf{F} = \mathbf{F}^{(\kappa)} + \mathbf{F}^{(\gamma)} + \boldsymbol{\Lambda}$ in (3.46), where $\mathbf{F}^{(\kappa)}$ and $\mathbf{F}^{(\gamma)}$ are defined to eliminate the perpendicular torque in $\partial_s \mathbf{N}$, then setting $\boldsymbol{\Lambda} = T \boldsymbol{\tau}$, which eliminates the perpendicular torque Lagrange multipliers that are the distinguishing feature of the Kirchhoff rod model. These two steps eliminate *all* perpendicular torques from the formulation, which leads to the Euler model.

In this section, we repeat this procedure on a fiber with intrinsic curvature and twist (e.g., a flagellum [85]) to derive the corresponding Euler model. In this case, we have the moment [50]

$$\begin{aligned}
\mathbf{N} &= \kappa \left(\left(\partial_s \mathbf{D}^{(2)} \cdot \mathbf{D}^{(3)} - \kappa_1 \right) \mathbf{D}^{(1)} + \left(\partial_s \mathbf{D}^{(3)} \cdot \mathbf{D}^{(1)} - \kappa_2 \right) \mathbf{D}^{(2)} \right) \\
&\quad + \gamma \left(\partial_s \mathbf{D}^{(1)} \cdot \mathbf{D}^{(2)} - \phi \right) \mathbf{D}^{(3)} \\
&= \kappa \left(\boldsymbol{\tau} \times \partial_s \boldsymbol{\tau} - \kappa_1 \mathbf{D}^{(1)} - \kappa_2 \mathbf{D}^{(2)} \right) + \gamma (\psi - \phi) \boldsymbol{\tau},
\end{aligned} \tag{3.61}$$

where κ_1 and κ_2 are the preferred curvatures, and ϕ is the preferred twist angle.

We illustrate the key steps here assuming constant κ_1 , κ_2 , and ϕ . Taking the s derivative of \mathbf{N} and separating the result into tangential and perpendicular parts, we have

$$\partial_s \mathbf{N} = \kappa \left(\boldsymbol{\tau} \times \partial_s^2 \boldsymbol{\tau} - \kappa_1 \partial_s \mathbf{D}^{(1)} - \kappa_2 \partial_s \mathbf{D}^{(2)} \right) + \gamma \left((\partial_s \psi) \boldsymbol{\tau} + (\psi - \phi) \partial_s \boldsymbol{\tau} \right) \tag{3.62}$$

$$\begin{aligned}
&= \boldsymbol{\tau} \times \left(\kappa \left(\partial_s^2 \boldsymbol{\tau} + \kappa_1 \left(\boldsymbol{\tau} \times \partial_s \mathbf{D}^{(1)} \right) + \kappa_2 \left(\boldsymbol{\tau} \times \partial_s \mathbf{D}^{(2)} \right) \right) - \gamma (\psi - \phi) \left(\boldsymbol{\tau} \times \partial_s \boldsymbol{\tau} \right) \right) \\
&\quad + \left(\kappa \left(-\kappa_1 \left(\partial_s \mathbf{D}^{(1)} \cdot \boldsymbol{\tau} \right) - \kappa_2 \left(\partial_s \mathbf{D}^{(2)} \cdot \boldsymbol{\tau} \right) \right) + \gamma \partial_s \psi \right) \boldsymbol{\tau}
\end{aligned} \tag{3.63}$$

To derive an Euler model, we now amend (3.45) and (3.46) to

$$\mathbf{F}^{(\kappa)} = -\kappa \left(\partial_s^2 \boldsymbol{\tau} - \kappa_1 \psi \mathbf{D}^{(1)} - \kappa_2 \psi \mathbf{D}^{(2)} \right), \quad \mathbf{F}^{(\gamma)} = \gamma (\psi - \phi) \left(\boldsymbol{\tau} \times \partial_s \boldsymbol{\tau} \right), \tag{3.64}$$

$$\mathbf{F} = \mathbf{F}^{(\kappa)} + \mathbf{F}^{(\gamma)} + T \boldsymbol{\tau}.$$

Just like in the case of intrinsically straight and untwisted fibers, substituting this representation into (3.40) gives the force and torque density applied to the fluid as

$$\mathbf{f} = \partial_s \mathbf{F}, \quad \mathbf{n} = n^\parallel \boldsymbol{\tau}, \quad n^\parallel = \kappa \left(\kappa_1 \left(\mathbf{D}^{(1)} \cdot \partial_s \boldsymbol{\tau} \right) + \kappa_2 \left(\mathbf{D}^{(2)} \cdot \partial_s \boldsymbol{\tau} \right) \right) + \gamma \partial_s \psi. \tag{3.65}$$

Free fiber boundary conditions are modified in the case of intrinsic curvature and twist.

As before, we require $\mathbf{N} = \mathbf{0}$ and $\mathbf{F} = \mathbf{0}$ at the free end, which implies that

$$\psi = \phi, \quad \partial_s \boldsymbol{\tau} = \kappa_2 \mathbf{D}^{(1)} - \kappa_1 \mathbf{D}^{(2)}, \quad \partial_s^2 \boldsymbol{\tau} = \kappa_1 \phi \mathbf{D}^{(1)} + \kappa_2 \phi \mathbf{D}^{(2)}, \quad \text{and} \quad T = 0 \quad (3.66)$$

at a free end. The equation for $\partial_s \boldsymbol{\tau}$ is derived by using (3.61) and substituting $\mathbf{D}^{(1)} = \mathbf{D}^{(2)} \times \boldsymbol{\tau}$ and $\mathbf{D}^{(2)} = \boldsymbol{\tau} \times \mathbf{D}^{(1)}$.

3.2.4 Comparing the Kirchhoff and Euler formulations

In this section, we analyze the difference between the Kirchhoff model of Section 3.2.2 and the Euler model of Section 3.2.3 for fibers without intrinsic curvature and twist. Because the Kirchhoff model accounts for the angular fluid velocity and has more degrees of freedom, we refer to the translational velocity obtained from (3.54) as the “exact” velocity \mathbf{U} , with corresponding tangent vector evolution $\partial_t \boldsymbol{\tau} = \partial_s \mathbf{U} = \boldsymbol{\Psi} \times \boldsymbol{\tau} = \boldsymbol{\Omega} \times \boldsymbol{\tau}$. We refer to the translational velocity from the Euler model (3.59) as \mathbf{U}_E and the parallel rotational velocity from the Euler model (3.60) as Ω_E^\parallel . Since the rotation of the tangent vector in the Euler model is obtained from $\partial_s \mathbf{U}_E$, the parallel rotational velocity Ω_E^\parallel represents the only remaining degree of freedom. To examine the difference between the models, we substitute the definitions (3.54), (3.59), (3.55), and (3.60) to define

$$\Delta \mathbf{U} = \mathbf{U} - \mathbf{U}_E = \mathcal{M}_{\text{tt}}(\boldsymbol{\lambda} - \boldsymbol{\lambda}_E) + \mathcal{M}_{\text{tr}}(\boldsymbol{\tau} \times \Delta \boldsymbol{\Lambda}), \quad (3.67)$$

$$\Delta \Psi^\parallel = \Omega^\parallel - \Omega_E^\parallel = (\mathcal{M}_{\text{rt}}(\boldsymbol{\lambda} - \boldsymbol{\lambda}_E) + \mathcal{M}_{\text{rr}}(\boldsymbol{\tau} \times \Delta \boldsymbol{\Lambda})) \cdot \boldsymbol{\tau} := \Delta \boldsymbol{\Psi} \cdot \boldsymbol{\tau} \quad (3.68)$$

The differences in velocity (3.67) and (3.68) depend on the difference in the constraint force

$$\Delta \boldsymbol{\Lambda} := \boldsymbol{\Lambda} - \boldsymbol{\Lambda}_E, \quad (3.69)$$

and its derivative, the force density $\boldsymbol{\lambda} - \boldsymbol{\lambda}_E$, where $\boldsymbol{\Lambda}_E = T\boldsymbol{\tau}$ and $\boldsymbol{\lambda}_E = \partial_s(T\boldsymbol{\tau})$ are the constraint force and force density from the Euler model.

To obtain $\Delta\boldsymbol{U}$ and $\Delta\Psi^\parallel$, we solve for the perturbation of the total velocity from the Euler velocity by solving the ‘‘mismatch’’ problem

$$\partial_s\Delta\boldsymbol{U} + \boldsymbol{\tau} \times \Delta\Psi = -\boldsymbol{m} \quad (3.70)$$

$$\boldsymbol{m} = \partial_s \left(\mathcal{M}_{\text{tt}} (\partial_s \boldsymbol{F}_E) + \mathcal{M}_{\text{tr}} (n^\parallel \boldsymbol{\tau}) \right) + \boldsymbol{\tau} \times \left(\mathcal{M}_{\text{rt}} (\partial_s \boldsymbol{F}_E) + \mathcal{M}_{\text{rr}} (n^\parallel \boldsymbol{\tau}) \right), \quad (3.71)$$

$$\text{where} \quad \boldsymbol{F}_E = \boldsymbol{F}^{(\kappa)} + \boldsymbol{F}^{(\gamma)} + \boldsymbol{\Lambda}_E, \quad (3.72)$$

is the total force in the Euler model. This problem, which we obtain by substituting (3.69) into the Kirchhoff model (3.53), gives a second order nonlocal BVP for $\Delta\boldsymbol{\Lambda}$ as a function of the inconsistency in the evolution of the tangent vector in the Euler formulation $\boldsymbol{m}(s)$. Because the Euler formulation does not enforce (3.39), the mismatch \boldsymbol{m} is nonzero in general. We therefore need the Lagrange multiplier $\Delta\boldsymbol{\Lambda}$ to correct for it and make the total Kirchhoff velocities $\boldsymbol{U} = \boldsymbol{U}_E + \Delta\boldsymbol{U}$ and $\boldsymbol{\Omega} = \boldsymbol{\Psi} = \boldsymbol{\Psi}_E + \Delta\Psi$ consistent with (3.39). This Lagrange multiplier can be split into two components: a perpendicular part $\Delta\boldsymbol{\Lambda}^\perp$ that contributes to the torque $\boldsymbol{\tau} \times \Delta\boldsymbol{\Lambda}$, and a parallel part $\Delta\boldsymbol{\Lambda}^\parallel$ which preserves the inextensibility of the velocity $\Delta\boldsymbol{U}$. Because the constraint force $\boldsymbol{\Lambda}_E = T\boldsymbol{\tau}$ in the Euler model is tangential, we have $\Delta\boldsymbol{\Lambda}^\perp = \boldsymbol{\Lambda}^\perp$ and $\Delta\boldsymbol{\Lambda}^\parallel = \boldsymbol{\Lambda}^\parallel - T\boldsymbol{\tau}$, where $T(s)$ is the tension in the Euler model.

In this section, we use asymptotic estimates and numerical calculations to affirm that

$$\|\partial_s\Delta\boldsymbol{U}\| / \|\partial_s\boldsymbol{U}_E\| \sim \epsilon^2 \log \epsilon \ll 1, \quad (3.73)$$

$$\|\Delta\boldsymbol{\Omega}^\parallel\| / \|\boldsymbol{\Omega}_E^\parallel\| \sim \begin{cases} \epsilon^4 (\log \epsilon)^2 & \gamma > 0 \\ \epsilon^2 \log \epsilon & \gamma = 0 \end{cases} \quad (3.74)$$

away from the endpoints, so that the difference between the Kirchhoff and Euler translational and rotational velocities is very small for sufficiently slender fibers.

3.2.4.1 Asymptotic estimates

To estimate the strength of $\Delta\mathbf{\Lambda}$, we use the local drag estimates of Chapter 2, repeated here for clarity

$$\mathcal{M}_{\text{tt}} \sim \log \epsilon, \quad \mathcal{M}_{\text{tr}} \sim \log \epsilon, \quad \mathcal{M}_{\text{rt}} \sim \log \epsilon, \quad \mathcal{M}_{\text{rr}} \sim \epsilon^{-2}, \quad (3.75)$$

Beginning with the perpendicular Lagrange multipliers $\Delta\mathbf{\Lambda}^\perp = \mathbf{\Lambda}^\perp$, we first observe that the mismatch \mathbf{m} in (3.71) is $\mathcal{O}(\log \epsilon)$ at most, since we don't expect \mathbf{F}_E as defined in (3.72) to change with ϵ . So $\Delta\mathbf{\Lambda}^\perp$ cannot exceed $\mathcal{O}(\epsilon^2 \log \epsilon)$, since otherwise the angular velocity from \mathcal{M}_{rr} would be too large to be consistent with the $\mathcal{O}(\log \epsilon)$ mismatch. If $\Delta\mathbf{\Lambda}^\perp \sim \epsilon^2 \log \epsilon$, then the rotational velocity $\Delta\mathbf{\Psi}$ is $\mathcal{O}(\log \epsilon)$, and therefore of the correct order to cancel the mismatch \mathbf{m} , while the translational velocity $\Delta\mathbf{U} = \mathcal{M}_{\text{tt}}(\partial_s \Delta\mathbf{\Lambda}) + \mathcal{M}_{\text{tr}}(\boldsymbol{\tau} \times \Delta\mathbf{\Lambda})$, is of the order $\Delta\mathbf{U} \sim \epsilon^2 (\log \epsilon)^2$ (if we assume that $\Delta\mathbf{\Lambda}$ and its derivative are of the same order, which should be true away from the endpoints).² This implies that in the fiber interior $\partial_s \Delta\mathbf{U} \sim \epsilon^2 (\log \epsilon)^2$, which gives (3.73).

The difference in parallel angular velocity $\Delta\Omega^\parallel$ is given asymptotically by

$$\Delta\Omega^\parallel = \boldsymbol{\tau} \cdot \Delta\mathbf{\Psi} = \boldsymbol{\tau} \cdot (\mathcal{M}_{\text{rt}}(\partial_s \Delta\mathbf{\Lambda}) + \mathcal{M}_{\text{rr}}(\boldsymbol{\tau} \times \Delta\mathbf{\Lambda})) \approx \boldsymbol{\tau} \cdot (\mathcal{M}_{\text{rt}}(\partial_s \Delta\mathbf{\Lambda})), \quad (3.76)$$

where the last approximate equality is because $\mathcal{M}_{\text{rr}}(\boldsymbol{\tau} \times \Delta\mathbf{\Lambda})$ is orthogonal to $\boldsymbol{\tau}$ to leading order by (2.47). In fact, in the fiber interior $\mathcal{M}_{\text{rr}}(\boldsymbol{\tau} \times \Delta\mathbf{\Lambda}) \cdot \boldsymbol{\tau} \sim \log \epsilon \|\Delta\mathbf{\Lambda}\| \sim \epsilon^2 (\log \epsilon)^2$. Adding the rot-trans contribution, since $\Delta\mathbf{\Lambda} \sim \epsilon^2 \log \epsilon$ and $\mathcal{M}_{\text{rt}} \sim \log \epsilon$, we get $\Delta\Omega^\parallel \sim \epsilon^2 (\log \epsilon)^2$. To non-dimensionalize this, we need to estimate the parallel rotational velocity Ω_E^\parallel defined in (3.60). When there is no twist, there is no parallel torque, so $\Omega_E^\parallel = \boldsymbol{\tau} \cdot (\mathcal{M}_{\text{rt}} \partial_s \mathbf{F}_E) \sim \log \epsilon$, and the relative difference should be $\|\Delta\Omega^\parallel\| / \|\Omega_E^\parallel\| \sim \epsilon^2 \log \epsilon$. When

²Recall that $\Delta\mathbf{\Lambda}^\parallel$ is a Lagrange multiplier that makes the perturbation velocity $\Delta\mathbf{U}$ inextensible. Since the interior velocity is at worst $\Delta\mathbf{U} \sim \epsilon^2 (\log \epsilon)^2$ and the translational mobility is $\mathcal{O}(\log \epsilon)$, the Lagrange multiplier $\Delta\mathbf{\Lambda}^\parallel \sim \epsilon^2 \log \epsilon$ at worst, just like $\Delta\mathbf{\Lambda}^\perp$.

there is twist, the disparity is even larger, since $\Omega_E^{\parallel} \sim \epsilon^{-2}$, which explains the two different cases in (3.74).

Summarizing, the mismatch \mathbf{m} in (3.70) is compensated for by $\Delta\Psi^{\perp}$, which grows as $\epsilon \rightarrow 0$, even as $\Delta\mathbf{\Lambda}$ decays. These are the components of the velocity that are discarded in the Euler model, so it does not matter if they differ from those of the Kirchhoff model. That said, because $\Delta\mathbf{\Lambda}$ must be zero at the endpoints by the boundary conditions (3.44), the perpendicular torques cannot generate angular velocity there, and the term $\Delta\Psi^{\perp}$ cannot compensate for the mismatch \mathbf{m} at the endpoints. Instead, the mismatch there has to be compensated for by terms involving $\partial_s\Delta\mathbf{\Lambda}$, meaning that the derivative of $\Delta\mathbf{\Lambda}$ must be large near the boundaries, as we see in Fig 3.1, where we solve it numerically. Because of the nonsmoothness of the mismatch problem, we will use a simple second-order numerical method to solve it, as described next. Once we switch to the Euler model, we can use more efficient spectral methods, as discussed in Chapter 6.

3.2.4.2 Solving numerically: method

We now solve the mismatch problem (3.70) numerically to establish that our asymptotic estimates do indeed hold. Following the method outlined in [47], we introduce a number of links N and set $\Delta s = L/N$. The discretization of the fiber is *staggered*: the “nodes” are located at $s_j = j\Delta s$, where $j = 0, \dots, N$, so that there are actually $N + 1$ nodes. The “links” are centered on $s_{j+1/2} = (j + 1/2)\Delta s$, where $j = 0, \dots, N - 1$, so that there are N links.

Euler model. In our second-order discretization of the Euler model (3.59), we define $\boldsymbol{\lambda}$, $\mathbf{f}^{(\kappa)}$, and $\mathbf{f}^{(\gamma)}$ at link centers $(j + 1/2)\Delta s$, for $j = 0, \dots, N - 1$. All mobility operators in (2.4) are then replaced with the obvious discretizations of the integrals which sum over the links with weight Δs . For example,

$$(\mathcal{M}_{\text{tt}}\mathbf{f})(s_{i+1/2}) \approx \sum_{j=0}^{N-1} \mathbb{M}_{\text{tt}}(\mathbb{X}_{i+1/2}, \mathbb{X}_{j+1/2}) \mathbf{f}_{j+1/2} \Delta s := (\mathbf{M}_{\text{tt}}\mathbf{f})_{i+1/2}. \quad (3.77)$$

Note that this discretization does not correspond to a more common discretization of an inextensible bead-link polymer chain, which would place forces on $N + 1$ blobs connected by N inextensible links, e.g., [52]. We use the staggered discretization instead for a more direct comparison to the Kirchhoff model.

To complete the second-order discretization of (3.59), we have to define the matrices \mathbf{K} and \mathbf{K}^* . In the second-order case, these matrices are exact transposes of each other, so we only have to define one of them. The matrix \mathbf{K} is a midpoint discretization of the continuum operator (3.7), with action given by

$$(\mathbf{K}\boldsymbol{\alpha})_{j+1/2} = \bar{U} + \sum_{k=0}^{j-1} \Delta s (\alpha_{1,k+1/2} \mathbf{n}_{1,k+1/2} + \alpha_{2,k+1/2} \mathbf{n}_{2,k+1/2}) + \frac{\Delta s}{2} (\alpha_{1,j+1/2} \mathbf{n}_{1,j+1/2} + \alpha_{2,j+1/2} \mathbf{n}_{2,j+1/2}). \quad (3.78)$$

The midpoint quadrature assigns a weight of Δs to all prior nodes $1/2, 3/2, \dots, j - 1/2$, and then a weight of $\Delta s/2$ to the current node $j + 1/2$. The matrix \mathbf{K}^* is the transpose of \mathbf{K} .

Finally, the parallel angular velocity (3.60) is obtained at the link centers by dotting the full angular velocity, $\mathcal{M}_{\text{rt}} \mathbf{f} + \mathcal{M}_{\text{tr}} (n^{\parallel} \boldsymbol{\tau})$ (computed at the link centers using the forces computed at the link centers), with the tangent vectors $\boldsymbol{\tau}$, which are also defined at the link centers.

Kirchhoff model. The discretization of the Kirchhoff model, which we take from [47], is more complicated because the forces $\boldsymbol{\Lambda}$ are defined at the *nodes*, and not the link centers. To communicate between the nodes and links, we let \mathbf{D}_{lb} be the $N \times (N + 1)$ differentiation matrix, which gives derivatives at the links from values at the nodes. We also define \mathbf{D}_{bl} as the $(N - 1) \times N$ differentiation matrix, which gives derivatives at interior nodes from values at the link centers. Likewise the $N \times (N + 1)$ averaging matrix \mathbf{A}_{lb} gives averages at links from nodes, and the $(N - 1) \times N$ averaging matrix \mathbf{A}_{bl} gives averages at interior nodes from

links. This gives the $(N - 1) \times (N + 1)$ system of equations discretizing (3.70)

$$[\mathbf{D}_{\text{bl}}(\mathbf{M}_{\text{tt}}\mathbf{D}_{\text{lb}} + \mathbf{M}_{\text{tr}}\mathbf{C}\mathbf{A}_{\text{lb}}) + \mathbf{A}_{\text{bl}}\mathbf{C}(\mathbf{M}_{\text{rt}}\mathbf{D}_{\text{lb}} + \mathbf{M}_{\text{rr}}\mathbf{C}\mathbf{A}_{\text{lb}})] \Delta\mathbf{\Lambda} = -\mathbf{m} \quad (3.79)$$

$$= -(\mathbf{D}_{\text{bl}}\mathbf{M}_{\text{tt}} + \mathbf{A}_{\text{bl}}\mathbf{C}\mathbf{M}_{\text{rt}}) \left(\mathbf{f}^{(\kappa)} + \mathbf{f}^{(\gamma)} + \boldsymbol{\lambda}_E \right) - (\mathbf{D}_{\text{bl}}\mathbf{M}_{\text{tr}} + \mathbf{A}_{\text{bl}}\mathbf{C}\mathbf{M}_{\text{rr}}) (n^{\parallel}\boldsymbol{\tau}) \quad (3.80)$$

which we solve for $\Delta\mathbf{\Lambda}$ at the $N - 1$ *interior blobs* away from the fiber boundaries. Because $\mathbf{\Lambda}$ is defined at $N + 1$ points, the other two equations are the BCs $\Delta\mathbf{\Lambda}_j = \mathbf{0}$ for $j = 0, N$, see (3.44). The matrix \mathbf{C} represents cross products with $\boldsymbol{\tau}_{j+1/2} = \boldsymbol{\tau}((j + 1/2)\Delta s)$, where in the second-order method $\boldsymbol{\tau} = \partial_s \mathbb{X}$ is evaluated analytically at the links and $(\mathbf{C}\mathbf{x})_{j+1/2} = \boldsymbol{\tau}_{j+1/2} \times \mathbf{x}_{j+1/2}$. Notice that the calculation of the mismatch (3.80) requires calculation of the Euler quantities $\mathbf{f}^{(\kappa)}$ and $\mathbf{f}^{(\gamma)}$ and $\boldsymbol{\lambda}_E$ at the N links, which we first obtain by solving the Euler problem to second order accuracy.

3.2.4.3 Numerical results

To numerically compare the Euler and Kirchhoff constitutive models, we will apply the second-order methods to a curved fiber with $L = 2$ and tangent vector

$$\boldsymbol{\tau}(s) = \frac{1}{\sqrt{2}} \left(\cos(s^3(s - L)^3), \sin(s^3(s - L)^3), 1 \right), \quad (3.81)$$

(see Fig. 6.8 for a visual). This fiber is chosen to satisfy the free fiber boundary conditions (3.47). We will also assume an instantaneous twist

$$\psi = \sin(2\pi s), \quad (3.82)$$

so that all of the free fiber boundary conditions (3.47) are satisfied. Because \mathbb{X} satisfies the free-fiber BCs, we can compute the elastic force to spectral accuracy simply by differentiating $\boldsymbol{\tau}$ three times analytically, and then setting $\mathbf{f}^{(\kappa)} = -\kappa \partial_s^3 \boldsymbol{\tau}$. For this section, we will compute

all geometric quantities, for example $\partial_s \boldsymbol{\tau}$, $\partial_s^2 \boldsymbol{\tau}$ and $\partial_s \psi$, analytically. We will use a viscosity of $\mu = 1$ and aspect ratios $\hat{\epsilon} = \hat{a}/L$ that are a function of the regularized singularity radius \hat{a} ; the corresponding value of the true fiber radius $a \approx \hat{a}/1.12$ (see (2.53)).

Bending but no twisting. We first solve the mismatch problem (3.70) with bending modulus $\kappa = 1$ and twist modulus $\gamma = 0$ using our second-order method with $N = 1/\hat{\epsilon}$ points. In Fig. 3.1, we plot \mathbf{m} , $\Delta \boldsymbol{\Lambda}$, $\Delta \mathbf{U}$, $\boldsymbol{\tau} \times \Delta \boldsymbol{\Psi}$, and $\partial_s \Delta \mathbf{U}$ in the basis of the local triad for $\hat{\epsilon} = \hat{a}/L = 0.005$. We observe several features: first, the mismatch \mathbf{m} is smooth until it reaches the endpoints, after which it jumps to a large value at the boundaries. In the fiber interior, the function $\Delta \boldsymbol{\Lambda}$ is approximately a constant multiple of \mathbf{m} , but at the endpoints it transitions to zero (as required by the boundary condition). Away from the fiber boundaries, the mismatch \mathbf{m} is almost entirely compensated for by the angular velocity perturbation $\boldsymbol{\tau} \times \Delta \boldsymbol{\Psi}$, and $\partial_s \Delta \mathbf{U}$ is small in the fiber interior, validating (3.73).

Figure 3.1 also shows that $\boldsymbol{\tau} \times \Delta \boldsymbol{\Psi}$ by itself does not compensate for the mismatch \mathbf{m} near the boundaries, since the dominant contribution from torque $\Delta \boldsymbol{\Lambda}$ must be zero there. The mismatch \mathbf{m} near the boundaries must therefore be compensated for by $\partial_s \Delta \mathbf{U}$, which means that $\Delta \mathbf{U}$ is highly nonsmooth at the boundaries. Since $\Delta \mathbf{U}$ comes from $\Delta \boldsymbol{\Lambda}$, it is not surprising that $\Delta \boldsymbol{\Lambda}$ is also not smooth at the boundaries, as it takes on its largest values near the endpoints, only to have to fall back to zero as required by the boundary conditions.

We will focus on the equivalence of the Euler and Kirchhoff models in the fiber *interior*, or points $\mathcal{O}(1)$ from the fiber endpoints. We therefore define an interior “norm” for all functions which is simply the L^2 norm on $0.05 \leq s/L \leq 0.95$. In Fig. 3.2(a), we plot the relative difference in the velocity and forces using this interior norm. We observe convergence of order $\hat{\epsilon}^2$ for all quantities. This verifies, and in fact exceeds, the predictions of (3.73) and (3.74) of $\hat{\epsilon}^2 \log \hat{\epsilon}$. The reason is that the mismatch is actually dominated by $\mathcal{O}(1)$ terms for the values of $\hat{\epsilon}$ we are interested in, and so our estimates are off by a factor of $\log \hat{\epsilon}$.

Twisting but no bending. We now proceed to the case of finite twist modulus, $\gamma = 1$,

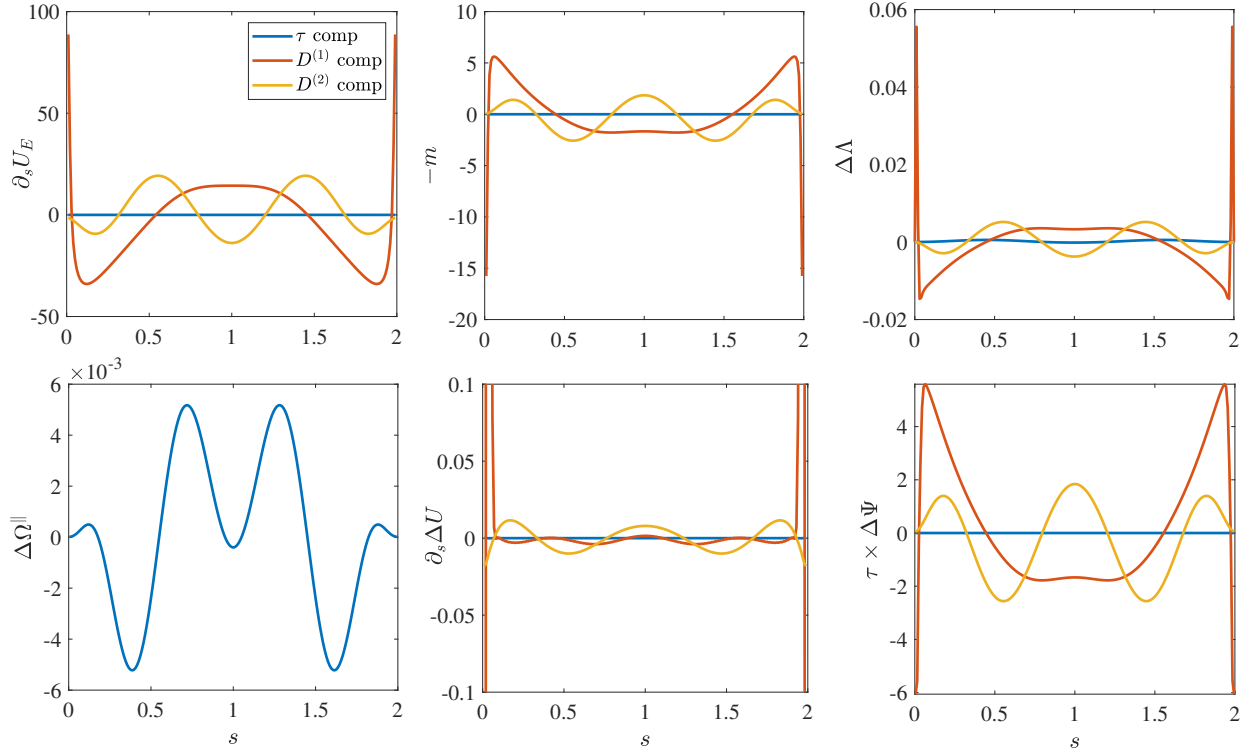


Figure 3.1: Solving the mismatch problem (3.70) with bending stiffness $\kappa = 1$, twist modulus $\gamma = 0$, and slenderness $\hat{\epsilon} = \hat{a}/L = 0.005$. All quantities are shown in the basis of the local triad, with the $\boldsymbol{\tau}$ component in blue and the $\mathbf{D}^{(1/2)}$ component in red/yellow. In the top row, from left to right, we show $\partial_s \mathbf{U}_E$ (the $\boldsymbol{\tau}$ component of this is zero because the fiber is inextensible), (negative) mismatch $-\mathbf{m}$, and the corresponding $\Delta \mathbf{\Lambda}$. Notice how the shapes of $\Delta \mathbf{\Lambda}$ and \mathbf{m} match, which implies that $\Delta \mathbf{\Lambda}(s) \approx c(\hat{\epsilon}) \mathbf{m}(s)$ away from the fiber boundaries. In the bottom row, from left to right, we show the resulting velocity perturbations $\Delta \Omega^{\parallel}$, $\partial_s \Delta \mathbf{U}$, and $\boldsymbol{\tau} \times \Delta \boldsymbol{\Psi}$. In the fiber interior, $\boldsymbol{\tau} \times \Delta \boldsymbol{\Psi} \approx -\mathbf{m}$, and $\partial_s \Delta \mathbf{U}(s) \approx \mathbf{0}$. At the endpoints, $\partial_s \Delta \mathbf{U}$ reaches values as negative as -10 , so we truncate the y axis to show the behavior in the fiber interior.

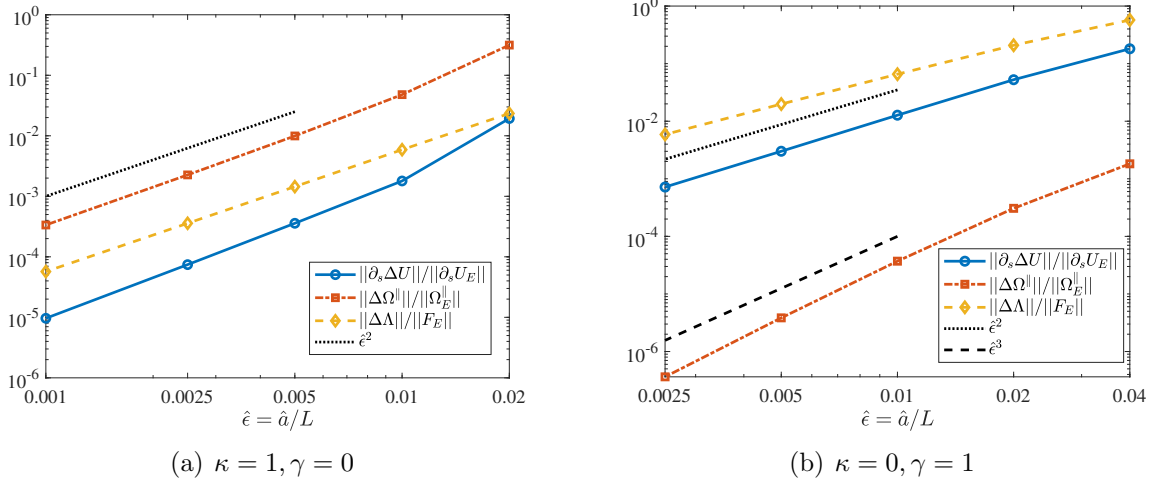


Figure 3.2: Decay of (normalized) magnitude of $\Delta\mathbf{\Lambda}$ (yellow), $\partial_s\Delta\mathbf{U}$ (blue), and $\Delta\Omega^\parallel$ (red) as slenderness $\hat{\epsilon} \rightarrow 0$ for: (a) $\gamma = 0$ and $\kappa = 1$ with $N = 1/\hat{\epsilon} + 1$ grid points; and, (b) $\gamma = 1$ and $\kappa = 0$ and $N = 4/\hat{\epsilon} + 1$ grid points. We use an interior “norm” (L^2 norm on $0.05 \leq s/L \leq 0.95$). All quantities scale like $\hat{\epsilon}^2$ when properly normalized, with the exception of $\Delta\Psi^\parallel = \Delta\Omega^\parallel$ when there is twist, which scales faster than $\hat{\epsilon}^3$. This establishes (3.73) and (3.74) numerically.

with zero bend modulus. We first plot the quantities related to the mismatch problem (3.70) in Fig. 3.3. We see similar behavior to the case without twist and with bending force. The velocity difference $\partial_s\Delta\mathbf{U}$ is again small, since most of the mismatch is compensated for by $\boldsymbol{\tau} \times \Delta\Psi$. A major difference from the non-twisted case is that $\Delta\mathbf{\Lambda}$ and $\Delta\mathbf{U}$ are smoother at the boundaries, since the force when $\kappa = 0$, given in (3.51), is zero at the fiber boundaries because of the free fiber boundary conditions (3.47). Because the force decays to zero there, the mismatch is much smaller at the boundaries (but is still not zero). This means that $\boldsymbol{\tau} \times \Delta\Psi$ can almost compensate for the entire mismatch everywhere. At the boundaries, whatever small mismatch there is must be compensated for by $\partial_s\Delta\mathbf{U}$, and so as $\hat{\epsilon}$ becomes smaller we do see some localized peaks in $\partial_s\Delta\mathbf{U}$ near the endpoints.

In Figure 3.2(b), we plot the convergence of the velocities and forces for $\gamma = 1$ and $\kappa = 0$. For $\Delta\mathbf{\Lambda}$ and $\Delta\mathbf{U}$, we observe a similar rate of convergence (approximately $\hat{\epsilon}^2$) to Fig. 3.2(a). In accordance with (3.74), the parallel angular velocity Ω_E^\parallel is larger when there is twist, and so as a result the relative rate at which $\Delta\Omega^\parallel$ approaches zero is faster. In

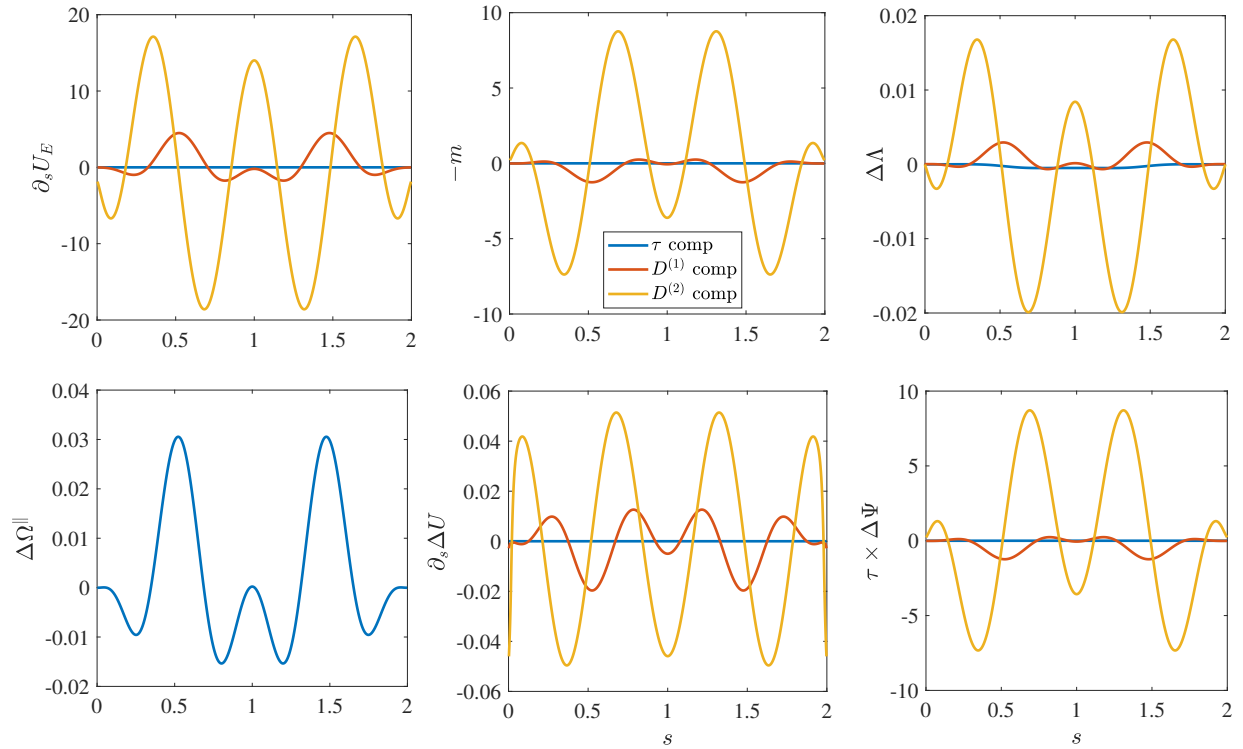


Figure 3.3: Solutions of the mismatch problem (3.70) with twist modulus $\gamma = 1$, bending modulus $\kappa = 0$, and slenderness ratio $\hat{\epsilon} = \hat{a}/L = 0.005$. The behavior is qualitatively the same as the mismatch problem without twist, with the perturbations in $\Delta\Omega^{\parallel}$ and $\partial_s \Delta U$ being small relative to the values obtained from the Euler method. Because the twist force $\mathbf{f}^{(\gamma)}$ is zero at the endpoints, the mismatch there is small, and $\partial_s \Delta U$ is smoother than seen in Fig. 3.1.

(3.74) we estimated $\hat{\epsilon}^4 (\log \hat{\epsilon})^2$; in Fig. 3.2(b) we simply show it is faster than $\hat{\epsilon}^3$, but the line isn't straight, which indicates some kind of log correction. This verifies (3.73) and (3.74) numerically for finite twist modulus. Combining with our asymptotic estimates in Section 3.2.4.1, we can conclude that the Euler and Kirchhoff models are indeed equivalent to roughly $\mathcal{O}(\epsilon^2)$ in the fiber interior. Our choice will be to use the Euler model going forward, as it displays smooth behavior at the endpoints and is consequently easier to design a (spectral) numerical method around.

Chapter 4

Conclusions and governing equations

This part of the dissertation was devoted to developing the governing equations that we will solve numerically in the next part. All of the content in this part was to some extent driven by the following question: how do we model the hydrodynamics of rotating inextensible slender fibers? The emphasis here is on rotation, as this drove the development of our RPY-based slender body theory in Chapter 2 and our comparison of the Kirchhoff and Euler models in Chapter 3.

In Chapter 1, we examined the deficiencies of existing slender body theories (SBTs) for rotating filaments. We showed that the singularity-based SBTs of Keller+Rubinow and Johnson [19, 20], which were designed with the resistance problem in mind, both suffer from the same issues for rotation-translation coupling in the mobility problem: when the fiber is driven by an applied torque at its base, their theories imply that torque induces an $\mathcal{O}(1)$ angular-dependent translational velocity on the fiber cross section, which violates the SBT boundary condition that the fiber only translates and rotates. Indeed, the recent work of Walker et al. [86] justifies Johnson's neglect of rot-trans coupling, but only when the translational velocity is $\mathcal{O}(\epsilon)$ times the rotational velocity, which is not an assumption we want to restrict ourselves to. As we show in Chapter 6, neglecting rotation-translation coupling can lead to overestimation of the torque required to induce large-scale filament

deformations. It is therefore important to derive an SBT which properly treats rotation-translation coupling.

In Chapter 2, we used a different definition of a filament as a line of infinitely many Rotne-Prager-Yamakawa (RPY) regularized singularities, similar to what has been done previously for translation by Cortez and Nicholas using regularized Stokeslets [70]. We then defined each component of the 2×2 mobility operator (translation-translation, rotation-translation, etc.) as a line integral of the corresponding fundamental regularized singularity given in (2.5). Because these regularized singularities, which define the mobility for a pair of “blobs,” give a symmetric 2×2 mobility, our mobility for fibers is automatically symmetric, which alleviates one of the two major problems with previous SBTs. The other issue, the $\mathcal{O}(1)$ angular dependence of the translational velocity on a cross section, is sidestepped because we use an isotropic regularized singularity function (a delta function on a sphere). This comes at a price: by using a series of spheres to represent the fiber centerline, we lose fidelity to the true three-dimensional geometry and the fluid flows near the filament surface. While this is not ideal, it gives a mobility which is physical and simple to compute, since all that results is a local drag term plus a finite part integral for the trans-trans, rot-trans, and trans-rot terms, for which special quadrature schemes have already been developed [80].

Using matched asymptotics on the RPY kernels, we found a 2×2 mobility relationship, which we then compared to SBT for a three-dimensional tubular fiber. For translation-translation and rotation-rotation, we found that the two mobilities take the same form, but with an $\mathcal{O}(1)$ difference in the constants. The discrepancy in the constants can be removed via a judicious choice of regularized singularity radius, which is different for each component of the mobility; $\hat{a} \approx 1.06a$ gives a perfect match in the rot-rot mobilities, while $\hat{a} \approx 1.12a$ matches the trans-trans mobilities. We used the correspondence between RPY asymptotics and SBT to posit the relationship (2.57) for the off-diagonal rotation-translation components in SBT. The relationship has a single unknown coefficient k , which we tried to

obtain in Section 2.1.6 from boundary integral calculations. Our results showed the value of the coefficient requires an extremely accurate numerical method the more-recently developed quadrature schemes of [79] (which gave $k = 1$) produced different results than the first-order numerical method we presented in [E] ($k \approx 2.85$). Because of this, more work (testing on hundreds or thousands of additional shapes) is required to definitively establish that (2.57) holds with the same value of $k = 1$ for a variety of filament shapes.

Using the RPY kernels for the mobility simplified a number of potential issues that arise for more rigorous models of the three-dimensional fiber geometry. First, we were able to easily derive mobility formulas near the fiber endpoints simply by modifying the domain of integration in the inner expansion. This is in contrast to the more rigorous three-dimensional case, where a nonsingular endpoint flow is obtained by assuming an ellipsoidally-decaying radius function and distributing the centerline singularities only between the generalized foci of the body [20]. Furthermore, the RPY mobility immediately generalizes to multiple filaments (see Section 2.2), since the flow induced on filament j can also be computed by an integral of the RPY kernel over filament i . In SBT, while the pointwise velocity induced by singularities might be constant on a cross section of fiber i , the same might not be true on a nearby fiber j , which leads to an ambiguity in how the induced flow on the centerline of fiber j is assigned. This is especially relevant when fibers are in near contact, and shows why only three-dimensional boundary integral approaches are precise when fibers are $\mathcal{O}(\epsilon)$ apart; such calculations are however very complex and expensive for suspensions of many fibers [87].

In Chapter 3, we looked at how the mechanical properties of the filament coupled to those of the fluid. We derived a new formulation of inextensibility based on rotating tangent vectors, which we coupled with the principle that constraint forces do no work to yield a saddle point system in two unknowns (the tangent vector rotation rates and constraint forces). We examined two different models of fluid-fiber coupling for twist elasticity, finding

that the dynamics of the ‘‘Kirchhoff’’ model, which constrains the perpendicular rotational fluid velocity on the fiber surface to be equal to that obtained from translational dynamics, are identical to those of the ‘‘Euler’’ model, which considers only parallel rotational dynamics and translational velocity, in the limit $\epsilon \rightarrow 0$. We choose to utilize the Euler model in the remainder of this dissertation, since the extra degree of rotational freedom (parallel torque and parallel rotational velocity) can be easily incorporated into the formulation without twist. To compute the additional force and parallel torque due to twist, we use the Bishop (twist-free) frame and evolved the angle of twist relative to that frame [84].

4.1 Summary of equations

We now summarize the equations we will solve in Part II of the dissertation. We recall that inextensible motions of the fiber centerline are parameterized by

$$(\boldsymbol{\kappa}\boldsymbol{\alpha})(s) = \mathbf{U}_{\text{MP}} + \int_{L/2}^s \boldsymbol{\Omega}(s') \times \boldsymbol{\tau}(s') ds'. \quad (4.1)$$

Letting $\boldsymbol{\alpha} = (\boldsymbol{\Omega}, \mathbf{U}_{\text{MP}})$, the dynamics of the fiber centerline(s) are obtained by solving the saddle point system

$$\begin{pmatrix} -\mathcal{M}_{\text{tt}} & \boldsymbol{\kappa} \\ \boldsymbol{\kappa}^* & \mathbf{0} \end{pmatrix} \begin{pmatrix} \boldsymbol{\lambda} \\ \boldsymbol{\alpha} \end{pmatrix} = \begin{pmatrix} \mathcal{M}_{\text{tt}} (\mathbf{f}^{(\kappa)} + \mathbf{f}^{(\gamma)}) + \mathcal{M}_{\text{tr}} (n^{\parallel} \boldsymbol{\tau}) + \mathbf{u}_0(\mathbb{X}) \\ \mathbf{0} \end{pmatrix}, \quad (4.2)$$

$$\mathbf{f}^{(\kappa)} = -\kappa \partial_s^4 \mathbb{X}, \quad \mathbf{f}^{(\gamma)} = \gamma \partial_s (\psi (\partial_s \mathbb{X} \times \partial_s^2 \mathbb{X})), \quad n^{\parallel} = \gamma \partial_s \psi,$$

$$\boldsymbol{\kappa}^* \boldsymbol{\lambda} := \begin{pmatrix} \int_0^L \boldsymbol{\lambda}(s) ds \\ \boldsymbol{\tau}(s) \times \int_s^L \boldsymbol{\lambda}(s') ds' \end{pmatrix},$$

subject to the free fiber BCs (3.47). Solving the saddle-point system (4.2) in the least squares sense yields the constraint forces $\boldsymbol{\lambda} = \partial_s (T\boldsymbol{\tau})$ and translational velocity

$$\partial_t \mathbb{X}(s) = (\boldsymbol{\mathcal{K}}\boldsymbol{\alpha})(s) = \left(\boldsymbol{\mathcal{K}} (\boldsymbol{\mathcal{K}}^* \boldsymbol{\mathcal{M}}_{\text{tt}}^{-1} \boldsymbol{\mathcal{K}})^{-1} \boldsymbol{\mathcal{K}}^* \right) \boldsymbol{\mathcal{M}}_{\text{tt}}^{-1} \boldsymbol{U}_{\text{Ex}} := \boldsymbol{\mathcal{K}} \boldsymbol{\mathcal{N}} \boldsymbol{\mathcal{K}}^* \boldsymbol{\mathcal{M}}_{\text{tt}}^{-1} \boldsymbol{U}_{\text{Ex}}, \quad (4.3)$$

where $\boldsymbol{U}_{\text{Ex}}$ denotes the right hand side in the saddle point system (4.2) and the projected mobility operator $\boldsymbol{\mathcal{N}}$ is self-adjoint with respect to the L^2 inner product, $\boldsymbol{\mathcal{N}} = \boldsymbol{\mathcal{N}}^*$. The saddle point system (4.2) as written can apply to one fiber or multiple fibers, in which case the operators $\boldsymbol{\mathcal{K}}$ and $\boldsymbol{\mathcal{K}}^*$ are block-diagonal operators with the diagonal having the corresponding operator for each individual fiber.

Because $\boldsymbol{\Omega} = \boldsymbol{\Omega}^{\parallel} \boldsymbol{\tau}$ is in the null space of $\boldsymbol{\mathcal{K}}$, it is not determined by (4.2), and is instead obtained from the post-processing step (3.60),

$$\boldsymbol{\Omega}^{\parallel} = (\boldsymbol{\mathcal{M}}_{\text{rt}} \boldsymbol{f} + \boldsymbol{\mathcal{M}}_{\text{rr}} (n^{\parallel} \boldsymbol{\tau})) \cdot \boldsymbol{\tau}, \quad (4.4)$$

where each component of the mobility in the foregoing equations is defined via a line integral of the RPY regularized singularity in (2.4). The evolution of the twist $\psi = \partial_s \theta$ is given by (3.37)

$$\partial_t \psi = \partial_s \boldsymbol{\Omega}^{\parallel} - (\boldsymbol{\Omega} \cdot \partial_s \boldsymbol{\tau}) = \partial_s \boldsymbol{\Omega}^{\parallel} - (\boldsymbol{\Omega}^{\perp} \cdot \partial_s \boldsymbol{\tau}) \quad (4.5)$$

where in the second equality we have used the formula $\boldsymbol{\tau} \cdot \partial_s \boldsymbol{\tau} = 0$, which means that the parallel part of $\boldsymbol{\Omega}$ makes no contribution to the dot product $\boldsymbol{\Omega} \cdot \partial_s \boldsymbol{\tau}$.

If needed (e.g., for visualization or intrinsic curvature), the material frame vectors $\boldsymbol{D}^{(1)}(s)$ and $\boldsymbol{D}^{(2)}(s)$ can be obtained by computing the Bishop frame from (3.26), then twisting it by the angle θ . The Bishop frame has to be chosen at one point on the fiber, and the angle θ is only determined up to a constant. Our choice is to keep track of the material vector $\boldsymbol{D}^{(1)}$

at the fiber midpoint by explicitly evolving it via rotation by $\boldsymbol{\Omega}(L/2)$,

$$\partial_t \mathbf{D}^{(1)}(L/2) = (\boldsymbol{\Omega}^\perp(L/2) + \Omega^\parallel(L/2)\boldsymbol{\tau}(L/2)) \times \mathbf{D}^{(1)}(L/2). \quad (4.6)$$

We then assign $\theta(L/2) = 0$, i.e., require that the Bishop frame be the same as the material frame at $L/2$

$$\mathbf{b}^{(1)}(L/2) = \mathbf{D}^{(1)}(L/2), \quad (4.7)$$

which provides a boundary condition for the Bishop ODE (3.26).

4.2 Future directions

Our work here is by no means the final say on SBT for twisting filaments, but is rather intended as a starting point to provoke future research in three main directions. First, there is the possibility of developing a singularity representation for the flow around a twisting filament that properly matches the velocity on the fiber cross section when the rotational velocity is much faster than the translational velocity. For an ellipsoidal filament, the velocity on a cross section due to a rotlet (1.10) can be expanded using (1.12) as

$$\begin{aligned} 8\pi\mu\mathbf{u}^{(\Psi)}\left(\widehat{\mathbb{X}}(s, \theta)\right) &= \frac{2n^\parallel}{a}(\boldsymbol{\tau} \times \widehat{\mathbf{r}}) + \frac{n^\parallel}{2}(\ln(\epsilon^{-2}) - 3)(\boldsymbol{\tau} \times \partial_s \boldsymbol{\tau}) + \mathbf{U}_n^{(\text{FP})} \\ &+ \frac{n^\parallel \kappa}{2}(\mathbf{e}_b \cos 2\theta - \mathbf{e}_n \sin 2\theta) + \mathcal{O}(\epsilon). \end{aligned} \quad (4.8)$$

Thus, the task is to find another singularity (or set of singularities) that can eliminate the $\mathcal{O}(1)$ dependence on 2θ in the second line, thereby making the cross-sectional velocity a pure rotation (the $\boldsymbol{\tau} \times \widehat{\mathbf{r}}$ term) plus a pure translation (the next two terms in the first line). Based on our empirical SBT (2.57), we expect such a singularity, if it exists, to eliminate the angular dependence while introducing a small correction to the $\mathcal{O}(1)$ term multiplying $\boldsymbol{\tau} \times \partial_s \boldsymbol{\tau}$.

There is also an opportunity to analyze what Stokes-related partial differential equation such a slender body theory is the solution of. For the case of translation, Mori, Ohm, and others [24–26] have shown that slender body theory is an $\mathcal{O}(\epsilon)$ approximation to a Stokes partial differential equation where the velocity on the fiber surface is required to be independent of cross sectional angle. Repeating this analysis for twist, where the velocity could be a strong function of angle, but only as a rigid body rotation, might offer insights into why it is so difficult to obtain a singularity solution.

Part II

Spectral methods for fiber dynamics

Chapter 5

Introduction

This part of the thesis is concerned with the numerical methods required to efficiently¹ simulate the dynamics of many filaments in viscous flow. While the equations we wrote in Section 4.1 apply to the deterministic bending and twisting dynamics of many filaments interacting hydrodynamically, in this part of the thesis we will only consider twisting dynamics for a single filament. That is, we will develop numerical methods for the following problems:

1. Filaments which resist bending *and* twist elasticity, but with mobility operators \mathcal{M}_{tr} , \mathcal{M}_{rt} , and \mathcal{M}_{rr} that only apply on one filament at a time (Chapter 6).
2. Filaments with free ends which resist bending, but not twisting, and interact hydrodynamically through the viscous medium (Chapter 7).
3. Filaments which exhibit thermal bending (but not twisting) fluctuations and interact hydrodynamically through the viscous medium (Chapter 8). This is an extension of #2 to *fluctuating* fibers, which were not discussed in our theoretical treatment of the problem. As we will see, the reason for this is that we first need to discretize the filaments in space prior to adding thermal bending fluctuations.

¹We ask the reader to forgive our use of this word, which is probably the most overused in all of the numerical analysis literature. Indeed, what do you call the method which is the worst out of all those labeled “efficient?” Hopefully the answer is not “the one described here!”

The key bridge from Chapter 6 to Chapter 7 is the demonstration in Section 6.5.1 that twist elasticity contributes negligibly to the dynamics of free-ended filaments. Building on this, in Chapter 7, we specialize our spatio-temporal discretizations to the case of free-ended filaments which only resist bending and interacting hydrodynamically. We then use this spatio-temporal discretization in Chapter 8 to construct a Langevin equation and corresponding temporal integrator for thermal fluctuations. This introduction discusses the literature for each of these chapters in order.

5.1 Single fiber hydrodynamics

Once we specify the model of the fiber mechanics, as we did in Chapter 3, a common numerical choice is to represent the fiber by a discrete set of regularized singularities or “blobs,” where each blob exerts a regularized force (coming from internal mechanics) on the fluid, and the blob velocity is obtained by interpolating the fluid velocity back onto the blob center, as described in Section 2.1. The regularized delta function we use to spread the force to the fluid and interpolate the velocity back onto the blobs makes the method: in this dissertation, we always use the RPY kernel, for which the regularized delta function is a surface delta function on a sphere. Still, other choices are possible, including the immersed boundary method [60, 74, 76], the method of regularized Stokeslets [62, 63], and the force coupling method [61, 73, 88].

No matter the method chosen, blob-based methods suffer from the same fundamental problem when trying to represent slender tubular bodies. If we want the fluid to see a tubular body rather than a series of disconnected blobs, then the fluid force and fiber velocity ought to be obtained from *integrals* of the regularized singularities (just as SBT is obtained from integrals of singularities), as we do in Chapter 2. Thus, in blob-based methods, the fluid will see a tubular body if there are sufficiently many points for the *sum* over the blobs to be indistinguishable from the integral. It has been shown [69, 76, 89] that this requires $1/\epsilon$

blobs, which means that blob based methods do not scale well in the slender limit. While this limitation can be partially overcome with adaptive mesh refinement [90], grid coarsening with local velocity correction [91–93], and kernel-independent fast multipole methods (to accelerate many-body sums) [94], the fact remains that to achieve controlled accuracy for dilute suspensions of many fibers, discretizing the fiber by a collection of marker points is much less efficient than treating it as a continuum. This is especially true for semi-flexible or stiff fibers, where the smooth fiber shapes are well represented in a spectral basis with rapidly decaying coefficients.

In order to circumvent this problem, we use the formulation of Chapter 2 to define the velocity as an *integral* of the RPY singularities over the continuum curve $\mathbb{X}(s)$, which we already showed is asymptotically equivalent to the more common Keller-Rubinow slender body theory [19], if the RPY regularization radius $\hat{a} = ae^{3/2}/4 \approx 1.12a$. The question then becomes how to numerically track the curve $\mathbb{X}(s)$. Since the fiber shapes are smooth, we choose to use a set of Chebyshev tracking points \mathbf{X} , which define a Chebyshev interpolant of the fiber centerline $\mathbb{X}(s)$. This gives us all the tools we need to define the hydrodynamic mobility.

A more difficult task is to evaluate the mobility efficiently in a manner independent of the fiber slenderness. One possibility is to use the asymptotic formulas of Chapter 2, which remove the \hat{a} lengthscale by evaluating those integrals asymptotically, and leave an $\mathcal{O}(1)$ finite part integral to be evaluated, just as in SBT. The finite part integral poses a serious challenge for numerical methods for two reasons: first, its integrand is singular, so it cannot be done with traditional direct quadrature schemes. Second, it turns out that the eigenvalues of the SBT continuum operator \mathcal{M}_{tt} become negative for high-frequency modes [25]. This unphysical behavior occurs when the finite part integral dominates over the positive definite local drag term in the mobility. For ellipsoidal fibers, the local drag term is even more dominant than for the cylindrical fibers we consider here, and Tornberg and Shelley [22,

Sec. B.1] showed that only 33 modes have positive eigenvalues when $\epsilon = 10^{-2}$. As the fibers become more cylindrical, the discrete asymptotic mobility \mathbf{M}_{tt} becomes very ill-conditioned even for a moderate number of points (see Fig. 6.9), and the constraint force $\boldsymbol{\lambda}$ fails to converge as the number of points increases, even in the weak sense.

There have been several attempts to date [22, 25, 95] to remedy this problem and keep the eigenvalues of the translational mobility away from zero. There are, broadly speaking, two classes of these methods: methods which regularize the original slender body equations, and methods which change the integral kernel used to evaluate the fiber velocity. In the first class, Tornberg and Shelley [22] regularized only the finite part integral by adding a function $\delta(s, \epsilon)$ in the denominator of each term. They showed that a proper choice of $\delta \sim \epsilon$ does not change the overall asymptotic accuracy of SBT. Mori and Ohm [25] justified this approach theoretically by analyzing the error obtained with respect to a more faithful model of the true three-dimensional dynamics, which they term the “slender body PDE.” They showed that both the Tornberg and Shelley regularization and spectral truncation (only including a finite number of modes) both preserve the asymptotic accuracy of SBT while keeping the eigenvalues of \mathbf{M}_{tt} positive (provided that, in the latter case, the number of modes is sufficiently small).

Recently, Andersson et al. recognized that another choice of regularization is to change the *original* integral kernel, which in slender body theory is the linear combination of a Stokeslet and a doublet. They postulated that expanding the SBT kernel a distance a away from the fiber centerline, then eliminating all angular-dependent terms, gives an integral kernel that eliminates the high wave number instabilities observed in the original slender body equations [95, Fig. 1]. They did not, however, develop an efficient quadrature scheme for their integral kernel, since the number of points required for a given accuracy still scales as $1/\epsilon$ [95, Fig. 3].

It is clear then, that while asymptotic evaluation of the RPY mobility gives a formula that

compares well with SBT, this is something which is not useful in practice, since SBT itself is not well-suited to numerical computations. Seeing as the original integral of the RPY kernel (2.63) *is* well-suited to numerical computations, in this chapter we will develop numerical methods which compute this integral to numerical, rather than asymptotic, accuracy. These methods turn out to be based on schemes previously designed for the SBT finite part integral [80]. Thus our schemes, which are described in Section 6.2, will immediately generalize to SBT integrals. These schemes form the numerical treatment of the mobility on a single fiber.

For inextensibility, we have already discussed (see Section 1.2.1) the line-tension based method of Tornberg and Shelley [22] in continuum. The discrete version of this method has since been used in applications with flexible (and sometimes fluctuating) filaments [96, 97], and was also extended to simulate falling rigid fibers, the novelty there being that many of the SBT-related integrals can be done analytically [98]. But in discrete simulations, all of these methods add a discrete penalty parameter to the line tension equation, which can set the temporal stiffness. The filament is then no longer discretely inextensible, which makes it similar to immersed boundary approaches with penalty terms [49, 50]. To our knowledge, only the recent approaches of Schoeller et al. [47] and Jabbarzadeh and Fu [41] enforce inextensibility rigorously with Lagrange multipliers. In the case of [41], the fiber is broken up into segments of regularized point forces, and each segment is updated via a rotation that preserves inextensibility exactly. Yet both [47] and [41] suffer from the same pitfall as all regularization methods already mentioned—the scaling of the number of required points with ϵ . There is therefore a need for a numerical method which enforces inextensibility rigorously *and* has a cost independent of the fiber slenderness. This is the method we will develop in Chapter 6.

Because SBT only reliably gives the trans-trans and rot-rot terms of the mobility (see Chapter 1), a disconnect has emerged in the numerical analysis literature in which immersed boundary methods [49], regularized Stokeslets [41, 50], and the force coupling method [47],

have accounted for the full hydrodynamics (including rot-trans and trans-rot coupling) of twisting fibers, while methods that use SBT [33, 52, 54] have ignored rot-trans coupling. Because our RPY-integral formulation in (2.63) includes fully nonlocal hydrodynamics, and our quadrature schemes allow us to compute the integrals with cost independent of ϵ , we can for the first time look at the error in excluding rot-trans coupling. To do this, we introduce a (constrained) backward Euler temporal integration scheme to solve dynamic problems involving twist. We apply this numerical method to two problems involving twisting fibers: a relaxing bent fiber, where we explore the feedback of twist on the fiber dynamics, and a twirling fiber, where we look at the critical frequency at which the fiber starts to whirl and then transitions to a stable “overwhirling” state [54, 99–101]. For fiber relaxation, we show that the assumption of twist being in quasi-equilibrium [57, 84] is satisfied for materials in which the twist and bending moduli are comparable, as is the case in most materials (for inextensible isotropic elastic cylinders, the ratio is 2/3 [52]). For a twirling fiber, we show how the systematic neglect of rot-trans coupling in prior studies based on the Euler model [54, 99] leads to an overestimation in the critical frequency of 20%. Our results explain half of the deviation between prior theoretical/computational studies and what is observed experimentally [102].

5.2 Nonlocal hydrodynamics

When we consider the dynamics of the fiber centerline, the disturbance flow from the other filaments can essentially be treated as a background flow. There is then an iterative process whereby we update the background flow, and then solve a local problem for the constraint forces, continuing this until we get convergence. The question becomes how to evaluate the nonlocal flows efficiently. We will only consider the trans-trans mobility \mathcal{M}_{tt} for this part.

In the original approach of Tornberg and Shelley [22], the nonlocal integrals are treated by quadratic brute-force summation, which makes the method ill-suited to scale to many

fibers. Thus Nazockdast et al. modified the approach by replacing the second-order spatial discretization of Tornberg and Shelley with a spectral spatial discretization and utilizing a kernel-independent FMM to accelerate sums. In doing this, they developed a parallel algorithm that makes it possible to simulate $\mathcal{O}(1000)$ fibers in linear time [15]. Our approach is ultimately similar to theirs: we too use a spectral discretization of the fiber centerline, but we treat the sums on a triply-periodic domain using fast Ewald summation [103, 104].

This summation ultimately boils down to discretizing each fiber by some number of (potentially oversampled) Chebyshev points, then performing a many-body RPY *blob* sum on those point-force pairs. This means that every nonlocal integral in (2.63) is discretized with the same number of points, implying either a loss of accuracy for nearly touching fibers, or a waste of computational power for well-separated ones. In Chapter 7, we explore the accuracy of direct quadrature and introduce a hybrid scheme whereby the RPY integrals for nearly-touching fibers are evaluated with a special quadrature method, while the others are treated with direct quadrature and Ewald splitting. There is then a computational trade-off where less direct quadrature points makes the Ewald step cheaper, but the “correction” (special quadrature) step more expensive. In our implementation, we found that doing the Ewald step on a GPU using the UAMMD library [105] is sufficiently fast for the optimal set-up to be upsampling for *all* pairs of fibers, with the special quadrature step skipped entirely. The upsampling factor is set in practice by using the minimum number such that results do not change when double the number of points is used.

For temporal integration, in deterministic methods we have the freedom to separate the *local* parts of the mobility, which are easy to invert, from the nonlocal parts, which are not. What we will do in practice is to write the (trans-trans) mobility as $\mathbf{M} = \mathbf{M}_L + \mathbf{M}_{NL}$, where the matrix \mathbf{M}_L is block-diagonal and represents the hydrodynamics of a single fiber, while the matrix \mathbf{M}_{NL} represents the contribution from all other fibers. We will then time-lag the nonlocal part, so that the number of nonlocal evaluations per time step is one, if

the suspension is sufficiently dilute to be dominated by local drag. When the suspension is more concentrated, the matrix \mathbf{M}_L becomes an efficient preconditioner for the full nonlocal system, which we solve with a few GMRES iterations, as discussed in Section 7.3.

5.3 Thermal fluctuations

So far, we have only introduced the numerical methods for simulating deterministic fibers in fluid. But what about Brownian fibers? There is already a large body of literature on this topic, which can be analyzed by asking the following two questions: is the chain extensible (with a penalty for stretching) or constrained to be exactly inextensible? And, is the chain being simulated discrete or continuous? That is, is the discrete chain considered the truth itself, or is it simply a discretization of a continuum equation which converges in the limit as the grid spacing goes to zero? This second question is quite difficult, so we begin our review of the literature with discrete chains.

Discrete extensible chains (known in the literature as bead-spring models) are the most commonly-used models for fluctuating filaments because of their simplicity [106–110]. In this case, the chain is represented by a series of beads or blobs connected by springs which penalize but do not prohibit extensibility and bending. There are therefore no constraints, and it is relatively straightforward to simulate the overdamped Langevin dynamics using standard algorithms [111, Sec. 3]. Actin filaments, which are nearly inextensible and therefore require a high stretching modulus to accurately simulate, push the boundaries of bead-spring models, as the high stretching and bending moduli restrict the maximum possible time step size, making long simulations of actin filament systems prohibitive [112].

It is therefore attractive to constrain the distance between each bead, which leads to a second class of methods known as bead-link or blob-link models [111, Sec. 4] (we will use the lesser-used “blob-link” terminology because of the coupling with “blob-based” hydrodynamics [113]). This approach, which in theory enables larger time step sizes and longer

simulations, is difficult in practice because it requires the formulation of an overdamped Langevin equation with the constraint that the chain is inextensible. In previous work, this has been done by defining a constrained Langevin equation for the *positions* of the blobs which has complicated stochastic drift terms and requires highly specialized algorithms to simulate [114, 115]. An alternative view, which is one we adopt here, is to view the degrees of freedom as the *link* orientations, as well as the position of the fiber center. For an inextensible chain, this reduces the problem to a series of connected rigid rods, for which we can make use of previous work by some of us on Langevin equations for rigid bodies [116].

As we have already discussed at length, however, the blob-link model breaks down in the limit of very slender fibers, such as actin filaments, which have a radius of $a = 4$ nm [117] and can have lengths L of a few microns in vivo and tens of microns in vitro [118, 119]. This set of parameters gives aspect ratios on the order of $\epsilon = a/L = 10^{-3}$, making the hydrodynamics multiscale in nature. Using a blob-link model with blobs spaced a lengthscale a apart often requires us to temporally resolve the fluctuations on that lengthscale; i.e., the time step size is constrained by the *smallest* lengthscale in the system. Thus, for slender filaments a blob-link discretization is prohibitively expensive in both the spatial and temporal variables using existing methods.

Since the hydrodynamics forces such a large number of degrees of freedom per fiber, our continuum framework, where the positions \mathbf{X} of the nodes become a function $\mathbb{X}(s)$ of a continuous variable s which describes the arclength of a curve, is an attractive choice [120, 121]. The problem with taking the continuum limit in the presence of thermal fluctuations is that it leads to ill-posed constrained Langevin partial-differential equations when dynamics are taken into account. From equilibrium statistical mechanics, it is known that for a free inextensible worm-like fiber the tangent vector $\boldsymbol{\tau}(s)$ performs Brownian motion in s on the unit sphere with diffusion coefficient ℓ_p^{-1} , where the persistence length $\ell_p = \kappa/k_B T$ is defined as the ratio of the bending stiffness κ to the thermal energy $k_B T$. On the other hand, we




Figure 5.1: Conceptual picture of our filament model. We begin with a blob-link chain at left, for which we need a large number of beads (second from left) to resolve the hydrodynamic interactions. A sensible model then becomes a continuum filament (second from right), shown as a solid curve, which we then replace by a fully discrete model based on Chebyshev collocation points (the gold nodes at right show 16 such points). In our approach, the model on the right is a coarse-grained approximation of the fully-resolved blob-link (second-from-left) model with the continuum model (second-from-right) used only in elasticity and hydrodynamic calculations. Figure 6.1 has a more detailed look at this discretization.

show in Chapter 8 that the overdamped Ito Langevin equation for a blob-link model with dynamics has stochastic drift terms which are required to obtain the correct equilibrium statistics (see Fig. 8.6, where we simulate the Langevin equation with and without the drift terms). These drift terms, which are a consequence of *both* the inextensibility constraint and the hydrodynamics, are only well-defined with respect to the *discrete* orientations of the tangent vectors, and do not converge in the continuum limit. This issue has been ignored in previous continuum methods for filament Brownian dynamics [39, 40], or been partially sidestepped by making the filament extensible through a penalty force [122, 123]. The latter approach, which has characterized a number of finite element methods for biopolymers [124–127] takes us back to an extensible chain (and its aforementioned temporal stiffness), and leaves unclear how to make sense of the continuum limit for constrained fluctuating filaments. And if there is no continuum limit, how is a spectral discretization possible?

While the continuum limit is an interesting object from the standpoint of applied stochastic analysis, for the purposes of simulation it has little relevance, since it requires an infinite number of degrees of freedom. The object that we simulate, like any other statistical mechanics model, must ultimately be discrete, and so we will propose a *fully discrete* or coarse-grained model of a fluctuating fiber based on a spectral representation of the chain (see Fig. 5.1). The important conceptual leap is then to think of the “continuum limit” only as a way to efficiently approximate the “true” hydrodynamics of $1/\epsilon$ blobs without actually needing to track that many degrees of freedom. As such, for a given value of \hat{a} and L we define our reference result to be a discrete blob-link chain with blobs spaced \hat{a} apart ($L/\hat{a} = 1/\hat{\epsilon}$ blobs). Our hope is then to approximate, to reasonable ($\approx 10\%$) accuracy, the equilibrium statistical mechanics and dynamics of the blob-link chain using a coarse-grained Chebyshev discretization with as few nodes as possible. The goal of this, which we show is achieved at least in part, is to extract the advantages of the blob-link and continuum discretizations into a single numerical method: by proposing a fully discrete “spectral” chain, we can write

a well-defined overdamped Langevin equation which samples from a well-defined *discrete* equilibrium Gibbs-Boltzmann distribution. But by using the mobility (2.63) to define the velocity at each point on the fiber, we avoid having to scale the number of collocation points with the fiber aspect ratio. In Chapter 8, we show that this approach does indeed give an effective way to simulate the Brownian hydrodynamics of a fluctuating filament, with a cost independent of the fiber aspect ratio.

Chapter 6

Numerics for a single filament

This chapter discusses our discretization of the hydrodynamics of a single bent, twisted fiber. This is a discrete version of the continuum theory we presented in Chapter 3, using the mobility (2.63) composed of integrals of the RPY kernel. In some sense, this chapter is meant to set the record straight. Like a politician who changes his opinion to match the winds of the age, our discretizations of the mobility \mathcal{M} , kinematic operators \mathcal{K} , and elastic force (density) $\mathbf{f}^{(\kappa)}$ have changed repeatedly over a few years of publications, as the problems we worked on became progressively more difficult. The reason for the multiple discretizations is that spectral methods, which help us achieve costs independent of slenderness, can be very fragile, as nonlinearities have to be handled carefully due to aliasing. In numerical methods with fluctuations, we cannot hide behind the assumption of filament smoothness, and so this chapter is the result of several years of trial and error in proper handling of nonlinearities.

6.1 Spectral spatial discretization

We will use a collocation discretization to discretize all quantities in the fluctuating filament problem. Because we use a collocation discretization, the fiber positions and tangent vectors are discretized on Chebyshev grids of nodes s_p , where $p = 1, \dots, N$ for the tangent vectors

and $p = 1, \dots, N_x$ for the positions. The Bishop and material frames (see Section 3.2.1) are then most logically defined on a grid of size N (they share the same collocation points as the tangent vectors). The twist angle ψ is defined on a grid of size N_ψ .

Our notation will shift slightly here to reflect the change from continuous to discrete. We use \mathbf{X} to refer to the $N_x \times 3$ matrix of fiber positions at the collocation points. The p th row of this matrix will be denoted by $\mathbf{X}_{\{p\}} = \mathbb{X}(s_p)$, while the p th entry of the corresponding $3N_x \times 1$ row-stacked vector will be denoted at \mathbf{X}_p . Likewise, $\boldsymbol{\tau}$ refers to the $N \times 3$ matrix of tangent vectors at the collocation points with $\boldsymbol{\tau}_{\{p\}} = \boldsymbol{\tau}(s_p)$. We will not try to distinguish between the unknown “true” fiber shape (which could have more than N Chebyshev modes) and its Chebyshev approximation $\mathbb{X}(s)$.

The tools we use for differentiation and integration are standard [128]. For differentiation, we use the Chebyshev differentiation matrix \mathbf{D}_{N_x} . By $\mathbf{D}_{N_x}\mathbf{X}$, we mean the linear operation that takes \mathbf{X} , computes the $N_x - 1$ degree Chebyshev polynomial representation $\mathbb{X}(s)$, differentiates it, and returns $\mathbb{X}'(s)$ on the same collocation grid. We also define \mathbf{D}^\dagger , the pseudo-inverse of the Chebyshev differentiation matrix, which gives the values of the indefinite integral of a function $f(s)$ modulo an unknown constant,

$$(\mathbf{D}^\dagger \mathbf{f})_p \approx \int_0^{s_p} f(s') ds' + C. \quad (6.1)$$

Importantly, on a grid of size N , this integration is *exact* when applied to polynomials of degree $N - 2$ or less (the result must be a polynomial of degree $N - 1$ or less). For definite integration, we use Clenshaw-Curtis quadrature with weight w_p associated with each collocation point,

$$\int_0^L f(s') ds' \approx \mathbf{w}^T \mathbf{f} := \sum_{p=1}^N f_p w_p. \quad (6.2)$$




Figure 6.1: Transferring the fiber discretization for the blob-link model (bottom) to a spectral grid (top). We show the tangent vector locations as magenta vectors visible through the “x-ray” inset – for the blob-link discretization these are the points exactly between the beads, while for the spectral discretization they are defined on a type 1 Chebyshev grid with N points ($N = 12$ here). The positions, which can be obtained from τ and \mathbf{X}_{MP} (red sphere), are shown as yellow spheres – for the blob-link discretization these are evenly spaced points, while for the spectral discretization they are defined on a type 2 Chebyshev grid with $N_x = N + 1$ points. Note the clustering of both τ and \mathbf{X} near the fiber endpoints in the spectral model, which is a characteristic of the Chebyshev grid. This means that the tangent vectors no longer connect individual beads, as they do for the blob-link discretization.

6.1.1 Inextensibility constraint

To properly track all of the modes of the fiber dynamics, we discretize an inextensible filament by transferring the more common blob-link discretization concept onto a non-uniform Chebyshev grid (see Fig. 6.1). We begin with a collection of N unit-length tangent vectors $\boldsymbol{\tau} = \{\boldsymbol{\tau}_{\{p\}}\}_{p=1}^N$, with $\boldsymbol{\tau}_{\{p\}} \cdot \boldsymbol{\tau}_{\{p\}} = 1$ for $p = 1, \dots, N$, on a type 1 Chebyshev grid (i.e., a Chebyshev grid that does not include the endpoints). To obtain the fiber position, we over-sample $\boldsymbol{\tau}$ to a (type 2, endpoints-included) Chebyshev grid of size $N + 1$ using the evaluation matrix $\mathbf{E}_{N \rightarrow N+1}$, then integrate the result exactly using the Chebyshev integration matrix \mathbf{D}_{N+1}^\dagger (pseudo-inverse of the differentiation matrix \mathbf{D}_{N+1}). This can be written as

$$\mathbf{X} = \left(\mathbf{D}_{N+1}^\dagger \mathbf{E}_{N \rightarrow N+1} \quad \mathbf{B} \right) \begin{pmatrix} \boldsymbol{\tau} \\ \mathbf{X}_{\text{MP}} \end{pmatrix} := \boldsymbol{\mathcal{X}} \bar{\boldsymbol{\tau}}. \quad (6.3)$$

The matrix \mathbf{B} is such that the midpoint of \mathbf{X} on the grid of size $N + 1$ is \mathbf{X}_{MP} (when N is even, \mathbf{X}_{MP} is an actual point on the $N + 1$ size Chebyshev grid, but N could also be odd, for which the midpoint is obtained via interpolation). Thus, the nodal points that track the fiber position \mathbf{X} are defined on a grid of size $N + 1$, whereas $\boldsymbol{\tau}$ is defined on a grid of size N , just as in the blob-link discretization. To go in the reverse direction, we apply the inverse of $\boldsymbol{\mathcal{X}}$, which is done by differentiating \mathbf{X} on the $N + 1$ point grid and then downsampling to the grid of size N via the matrix $\mathbf{E}_{N+1 \rightarrow N}$. The midpoint is determined from the $N + 1$ nodes via sampling the $N + 1$ -term interpolating polynomial at the midpoint using the matrix $\mathbf{E}_{N+1 \rightarrow \text{MP}}$. Together this gives,

$$\boldsymbol{\mathcal{X}}^{-1} = \begin{pmatrix} \mathbf{E}_{N+1 \rightarrow N} \mathbf{D}_{N+1} \\ \mathbf{E}_{N+1 \rightarrow \text{MP}} \end{pmatrix}, \quad (6.4)$$

which is the actual inverse of $\boldsymbol{\mathcal{X}}$ because we carefully handle the integration of N -term

Chebyshev polynomials by using a grid of size $N + 1$. We contrast this with handling everything on a grid of size N , in which case information is lost when converting the unit-length tangent vectors to the positions \mathbf{X} . For the deterministic examples we study here, the fibers are relatively smooth, and so the lost (high-frequency) information has a negligible effect. But when we extend to Brownian filaments in Chapter 8, the method here is required to correctly track the high-frequency modes. It is important to note, however, that this formulation only applies to open, two-ended, filaments, and *not* looped filaments, which require further (continuity) conditions on τ .

While we motivated (6.3) using a spectral discretization, these equations also hold for the standard blob-link discretization [47]; in that case the map \mathcal{X}^{-1} is defined by taking finite differences of nodal points to give values at the links, and the product \mathcal{X} is defined by summing the values of the links to obtain values at the nodes. Thus, the equations we write in this section are general once the linear maps \mathcal{X} and \mathcal{X}^{-1} have been specified.

6.1.2 Spectral fibers from blob-link fibers

Throughout this part of the dissertation, and especially in Chapter 8, our goal will always be for the spectral method to give an inexpensive approximation to the blob-link model, where the fiber is discretized by a series of blobs spaced ϵ apart. We quantify this by computing static and dynamic statistics for both models, and checking that they are the same. A worthwhile question to ask before doing any of this is the following: supposing we have a blob-link discretization of the fiber, how can we obtain the “best” spectral approximation? Because the spectral discretization has far fewer degrees of freedom than the blob-link model, this is fundamentally a coarse-graining (CG) operation. This section contains a few ideas about how to coarse grain blob-link fibers into spectral ones. Throughout the section, we denote $\mathbf{X}^{(B)}$ as the configuration for the blob-link model, with \mathbf{X} continuing to stand for the positions of the Chebyshev nodes in the spectral model.

For a given persistence length $\ell_p = \kappa/k_B T$, where κ is the bending stiffness and $k_B T$ is the thermal energy, we have an exact sampler for a blob-link chain, so that we can generate a blob-link configuration $\mathbf{X}^{(B)}$ on a grid with spacing $\Delta s = \hat{a}$. Our goal is then to construct a “coarse-grained” spectral discretization, which is (usually) defined by the set of tangent vectors $\boldsymbol{\tau}_{\{p\}}$ that are of unit length on a type 1 Chebyshev grid with N nodes. To obtain a set of nodal positions \mathbf{X} on a type 2 Chebyshev grid of $N_x = N + 1$ nodes, we apply the map $\boldsymbol{\mathcal{X}}$ in (6.3). A well-defined operation is then to use the polynomial interpolant $\mathbb{X}(s)$ that is defined by the points \mathbf{X} to get an approximation on the blob-link grid. We define this interpolation via $\mathbf{X}^{(\text{SB})} = \mathbf{E}_{S \rightarrow B} \mathbf{X}$. We can then compute the (dimensionless) discrete L^2 norm

$$\frac{1}{L} \left\| \mathbf{X}^{(\text{SB})} - \mathbf{X}^{(B)} \right\|_2 = \frac{1}{L} \sqrt{\sum_{k=1}^{N_B} \left\| \mathbf{X}_{\{p\}}^{(\text{SB})} - \mathbf{X}_{\{p\}}^{(B)} \right\|^2 \frac{\Delta s}{L}} \quad (6.5)$$

and use this as a metric for the error between the two discretizations. We now outline a few different approaches for doing so.

6.1.2.1 Minimization problem

The first CG approach is to solve the constrained minimization problem

$$\begin{aligned} \begin{pmatrix} \boldsymbol{\tau} \\ \mathbf{X}_{\text{MP}} \end{pmatrix} &= \underset{\boldsymbol{\tau}, \mathbf{X}_{\text{MP}}}{\text{argmin}} \left\| \mathbf{X}^{(\text{SB})} - \mathbf{X}^{(B)} \right\|_2^2 = \left\| \mathbf{E}_{S \rightarrow B} \boldsymbol{\mathcal{X}} \begin{pmatrix} \boldsymbol{\tau} \\ \mathbf{X}_{\text{MP}} \end{pmatrix} - \mathbf{X}^{(B)} \right\|_2^2 \\ \boldsymbol{\tau}_{\{p\}}^T \boldsymbol{\tau}_{\{p\}} &= 1 \quad \text{for } p = 1, \dots, N \end{aligned} \quad (6.6)$$

for $(\boldsymbol{\tau}, \mathbf{X}_{\text{MP}})$, using built-in methods (`fmincon`) in Matlab. This computes the configuration with unit-length tangent vectors that minimizes the error (6.5). In this sense, it is the optimal configuration.




Figure 6.2: Comparing the minimization problem with the ad-hoc approach for a fiber of persistence length $\ell_p/L = 1/2$. For each N , the solution of the minimization problem is shown as a solid line, while the ad-hoc approach is shown as a dotted line. The blob-link fiber $\mathbf{X}^{(B)}$ is shown as a solid black line.

6.1.2.2 Ad-hoc approach

The approach in the previous section is quite expensive, since it requires a constrained quadratic programming problem at every time step. A more *ad hoc* approach is to just sample the blob-link chain at the type-2 Chebyshev node locations using linear interpolation. Then apply the map \mathcal{X}^{-1} to obtain candidate tangent vectors (which will not have norm one). We then normalize those tangent vectors, and apply the map \mathcal{X} again to obtain $\mathbf{X}^{(SB)}$. Figure 6.2 shows a visual comparison of these two approaches for various N with a fiber of persistence length $\ell_p/L = 1/2$ (bottom). We can only see a visual difference between the minimization and ad-hoc approaches when $N = 8$. Even for $N = 8$, the ad-hoc approach generates a reasonable fiber shape.

6.1.2.3 BVP approach

In the last approach, which is inspired by [129, Eq. (7)], we use a smoothing strategy to generate $\mathbf{X}^{(\text{SB})}$. Beginning with the blob-link configuration $\mathbf{X}^{(B)}$, we fit a natural spline $\mathbb{X}^{(\text{spl})}$ through it, which we sample at the N_x Chebyshev points. We use these sampled locations to solve the BVP

$$-b^2 \partial_s^2 \mathbb{X} + \mathbb{X} = \mathbb{X}^{(\text{spl})} \quad (6.7)$$

by enforcing this equation at nodes $2, \dots, N_x - 1$ and the BCs $\mathbb{X}^{(\text{spl})}(0, L) = \mathbb{X}(0, L)$ at the endpoints. Unlike our previous two approaches, this approach does not enforce $\|\partial_s \mathbb{X}\| = 1$ anywhere, but assumes the tangent vectors will have nearly norm one because they come from an underlying blob-link discretization.

In Figs. 6.3 and 6.4, we show samples of applying this procedure to fibers with $\ell_p = L = 1$ and $\ell_p/L = 1/2$, respectively. In both plots, the left plot shows fiber samples, while the right column shows the signed error in the norm of the tangent vectors on the type 1 grid of size N (where we are used to constraining them). We see that the fiber gets smoother and smoother as b gets larger, and that smaller N is needed to resolve the shape for larger b . As b gets smaller, the tangent vectors approach norm 1. We notice that they typically (but not always) approach 1 from below, so that for finite b the tangent vectors always have norm less than 1.

6.1.2.4 Choosing the value of b

We want the parameter b to be a proxy for the error (6.5). Thus, to choose b , we first perform an experiment where we vary b , L , and ℓ_p , and look at the dimensionless L^2 error given in (6.5). For small b , there isn't any smoothing at all, and the error is from approximating the curve by a polynomial of finite degree (we choose $N = 36$ so that this is small). Large b just gives a straight fiber, and the error plateaus to the value for a straight filament connecting the two endpoints. We are interested in the error in the intermediate regime (between “large”

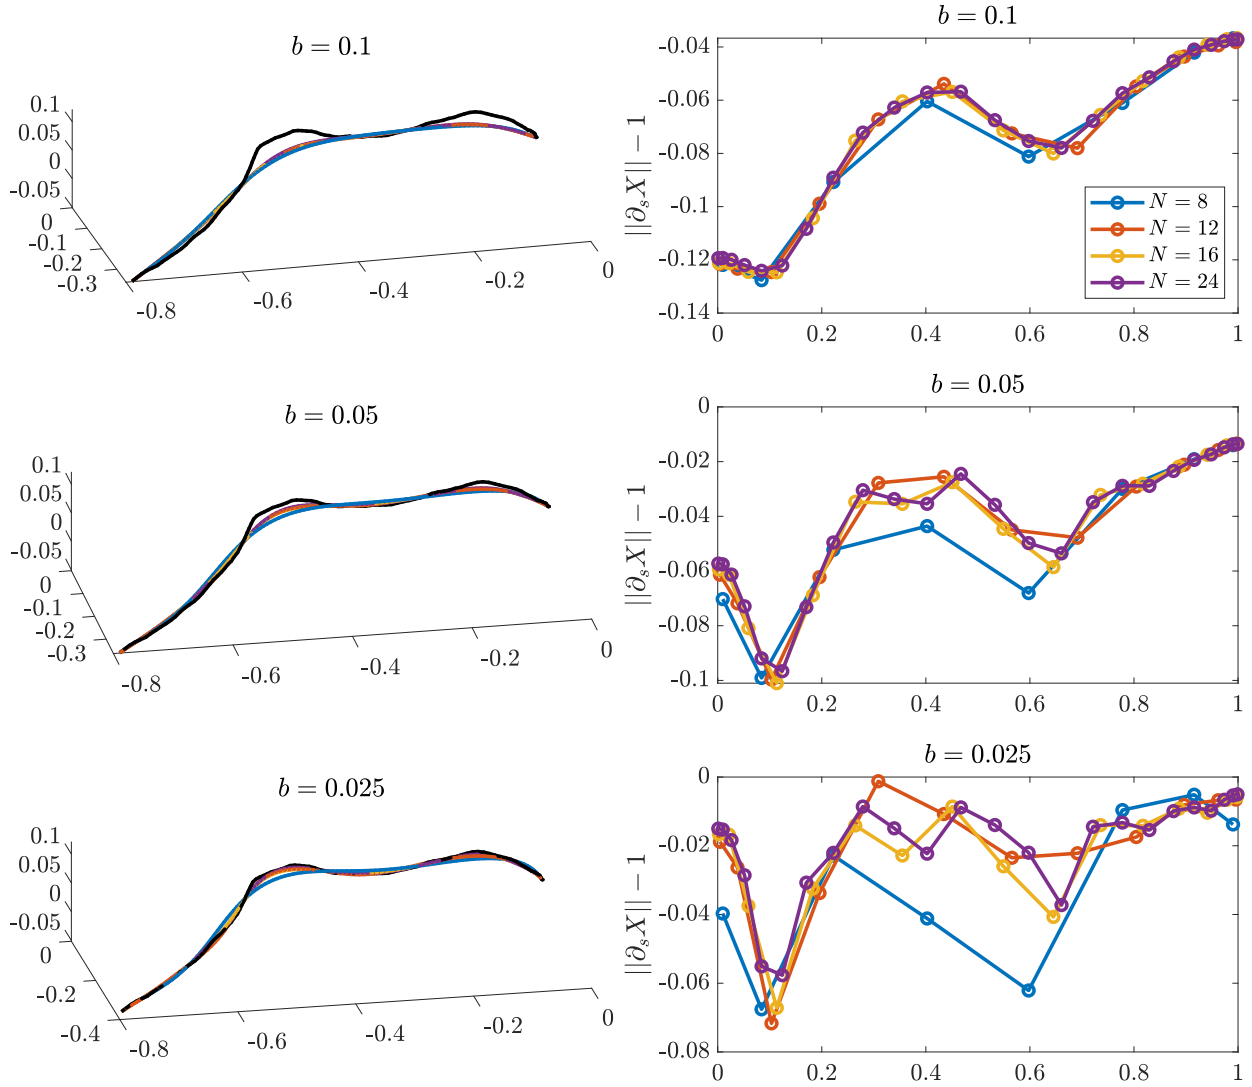


Figure 6.3: Coarse-graining a blob-link fiber (black) with $\ell_p/L = 1$ using the BVP (6.7). We show a sample on the left, and the norm of the tangent vectors ($\|\partial_s \mathbf{X}\| - 1$) on the type 1 Chebyshev grid at the right.

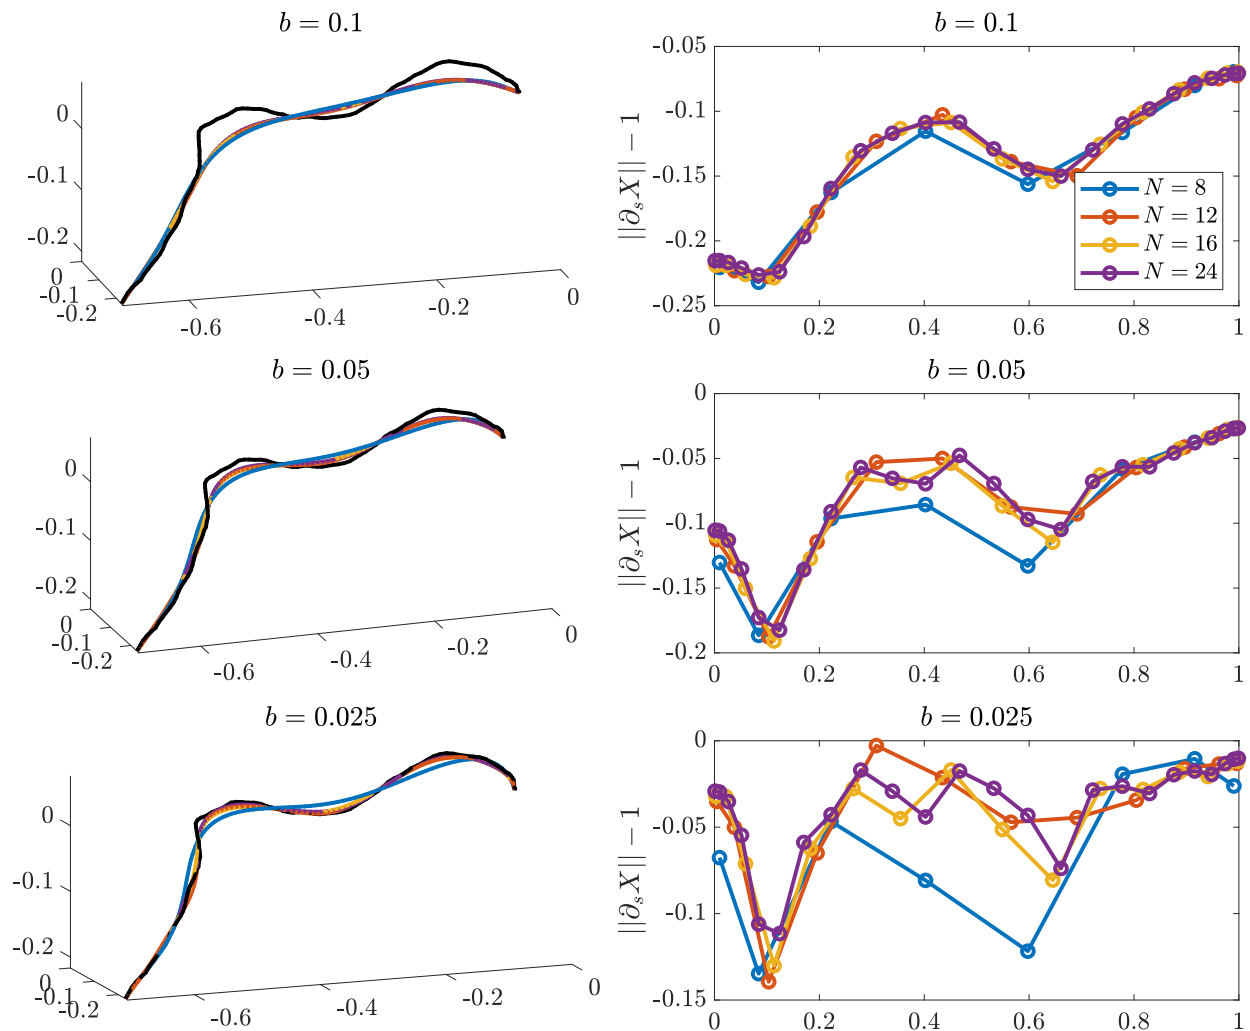


Figure 6.4: Coarse-graining a blob-link fiber (black) with $\ell_p/L = 1/2$ using the BVP (6.7). We show a sample on the left, and the norm of the tangent vectors ($\|\partial_s \mathbf{X}\| - 1$) on the type 1 Chebyshev grid at the right.




Figure 6.5: The L^2 error (6.5) when using the BVP procedure of this section with various b , ℓ_p and L . Using the rescaling $\bar{b} = b/\sqrt{L\ell_p}$ collapses curves with the same ℓ_p/L (blue and purple, red and cyan, yellow and green) onto the same curve. There is then a weak dependence on ℓ_p/L , another dimensionless parameter.

and “small” b).

Figure 6.5 shows the data for the error as a function of $\bar{b} = b/\sqrt{L\ell_p}$. When we use this rescaling, the error curves for the same ratio ℓ_p/L are on top of each other, which tells us we have removed the dimensions from the problem and have $\text{Error} = f(\bar{b}, \ell_p/L)$. The dependence on ℓ_p/L is quite weak in the area we consider, and we see that the error is roughly given by $\text{Error} = 0.5\bar{b}^{4/3}$, although the leading constant is subject to some variation for the different values of ℓ_p/L .

6.1.2.5 Least squares approximation ($b = 0$)

If we do not enforce inextensibility of the tangent vectors, another obvious approach is to find the Chebyshev polynomial that gives the best approximation to the blob-link chain. This is a continuum-style approach to solving (6.7) with $b = 0$, and amounts to solving $\mathbf{E}_{S \rightarrow B} \mathbf{X} - \mathbf{X}^{(B)}$ in the least squares sense. We do this by setting

$$\mathbf{X} = (\mathbf{E}_{S \rightarrow B}^T \mathbf{W}_B \mathbf{E}_{S \rightarrow B})^{-1} \mathbf{E}_{S \rightarrow B}^T \mathbf{W}_B \mathbf{X}^{(B)}, \quad (6.8)$$

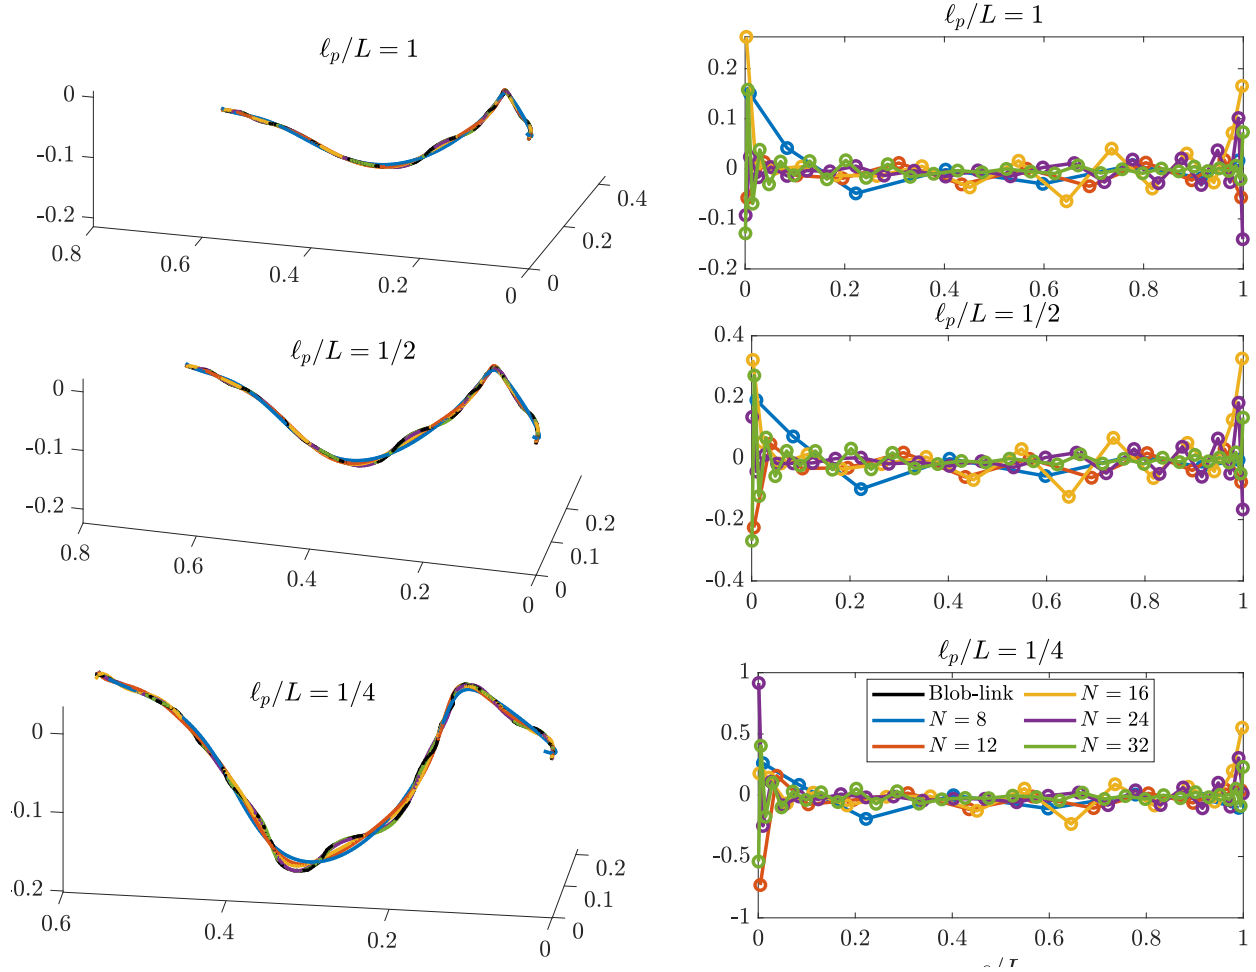


Figure 6.6: Fiber approximations using (6.8) as the least-squares fit. We show the fits and the error in the tangent vector norms for different ℓ_p/L .

where $\mathbf{W}_B = \text{diag}(1/2, 1, \dots, 1, 1/2)$ is the matrix of L^2 weights on the blob-link grid. Figure 6.6 shows the results in this case. Comparing to our previous approaches, we notice that the polynomials, especially $N = 8$, appear to be better approximations to the blob-link chain. However, when we look at the norm of the tangent vectors, we see disastrous effects at the boundary. This boundary layer comes from the fact that $b = 0$ is a singular limit of (6.7), since $b = 0$ converts (6.7) from a differential equation (with BCs) to an algebraic one. Thus this approach is *not* to be used, since the tangent vectors it generates are not smooth and have large errors near the boundary.

6.1.3 Discretizing the internal fiber force and torque (density)

Now that we have defined a spectral discretization of the fiber centerline, we can go about computing a force/torque and force/torque density associated with bending and twisting. This section will again present two ways of doing so: in the first case, which we use for free-ended filaments which resist bending, we begin with an energy functional, which we discretize using the discrete L^2 inner product. We then compute the force and force density from this energy, which enforces the natural free-ended boundary conditions. In the second case, which can be more generally adapted to different boundary conditions, our approach is to work with the continuum force density (e.g., $\mathbf{f}^{(\kappa)} = -\kappa \partial_s^4 \mathbf{X}$), which we discretize using rectangular spectral collocation [130, 131]. The former discretization is essential for fluctuating filaments; since the force comes from an energy, the correct Gibbs-Boltzmann distribution is automatically implied (see Chapter 8). But the latter discretization is useful when we compute twist elasticity and have different boundary conditions, so it is worth presenting here.

6.1.3.1 A discrete energy-based approach

In continuum, the fiber bending or curvature energy is given by the squared L^2 norm of the fiber curvature vector,

$$\mathcal{E}_{\text{bend}}[\mathbb{X}(\cdot)] = \frac{\kappa}{2} \int_0^L \partial_s^2 \mathbb{X}(s) \cdot \partial_s^2 \mathbb{X}(s) ds. \quad (6.9)$$

It is useful to discretize the energy functional (6.9) directly, then use the matrix that results to compute forces, implementing the boundary conditions naturally. On a spectral grid, we can put the energy functional (6.9) in the form $\mathcal{E}_{\text{bend}} = \frac{1}{2} \mathbf{X}^T \mathbf{L} \mathbf{X}$, where

$$\mathbf{L} = \kappa \left(\mathbf{E}_{N_x \rightarrow 2N_x} \mathbf{D}^2 \right)^T \mathbf{W}_{2N_x} \left(\mathbf{E}_{N_x \rightarrow 2N_x} \mathbf{D}^2 \right) := \kappa \left(\mathbf{D}^2 \right)^T \widetilde{\mathbf{W}} \mathbf{D}^2. \quad (6.10)$$

is a matrix that takes two derivatives of \mathbf{X} on a grid of size N_x , then does an inner product of those derivatives by upsampling to a $2N_x$ grid with extension matrix \mathbf{E} and Clenshaw-Curtis weights matrix \mathbf{W}_{2N_x} . The reason for this upsampling is to compute the inner product of the Chebyshev polynomials representing $\mathbf{D}^2\mathbf{X}$ *exactly* (this could be done on a slightly smaller grid because the two derivatives make the polynomial representation $\mathbf{D}^2\mathbf{X}$ have degree at most $N_x - 3$). The L^2 inner product on the upsampled grid corresponds to an inner product on \mathbb{R}^{N_x} with the weights matrix

$$\widetilde{\mathbf{W}} = \mathbf{E}_{N_x \rightarrow 2N_x}^T \mathbf{W}_{2N_x} \mathbf{E}_{N_x \rightarrow 2N_x}. \quad (6.11)$$

Note that an equivalent way to formulate the energy is to use $\bar{\boldsymbol{\tau}}$ as the degrees of freedom, so that

$$\mathcal{E}_{\text{bend}}(\bar{\boldsymbol{\tau}}) = \frac{1}{2} \bar{\boldsymbol{\tau}}^T \boldsymbol{\mathcal{X}}^T \mathbf{L} \boldsymbol{\mathcal{X}} \bar{\boldsymbol{\tau}} := \frac{1}{2} \bar{\boldsymbol{\tau}}^T \mathbf{L}_\tau \bar{\boldsymbol{\tau}}. \quad (6.12)$$

The *force* on the fiber nodes is then obtained by differentiating the discrete energy

$$\mathbf{F} = -\frac{\partial \mathcal{E}_{\text{bend}}}{\partial \mathbf{X}} = -\mathbf{L}\mathbf{X} = -\mathbf{L}\boldsymbol{\mathcal{X}}\bar{\boldsymbol{\tau}}. \quad (6.13)$$

6.1.3.2 The map between force and force density

When we compute the fiber hydrodynamics, we will need to input the force *density* on the centerline. We can obtain this at the nodes from the force in (6.13) by

$$\mathbf{f} = \widetilde{\mathbf{W}}^{-1} \mathbf{F}, \quad (6.14)$$

and use the Chebyshev interpolant $\mathbf{f}(s)$ as a continuum force density. Note that the matrix $\widetilde{\mathbf{W}}$, rather than the diagonal matrix \mathbf{W} of weights on the N_x point Chebyshev grid, must be used for the force density to converge as the spatial discretization is refined. The reason for this is that the weights matrix (6.11) enters in the discrete L^2 inner product, and the

force density is the representation of the derivative of energy with respect to that inner product [132, Sec. 6], i.e., the function that satisfies $\langle \mathbf{X}, \mathbf{f} \rangle = \mathbf{X}^T \mathbf{F}$. As a more concrete justification of this, suppose that the filament satisfied the free fiber boundary conditions $\partial_s^2 \mathbb{X}(0, L) = \partial_s^3 \mathbb{X}(0, L) = \mathbf{0}$. Then in continuum we could equivalently write the energy as $U = \kappa/2 \int_0^L \mathbb{X}(s) \cdot \partial_s^4 \mathbb{X}(s) ds$, which we would discretize as $U = \kappa/2 \mathbf{X}^T \widetilde{\mathbf{W}} \mathbf{D}^4 \mathbf{X}$. Taking the derivative, we have the force $\mathbf{F} = -\kappa \widetilde{\mathbf{W}} \mathbf{D}^4 \mathbf{X}$. But we know that force density should be $\mathbf{f} = -\kappa \mathbf{D}^4 \mathbf{X}$, which can only be obtained from \mathbf{F} using (6.14)

6.1.3.3 A continuum approach

As an alternative approach that works for other boundary conditions, let us now look at how we can use rectangular spectral collocation [130, 131] to discretize the bending and twisting force and torque *densities* to high accuracy while respecting the boundary conditions (3.47). In this discretization, it is vital that the matrix \mathbf{X} gives the positions of the fiber on an N_x point type 1 Chebyshev grid that does *not* include the boundaries. Whenever we use this discretization of the force, we therefore switch the \mathbf{X} nodes to a type 1 grid.

The idea of rectangular spectral collocation as proposed in [130, 131] is to compute an upsampled representation $\widetilde{\mathbf{X}}$ of \mathbf{X} on a type 2 (including the endpoints) Chebyshev grid that satisfies the boundary conditions exactly, and gives \mathbf{X} when downsampled to the original (type-1) Chebyshev grid. Since there are four boundary conditions, the upsampled representation is on a $N + 4$ point type 2 Chebyshev grid that includes the endpoints. In this section, we will assume that the $s = 0$ end is clamped, i.e., the boundary conditions at $s = 0$ are

$$\mathbb{X}(s = 0) = \mathbf{X}_0 \quad \boldsymbol{\tau}(s = 0) = \boldsymbol{\tau}_0 \quad \Omega^{\parallel}(0) = \omega, \quad (6.15)$$

while at $s = L$ the end is free with boundary conditions

$$\partial_s^2 \mathbb{X}(s = 0) = \mathbf{0} \quad \partial_s^3 \mathbb{X}(s = 0) = \mathbf{0} \quad \psi(L) = 0. \quad (6.16)$$

(recall that ψ represents the twist density). These equations will show how to treat one free end and one clamped end, and are straightforward to modify when we have two free ends.

Beginning with the bending force, the unique configuration $\widetilde{\mathbf{X}}$ on the $\tilde{N} = N_x + 4$ point grid can be obtained by solving

$$\begin{pmatrix} \mathbf{R}_X \\ \mathbf{B} \end{pmatrix} \widetilde{\mathbf{X}} =: \begin{pmatrix} \mathbf{R}_X \\ \mathbf{I}(0) \\ \mathbf{D}_{N_x+4}(0) \\ \mathbf{D}_{N_x+4}^2(L) \\ \mathbf{D}_{N_x+4}^3(L) \end{pmatrix} \widetilde{\mathbf{X}} = \begin{pmatrix} \mathbf{X} \\ \mathbf{X}_0 \\ \boldsymbol{\tau}_0 \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix} := \begin{pmatrix} \mathbf{X} \\ \mathbf{b}_X \end{pmatrix}, \quad (6.17)$$

where \mathbf{R}_X is the downsampling matrix that evaluates the interpolant $\widetilde{\mathbf{X}}(s)$ on the original N_x point grid, $\mathbf{I}(0)$ encodes evaluation at $s = 0$, and $\mathbf{D}_M^n(s)$ represents evaluation of the n th derivative at the point s on a grid of size M . In (6.17), \mathbf{B} is the $4 \times (N_x + 4)$ matrix that encodes the boundary conditions and \mathbf{b}_X is the values of the BCs. We now write the affine transformation from \mathbf{X} to $\widetilde{\mathbf{X}}$ as

$$\widetilde{\mathbf{X}} = \begin{pmatrix} \mathbf{R}_X \\ \mathbf{B} \end{pmatrix}^{-1} \left(\begin{pmatrix} \mathbf{X} \\ \mathbf{0} \end{pmatrix} + \begin{pmatrix} \mathbf{0} \\ \mathbf{b} \end{pmatrix} \right) := \mathbf{E}_X \mathbf{X} + \boldsymbol{\beta}_X. \quad (6.18)$$

Note that $\boldsymbol{\beta}_X$ is independent of time in this case.

We then use the configuration $\widetilde{\mathbf{X}}$ to compute any quantities with more than one s derivative in the Euler equations (4.2). Thus the elastic force density is computed on the upsampled grid as $\widetilde{\mathbf{f}}^\kappa = -\kappa \mathbf{D}_{N_x+4}^4 \widetilde{\mathbf{X}}$ and is then downsampled to the original N_x point type 1 grid to give the bending elasticity discretization

$$\mathbf{f}^\kappa(\mathbf{X}) = \mathbf{R}_X \widetilde{\mathbf{f}}^\kappa = -\kappa \mathbf{R}_X \mathbf{D}_{N_x+4}^4 (\mathbf{E}_X \mathbf{X} + \boldsymbol{\beta}_X) := \mathbf{F}_\kappa \mathbf{X} + \boldsymbol{\beta}_\kappa \quad (6.19)$$

In all cases in which derivatives of $\boldsymbol{\tau}$ enter in the forcing in (3.59), we compute the derivatives on the $N + x + 4$ point grid from the configuration $\widetilde{\mathbf{X}}$, then downsample them to evaluate the derivatives on the original N point grid, e.g., we use $\mathbf{R}_{\mathbf{X}}\mathbf{D}_{N+4}^2\widetilde{\mathbf{X}}$ as the discretization of $\partial_s\boldsymbol{\tau}$. This helps by smoothing the higher order derivatives due to the imposition of smoothing BCs at the endpoints. Note however that $\widetilde{\boldsymbol{\tau}} = \partial_s\widetilde{\mathbf{X}}$ is not of unit length at the collocation points, so we use $\boldsymbol{\tau}$ on the original N point grid instead when a vector of unit length is required (in the kinematic matrix \mathbf{K}).

6.1.3.4 A continuum approach for twist

Let us now consider using rectangular spectral collocation to solve for the twist density ψ . Recall that the degrees of rotation $\theta(s)$ are defined on the same grid as the tangent vectors – that is, a grid of size N . Thus, we compute $\psi = \partial_s\theta$ on a type 1 Chebyshev grid of size $N_\psi = N - 1$ (one less point is needed). In exactly the same manner as we did for \mathbf{X} , we define an affine operator which gives an upsampled representation $\tilde{\psi}$ on an $N_\psi + 2$ point type 2 Chebyshev grid that satisfies the boundary conditions (6.15) and (6.16),

$$\begin{pmatrix} \mathbf{R}_\psi \\ \tilde{\Omega}_n^{\parallel}(0) \\ \mathbf{I}(L) \end{pmatrix} \tilde{\psi} = \begin{pmatrix} \psi \\ \omega - \Omega_f^{\parallel}(0) \\ 0 \end{pmatrix} = \begin{pmatrix} \psi \\ \mathbf{b}_\psi \end{pmatrix}, \quad (6.20)$$

$$\text{where } \tilde{\Omega}_n^{\parallel}\tilde{\psi} = \mathbf{M}_{\text{rr}}\gamma\mathbf{D}_{N_\psi+2}\tilde{\psi}, \quad (6.21)$$

is the angular velocity due to torque on the $N_\psi + 2$ point grid (throughout this chapter, the matrix \mathbf{M}_{rr} is a diagonal matrix on the $N_\psi + 2$ point grid). The analogue of (6.18) in this case is the affine operator that gives $\tilde{\psi}$ from ψ on the N_ψ point grid, which we define as $\tilde{\psi} = \mathbf{E}_\psi\psi + \boldsymbol{\beta}_\psi$.

In (6.20), we use the angular velocity at $s = 0$ as the boundary condition for the clamped

end. Another possibility is to prescribe the parallel torque at the clamped end,

$$\mathbf{N}(s=0) \cdot \boldsymbol{\tau}(0) = \gamma\psi(0) = N_0 \quad (6.22)$$

which gives a boundary condition $\psi(0) = N_0/\gamma$ that is easy to handle; we therefore focus on the more challenging (6.15). See [52, Sec II(A)] for a discussion of the difference between the constant torque vs. constant angular velocity boundary conditions.

To calculate the torque on an N_ψ point grid, we take derivatives of $\tilde{\psi}$ on the $N+2$ point grid, and evaluate the result on the N_ψ point grid using \mathbf{R}_ψ ,

$$n^\parallel(\psi) = \gamma \mathbf{R}_\psi \mathbf{D}_{N+2} (\mathbf{E}_\psi \psi + \boldsymbol{\beta}_\psi) := \mathbf{R}_\psi \left(\widetilde{\mathbf{N}}_\psi \psi + \widetilde{\boldsymbol{\beta}}_n \right), \quad (6.23)$$

where $\widetilde{\boldsymbol{\beta}}_n$ is time-dependent in general. The twist force $\mathbf{f}^{(\gamma)}$ as defined in (4.2) is slightly more complex, given that it is a nonlinear function of \mathbf{X} and involves interaction between \mathbf{X} and $\psi = \partial_s \theta$. To reduce aliasing on the spectral grid, we use the product rule to separate $\mathbf{f}^{(\gamma)}$ into two terms, then use a common upsampled grid to compute each term with proper boundary conditions,

$$\mathbf{f}^{(\gamma)}(\mathbf{X}, \psi) = \left(\gamma \left(\partial_s \tilde{\psi} \left(\partial_s \widetilde{\mathbf{X}} \times \partial_s^2 \widetilde{\mathbf{X}} \right) + \tilde{\psi} \left(\partial_s \widetilde{\mathbf{X}} \times \partial_s^3 \widetilde{\mathbf{X}} \right) \right) \right) \quad (6.24)$$

We perform the multiplications of ψ and derivatives of \mathbf{X} on the upsampled grid, then downsample the result to the N_x point grid to obtain the final force density $\mathbf{f}^{(\gamma)}$ at the collocation points. For the upsampling factor, we find $N_x + 4$ to give identical results to $2N_x$, so we will use $N_x + 4$ for the upsampled grid size, which means we do not need any upsampling of $\widetilde{\mathbf{X}}$. Additionally, when we compute the trans-rot velocity $\mathbf{M}_{\text{tr}} n^\parallel$, we upsample n^\parallel in (6.23) from the N_ψ point grid to the grid of size N_x .

6.1.4 Discretization of kinematics

We now discuss the discretization of the kinematic operator \mathcal{K} . The discretization we present here, which is the only one that should be used, is based on the “strong form” of \mathcal{K} in Section 3.1.2.1. Other discretizations based on the weak form in Section 3.1.2.3 are subject to strong aliasing errors and consequently should not be used.

Let us first consider the evolution of the fiber tangent vectors. Since, for any p , $\boldsymbol{\tau}_{\{p\}} \cdot \boldsymbol{\tau}_{\{p\}} = 1$ for all time, it follows that the evolution of the tangent vectors can be described by (3.5), which we denote in matrix form as

$$\partial_t \boldsymbol{\tau} = \boldsymbol{\Omega} \times \boldsymbol{\tau} := -\mathbf{C}[\boldsymbol{\tau}] \boldsymbol{\Omega}, \quad (6.25)$$

where the matrix $\mathbf{C}[\boldsymbol{\tau}]$ is such that $\mathbf{C}[\boldsymbol{\tau}] \boldsymbol{\Omega} = \boldsymbol{\tau} \times \boldsymbol{\Omega}$; we will drop the explicit $\boldsymbol{\tau}$ dependence when clear from context. The matrix \mathbf{C} satisfies the following properties, which will be useful when we formulate the overdamped Langevin equation for the fiber evolution in Chapter 8:

$$\mathbf{C} \boldsymbol{\tau} = -\mathbf{C}^T \boldsymbol{\tau} = \mathbf{0}, \quad \partial_{\boldsymbol{\tau}} \cdot \mathbf{C}^T = \mathbf{0}. \quad (6.26)$$

The first equation (6.26) follows from the definition of \mathbf{C} as a cross product with $\boldsymbol{\tau}$, while the second, which is a divergence with $\boldsymbol{\tau}$, has j th entry

$$(\partial_{\boldsymbol{\tau}} \cdot \mathbf{C}^T)_j := \frac{\partial}{\partial \tau_k} \mathbf{C}_{jk}^T = \frac{\partial}{\partial \tau_k} \mathbf{C}_{kj} = 0 \quad (6.27)$$

because the k th row of \mathbf{C} has no entries that depend on τ_k ; here and throughout this dissertation, repeated indices are summed over using Einstein’s convention. Based on these two properties, we can conclude that the time evolution of the tangent vector (6.25) is analogous to the time evolution of the unit quaternion describing a rigid body’s orientation [116, Eqs. (5–8)], which can be represented as a unit vector on the unit 4 sphere.

The evolution of the tangent vectors according to (6.25) automatically implies that the evolution of the positions \mathbf{X} is given by

$$\partial_t \mathbf{X} = \mathcal{X} \begin{pmatrix} -\mathbf{C} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \boldsymbol{\Omega} \\ \mathbf{U}_{\text{MP}} \end{pmatrix} := \mathcal{X} \bar{\mathbf{C}} \boldsymbol{\alpha} := \mathbf{K} \boldsymbol{\alpha}. \quad (6.28)$$

Note the analogy with (6.3), but here we transform *velocity* instead of position using a square $3(N+1) \times 3(N+1)$ non-invertible matrix \mathbf{K} . We define a pseudo-inverse of \mathbf{K} as

$$\mathbf{K}^{-1} = \begin{pmatrix} \mathbf{C} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \boldsymbol{\alpha}^{-1} = \bar{\mathbf{C}}^T \boldsymbol{\alpha}^{-1}. \quad (6.29)$$

Note that the matrices \mathbf{K} and \mathbf{K}^{-1} have rank $2N+3$, since the N tangent vectors $\boldsymbol{\tau}$ live in the null space¹ of \mathbf{C} . Thus \mathbf{K}^{-1} is not a true inverse of \mathbf{K} , but rather

$$\mathbf{K}^{-1} \mathbf{K} = \begin{pmatrix} -\mathbf{C}^2 & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \rightarrow \bar{\mathbf{C}} \mathbf{K}^{-1} \mathbf{K} = \bar{\mathbf{C}}, \quad (6.30)$$

i.e., $\mathbf{K}^{-1} \mathbf{K}$ acts like the identity when applied to $\bar{\mathbf{C}}$ from the right, since the matrix \mathbf{C}^2 is block diagonal with p th block $\mathbf{I}_3 - \boldsymbol{\tau}_{\{p\}} \boldsymbol{\tau}_{\{p\}}$, and when multiplied by \mathbf{C} the term involving $\boldsymbol{\tau}_{\{p\}}$ becomes zero. What (6.30) also means in practice is that

$$\mathbf{K}^{-1} \mathbf{K} \boldsymbol{\alpha} = \mathbf{K}^{-1} \mathbf{K} \begin{pmatrix} \boldsymbol{\Omega} \\ \mathbf{U}_{\text{MP}} \end{pmatrix} = \begin{pmatrix} \boldsymbol{\Omega}^\perp \\ \mathbf{U}_{\text{MP}} \end{pmatrix}, \quad (6.31)$$

where $\boldsymbol{\Omega}_{\{p\}}^\perp = (\mathbf{I} - \boldsymbol{\tau}_{\{p\}} \boldsymbol{\tau}_{\{p\}}) \boldsymbol{\Omega}_{\{p\}}$. Thus, because the fiber evolves according to (6.25), it

¹This is true even for straight fibers. In continuum, the only nontrivial solutions of $\left(\int_{L/2}^s \boldsymbol{\Omega}(s') ds'\right) \times \boldsymbol{\tau} + \mathbf{U}_{\text{MP}} = \mathbf{0}$ are those with $\boldsymbol{\Omega} = \boldsymbol{\Omega}^\parallel \boldsymbol{\tau}$. Our discretization preserves this property; since in our method $\int \boldsymbol{\Omega}(s) ds$ (and, for non-straight fibers $\int \boldsymbol{\Omega}(s) \times \boldsymbol{\tau}(s) ds$) is a polynomial with $N+1$ terms (degree N), it can only be exactly equal to $-\mathbf{U}_{\text{MP}}$ at N points (and not $N_x = N+1$). This shows why careful handling of high-frequency modes (on a grid of size $N+1$) is necessary.

makes no difference whether we use $\boldsymbol{\alpha}$ or $\mathbf{K}^{-1}\mathbf{K}\boldsymbol{\alpha}$ for the fiber velocities. That is, *the rotation rate $\boldsymbol{\alpha}$ is the actual rotation rate of the tangent vectors on the N point grid.*

To obtain the discretization of \mathbf{K}^* , we simply use the discrete L^2 inner product. The virtual work constraint in Section 3.1.3 says that for all $\boldsymbol{\alpha}$

$$0 = \langle \mathbf{K}\boldsymbol{\alpha}, \boldsymbol{\lambda} \rangle_{L^2} = \boldsymbol{\alpha}^T \mathbf{K}^T \widetilde{\mathbf{W}} \boldsymbol{\lambda}. \quad (6.32)$$

For this to be zero for any $\boldsymbol{\alpha}$, it follows that $\mathbf{K}^T \widetilde{\mathbf{W}} \boldsymbol{\lambda} = \mathbf{0}$, which gives a simple discretization of \mathbf{K}^* as

$$\mathbf{K}^* = \mathbf{K}^T \widetilde{\mathbf{W}}. \quad (6.33)$$

6.2 Discretization of mobility for a single fiber

Now that we have discretized the kinematic operators and forces, the last ingredient we need to solve (4.2) is to evaluate the action of the mobility operators $\mathcal{M}_{\text{tt}}, \mathcal{M}_{\text{rt}}, \mathcal{M}_{\text{tr}}$, and \mathcal{M}_{rr} at each point on the Chebyshev grid, giving discrete mobility matrices $\mathbf{M}_{\text{tt}}, \mathbf{M}_{\text{rt}}, \mathbf{M}_{\text{tr}}$, and \mathbf{M}_{rr} . We discuss how to do this in this section.

Our approach is to use the mobility operators given by the RPY integral (2.4). For translational mobility, if the integral is computed to sufficient accuracy, then the eigenvalues of the discrete mobility matrix \mathbf{M}_{tt} are guaranteed to be positive since the RPY kernel is SPD [65]. While (2.4) is still a “first kind” integral operator, and therefore has eigenvalues that cluster around zero, unlike the “second kind” integral operator proposed in [95] to regularize SBT, we will demonstrate that the ill-conditioning of \mathbf{M}_{tt} is *not* a problem in practice. Furthermore, unlike [95], we develop a nearly-singular quadrature scheme so that the RPY integrals for smooth fibers can be evaluated to 3 digits of accuracy with 10 – 40 collocation points, regardless of the slenderness $\hat{\epsilon}$.

Our nearly singular “slender body” quadrature schemes are based on an input position

function $\mathbb{X}(s)$ and input force density $\mathbf{f}(s)$ or torque density $n^{\parallel}(s)$. The goal is to compute the integrals (2.4) to high accuracy from these inputs. The schemes we present in this section do *not* assume inextensible fibers; that is, s does not have to be arclength parameterization for these schemes to accurately resolve the integrals. We will, however, assume that the curves are approximately inextensible, so that we can make the approximation $R(s, s') \approx |s - s'|$ when we separate the integral into a piece over $R < 2\hat{a}$ and $R \geq 2\hat{a}$.

6.2.1 Translational mobility

We define the translational velocity of the fiber centerline due to a force density $\mathbf{f}(s)$ using (2.4),

$$\begin{aligned} \mathbf{U}_f(s) &= (\mathcal{M}_{\text{tt}} \mathbf{f})(s) = \int_0^L \mathbb{M}_{\text{tt}}(\mathbb{X}(s), \mathbb{X}(s')) \mathbf{f}(s') ds' \\ &= \int_{D(s)} \left(\mathbb{S}(\mathbb{X}(s), \mathbb{X}(s')) + \frac{2\hat{a}^2}{3} \mathbb{D}(\mathbb{X}(s), \mathbb{X}(s')) \right) \mathbf{f}(s') ds' \\ &\quad + \frac{1}{8\pi\mu} \int_{D^c(s)} \left(\left(\frac{4}{3\hat{a}} - \frac{3R(\mathbb{X}(s), \mathbb{X}(s'))}{8\hat{a}^2} \right) \mathbf{I} + \frac{\mathbf{R}\hat{\mathbf{R}}(\mathbb{X}(s), \mathbb{X}(s'))}{8\hat{a}^2} \right) \mathbf{f}(s') ds'. \end{aligned} \quad (6.34)$$

Here we have used the definition (2.6) to split the integral into a region $D(s)$ for $R(\mathbf{X}(s), \mathbf{X}(s')) > 2\hat{a}$ and $D^c(s)$, which uses the RPY kernel for $R \leq 2\hat{a}$. We use the (approximate) fiber inextensibility to make the approximation $R \approx |s' - s|$ when $R \lesssim 2\hat{a}$, so that

$$D(s) = \begin{cases} (0, s - 2\hat{a}) \cup (s + 2\hat{a}, L) & 2\hat{a} \leq s \leq L - 2\hat{a} \\ (s + 2\hat{a}, L) & s < 2\hat{a} \\ (0, s - 2\hat{a}) & s > L - 2\hat{a} \end{cases}, \quad (6.35)$$

with the complement $D^c(s) = [0, L] \setminus D(s)$. In making the approximation $R \approx |s' - s|$, we also assume that the fiber never re-encroaches itself, which can be accomplished in practice by including steric interactions in simulations (see Chapter 9). We now discuss each of the

integrals in (6.34) in more detail, separating out the Stokeslet and doublet into two separate integrals.

6.2.1.1 Stokeslet integral

For the integral of the Stokeslet in (6.34), we use a singularity subtraction technique which is closely tied with the asymptotics of the Stokeslet. In particular, we subtract from the integrand the leading order singular behavior and perform that integral separately, which gives

$$\begin{aligned}
\mathbf{U}_f^{(S)} &= \int_{D(s)} \mathbb{S}(\mathbb{X}(s), \mathbb{X}(s')) \mathbf{f}(s') ds' \tag{6.36} \\
&= \frac{1}{8\pi\mu} \int_{D(s)} \left(\frac{\mathbf{I} + \widehat{\partial_s \mathbb{X}(s)} \widehat{\partial_s \mathbb{X}(s)}}{\|\partial_s \mathbb{X}(s)\| |s - s'|} \right) \mathbf{f}(s) ds' \\
&\quad + \int_{D(s)} \left(\mathbb{S}(\mathbb{X}(s), \mathbb{X}(s')) \mathbf{f}(s') - \frac{1}{8\pi\mu} \left(\frac{\mathbf{I} + \widehat{\partial_s \mathbb{X}(s)} \widehat{\partial_s \mathbb{X}(s)}}{|s - s'|} \right) \mathbf{f}(s) \right) ds' \\
&:= \mathbf{U}_f^{(\text{inner}, S)}(s) + \mathbf{U}_f^{(\text{int}, S)}(s) \\
\text{where } \mathbf{U}_f^{(\text{inner}, S)}(s) &= \frac{a_L(s)}{8\pi\mu} \frac{\left(\mathbf{I} + \widehat{\partial_s \mathbb{X}(s)} \widehat{\partial_s \mathbb{X}(s)} \right) \mathbf{f}(s)}{\|\partial_s \mathbb{X}(s)\|},
\end{aligned}$$

where $a_L(s)$ is defined in (2.20). This leaves a nearly-singular integral that represents the Stokeslet minus the leading order singularity,

$$\mathbf{U}_f^{(\text{int}, S)}(s) = \int_{D(s)} \left(\mathbb{S}(\mathbb{X}(s), \mathbb{X}(s')) \mathbf{f}(s') - \frac{1}{8\pi\mu} \left(\frac{\mathbf{I} + \widehat{\partial_s \mathbb{X}(s)} \widehat{\partial_s \mathbb{X}(s)}}{\|\partial_s \mathbb{X}(s)\| |s - s'|} \right) \mathbf{f}(s) \right) ds', \tag{6.37}$$

which has the same integrand as the finite part integral in slender body theory, but with the different domain of integration $D(s)$. Because of our singularity subtraction scheme, the second term in the integral (6.37) cancels the leading order $1/|s' - s|$ singularity in the first. The next singularity is $\text{sign}(s' - s)$, which means that the near singular integral (6.37) can

be written as

$$\mathbf{U}_f^{(\text{int}, \mathbb{S})}(s) = \int_{D(s)} \mathbf{g}_{\text{tt}}(s, s') \frac{s' - s}{|s' - s|} ds' = \frac{L}{2} \int_{D'(\eta)} \mathbf{g}_{\text{tt}}(\eta, \eta') \frac{\eta' - \eta}{|\eta' - \eta|} d\eta', \quad \text{where} \quad (6.38)$$

$$\mathbf{g}_{\text{tt}}(s, s') = \left(\mathbb{S}(\mathbb{X}(s), \mathbb{X}(s')) \mathbf{f}(s') |s' - s| - \frac{1}{8\pi\mu} \frac{\left(\mathbf{I} + \widehat{\partial_s \mathbb{X}(s)} \widehat{\partial_s \mathbb{X}(s)} \right) \mathbf{f}(s)}{\|\partial_s \mathbb{X}(s)\|} \right) \frac{1}{s' - s}$$

is a smooth function, $\eta = -1 + 2s/L$ is a rescaled arclength coordinate on $[-1, 1]$, and $D'(\eta) = [0, \eta_\ell] \cup [\eta_h, L]$ is defined from $D(s)$. The function \mathbf{g}_{tt} is nonsingular at $s = s'$ with the finite limit

$$\lim_{s' \rightarrow s} \mathbf{g}_{\text{tt}}(s, s') = \frac{1}{8\pi\mu} \left[\frac{1}{2 \|\partial_s \mathbb{X}(s)\|^3} \left(\partial_s \mathbb{X}(s) \partial_s^2 \mathbb{X}(s) + \partial_s^2 \mathbb{X}(s) \partial_s \mathbb{X}(s) \right. \right. \quad (6.39)$$

$$\left. \left. - (\partial_s \mathbb{X}(s) \cdot \partial_s^2 \mathbb{X}(s)) \left(\mathbf{I} + 3 \widehat{\partial_s \mathbb{X}(s)} \widehat{\partial_s \mathbb{X}(s)} \right) \right) \mathbf{f}(s) + \left(\frac{\mathbf{I} + \widehat{\partial_s \mathbb{X}(s)} \widehat{\partial_s \mathbb{X}(s)}}{\|\partial_s \mathbb{X}(s)\|} \right) \partial_s \mathbf{f}(s) \right]$$

Furthermore, \mathbf{g}_{tt} is smooth, so we can express it in a truncated Chebyshev series on $[-1, 1]$,

$$\frac{L}{2} \mathbf{g}_{\text{tt}}(\eta, \eta') \approx \sum_{k=0}^{N-1} \mathbf{c}_k(\eta) T_k(\eta'), \quad (6.40)$$

where \mathbf{c}_k is a vector of 3 coefficients for each η . Substituting the Chebyshev expansion (6.40) into the integrand (6.38), we obtain

$$\mathbf{U}_f^{(\text{int}, \mathbb{S})}(\eta) = \sum_{k=0}^{N-1} \mathbf{c}_k(\eta) \int_{D'(\eta)} T_k(\eta') \frac{\eta' - \eta}{|\eta' - \eta|} d\eta' = \sum_{k=0}^{N-1} \mathbf{c}_k(\eta) q_k^{(S)}(\eta) = \mathbf{c}^T(\eta) \mathbf{q}^{(S)}(\eta), \quad (6.41)$$

$$\text{where} \quad q_k^{(S)}(\eta) = \int_{D'(\eta)} T_k(\eta') \frac{\eta' - \eta}{|\eta' - \eta|} d\eta' = - \int_{-1}^{\eta_\ell} T_k(\eta') d\eta' + \int_{\eta_h}^L T_k(\eta') d\eta', \quad (6.42)$$

are integrals that can be precomputed to high accuracy for each η on the Chebyshev collocation grid. The precomputed integrals in (6.42) can then be used in an adjoint method

to accelerate the repeated evaluation of (6.41) [80]. Specifically, the coefficients of $\mathbf{g}_{\text{tt}}(\eta)$ for each η can be written as $\mathbf{c}(\eta) = \mathbf{V}^{-1}\mathbf{g}_{\text{tt}}(\eta, \eta')$, where \mathbf{V} is the matrix that maps coefficients of the basis functions to values on the collocation grid. Then $\mathbf{c}^T(\eta) = \mathbf{g}^T(\eta, \eta')\mathbf{V}^{-T}$, and $\mathbf{U}_{\text{f}}^{(\text{int}, \text{S})}(\eta) = \mathbf{g}^T(\eta, \eta')\mathbf{V}^{-T}\mathbf{q}^{(\text{S})}(\eta)$. The matrix vector product $\mathbf{V}^{-T}\mathbf{q}^{(\text{S})}(\eta)$ can then be pre-computed for each η , and so at each time step all that is needed to compute the integral (6.37) for each η are the values of $\mathbf{g}_{\text{tt}}(\eta, \eta')$ on the collocation grid for η' being another collocation point. Evaluating $\mathbf{U}_{\text{f}}^{(\text{int}, \text{S})}$ therefore has a cost of $\mathcal{O}(N^2)$.

We note that our use of Chebyshev polynomials differs from the approach of [80], which uses a monomial basis for the expansion (6.40). Although the integrals (6.42) with monomials can be computed analytically, using the adjoint method with monomial coefficients requires inverting the Vandermonde matrix \mathbf{V} , which is too ill-conditioned with $N \gtrsim 40$ points in double precision [133]. The Chebyshev expansion is an improvement in the sense that computing coefficients from values on the Chebyshev grid is a well-conditioned problem, and so the technique we use to evaluate finite part integrals does not limit the number of points on the fiber or require the use of panels (composite quadrature).

6.2.1.2 Modifications to discretize asymptotic mobility

Let us now suppose that instead of numerically integrating the RPY mobility, we want to use the asymptotic, SBT-style mobility given by (2.24). In this case, we still have a local drag matrix, and the only nonlocal integral is again given by (6.37), but with $D(s) = [0, L]$. The calculation then proceeds in exactly the same way. It is possible, however, to compute the integrals (6.42) analytically using a monomial basis, which gives

$$q_k^{(\text{S})}(\eta) = \int_{-1}^1 (\eta')^k \frac{\eta' - \eta_p}{|\eta' - \eta_p|} d\eta' = \frac{1 + (-1)^{k+1} - 2\eta_p^{k+1}}{k + 1}. \quad (6.43)$$

The rest of the calculation for the finite part integral is precisely the same. Notice here that we do not have any additional nonlocal integrals to compute; thus, using the asymptotic for-

mulation is indeed simpler and faster than doing the full RPY integral numerically. However, we will see in Section 6.3 that using the asymptotic mobility produces negative eigenvalues in the trans-trans mobility matrix \mathbf{M}_{tt} , making the Euler problem (4.2) ill-posed.

6.2.1.3 Doublet integral

The next part of the translational mobility (6.34) is the integral of the doublet kernel. We begin by applying the same singularity subtraction technique we used in (6.36) for the Stokeslet integral

$$\begin{aligned}
\mathbf{U}_f^{(D)}(s) &= \int_{D(s)} \mathbb{D}(\mathbf{X}(s), \mathbf{X}(s')) \mathbf{f}(s') ds' \\
&= \frac{1}{8\pi\mu} \int_{D(s)} \left(\frac{\mathbf{I} - 3\widehat{\partial_s \mathbb{X}}(s)\widehat{\partial_s \mathbb{X}}(s)}{\|\partial_s \mathbb{X}(s)\|^3 |s - s'|^3} \right) \mathbf{f}(s) \\
&\quad + \int_{D(s)} \left(\mathbb{D}(\mathbf{X}(s), \mathbf{X}(s')) \mathbf{f}(s') - \frac{1}{8\pi\mu} \left(\frac{\mathbf{I} - 3\widehat{\partial_s \mathbb{X}}(s)\widehat{\partial_s \mathbb{X}}(s)}{\|\partial_s \mathbb{X}(s)\|^3 |s - s'|^3} \right) \mathbf{f}(s) \right) ds' \\
&= \mathbf{U}_f^{(\text{inner}, D)}(s) + \mathbf{U}_f^{(\text{int}, D)}(s).
\end{aligned} \tag{6.44}$$

The term $\mathbf{U}_f^{(\text{inner}, D)}(s)$ is given for all s by

$$\mathbf{U}_f^{(\text{inner}, D)}(s) = \frac{(\mathbf{I} - 3\boldsymbol{\tau}(s)\boldsymbol{\tau}(s)) a_D(s) \mathbf{f}(s)}{8\pi\mu \|\partial_s \mathbb{X}(s)\|^3}, \tag{6.45}$$

where $a_D(s)$ is defined in (2.20).

The nearly-singular doublet integral is handled in exactly the same way as the Stokeslet one, but the singular behavior is different, namely,

$$\mathbf{U}_f^{(\text{int}, D)}(s) = \int_{D(s)} \frac{(s' - s)}{|s' - s|^3} \mathbf{g}_D(s', s) ds', \tag{6.46}$$

where $\mathbf{g}_D(s', s) = \left(\mathbb{D}(\mathbf{X}(s), \mathbf{X}(s')) \mathbf{f}(s') - \frac{1}{8\pi\mu} \left(\frac{\mathbf{I} - 3\widehat{\partial_s \mathbb{X}}(s)\widehat{\partial_s \mathbb{X}}(s)}{\|\partial_s \mathbb{X}(s)\|^3 |s - s'|^3} \right) \mathbf{f}(s) \right) \frac{|s' - s|^3}{(s' - s)}$

is a smooth function with the finite limit

$$\begin{aligned} \lim_{s' \rightarrow s} \mathbf{g}_D(s', s) &= \frac{1}{8\pi\mu} \left[\frac{1}{2 \|\partial_s \mathbb{X}(s)\|^5} \left(-3 (\partial_s \mathbb{X}(s) \partial_s^2 \mathbb{X}(s) + \partial_s^2 \mathbb{X}(s) \partial_s \mathbb{X}(s)) \right. \right. \\ &\quad \left. \left. - (\partial_s \mathbb{X}(s) \cdot \partial_s^2 \mathbb{X}(s)) \left(3\mathbf{I} - 15 \widehat{\partial_s \mathbb{X}(s)} \widehat{\partial_s \mathbb{X}(s)} \right) \right) \mathbf{f}(s) + \left(\frac{\mathbf{I} - 3 \widehat{\partial_s \mathbb{X}(s)} \widehat{\partial_s \mathbb{X}(s)}}{\|\partial_s \mathbb{X}(s)\|^3} \right) \partial_s \mathbf{f}(s) \right]. \end{aligned}$$

If we expand \mathbf{g}_D in terms of Chebyshev polynomials, the efficient evaluation of (6.44) requires precomputing integrals of the form

$$q_k^{(D)}(\eta) = \int_{D'(\eta)} \frac{(\eta' - \eta)}{|\eta' - \eta|^3} T_k(\eta') d\eta', \quad (6.47)$$

for all η on the Chebyshev grid. Notice that the integrals (6.47) are not defined if $\eta' = \eta$ is included in the integration domain, but $D'(\eta)$ does *not* contain η .

6.2.1.4 Integral on $R < 2\hat{a}$

After discretizing the Stokeslet and doublet integrals, we are left with the integral over $D^c(s)$ in the third line of (6.34). The integrand is nonsingular, but behaves like $|s - s'|$, and so for each s we split the domain D^c into $(s - 2\hat{a}, s)$ and $(s, s + 2\hat{a})$ (with appropriate modifications at the endpoints), and use $N_2/2$ Gauss-Legendre quadrature points to sample the fiber and force density (i.e., sample the Chebyshev interpolant of each) and evaluate the integral on each of the two subdomains. We use $N_2/2$ points for each of these integrals so that there are a total of N_2 additional (local) quadrature nodes per collocation point.

6.2.2 Rot-trans mobility

For the rotational velocity from force, we have the RPY integrals

$$\Omega_{\mathbf{f}}^{\parallel}(s) := (\mathcal{M}_{\text{rt}} \mathbf{f})(s) = \partial_s \mathbb{X}(s) \cdot \int_0^L \mathbb{M}_{\text{rt}}(\mathbb{X}(s), \mathbb{X}(s')) \mathbf{f}(s') ds' \quad (6.48)$$

$$\begin{aligned} &= \frac{1}{8\pi\mu} \int_{D(s)} \frac{(\mathbf{f}(s') \times \mathbf{R}(\mathbb{X}(s'), \mathbb{X}(s))) \cdot \partial_s \mathbb{X}(s)}{R(\mathbb{X}(s'), \mathbb{X}(s))^3} ds' \\ &+ \frac{1}{8\pi\mu} \int_{D^c(s)} \frac{1}{2\hat{a}^2} \left(\frac{1}{\hat{a}} - \frac{3R(\mathbb{X}(s'), \mathbb{X}(s))}{8\hat{a}^2} \right) (\mathbf{f}(s') \times \mathbf{R}(\mathbb{X}(s'), \mathbb{X}(s)) \cdot \partial_s \mathbb{X}(s)) ds', \end{aligned} \quad (6.49)$$

which we evaluate to spectral accuracy using singularity subtraction and slender-body quadrature for the rotlet integral over $D(s)$, and direct Gauss-Legendre quadrature for the integral over $D^c(s)$ (split into two pieces).

6.2.2.1 Slender body quadrature

The long-range integral for the rot-trans mobility (6.48) is given by an integral of the rotlet kernel over the domain $D(s)$. Using singularity subtraction, we compute this integral as

$$\begin{aligned} \Omega_{\mathbf{f}}^{(\text{R})}(s) &= \int_{D(s)} \frac{(\mathbf{f}(s') \times \mathbf{R}(s')) \cdot \partial_s \mathbb{X}(s)}{R(s')^3} ds' \\ &= \int_{D(s)} \left(\frac{(\partial_s^2 \mathbb{X}(s) \times \mathbf{f}(s)) \cdot \partial_s \mathbb{X}(s)}{2|s - s'| \|\partial_s \mathbb{X}(s)\|^3} \right) ds' \\ &+ \int_{D(s)} \left(\frac{(\mathbf{f}(s') \times \mathbf{R}(s')) \cdot \boldsymbol{\tau}(s)}{R(s')^3} - \frac{(\partial_s \boldsymbol{\tau}(s) \times \mathbf{f}(s)) \cdot \boldsymbol{\tau}(s)}{2|s - s'| \|\partial_s \mathbb{X}(s)\|^3} \right) ds' \\ &:= \Omega_{\mathbf{f}}^{(\text{inner, R})}(s) + \Omega_{\mathbf{f}}^{(\text{int, R})}(s), \end{aligned} \quad (6.50)$$

where $\mathbf{R}(s') = \mathbf{X}(s) - \mathbf{X}(s')$ and $R = \|\mathbf{R}\|$. Since the leading order singularity in $\Omega_{\mathbf{f}}^{(\text{inner, R})}$ of $1/|s - s'|$ is the same as that of the Stokeslet,

$$\Omega_{\mathbf{f}}^{(\text{inner, R})}(s) = a_L(s) \left(\frac{(\partial_s^2 \mathbb{X}(s) \times \mathbf{f}(s)) \cdot \partial_s \mathbb{X}(s)}{2 \|\partial_s \mathbb{X}(s)\|^3} \right), \quad (6.51)$$

where $a_L(s)$ is defined in (2.20). The nearly singular integral $\Omega_f^{(\text{int}, R)}(s)$ in (6.50) can then be put in the form (6.38) (with scalar $g_{\text{rt}}(s, s')$) by setting

$$g_{\text{rt}}(s, s') = \left(\frac{|s - s'|}{R(s')^3} (\mathbf{R}(s') \times \partial_s \mathbb{X}(s)) \cdot \mathbf{f}(s') - \frac{(\partial_s \mathbb{X}(s) \times \partial_s^2 \mathbb{X}(s)) \cdot \mathbf{f}(s)}{2 \|\partial_s \mathbb{X}(s)\|^3} \right) \frac{1}{s' - s}, \quad (6.52)$$

which has the finite limit

$$\begin{aligned} \lim_{s' \rightarrow s} g_{\text{rt}}(s, s') &= -\frac{1}{2 \|\partial_s \mathbb{X}(s)\|^3} \left(\partial_s \mathbf{f}(s) \cdot (\partial_s^2 \mathbb{X}(s) \times \partial_s \mathbb{X}(s)) + \frac{1}{3} \mathbf{f}(s) \cdot (\partial_s^3 \mathbb{X}(s) \times \partial_s \mathbb{X}(s)) \right) \\ &\quad - \frac{3}{4 \|\partial_s \mathbb{X}(s)\|^5} (\partial_s \mathbb{X}(s) \cdot \partial_s^2 \mathbb{X}(s)) (\partial_s \mathbb{X}(s) \times \partial_s^2 \mathbb{X}(s)) \cdot \mathbf{f}(s) \end{aligned} \quad (6.53)$$

Now the scalar function $g_{\text{rt}}(s, s')$ can be expanded in a Chebyshev series as in (6.41), and the precomputed integrals $q_k^{(S)}(\eta)$ in (6.42) used once again to evaluate the second integral in (6.50).

6.2.3 Trans-rot mobility

The trans-rot mobility is the adjoint of the rot-trans mobility,

$$\mathbf{U}_n(s) := (\mathcal{M}_{\text{tr}} n^\parallel)(s) = \int_0^L \mathbb{M}_{\text{rt}}(\mathbf{X}(s), \mathbf{X}(s')) \boldsymbol{\tau}(s') n^\parallel(s') ds' \quad (6.54)$$

$$\begin{aligned} &= \frac{1}{8\pi\mu} \int_{D(s)} \frac{(\boldsymbol{\tau}(s') \times \mathbf{R}(\mathbf{X}(s'), \mathbf{X}(s))) n^\parallel(s')}{R(\mathbf{X}(s'), \mathbf{X}(s))^3} ds' \\ &+ \frac{1}{8\pi\mu} \int_{D^c(s)} \frac{1}{2\hat{a}^2} \left(\frac{1}{\hat{a}} - \frac{3R(\mathbf{X}(s'), \mathbf{X}(s))}{8\hat{a}^2} \right) (\boldsymbol{\tau}(s') \times \mathbf{R}(\mathbf{X}(s'), \mathbf{X}(s)) n^\parallel(s')) ds'. \end{aligned} \quad (6.55)$$

As in the rot-trans mobility, we compute the first (rotlet) integral using slender-body quadrature, and the second using direct Gauss-Legendre quadrature.

6.2.3.1 Slender body quadrature

For the trans-rot mobility (6.54), the singularity subtraction approach is to set

$$\begin{aligned}
\mathbf{U}_n^{(\text{R})}(s) &= \int_{D(s)} \frac{(\partial_s \mathbb{X}(s') \times \mathbf{R}(s')) n^\parallel(s')}{R(s')^3} ds' \\
&= \int_{D(s)} \left(\frac{(\partial_s \mathbb{X}(s) \times \partial_s^2 \mathbb{X}(s)) n^\parallel(s)}{2|s-s'| \|\partial_s \mathbb{X}(s)\|^3} \right) ds' \\
&\quad + \int_{D(s)} \left(\frac{(\partial_s \mathbb{X}(s') \times \mathbf{R}(s')) n^\parallel(s')}{R(s')^3} - \frac{(\partial_s \mathbb{X}(s) \times \partial_s^2 \mathbb{X}(s)) n^\parallel(s)}{2|s-s'| \|\partial_s \mathbb{X}(s)\|^3} \right) ds' \\
&:= \mathbf{U}_n^{(\text{inner, R})}(s) + \mathbf{U}_n^{(\text{int, R})}(s)
\end{aligned} \tag{6.56}$$

$$\mathbf{U}_n^{(\text{inner, R})}(s) = a_L(s) \left(\frac{(\partial_s \mathbb{X}(s) \times \partial_s^2 \mathbb{X}(s)) \cdot n(s)}{2 \|\partial_s \mathbb{X}(s)\|^3} \right), \tag{6.57}$$

where $a_L(s)$ is defined in (2.20). The nearly singular integral $\mathbf{U}_n^{(\text{int, R})}(s)$ can then be put in the form (6.38) by setting

$$\mathbf{g}_{\text{tr}}(s, s') = \left(\frac{|s-s'|}{R(s')^3} (\boldsymbol{\tau}(s') \times \mathbf{R}(s')) n^\parallel(s') - \left(\frac{\boldsymbol{\tau}(s) \times \partial_s \boldsymbol{\tau}(s)}{2 \|\partial_s \mathbb{X}(s)\|^3} \right) n^\parallel(s) \right) \frac{1}{s' - s}, \tag{6.58}$$

$$\begin{aligned}
\lim_{s' \rightarrow s} \mathbf{g}_{\text{tr}}(s, s') &= -\frac{\partial_s n^\parallel(s)}{2 \|\partial_s \mathbb{X}(s)\|^3} (\partial_s^2 \mathbb{X}(s) \times \partial_s \mathbb{X}(s)) - \frac{n^\parallel(s)}{3 \|\partial_s \mathbb{X}(s)\|^3} (\partial_s^3 \mathbb{X}(s) \times \partial_s \mathbb{X}(s)) \\
&\quad - \frac{3}{4 \|\partial_s \mathbb{X}(s)\|^5} (\partial_s \mathbb{X}(s) \cdot \partial_s^2 \mathbb{X}(s)) (\partial_s \mathbb{X}(s) \times \partial_s^2 \mathbb{X}(s)) n^\parallel(s)
\end{aligned} \tag{6.59}$$

6.2.4 Rot-rot mobility

We now consider the rot-rot mobility. Substituting the tensor \mathbb{M}_{rr} from (2.9) into the integral (2.4) and using the approximation $\|\mathbb{X}(s) - \mathbb{X}(s')\| \approx s - s'$ when $s - s' = \mathcal{O}(\hat{a})$, the mobility is given by

$$\begin{aligned}
8\pi\mu \boldsymbol{\Psi}_n(s) &= -\frac{1}{2} \int_{D(s)} \mathbb{D}(\mathbb{X}(s), \mathbb{X}(s')) \mathbf{n}(s') ds' \\
&\quad + \frac{1}{\hat{a}^3} \int_{D^c(s)} \left(\left(1 - \frac{27R(s')}{32\hat{a}} + \frac{5R(s')^3}{64\hat{a}^3} \right) \mathbf{I} + \left(\frac{9}{32\hat{a}R(s')} - \frac{3R(s')}{64\hat{a}^3} \right) (\mathbf{R}\mathbf{R})(s') \right) \mathbf{n}(s') ds'.
\end{aligned} \tag{6.60}$$

For the rot-rot mobility, we will always use the asymptotic result (2.47). In the fiber interior when $\|\partial_s \mathbb{X}(s)\| = 1$, this reduces to (2.49),

$$\Omega_n^{\parallel}(s) = (\mathcal{M}_{\text{rr}} n^{\parallel})(s) = \frac{1}{8\pi\mu} \frac{9n^{\parallel}(s)}{4\hat{a}^2}. \quad (6.61)$$

We use the asymptotic result here because the $1/\hat{a}^2$ term is so dominant as to render calculation of the full integral, which involves the rapidly-decaying doublet kernel, unnecessary. More importantly, the local operator \mathcal{M}_{rr} (and the diagonal discrete matrix \mathbf{M}_{rr}) are positive definite and well-behaved up to the end points. In our implementation, we derive corresponding expressions for (2.47) when $\|\partial_s \mathbb{X}(s)\| \neq 1$ in a straightforward way by substituting $R^p \approx |s - s'|^p \|\partial_s \mathbb{X}(s)\|^p$ into the expansions of Section 2.1.4.

6.2.4.1 Slender body quadrature (if desired)

Although we use asymptotic evaluation for the rot-rot mobility, it is straightforward to derive a nearly singular quadrature scheme to evaluate the RPY integral for \mathcal{M}_{rr} . The first of the integrals (6.60) can be evaluated to high accuracy using the doublet scheme of Section 6.2.1.3, while the second integral can be computed using direct Gauss-Legendre quadrature on $D^c(s)$ (split into two pieces). The parallel angular velocity Ψ_n^{\parallel} can then be obtained by taking the inner product of Ψ_n with $\boldsymbol{\tau}$.

6.2.5 Convergence of quadratures

We next test the accuracy of the (action of the) mobility matrices \mathbf{M}_{tt} , \mathbf{M}_{tr} , \mathbf{M}_{rt} , and \mathbf{M}_{rr} by comparing against upsampled direct quadrature. To illustrate that our quadrature




Figure 6.7: Convergence of the nearly singular quadrature schemes. We use the fiber and force/torque density in (6.62) and compute the velocities \mathbf{U}_f , \mathbf{U}_n , Ω_f^{\parallel} , and Ω_n^{\parallel} on the collocation grid by applying the discrete mobility matrices. In all cases, we compute a reference solution by upsampled direct quadrature and report the *relative* L^2 errors in the velocity.

scheme does *not* depend on the fiber being inextensible, we apply it to a fiber with

$$\mathbb{X}(s) = (\cos 3s, \sin 3s, \cos s), \quad (6.62)$$

$$\mathbf{f}(s) = (\cos 2s^3, s^2, \sin 2s) \quad (6.63)$$

$$n(s) = \cos(4s^2) \quad (6.64)$$

In each case, we compute the error relative to upsampling the integrals with a large number of Clenshaw-Curtis quadrature nodes. Here by “the integrals,” we mean the corresponding integral *after* dividing the domain into $D(s)$ and $D^c(s)$, so that we are testing the quadrature scheme and not our assumption that $R \approx |s - s'|$. The relative error is normalized by the L^2 norm of the reference solution.

Figure 6.7 shows the results for the trans-trans mobility (6.34) and trans-rot mobility (6.54) (top row), and rot-trans mobility (6.48) and rot-rot mobility (6.61) (bottom row) for two different values of $\hat{\epsilon}$, and various values of N and N_2 . In the case of the trans-trans, rot-trans, and trans-rot mobilities, we obtain spectral convergence, with an error that saturates when N_2 is too small. Increasing N_2 gives spectral convergence to machine precision, at a rate and saturated error that is independent of $\hat{\epsilon}$.

In Fig. 6.7, we also show the accuracy of the rot-rot mobility as defined in (6.61) relative to the actual RPY integral of doublets as defined in (2.9). We do this in two ways: we again perform special quadrature on the integral (2.37) of doublets, finding an accuracy that depends on N and (most especially) N_2 and is essentially independent of $\hat{\epsilon}$. The strong dependence on N_2 tells us that the value of the integral is controlled by resolution of scales $R \sim \hat{a}$. As such, we compare the numerical results to the *asymptotic* local mobility (6.61) as dotted lines. There we see that for $\hat{\epsilon} = 10^{-2}$, we get about one digit of accuracy, while for $\hat{\epsilon} = 10^{-3}$ we obtain about three digits. The error decrease reflects the fact that the full asymptotic formula for the rot-rot mobility (2.47) is $\mathcal{O}(\hat{\epsilon}^2)$ in the fiber interior ($2\hat{a} \leq s \leq$

$L - 2\hat{a}$) and $\mathcal{O}(\hat{\epsilon})$ at the fiber endpoints. This is sufficient for our purposes, so we use the asymptotic mobility going forward.

6.3 Well-posedness/convergence of the static problem

In this section, we study the convergence of our numerical method for the case of a static filament (solving for velocity at $t = 0$) in two steps. First, we show that our discrete mobility matrix \mathbf{M}_{tt} is indeed positive definite for larger $\hat{\epsilon}$, in contrast to \mathbf{M}_{tt} in the asymptotic formula (2.24), which has negative eigenvalues for as few as 15 collocations points on the fiber centerline when $\hat{\epsilon} = 10^{-2}$. Next, we solve for the velocity \mathbf{U} and constraint force $\boldsymbol{\lambda}$ in the Euler model equations (4.2) and examine their convergence. This is where our method hits a snag: if we maintain a cylindrical fiber shape, there is a physical singularity that develops at the fiber endpoints, where the force required to produce a uniform motion becomes very oscillatory [134, Sec. 4.3.2]. This means that $\boldsymbol{\lambda}$ fails to converge pointwise, which calls into question the convergence of \mathbf{U} as well. Nevertheless, we show empirically that the spectral method *does* give a convergent velocity and a weakly-convergent constraint force $\boldsymbol{\lambda}$ (i.e., the moments of $\boldsymbol{\lambda}$ converge). Furthermore, the accuracy we obtain in the spectral method with just 16 – 30 points on the fiber centerline is better than with $1/\hat{\epsilon}$ points in a second-order method.

Throughout this section, we will consider free fibers with tangent vectors of the form

$$\boldsymbol{\tau}(s) = \frac{1}{\sqrt{2}} (\cos(qs^3(s-L)^3), \sin(qs^3(s-L)^3), 1), \quad (6.65)$$

where we set $L = 2$, and q is a parameter that determines the number of helical turns and fiber curvature and smoothness. Fibers with $q = 1$ and $q = 7$ are shown in blue and red, respectively, in Fig. 6.8.

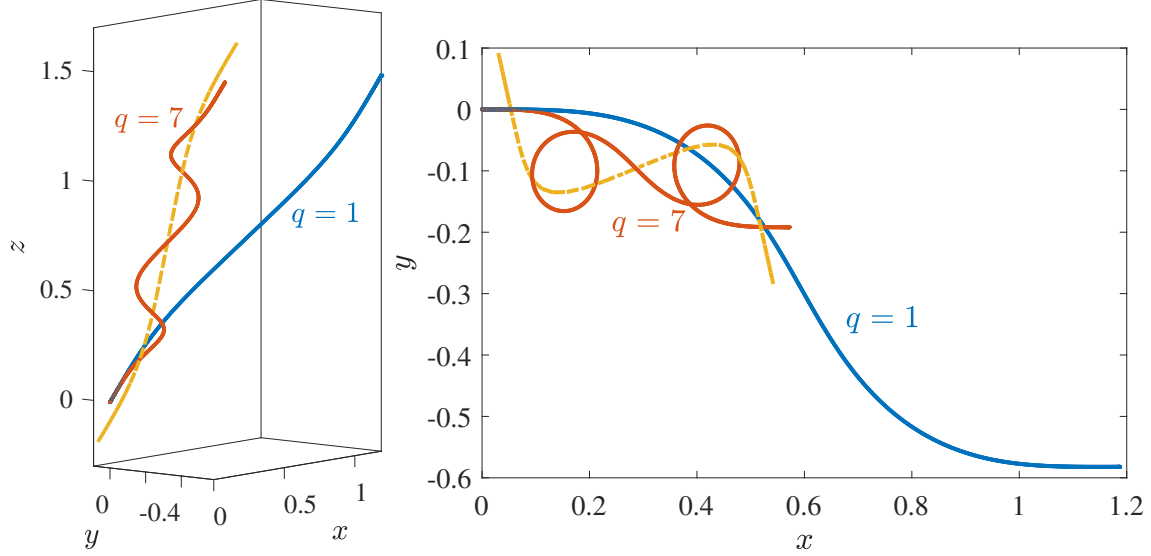


Figure 6.8: Fibers we use for the static convergence study, shown from two different views. Here $L = 2$ and the fiber tangent vector is given by (6.65), with $q = 1$ for the blue curve and $q = 7$ for the red curve. The dotted yellow curve shows the $q = 7$ fiber after $t = 0.01$ seconds of relaxation (when $\hat{\epsilon} = 10^{-2}$, $\mu = 1$, $\kappa = 1$ and $\gamma = 0$; see Section 6.5.1).

6.3.1 Eigenvalues of \mathbf{M}_{tt}

To illustrate that the RPY mobility (6.34) alleviates the problem of negative eigenvalues that plagues SBT, we compute the eigenvalues of the translation-translation matrix \mathbf{M}_{tt} numerically. We use $q = 7$ in the free fiber (6.65), although the mobility does not change substantially with the fiber shape (specifically, a straight fiber and four-turn helix give similar results).

There are two parameters associated with the quadrature for (6.34): the number of collocation points N that we use for the fiber centerline, which we also use for the Stokeslet and doublet integrals, and the total number of additional points N_2 that we use for the integrals in the region $|s - s'| \leq 2\hat{a}$ (recall that we use $N_2/2$ Gauss-Legendre points on either side of $s' = s$ for these integrals). In Fig. 6.9, we report the eigenvalues of the mobility matrix for $\hat{\epsilon} = 10^{-2}$ as a function of N (left plot) and N_2 (right plot). In the left plot, we also show the eigenvalues of \mathbf{M}_{tt} when we use the asymptotic SBT of (2.24) for translation. We see that the eigenvalues with SBT quadrature become negative even using $N = 15$ collocation




Figure 6.9: Absolute value of the eigenvalues of the discrete mobility matrix \mathbf{M}_{tt} for $\hat{\epsilon} = 10^{-2}$ with slender-body quadrature for the RPY integrals (positive eigenvalues shown using a solid line and negative using circles), as well as the RPY-based slender body theory (SBT) of (2.24) (positive eigenvalues shown using a dotted line and negative using squares). Left: the eigenvalues as a function of the number of collocation points N with $N_2 = 6$ fixed. A result for the RPY integrals (6.34) computed with oversampled quadrature is shown as a black line for $N = 60$. There are a few small negative eigenvalues (yellow circles) for $N = 45$ due to numerical errors. Right: eigenvalues for $k \geq 90$ for fixed $N = 45$ and changing N_2 (eigenvalues for $k < 90$ are visually indistinguishable for the lines plotted). Using $N_2 = 4$ and $N_2 = 6$ still gives some negative eigenvalues (circles), while $N_2 = 8$ gives all positive eigenvalues that accurately approximate the true eigenvalues (solid black line).

points, while the eigenvalues with the RPY slender-body quadrature remain positive up to $N \approx 35$. When we increase N_2 from 6 to 8 (right plot), we match the true eigenvalues of the RPY mobility almost perfectly, and all eigenvalues remain positive. We emphasize that the negative eigenvalues for the RPY quadrature are the result of numerical quadrature errors, while for SBT negative eigenvalues are inherent to the continuum operators. The discrepancy gets better as $\hat{\epsilon}$ decreases; for $\hat{\epsilon} = 10^{-3}$ there is close agreement in the eigenvalues between SBT and RPY for $N \lesssim 50$, and, even for $N_2 = 2$, there are no negative eigenvalues until $N \gtrsim 60$ for SBT quadrature and $N \gtrsim 80$ for RPY quadrature (results not shown).

6.3.2 Strong convergence of velocity

Before considering time-dependent problems, we examine the convergence of the solution to the static problem (4.2). The main object of study here is the constraint force density $\boldsymbol{\lambda}(s)$. Even though our discrete translational mobility \mathbf{M}_{tt} (with sufficiently large N_2) does not have negative eigenvalues (see Fig. 6.9), it still has very small positive eigenvalues and the resulting $\boldsymbol{\lambda}$ changes rapidly near the endpoints. Therefore, we cannot expect *pointwise* spectral convergence for the constraint force $\boldsymbol{\lambda}$.

The key result of this section is that, while the constraint force $\boldsymbol{\lambda}$ is not smooth, the velocity for $\hat{\epsilon} = 10^{-3}$ to 10^{-2} is sufficiently smooth to be resolved by the spectral method. Furthermore, while the constraint force does not converge pointwise at the fiber endpoints, moments of it, which are the physical observables in the problem (e.g., stress $\sim \int \boldsymbol{\lambda}(s)\mathbf{X}(s) ds$), do converge, and can be captured by the spectral method to reasonably high accuracy.

To define a reference solution, we utilize the second-order blob-link discretization described in Section 3.2.4.2. This discretization is more robust because the values of functions at the fiber endpoints do not affect derivatives at the fiber interior, which ensures that $\boldsymbol{\lambda}$ does not contain spurious oscillations in the fiber interior. Since the mobility in the second-order method converges to the integrals we compute in Section 6.2 as the number of blobs goes to

infinity, we will utilize Richardson extrapolation to form a reference solution and compute error with respect to that solution. Then, to compare the accuracy of our method with that of the second-order method, we also compute a solution using $1/\hat{\epsilon}$ points. According to [76], this is (approximately) the minimum required resolution such that the filament is treated hydrodynamically as a solid cylinder, rather than a series of disconnected blobs.² In the second-order method we compute the forces and torques due to bending and twisting analytically, thus focusing our investigation on the saddle point solve (4.2), and not the accuracy of computing the right hand side.

Let us begin by examining the function $\boldsymbol{\lambda}(s)$ in the solution of (4.2). For this test, we begin with the fiber (6.65) with $q = 1$, which is a smoother problem for which we could get high accuracy with on the order of 10 points (and enter the asymptotic regime of the second-order method). We perform Richardson extrapolation on the second-order solutions, and in Fig. 6.10 compare the result to the spectral method with $N = 40$ points. We consider two different values of N_2 in both cases. For $\hat{\epsilon} = 10^{-2}$, we start with $N_2 = 6$, for which we observe uncontrolled oscillations in $\boldsymbol{\lambda}$ that grow near the fiber endpoints. This is consistent with our observation in Fig. 6.9 that negative eigenvalues exist when N_2 is too small. When we increase to $N_2 = 8$, we again see oscillations at the endpoints, but the magnitude of these is closer to that of the reference constraint forces we obtain from the second-order method.

Quite surprisingly, despite the misbehavior of $\boldsymbol{\lambda}$, the translational velocity \boldsymbol{U} and parallel rotational velocity Ω^{\parallel} appear to converge pointwise, albeit with rapidly-changing behavior at the fiber endpoints (see Fig. 3.1, top left plot for an example). To quantify the errors in \boldsymbol{U} and Ω^{\parallel} , we perform a self-convergence study, in which the errors are computed by successive refinements, then confirm that the converged spectral solution is close to that of the second-order method (see Fig. 6.11 caption). The error as a function of N is shown in

²It might be possible to reduce the required resolution of the second-order method by designing second-order quadrature schemes for the integrals in the RPY mobility instead of evaluating direct sums of the RPY kernel. However, since our slender-body quadrature is based on *global* interpolants of the fiber centerline, it is not straightforward to extend these to second-order discretizations.

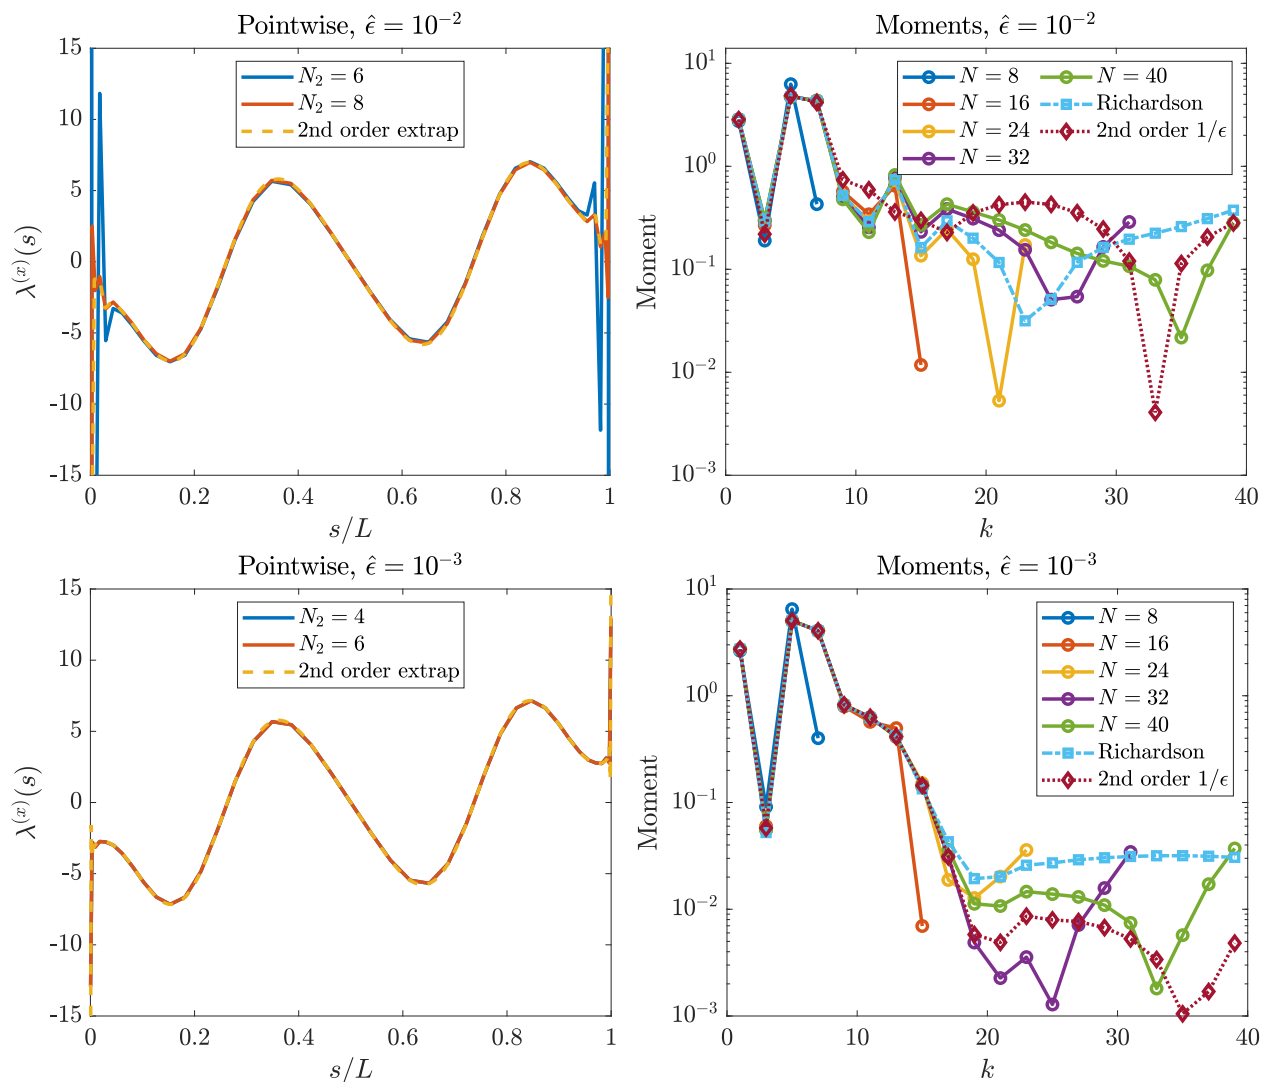


Figure 6.10: Comparison of the pointwise values (left) and moments (right) of λ in the spectral and second-order methods using the fiber (6.65) with $q = 1$ and slenderness (top) $\hat{\epsilon} = 10^{-2}$ and (bottom) $\hat{\epsilon} = 10^{-3}$. In the panels on the left, we show the x component, which is the least smooth, for $N = 40$ and various N_2 , and compare to the reference solution obtained by Richardson extrapolation of the second-order discretization. In the panels on the right, we show the magnitude of the moments of $\lambda^{(x)}(s)$ against the Chebyshev polynomials $T_k(s)$ with $N_2 = 8$ for $\hat{\epsilon} = 10^{-2}$ and $N_2 = 4$ for $\hat{\epsilon} = 10^{-3}$. The spectral method with $N \approx 16$ gives errors in the moments of λ that are smaller than or comparable to the second-order method with $1/\hat{\epsilon}$ points for both values of $\hat{\epsilon}$.




Figure 6.11: Convergence of the translational and (parallel) rotational velocities for the static Euler model equations (4.2) with $\hat{\epsilon} = 10^{-2}$ ($N_2 = 8$ and $N_2 = 12$, blue) and $\hat{\epsilon} = 10^{-3}$ ($N_2 = 4$, red). For each $\hat{\epsilon}$, the solid lines (and circles) show the self-convergence of the spectral method (error is the relative L^2 difference between the solution with N points and the solution with $N + 8$ points). The dotted (for $\hat{\epsilon} = 10^{-2}$) and dashed (for $\hat{\epsilon} = 10^{-3}$) show the error in the second-order method using $1/\hat{\epsilon}$ points, relative to Richardson extrapolation. For \mathbf{U} , the difference between the Richardson-extrapolated solutions and spectral solutions with $N = 40$ is 6.6×10^{-4} for $\hat{\epsilon} = 10^{-2}$ and 0.002 for $\hat{\epsilon} = 10^{-3}$. For Ω^{\parallel} , the Richardson-extrapolated solution differs from the spectral solution with $N = 40$ by 0.02 for $\hat{\epsilon} = 10^{-2}$ and 0.01 for $\hat{\epsilon} = 10^{-3}$. Here Ω^{\parallel} is computed on an N point grid by applying \mathbf{R}_{ψ} to (6.21).

Fig. 6.11 for both $\hat{\epsilon} = 10^{-2}$ and $\hat{\epsilon} = 10^{-3}$. We obtain about 3 digits of accuracy in both velocities when $N = 32$ and $\hat{\epsilon} = 10^{-2}$. When we increase to $N = 40$, the endpoints become more resolved, and the error in the translational velocity using $N_2 = 8$ points stagnates. By increasing to $N_2 = 12$, we obtain a smaller error when $N = 40$ (dashed-dotted squares in Fig. 6.11). In theory, this tells us that N_2 should increase with N , especially for large $\hat{\epsilon}$, which is when the contributions from the integrals on $|s - s'| < 2\hat{a}$ are nontrivial. However, for $\hat{\epsilon} = 10^{-2}$, using $N \gtrsim 50$ makes the number of degrees of freedom in the spectral method comparable to the second-order one with $1/\hat{\epsilon} = 100$ points, and in that case we might as well use the more robust low-order method.

When we decrease the slenderness to $\hat{\epsilon} = 10^{-3}$, Fig. 6.11 shows that we obtain between 2–3 digits in both \mathbf{U} and Ω^{\parallel} when $N = 40$ (the errors saturate at $N \approx 64$). The error in this case is larger than for $\hat{\epsilon} = 10^{-2}$ since the problem is less smooth at the fiber endpoints. While the convergence is slow due to this lack of smoothness, our spectral method still represents a great improvement over the lower-order discretization with $1/\hat{\epsilon}$ points, as we show in Fig. 6.11 with dashed-dotted lines. Thus, if we define the limit of infinitely many blobs as the reference solution, the spectral method provides a cheap, efficient way to approximate that result. In particular, we can obtain the same accuracy with 20 Chebyshev points on the fiber as we do with $1/\hat{\epsilon}$ blobs. Unsurprisingly, increasing the fiber shape to $q = 7$ in (6.65) makes the required number of Chebyshev points larger; in that case we find that 40 points is sufficient to give a translational velocity with lower error than $1/\hat{\epsilon}$ blobs, while about 24 points is sufficient for angular velocity (not shown). The translational, but not angular velocity, requires more points as curvature increases because \mathcal{M}_{tt} is the most nonlocal operator.

That said, it is still clear from the left panel in Fig. 6.10 that $\boldsymbol{\lambda}$ at the endpoints is not converging *pointwise*. Indeed, even the second-order method shows large jumps in $\boldsymbol{\lambda}$ near the endpoints, which suggests the problem is with the model physics, and not the spectral numerical method. Nevertheless, we see from Fig. 6.10(b) that there is still a surprisingly

good match between the spectral and second-order solutions for $\boldsymbol{\lambda}$ at the endpoints for $\hat{\epsilon} = 10^{-3}$, despite the lack of smoothness in $\boldsymbol{\lambda}$. We also wish to emphasize that, while $\boldsymbol{\lambda}$ is not smooth at the endpoints, when multiplied by \mathbf{M}_{tt} it gives *smooth enough* (to resolve with a spectral method to 2–3 digits) contributions to the translational and rotational velocities, which are the important quantities for the evolution of the fiber centerline. The smooth velocity maintains the smooth fiber shape, which makes our spectral method viable.

6.3.3 Weak convergence of $\boldsymbol{\lambda}$

Given that $\boldsymbol{\lambda}$ is not smooth at the endpoints, we do not expect pointwise convergence of $\boldsymbol{\lambda}$, and we cannot a priori expect spectral accuracy from the Chebyshev collocation discretization. We can, however, hope for *weak* convergence of $\boldsymbol{\lambda}$, which we study by computing the moments $\int_0^L \boldsymbol{\lambda}(s)T_k(s) ds$ for increasing k .

We study this in the following way: for each N , we solve for $\boldsymbol{\lambda}$ and compute the integral against the Chebyshev polynomial $T_k(s)$ for each k on a fine, upsampled grid. We do the same in the second-order method, except we compute the integral directly on the second-order grid without any upsampling. We use Richardson extrapolation of the second-order solutions to get a reference solution for the moments, then compare to the spectral moments in Fig. 6.10. The goal for our spectral discretization is then to obtain the first N moments of $\boldsymbol{\lambda}$ with greater accuracy than the second-order method with $1/\hat{\epsilon}$ points.

If we define the extrapolated moments from the second-order moments as the “true” moments of $\boldsymbol{\lambda}$, we see that the spectral method performs better than the second-order method with $1/\hat{\epsilon}$ points in approximating those moments, especially for $\hat{\epsilon} = 10^{-2}$. For $\hat{\epsilon} = 10^{-2}$, we obtain about 2 – 3 digits of accuracy in the first 20 moments using $N \geq 24$, whereas the second-order discretization with 100 points gives only 1 – 2 digits in each moment, which is comparable (but worse) than the accuracy of the spectral method with only 16 points. For $\hat{\epsilon} = 10^{-3}$, the error in the moments is comparable for $N \geq 16$ in the spectral method

and the second-order code with $1/\hat{\epsilon} = 10^3$ points. For both values of $\hat{\epsilon}$, the moments of the spectral approximation to $\boldsymbol{\lambda}$ start to increase for larger N when $k \gtrsim 24$. This comes from the non-smoothness at the endpoints, but has little effect in practice on the accuracy of physical quantities of interest like stress.

6.4 Temporal discretization

We now develop a temporal integrator for the Euler model (4.2)–(4.5) on a single fiber (Chapter 7 will cover integrators for multiple filaments). Our temporal integrator is based on a linearly-implicit, first-order, backward Euler discretization of the equations. While using higher-order schemes is certainly possible [135, 136], this is complicated by the fact that any scheme we use must be L -stable since twist equilibrates on a much faster timescale than bending, and the higher-order bending modes also equilibrate on a fast timescale. Since \mathbf{M}_{tt} is a nonlinear function of \mathbf{X} , linearizing the mobility as in [136, Eq. (25)] does not guarantee stability, which is why prior studies [47] have used nonlinearly implicit BDF formulas to obtain higher-order accuracy. Designing a temporal integrator for twist suitable for dense suspensions of fibers is a question we leave open for future work; here we only consider a single fiber to first order in time.

At the k th time step, we solve the *linear* system

$$\begin{aligned} \mathbf{K}^k \boldsymbol{\alpha}^k &= \mathbf{M}_{\text{tt}}^k \left(\mathbf{f}^{(\kappa)}(\mathbf{X}^{k+1,*}) + \mathbf{f}^{(\gamma)}(\mathbf{X}^k, \psi^k) + \boldsymbol{\lambda}^k \right) + \mathbf{M}_{\text{rt}}^k n^{\parallel}(\psi^k), \\ (\mathbf{K}^*)^k \boldsymbol{\lambda}^k &= \mathbf{0}, \end{aligned} \quad (6.66)$$

where $\mathbf{K}^k = \mathbf{K}(\mathbf{X}^k)$, and likewise for all other matrices. Here $\mathbf{X}^{k+1,*} = \mathbf{X}^k + \Delta t \mathbf{K}^k \boldsymbol{\alpha}^k$ is a linear approximation to the position at the next time step (the actual position is obtained by rotating the tangent vectors; see Section 6.4.1). The elastic and twist forces are obtained using (6.19) and (6.24), respectively, and the parallel torque is calculated from ψ using

(6.23). Substituting the formula for $\mathbf{X}^{k+1,*}$ and the discretization of $\mathbf{f}^{(\kappa)}$ from (6.19) into (6.66) gives the saddle point system

$$\begin{aligned} & \begin{pmatrix} -M_{\text{tt}}^k & \mathbf{K}^k - \Delta t M_{\text{tt}}^k \mathbf{F}_\kappa \mathbf{K}^k \\ (\mathbf{K}^*)^k & \mathbf{0} \end{pmatrix} \begin{pmatrix} \boldsymbol{\lambda}^k \\ \boldsymbol{\alpha}^k \end{pmatrix} \\ &= \begin{pmatrix} M_{\text{tt}}^k \left(\mathbf{f}^{(\kappa)}(\mathbf{X}^k) + \mathbf{f}^{(\gamma)}(\mathbf{X}^k, \psi^k) \right) + M_{\text{rt}}^k n^\parallel(\psi^k) \\ \mathbf{0} \end{pmatrix} \end{aligned} \quad (6.67)$$

which we solve for $\boldsymbol{\alpha}^k$ and $\boldsymbol{\lambda}^k$ using dense pseudoinverses via the Schur complement.

Solving (6.67) using the psuedo-inverse gives the *perpendicular* rate of rotation $\boldsymbol{\Omega}^{\perp,k}$ due to the forcing and torque. The parallel rate of rotation from the force can then be computed via (4.4),

$$\Omega_{\text{f}}^{\parallel,k} = \mathbf{R}_{N_x \rightarrow N_\psi} M_{\text{rt}}^k \left(\mathbf{f}^{(\kappa)}(\mathbf{X}^{k+1,*}) + \mathbf{f}^{(\gamma)}(\mathbf{X}^k, \psi^k) + \boldsymbol{\lambda}^k \right). \quad (6.68)$$

The parallel rate of rotation due to torque is then treated implicitly via a backward Euler discretization of (4.5), so that the evolution equation for the twist density $\psi = \partial_s \theta$ is given by

$$\psi^{k+1} = \psi^k + \Delta t \left(-\mathbf{R}_{N \rightarrow N_\psi} (\boldsymbol{\Omega}^{\perp,k} \cdot \mathbf{D}_N \boldsymbol{\tau}^k) + \mathbf{D}_{N_\psi} \Omega_{\text{f}}^{\parallel,k} + \partial_s \Omega_{\text{n}}^\parallel(\psi^{k+1}) \right) \quad (6.69)$$

To compute $\partial_s \Omega_{\text{n}}^\parallel$, we differentiate (6.21) on the $N_\psi + 2$ point grid, then downsample by applying \mathbf{R}_ψ . This converts (6.69) into the linear system

$$\begin{aligned} & \left(\mathbf{I} - \Delta t \mathbf{R}_\psi \mathbf{D}_{N_\psi+2} \mathbf{M}_{\text{rr}} \widetilde{\mathbf{N}}_\psi \right) \psi^{k+1} \\ &= \psi^k + \Delta t \left(-\mathbf{R}_{N \rightarrow N_\psi} (\boldsymbol{\Omega}^{\perp,k} \cdot \mathbf{D}_N \boldsymbol{\tau}^k) + \mathbf{D}_{N_\psi} \Omega_{\text{f}}^{\parallel,k} + \mathbf{R}_\psi \mathbf{D}_{N_\psi+2} \mathbf{M}_{\text{rr}} \widetilde{\boldsymbol{\beta}}_n \right), \end{aligned} \quad (6.70)$$

where $\widetilde{\mathbf{N}}_\psi$ and \mathbf{R}_ψ are defined in (6.20) and (6.23).

6.4.1 Updating $\boldsymbol{\tau}$ and \mathbf{X}

While the evolution (6.25) describes how the tangent vectors evolve in continuous time, in discrete time the update $\boldsymbol{\tau}^{k+1} = \boldsymbol{\tau}^k - \Delta t \mathbf{C} \boldsymbol{\Omega}$ does *not* preserve the unit-length constraint. Thus, in order to keep the dynamics on the constraint in discrete time, we will use $\boldsymbol{\Omega}^k$ from the solution of (6.67) (recall that $\boldsymbol{\alpha} = (\boldsymbol{\Omega}, \mathbf{U}_{\text{MP}})$), then evolve each tangent vector $\boldsymbol{\tau}_{\{p\}}^k$ by rotating by $\boldsymbol{\Omega}_{\{p\}}^k$,

$$\boldsymbol{\tau}_{\{p\}}^{k+1} = \text{rotate}(\boldsymbol{\tau}_{\{p\}}^k, \boldsymbol{\Omega}_{\{p\}}^k \Delta t), \quad (6.71)$$

the explicit formula for this rotation is the Rodrigues rotation formula [137]

$$\begin{aligned} \boldsymbol{\tau}_p^{k+1} = & \boldsymbol{\tau}_p^k \cos(\Omega_p^k \Delta t) + \left(\widehat{\boldsymbol{\Omega}}_p^k \times \boldsymbol{\tau}_p^k \right) \sin(\Omega_p^k \Delta t) + \\ & \widehat{\boldsymbol{\Omega}}_p^k \left(\widehat{\boldsymbol{\Omega}}_p^k \cdot \boldsymbol{\tau}_p^k \right) (1 - \cos(\Omega_p^k \Delta t)), \end{aligned} \quad (6.72)$$

where $\Omega = \|\boldsymbol{\Omega}\|$ and $\widehat{\boldsymbol{\Omega}} = \boldsymbol{\Omega}/\Omega$. We then update the midpoint \mathbf{X}_{MP} via $\mathbf{X}_{\text{MP}}^{n+1} = \mathbf{X}_{\text{MP}}^n + \Delta t \mathbf{U}_{\text{MP}}^n$, and compute \mathbf{X}^{n+1} from $\boldsymbol{\tau}^{n+1}$ and $\mathbf{X}_{\text{MP}}^{n+1}$ by applying the mapping $\boldsymbol{\mathcal{X}}$ defined in (6.3). Rotating the tangent vectors in this way keeps the dynamics on the constraint without needing to introduce penalty parameters [22] that lead to additional stiffness in temporal integration.

6.4.2 Updating the Bishop frame

In the case of intrinsic curvature and twist, the material frame is involved in the forces and torques, and so it must be computed from the Bishop frame and θ using (3.27). Computing the Bishop frame requires the solution of the Bishop ODE (3.26), which is subject to the

condition (4.7),

$$\partial_s \mathbf{b}^{(1)} = (\boldsymbol{\tau} \times \partial_s \boldsymbol{\tau}) \times \mathbf{b}^{(1)} \quad (6.73)$$

$$\mathbf{b}^{(1)}(L/2) = \mathbf{D}^{(1)}(L/2).$$

Thus in our temporal integration scheme, we will obtain $\boldsymbol{\tau}^{k+1}$ via (6.72), and ψ^{k+1} via (6.70). We then integrate ψ^{k+1} to obtain θ^{k+1} , setting $\theta(L/2) = 0$. With θ^{k+1} and $\left(\mathbf{D}^{(1)}(L/2)\right)^{k+1} = \text{rotate}\left(\left(\mathbf{D}^{(1)}(L/2)\right)^k, \Delta t (\boldsymbol{\Omega}^\perp(L/2) + \boldsymbol{\Omega}^\parallel(L/2)\boldsymbol{\tau})^k\right)$, we can then solve (6.73) numerically as follows.

To discretize the Bishop ODE (6.73), we first integrate both sides of (6.73) to obtain an integral equation for $\mathbf{b}^{(1)}$. Since each fiber has at most 40 points on its centerline, we will use dense linear algebra to solve for $\mathbf{b}^{(1)}$ at all of the collocation points. The discrete integral form of (6.73) is

$$\mathbf{b}^{(1)} + \mathbf{c} = \mathbf{D}_N^\dagger (\boldsymbol{\tau} \times \partial_s \boldsymbol{\tau}) \times \mathbf{b}^{(1)}. \quad (6.74)$$

where \mathbf{c} is an unknown integration constant and \mathbf{D}_N^\dagger is the pseudo-inverse of the Chebyshev differentiation matrix on the grid of size N (tangent vector grid). Combining with the boundary condition in (6.73), and denoting with $\mathbf{I}(s)$ evaluation at s , we get the system

$$\begin{pmatrix} \mathbf{I} - \mathbf{D}^\dagger ((\boldsymbol{\tau} \times \partial_s \boldsymbol{\tau}) \times) & \mathbf{I} \\ \mathbf{I}(L/2) & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{b}^{(1)} \\ \mathbf{c} \end{pmatrix} = \begin{pmatrix} \mathbf{0} \\ \mathbf{D}^{(1)}(L/2) \end{pmatrix}, \quad (6.75)$$

which we solve for $\mathbf{b}^{(1)}$.

In Fig. 6.12, we plot the errors in $\mathbf{b}^{(1)}(s)$ as a function of the grid size for the fiber (6.65) with $q = 7$ that satisfies the free fiber BCs (the refined solution is generated with $N = 100$, in which all vectors have unit length to tolerance 10^{-13}). We see that the errors in the parallel and perpendicular direction are the same, which means that choosing to renormalize $\mathbf{b}^{(1)}$ will not increase the overall error (in Fig. 6.12 we do not renormalize $\mathbf{b}^{(1)}$).




Figure 6.12: The L^2 function errors in the Bishop vector $\mathbf{b}^{(1)}(s)$. We use the fiber (6.65) with $q = 7$ and generate a reference solution using 100 collocation points (these vectors have unit length to 13 digits). We then solve (6.75) using N points and measure the L^2 function error in $\mathbf{b}^{(1)}(s)$ on a common grid. We separate the total error (blue circles) into a component parallel to $\mathbf{b}^{(1)}(s)$ (red squares) and a component perpendicular to $\mathbf{b}^{(1)}(s)$ (yellow diamonds).

6.5 Results

We now run our temporal integrator on two examples. In the first example, a relaxing bent fiber, twist elasticity contributes minimally to the dynamics. In the second, a clamped spinning fiber, twist elasticity is essential to generate the correct “overwhirling” [100] dynamics.

6.5.1 Role of twist in the relaxation of a bent filament

The first example we consider is a relaxing fiber that is initially bent. Previous works [15, 84], have neglected the twist elasticity of free filaments, with the justification that twist equilibrates much faster than bend, and therefore is always in a quasi-equilibrium state as the filament deforms. Our goal here is to test this assumption by simulating the relaxation of a filament with and without twist elasticity.

The fiber has length $L = 2$ and the initial tangent vector is given by (6.65) with $q = 7$. We start the fiber with no twist, $\psi(s, t = 0) = 0$, and simulate until $t = 0.01$ with $\mu = 1$,




Figure 6.13: L^2 error in the final position of a relaxing bent filament when we do not consider twist elasticity. We measure the error as a function of γ and consider $\hat{\epsilon} = 10^{-2}$ (blue circles), $\hat{\epsilon} = 10^{-2}$ without rot-trans coupling (red squares), and $\hat{\epsilon} = 10^{-3}$ (yellow diamonds).

$\kappa = 1$, and $\gamma = 0$. This is long enough for the fiber to make a non-negligible shape change, as shown in Fig. 6.8. We then repeat the same simulations using a varying twist modulus γ and in Fig. 6.13 report the L^2 error in the fiber position when twist elasticity is ignored (equivalently, $\gamma = 0$). We report results using $N = 36$ and $\Delta t = 10^{-4}$ for $\hat{\epsilon} = 10^{-2}$ and $N = 40$ and $\Delta t = 1.25 \times 10^{-5}$ for $\hat{\epsilon} = 10^{-3}$, having verified that increasing N and decreasing Δt makes little difference in the error curves shown in Fig. 6.13.

Figure 6.13 shows how the position of the filament changes when we add twist elasticity. We recall that the timescale of twist equilibration is $1/\gamma$, and the forcing that results from twist is proportional to γ . Thus, when γ is small, the timescale of twist equilibration is larger than that considered here, and there is an $\mathcal{O}(\gamma)$ correction to the position of the fiber. As γ increases, the $1/\gamma$ timescale goes to zero, but the contribution of $\mathbf{f}^{(\gamma)} \propto \gamma$ goes to infinity, so in total there is a constant difference in the position which scales as $\hat{\epsilon}^2$ in the limit as $\gamma \rightarrow \infty$ (in which twist is always in quasi-equilibrium).

The scaling of the curves with $\hat{\epsilon}$ suggests that the rot-rot mobility controls the equilibration of twist, which feeds back to the fiber shape through the force $\mathbf{f}^{(\gamma)}$ (see (4.2)) and

trans-trans mobility. Indeed, as shown in Fig. 6.13, dropping all rot-trans coupling from the dynamics, so that the *only* coupling between the position and twist angle comes through the term $\mathbf{M}_{\text{tt}}\mathbf{f}^{(\gamma)}$ in (4.2), gives the same behavior for $\hat{\epsilon} = 0.01$. Thus, in the case of a relaxing fiber, the rot-trans coupling is negligible. Furthermore, in most materials, $\gamma/\kappa \approx 1$ [33], which according to Fig. 6.13 falls within the plateau regime of equilibrated twist, especially for $\hat{\epsilon} \lesssim 10^{-3}$.

6.5.2 Twirling to whirling instability

Having considered an example for which twist elasticity is negligible, we now consider an example in which it is vital: the instability in a twirling fiber [54, 99, 100, 102]. In this case, the instability results when the torque due to twist becomes larger than the filament buckling torque, and the filament transitions from a straight twirling state to a curved whirling state [99].

To initialize a small perturbation of a straight fiber, we set

$$\boldsymbol{\tau}(s) = \frac{1}{\sqrt{1+\delta^2}} \begin{pmatrix} \cos(\tan^{-1}(\delta)) & -\sin(\tan^{-1}(\delta)) & 0 \\ \sin(\tan^{-1}(\delta)) & \cos(\tan^{-1}(\delta)) & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \delta \cos(s(s-L)^3) \\ 1 \\ \delta \sin(s(s-L)^3) \end{pmatrix} \quad (6.76)$$

at $t = 0$. Integrating (6.76) numerically (with the pseudo-inverse of the Chebyshev differentiation matrix) and setting $\mathbf{X}(0) = \mathbf{0}$ then gives a filament that satisfies free boundary conditions at $s = L$ and the clamped boundary conditions $\mathbf{X}(0) = \mathbf{0}$ and $\boldsymbol{\tau}(0) = (0, 1, 0)$. At the clamped end, there is a delta-like singularity in the perpendicular component of $\boldsymbol{\lambda}(s)$, which enforces the clamped boundary condition. As a result of this, the velocity has a boundary layer that requires more collocation points to resolve for smaller ϵ . We will therefore use $\hat{\epsilon} = 10^{-2}$ throughout this section, in addition to $\delta = 0.01$, $L = 2$, $\mu = 1$, $\kappa = 1$, and $\gamma = 1$.

To trigger the instability, we prescribe the rate of twist at the $s = 0$ end as $\Omega^{\parallel}(0) = \omega$

(see Section 6.1.3.3). Note that we could also simulate with the constant torque BC (6.22), in which case the applied torque N_0 is equal to the torque required to spin a straight fiber at rate ω ,

$$N_0 := - \int_0^L \frac{\omega}{m_{\text{rr}}(s)} ds, \quad (6.77)$$

where m_{rr} is the local rot-rot mobility coefficient, $\Omega_{\text{n}}^{\parallel}(s) = m_{\text{rr}}(s)n^{\parallel}(s)$. The results in this section are unchanged regardless of the BC we use.

We start the twist profile in a state that satisfies the boundary conditions with no rot-trans coupling, $\psi(s) = \omega/\gamma \int_L^s m_{\text{rr}}^{-1}(s') ds'$. In the case of an ellipsoidally-tapered filament we use $m_{\text{rr}}(s) \equiv 9/(32\pi\mu\hat{a}^2)$, and the steady state twist profile reduces to $\psi(s) = \omega(s - L)/(\psi m_{\text{rr}})$ [99]. When the fiber is cylindrical, we use the formulas in (2.47) to obtain $m_{\text{rr}}(s)$.

In the case of an ellipsoidal fiber with local drag and no rot-trans coupling, Wolgemuth et. al [99] performed a linear stability analysis to show that the spinning of the fiber about its axis is unstable at critical frequency [33, Eq. (74)]

$$\omega_c^{(\text{ELD})} \approx 8.9 \frac{m_{\text{rr}}\kappa}{L^2} = 8.9 \frac{9\kappa}{32\pi\mu L^2 \hat{a}^2} \quad (6.78)$$

with the centerline oscillating at a significantly smaller frequency [33, Eq. (75)]

$$\chi_c^{(\text{ELD})} \approx 22.9 \frac{m_{\text{tt}}^{\perp}\kappa}{L^4} = 22.9 \frac{\kappa (\log(\hat{\epsilon}^{-2}/16) + 4)}{8\pi\mu L^4}. \quad (6.79)$$

Here m_{tt}^{\perp} is the mobility coefficient for force perpendicular to the fiber centerline for an ellipsoidally-tapered filament. We will therefore report frequency in units $\bar{\omega} = \omega/\omega_c^{(\text{ELD})}$ and time in units $\bar{t} = t\chi_c^{(\text{ELD})}/(2\pi)$.

Since the publication of [99], there have been a number of studies on the whirling instability that leave a few questions open. A study by Lim and Peskin, for instance [100], found

a critical frequency of $\bar{\omega}_c = 0.3$, while recent experimental work [102] has put the critical frequency at about $\bar{\omega}_c = 0.6$. The discrepancy in Lim and Peskin’s work [100] could be due to a number of factors, including their use of nonlinear fluid equations, a periodic domain, and a larger initial deflection. It is not, however, due to nonlocal trans-trans interactions, as Wada and Netz [54] included those and obtained $\bar{\omega}_c = 1$. Wada and Netz also showed, however, that the frequency required to induce the overwhirling state drops as the initial perturbation increases (and the nonlinear regime is entered). It could be possible, therefore, that the simulations of [100] fall in this nonlinear regime. This observation was used in [102] to explain the deviations from the theory as well.

An alternative explanation for the different values of ω_c is the influence of rot-trans coupling. This is neglected in computational studies which use the RPY tensor or SBT [54, 99], but is included by necessity in simulations that use the IB method [100], and of course in experimental systems [102]. Here we will study the role of rot-trans coupling to see if we can understand the discrepancies in the literature. We will approach the problem in two steps. First, we will confirm the relationships (6.78) and (6.79) hold for an ellipsoidal filament with local drag and a cylindrical filament with all hydrodynamic terms included except rot-trans. Then, we will examine the influence of the rot-trans coupling in the linear regime.

6.5.2.1 Critical frequency as a function of hydrodynamic model

Figure 6.14 shows our study of the behavior close to the critical frequency ω_c . In the top panel, we simulate at $\bar{\omega} = 1$ with four different hydrodynamic models: local drag with ellipsoidal fibers (no rot-trans coupling), local drag with cylindrical fibers (endpoint formulas in (2.20), no rot-trans coupling), and the RPY integral mobility with and without rot-trans coupling. We see that using ellipsoidal local drag gives a trajectory which neither grows nor decays over five periods, and the period is exactly equal to that predicted from (6.79).




Figure 6.14: Simulating the whirling instability in the linear regime, $\delta = 0.01$ in (6.76). We show the x coordinate of the $s = L$ endpoint as a function of dimensionless time $t^* = t\chi_c^{(\text{ELD})}/(2\pi)$ for $\hat{\epsilon} = 10^{-2}$ (results for $\hat{\epsilon} = 10^{-3}$ are very similar, but require larger N to capture the changes in velocity near the endpoints, and are not shown). We consider ellipsoidal (blue) and cylindrical (dashed red) fibers with local drag, and cylindrical fibers (dashed-dotted yellow), where in all three of these cases we neglect rot-trans and trans-rot coupling. We then add rot-trans and trans-rot coupling and show the results as dotted purple lines. At the top, we show the endpoint x coordinate at $\omega = \omega_c^{(\text{ELD})}$ using 1000 time steps per period $2\pi/\chi_c^{(\text{ELD})}$, $N = 16$ for ellipsoidal fibers, and $N = 32$ for all other fibers. At bottom, we tune $\bar{\omega} = \omega/\omega_c^{(\text{ELD})}$ to be as close as possible to the frequency at which the perturbation neither grows nor decays (each line represents the same mobility as the top plot), and use 2000 time steps per period with $N = 40$ for all fibers.

When we switch to the cylindrical local drag formulas, the mobility at the fiber endpoints decreases, and so the period $T_c \sim \chi_c^{-1} \sim m_{tt}^{-1}$ increases (see (6.79)). Likewise, the rot-rot mobility m_{tr} at the fiber endpoints is also smaller when we use cylindrical fibers, and so the critical frequency ω_c ought to be smaller (see (6.78)). As shown in Fig. 6.14(b), this is indeed the case, with the critical frequency for cylindrical filaments coming in lower than that for ellipsoidal filaments, but with a small difference of only 2%. The same can be said for the RPY mobility without rot-trans coupling, since in this case the difference in the critical frequencies is at most 4%. Thus, nonlocal trans-trans hydrodynamics, as well as the shape of the fiber radius function, make little difference for the critical frequency, and cannot explain the previously-observed large deviations from $\omega_c^{(\text{ELD})}$.

It is only when we account for rot-trans coupling that we see a substantial difference in the critical frequency. Indeed, as shown in Fig. 6.14, including trans-rot dynamics gives a large growth rate of the oscillations in the fiber endpoints, which eventually leads to the “overwhirling” behavior that has been documented previously [33, 100] (see Fig. 6.15, but note that the same overwhirling behavior can be obtained without rot-trans and trans-rot dynamics for $\omega > \omega_c$ [54]). Figure 6.14 shows that the critical frequency with rot-trans coupling is about $\bar{\omega}_c = 0.8$, which explains about half of the observed experimental value of 0.6. The critical frequency of $\bar{\omega}_c = 0.8$ is also unchanged within 2% when we place the clamped end $2\hat{a} = 0.04$ above a bottom wall and use the wall-corrected RPY mobility derived in [138], which we discretize using oversampled integrals at every collocation point. Thus, confinement, at least in the direction perpendicular to the centerline, is not the cause of the other 20%, since the wall corrections decay too fast away from the wall to have a noticeable effect (the total effect is $\mathcal{O}(\hat{\epsilon})$ [139, Sec. 2.3–4]).

The rest of the deviation could be due to the initial fiber shape, as is suggested in [102]; however, our observations are the same up to a deflection of $\delta \approx 0.2$ in (6.76), and Wada and Netz showed that starting the filament in a circular configuration gives about 25% reduction




Figure 6.15: Nonlinear overwhirling dynamics: Fiber trajectory when $\omega = \omega_c^{(\text{ELD})} = 497$ rad/s $\approx 1.25\omega_c$ and we include rot-trans coupling in the RPY mobility (dotted purple curve in Fig. 6.14). The fiber endpoint drops below $y = 0$ and begins periodic crankshafting motion around $\bar{t} \approx 2.5$.

in ω_c [54, Fig. 1(g)]. Seeing as the initial fiber shapes in [102] are far more straight than circular, it seems unlikely that a reduction as large as 20% comes from the fiber shape. It is possible that the confinement inside a PVC pipe could reduce the critical frequency, although the authors of [102] say that the flow is fully developed around the filament. Other possibilities include the experimental set-up; for instance, the supplementary videos of [102] show that the shaft also slightly translates the fiber endpoint in addition to spinning it. This provides additional perturbations that could lower the critical frequency.


6.5.2.2 Nonlinear overwhirling dynamics

As ω increases, the twisting torque overcomes the bending moment of the filament, and the filament dynamics eventually reach a steady state in which the fiber takes the shape of a shepherd's crook and spins in a steady motion, as illustrated in Fig. 6.15. These dynamics, which are called “overwhirling” [100], have been studied in more detail in [54, 101]. In particular, it was shown in [101] that there is a subcritical Hopf bifurcation: if ω is


increased from below, the overwhirling frequency is $\omega = \omega_c$, but if ω is decreased once the fiber is already in the overwhirling state, overwhirling is maintained for $\omega < \omega_c$. In addition to determining the boundaries of the bistable region, in this section we also study the amplitude and frequency of the periodic overwhirling motion as a function of the motor twirling frequency ω .

To study the stability and behavior of overwhirling at various frequencies, we begin by getting the fiber to the steady overwhirling state using $\omega = 1.1\omega_c$ (by steady, we mean the crankshafting frequency is unchanged to 0.2% over the last 5 periods). Following this, we use this steady state configuration to initialize another simulation with $\omega = 1.2\omega_c$, then use that steady state to initialize $\omega = 1.3\omega_c$, and so forth. We do the same for decreasing frequency: starting from $\omega = 1.1\omega_c$, we decrease the frequency until the fiber comes out of the overwhirling state. As discussed in [101], this happens at a *lower* frequency than ω_c , which we find to be somewhere between $0.8\omega_c$ (stable overwhirling) and $0.72\omega_c$ (stable twirling). Thus, the bistable region in the bifurcation diagram is between $\omega/\omega_c = 0.8$ and $\omega/\omega_c = 1$, which is (approximately) consistent with the results of [101, Fig. 5], where the lower boundary appears to be $0.85 - 0.9\omega_c$.

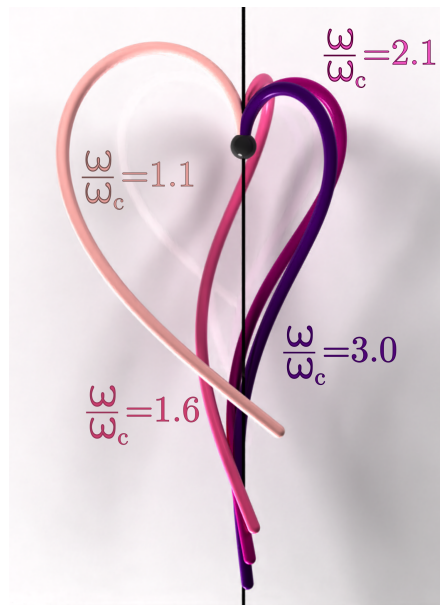
The overwhirling steady state dynamics are analyzed in Fig. 6.16. In Fig. 6.16(a), we show the trajectory of the $s = L$ endpoint when we use $\omega/\omega_c = 1.1$ and initialize the fiber to be nearly straight using (6.76). As depicted in Fig. 6.15, the endpoint dips below $y = 0$, and the whole fiber begins a steady crankshafting rigid rotation around the y axis, such that $Y(s)$ and $\sqrt{X(s)^2 + Z(s)^2}$ are constant for all s . In Fig. 6.16, we show how the amplitude $R = \sqrt{X(L)^2 + Z(L)^2}$ (Fig. 6.16(b)) and frequency χ (inset of Fig. 6.16(a)) of the crankshafting motion depend on the driving motor frequency. For $\omega/\omega_c \lesssim 1.6$, the amplitude contracts rapidly to about 5% of the fiber length, while the crankshafting frequency is roughly constant at $1 - 2\%$ of the motor frequency. On $\omega/\omega_c \gtrsim 1.6$, however, the fiber amplitude function is roughly constant; indeed, Fig. 6.16(c) shows that the steady state overwhirling



(a) Dynamics of fiber endpoint



(b) Steady state amplitude



(c) Steady state shapes

Figure 6.16: Steady state overwhirling dynamics. (a) Transition to the overwhirling state for $\omega = 1.1\omega_c$; see Fig. 6.15 for a 3D illustration for a different $\omega/\omega_c = 1.25$. The fiber endpoint dips below $y = 0$ and begins cranksafting motion with frequency $\chi(\omega)$. The top inset zooms in on the black box, while the bottom inset shows $\chi(\omega)$. (b) Steady state amplitudes for varying ω/ω_c . The left axis and black line shows the amplitude $R = \sqrt{X(s)^2 + Z(s)^2}$, while the right axis and red line show the steady state y coordinate. (c) Steady state overwhirling configurations for given ω/ω_c .

shapes approach a fiber that is essentially twirling upside down in this case. Instead of bringing the endpoint closer to the y axis, the extra work done by the motor is used to increase the frequency of the crankshafting motion at a higher rate, so that it transitions to about 10% of the motor frequency. The mean χ/ω of about 0.05 is similar to what Wada and Netz found for a thermally fluctuating fiber [54, Fig. 1(c)]. Our frequency results are also similar to those presented previously in [101, Fig. 6(a)], where nonlinear behavior is observed for $\omega/\omega_c \gtrsim 1.6$.

Interestingly, while we find that the overwhirling state can be stable for frequencies as large as $3\omega_c$, we also find that the ability to reach it depends on the initial fiber configuration. For instance, when ω/ω_c is as small as 2.1, the fiber intersects itself on its way to the overwhirling state. This suggests a mechanism for the formation of plectonemes that were observed at high ω in [102], and that the full nonlinear dynamics are likely more complex than previously thought.

Chapter 7

Free fibers with nonlocal hydrodynamics

The numerical tests in the previous chapter (in particular Section 6.5.1) showed that twist elasticity makes a negligible contribution to the dynamics of slender filaments with two free ends. Since these are precisely the type of filaments we are interested in moving forward, we will neglect twist elasticity in the rest of this dissertation. That is, the schemes we develop in this chapter solve the Euler model equations (4.2) with $\gamma = 0$ and only the \mathcal{M}_{tt} component of the mobility. While this is certainly a simpler problem than the full twist-bend elasticity problem, we make it more challenging by introducing multiple filaments into the system. In this case, we need *fast* (linear-scaling) and robust methods to compute the hydrodynamic interactions between filaments. That is the focus of this chapter.

We begin this chapter in Section 7.1 by reformulating the equations (4.2) in terms of the discrete *force* (not force density) on the fiber centerline, so that the mapping between fiber forces and velocities is symmetric positive definite. The question then becomes how to discretize it. Thus, in Section 7.2 we discuss how we can apply \mathcal{M}_{tt} in linear time using Ewald splitting [103]. Following this, we develop temporal integrators that reduce the number of nonlocal velocity evaluations, then test those integrators on three examples

of fiber suspensions. The main idea behind the temporal integrators, which only works in the deterministic context, is to time-lag the nonlocal parts of the mobility, so that only the local parts of the mobility need to be inverted. This saves on costly nonlocal hydrodynamic evaluations.

7.1 Reformulation in terms of force

If we no longer consider twist elasticity and restrict ourselves to clamped boundary conditions, then it becomes useful to formulate the equations in terms of *forces*, so that the mobility matrix $\widetilde{\mathbf{M}}$ mapping forces to velocities is symmetric positive definite (this will aid us in the next chapter when we introduce Brownian fluctuations). The force due to bending elasticity is given in (6.13), and can be written as $\mathbf{F} = -\mathbf{L}\mathbf{X}$

Let us now consider how the matrix $\widetilde{\mathbf{M}}$ is related to the matrix \mathbf{M} which maps force *densities* to forces. The work dissipated in the fluid, which is always positive, is given by

$$\mathcal{P} = \int_0^L \mathbf{f}(s) \cdot \mathbf{U}(s) ds = \mathbf{f}^T \widetilde{\mathbf{W}} \mathbf{U} = \mathbf{f}^T \widetilde{\mathbf{W}} \mathbf{M} \mathbf{f} = \mathbf{F}^T \mathbf{M} \widetilde{\mathbf{W}}^{-1} \mathbf{F}, \quad (7.1)$$

$$\mathcal{P} = \mathbf{F}^T \mathbf{U} = \mathbf{F}^T \widetilde{\mathbf{M}} \mathbf{F}.$$

The first line gives the work using the L^2 inner product on force density, and subsequent conversion via (6.14). The second line uses the definition of force $\langle \mathbf{f}, \mathbf{U} \rangle_{L^2} = \mathbf{F}^T \mathbf{U}$. Thus, we have that the SPD matrix $\widetilde{\mathbf{M}}$ is given in terms of \mathbf{M} by

$$\widetilde{\mathbf{M}} = \mathbf{M} \widetilde{\mathbf{W}}^{-1}. \quad (7.2)$$

Then, in a deterministic method, the rotation rates and midpoint velocities $\boldsymbol{\alpha} = (\boldsymbol{\Omega}, \mathbf{U}_{\text{MP}})^T$ are given by the solution of the saddle point system (4.2), which we reformulate in terms of




Figure 7.1: Summary of how we use the fiber configuration to compute the forces \mathbf{F} and mobility $\widetilde{\mathbf{M}}$. Given $(\boldsymbol{\tau}, \mathbf{X}_{\text{MP}})$, we apply the map $\boldsymbol{\chi}$ from (6.3) to obtain the Chebyshev polynomial $\mathbb{X}(s)$, which we then use to compute the bending energy $\mathcal{E}_{\text{bend}}$ in (6.9), force \mathbf{F} in (6.13), and force density \mathbf{f} in (6.14). The velocity $\widetilde{\mathbf{M}}\mathbf{F}$ is computed by applying special quadrature to the force density \mathbf{f} (because the quadrature matrix might have negative eigenvalues, we form $\widetilde{\mathbf{M}}$ explicitly to compute and truncate its eigenvalues).

force using (6.13) and (6.33)

$$\mathbf{K}\boldsymbol{\alpha} = \widetilde{\mathbf{M}}(-\mathbf{L}\mathbf{X} + \boldsymbol{\Lambda}) \quad (7.3)$$

$$\mathbf{K}^T \boldsymbol{\Lambda} = \mathbf{0},$$

where \mathbf{K} , $\widetilde{\mathbf{M}}$, \mathbf{L} and \mathbf{K}^T are all matrices of size $3N_x = 3(N + 1)$, and $\boldsymbol{\Lambda} = \widetilde{\mathbf{W}}\boldsymbol{\lambda}$ is the constraint force. Figure 7.1 summarizes the procedure we use to calculate forces and apply the mobility matrix $\widetilde{\mathbf{M}}$.

7.1.1 Saddle point system as gradient descent dynamics

We now show that the deterministic dynamics (7.3) take the form of a constrained gradient descent. Let us consider a backward Euler integrator for (7.3) and the Lagrangian (see (8.8))

for the Lagrangian with random forcing)

$$\begin{aligned} \mathcal{L}[\mathbf{X}, \boldsymbol{\alpha}, \boldsymbol{\lambda}] = & \frac{1}{2} \mathbf{X}^T \mathbf{L} \mathbf{X} + \frac{1}{2\Delta t} (\mathbf{X} - \mathbf{X}^n)^T \left(\widetilde{\mathbf{M}}^n \right)^{-1} (\mathbf{X} - \mathbf{X}^n) \\ & + \boldsymbol{\lambda}^T \widetilde{\mathbf{W}} \left(\mathbf{K}^n \boldsymbol{\alpha} - \frac{1}{\Delta t} (\mathbf{X} - \mathbf{X}^n) \right). \end{aligned} \quad (7.4)$$

Here the matrices \mathbf{K} and $\widetilde{\mathbf{M}}$ are evaluated at time n , and consequently are constant with respect to \mathbf{X} . In (7.4), the second term is $\Delta t/2$ times the rate of dissipation in the fluid (velocity \times force), and $\boldsymbol{\lambda}$ represents a force density, which is a Lagrange multiplier for the inextensibility constraint. We formulate that constraint as an L^2 inner product using the discrete weights matrix $\widetilde{\mathbf{W}}$, since this gives the work done by the constraint force density on the fluid.

We now differentiate the Lagrangian with respect to the three inputs to arrive at our equations of motion

$$\frac{\delta \mathcal{L}}{\delta \mathbf{X}} = \mathbf{0} \rightarrow -\mathbf{L} \mathbf{X} + \widetilde{\mathbf{W}} \boldsymbol{\lambda} = \frac{1}{\Delta t} \left(\widetilde{\mathbf{M}}^n \right)^{-1} (\mathbf{X} - \mathbf{X}^n) \quad (7.5)$$

$$\frac{\delta \mathcal{L}}{\delta \boldsymbol{\lambda}} = \mathbf{0} \rightarrow \mathbf{K}^n \boldsymbol{\alpha} = \frac{1}{\Delta t} (\mathbf{X} - \mathbf{X}^n) \quad (7.6)$$

$$\frac{\delta \mathcal{L}}{\delta \boldsymbol{\alpha}} = \mathbf{0} \rightarrow (\mathbf{K}^n)^T \widetilde{\mathbf{W}} \boldsymbol{\lambda} = \mathbf{0} \quad (7.7)$$

Combining (7.5) and (7.6) and taking $\Delta t \rightarrow 0$, we arrive at the deterministic saddle point system (7.3), where we set the force $\boldsymbol{\Lambda} = \widetilde{\mathbf{W}} \boldsymbol{\lambda}$, which represents another instance of the conversion between force and force density defined in (6.14). In this case it is easy to see how this conversion arises from the L^2 inner product. The second equation in (7.3) is the principle of virtual work introduced in Section 3.1.3.

7.1.1.1 Proximal gradient descent and boundary conditions

In continuum (in the spatial variable s), we can see why gradient descent on the Lagrangian (7.4) preserves the natural free fiber boundary conditions. In this regard, let us consider the simpler Lagrangian

$$\mathcal{L}[\mathbb{X}(\cdot)] = \frac{1}{2} \langle \partial_s^2 \mathbb{X}, \partial_s^2 \mathbb{X} \rangle + \frac{1}{2\Delta t} \langle \mathbb{X} - \mathbb{X}^n, \mathbb{X} - \mathbb{X}^n \rangle, \quad (7.8)$$

where the inner product is in L^2 . To differentiate this Lagrangian in continuum, we compute

$$\begin{aligned} \frac{\delta \mathcal{L}}{\delta \mathbb{X}}(\tilde{\mathbb{X}}) &= \left. \frac{d}{d\varepsilon} \mathcal{L}[\mathbb{X} + \varepsilon \tilde{\mathbb{X}}] \right|_{\varepsilon=0} \\ &= \langle \partial_s^2 \tilde{\mathbb{X}}, \partial_s^2 \mathbb{X} \rangle + \frac{1}{\Delta t} \langle \tilde{\mathbb{X}}, \mathbb{X} - \mathbb{X}^n \rangle \\ &= \left(\partial_s \tilde{\mathbb{X}} \cdot \partial_s^2 \mathbb{X} - \tilde{\mathbb{X}} \cdot \partial_s^3 \mathbb{X} \right) \Big|_0^L + \langle \tilde{\mathbb{X}}, \partial_s^4 \mathbb{X} \rangle + \frac{1}{\Delta t} \langle \tilde{\mathbb{X}}, \mathbb{X} - \mathbb{X}^n \rangle = 0 \end{aligned} \quad (7.9)$$

which must hold for all $\tilde{\mathbb{X}}$. This implies the backward Euler evolution $\mathbb{X} - \mathbb{X}^n = -\Delta t \partial_s^4 \mathbb{X}$, plus the free fiber boundary conditions $\partial_s^2 \mathbb{X}(0, L) = \mathbf{0}$ and $\partial_s^3 \mathbb{X}(0, L) = \mathbf{0}$. Thus, in continuum, proximal gradient descent automatically enforces the free fiber boundary conditions.

7.2 Hydrodynamic interactions in linear time

Now that we have a formulation of the equations in terms of force, the question is how to apply the mobility matrix $\tilde{\mathbf{M}} = \mathbf{M}\tilde{\mathbf{W}}^{-1}$. The application of this matrix ultimately reduces to discretizing integrals of the RPY kernel (2.63) along the fiber centerlines, so here by $\tilde{\mathbf{M}}$ we mean the matrix which maps *all* the fiber forces to *all* the fiber velocities. The most straightforward definition of these integrals is to use oversampled quadrature on an upsampled grid. When we do this, we obtain a mobility $\tilde{\mathbf{M}}$ which is *guaranteed* to be SPD, and for which we can write the Cholesky factor immediately. In particular, upsampling to a

fine grid gives us a reference SPD mobility

$$\widetilde{\mathbf{M}}_{\text{ref}} = \widetilde{\mathbf{W}}^{-1} \mathbf{E}_u^T \mathbf{W}_u \widetilde{\mathbf{M}}_{\text{RPY},u} \mathbf{W}_u \mathbf{E}_u \widetilde{\mathbf{W}}^{-1}, \quad (7.10)$$

where the subscript u denotes a matrix on the fine grid. Walking through the steps of this calculation, we first convert force to force density by applying $\widetilde{\mathbf{W}}^{-1}$. We then extend the force density to the upsampled grid (multiplying by \mathbf{E}_u) and integrate it against the RPY kernel there (the matrix $\widetilde{\mathbf{M}}_{\text{RPY},u}$ describes the pairwise RPY kernel on the upsampled grid, and \mathbf{W}_u gives the integration weights). The first three matrices, which follow from the symmetry of $\widetilde{\mathbf{M}}_{\text{ref}}$, downsample the velocity \mathbf{U}_u on the upsampled grid to the N_x point grid by minimizing the L^2 difference between $\mathbf{E}_u \mathbf{U}_{N_x}$ and \mathbf{U}_u . If we want to apply the matrix $\widetilde{\mathbf{M}}_{\text{ref}}$ in linear time, we need a fast summation method for the RPY matrix $\widetilde{\mathbf{M}}_{\text{RPY},u}$. We will treat this in Section 7.2.3.

The mobility (7.10) can also be justified in terms of a Galerkin formulation, where we let \mathbf{V} be a matrix mapping coefficients of basis functions ϕ_i to values on the collocation grid, thus setting $\mathbf{f} = \mathbf{V}\widehat{\mathbf{f}}$ and $\mathbf{U} = \mathbf{V}\widehat{\mathbf{U}}$. In this approach the mobility operator and matrix discretization we use to discretize the continuum equation $\widehat{\mathbf{U}} = \mathcal{M}_g \widehat{\mathbf{f}}$ is

$$(\mathcal{M}_g)_{ij} = \int_0^L \int_0^L \phi_i(s) \mathbf{M}_{\text{RPY}}(s, s') \phi_j(s') ds ds' \quad (7.11)$$

$$\mathbf{M}_g = \mathbf{V}^T \mathbf{E}_u^T \mathbf{W}_u \widetilde{\mathbf{M}}_{\text{RPY},u} \mathbf{W}_u \mathbf{E}_u \mathbf{V}.$$

A vital point in (7.11) is that, because we are getting the ij component of \mathcal{M}_g by taking inner products with $\phi_i(s)$ and $\phi_j(s)$, we are implicitly assuming the orthonormality of the basis functions in L^2 , i.e., the stiffness matrix is the identity,

$$\mathbf{S} = \mathbf{V}^T \widetilde{\mathbf{W}} \mathbf{V} = \mathbf{I} \rightarrow \mathbf{V}^{-1} = \mathbf{V}^T \widetilde{\mathbf{W}}. \quad (7.12)$$

Substituting into (7.11), we get a new expression of the Galerkin matrix which involves the matrix $\widetilde{\mathbf{M}}$ from (7.10),

$$\begin{aligned} \mathbf{M}_g &= \mathbf{V}^{-1} \widetilde{\mathbf{W}}^{-1} \mathbf{E}_u^T \mathbf{W}_u \mathbf{M}_{\text{RPY}} \mathbf{W}_u \mathbf{E}_u \widetilde{\mathbf{W}}^{-1} \\ &= \mathbf{V}^{-1} \widetilde{\mathbf{M}}_{\text{ref}} \mathbf{V}^{-T}. \end{aligned} \quad (7.13)$$

Here the $\mathbf{V}^{-T} = \widetilde{\mathbf{W}} \mathbf{V}$ matrix maps the Galerkin coefficients of the force density to the values of the forces on the collocation grid, the matrix $\widetilde{\mathbf{M}}_{\text{ref}}$ gives values of velocity on the collocation grid, and \mathbf{V}^{-1} maps to coefficients of velocity. Thus (7.10) could have been obtained using this Galerkin approach.

We want to emphasize that, for the rest of this dissertation, the goal of our numerical methods will be to capture the mobility (7.10). When there are thermal fluctuations with multiple filaments, further work is required to use a mobility other than (7.10). However, when there is only a single filament, or when we do not need to compute $\widetilde{\mathbf{M}}^{1/2}$ (deterministic methods), we have some freedom in making an approximation to (7.10) which is not as expensive. We will focus on the diagonal (self) terms of $\widetilde{\mathbf{M}}_{\text{ref}}$, which in theory require large oversampling to get correct when $\epsilon \rightarrow 0$. It is here that our special quadrature scheme helps us.

7.2.1 Treating the self term differently

In some cases, it might be useful to use special quadrature for the self term while using upsampling for the other RPY integrals. Indeed, in the case when the same oversampling factor is used in (7.10) for all fibers, the accuracy is controlled by the single fiber, which sets the required number of points to scale as $1/\epsilon$. We can avoid this by subtracting the self term from (7.10), then using a special quadrature scheme to obtain the correct result.

When we use our efficient quadrature scheme of Section 6.2 to approximate \mathbf{M} on a single

fiber, the matrix that results is $\widetilde{\mathbf{M}} = \mathbf{M}\widetilde{\mathbf{W}}^{-1}$, which is not guaranteed to be symmetric positive definite. Indeed, in Section 6.3.1, we showed that the mobility matrix \mathbf{M} obtained from quadrature can have negative eigenvalues due to numerical errors, especially for larger N and $\hat{\epsilon}$. The negative eigenvalues for large N and $\hat{\epsilon}$ are not altogether surprising; we know that lengthscales in the forcing on the order $\hat{\epsilon}$ will be filtered by our RPY regularization, resulting in eigenvalues close to zero. Putting more Chebyshev points, which are clustered together near the boundary, brings these lengthscales into play. Combining this with the imperfect accuracy of our quadrature scheme, which is designed for smooth forces, it is not hard to understand why negative eigenvalues result.

To work around this problem, we can replace the diagonal parts of the mobility (7.10) with the symmetric matrix

$$\widetilde{\mathbf{M}}_{\text{SQS}} = \frac{1}{2} \left(\mathbf{M}\widetilde{\mathbf{W}}^{-1} + \widetilde{\mathbf{W}}^{-1}\mathbf{M}^T \right). \quad (7.14)$$

Then, we compute an eigenvalue decomposition of this matrix and set all eigenvalues less than a threshold σ to be equal to σ . The choice of σ for a given discretization comes from the smallest eigenvalue of the reference mobility (7.10). This is a precomputation done by the code, where the input is a straight fiber (the smallest eigenvalue is not sensitive to the fiber shape; it changes by at most 10% when we change the shape), and the oversampling for the precomputation is done with $1/\epsilon$ points. This gives us a first option for the mobility, where the nonlocal parts are given by the upsampling (7.10) and the local parts are given by the special quadrature matrix (7.14),

$$\widetilde{\mathbf{M}} = \widetilde{\mathbf{M}}_{\text{ref}} - \text{BDiag} \left\{ \widetilde{\mathbf{M}}_{\text{ref}} \right\} + \text{BDiag} \left\{ \widetilde{\mathbf{M}}_{\text{SQS}} \right\}, \quad (7.15)$$

where BDiag stands for the diagonal blocks (self interaction terms) of the matrix $\widetilde{\mathbf{M}}$.

7.2.2 A hybrid approach: near fiber quadrature

If we are able to reduce the number of oversampled points by treating the self term with special quadrature, then the question remains: what about nearly touching fibers? If we want to obtain an accuracy guarantee for the integrals (2.63), then the oversampling in (7.10) will be controlled by the number of nearly touching fibers. In this section, we outline a scheme that can remove this restriction by treating nearly touching fibers with special quadrature. Our goal is to accurately compute the velocity

$$\mathbf{v}(\mathbf{x}) = \frac{1}{8\pi\mu} \int_0^L \left(\mathbb{S}(\mathbf{x}, \mathbb{X}(s)) + \frac{2\hat{a}^2}{3} \mathbb{D}(\mathbf{x}, \mathbb{X}(s)) \right) ds \quad (7.16)$$

where \mathbf{x} is a target point and $\mathbb{X}(s)$ is the Chebyshev interpolant of the centerline of a fiber. Here we have restricted ourselves to the case $\|\mathbf{x} - \mathbb{X}(s)\| > 2\hat{a}$, so that we only need to consider the Stokeslet and doublet combination of the RPY kernel (2.6). There are several components in our scheme to compute the interaction velocities $\mathbf{v}(\mathbf{x})$ to a guaranteed tolerance regardless of how close the target point \mathbf{x} is to the centerline of fiber $\mathbb{X}(s)$. We need to:

1. Understand how far \mathbf{x} can be from $\mathbb{X}(s)$ for direct quadrature on (7.16) to remain sufficiently accurate.
2. Obtain a reliable metric to compute or bound the minimum distance between the target \mathbf{x} and fiber $\mathbb{X}(s)$,

$$d := \min_s \|\mathbf{x} - \mathbb{X}(s)\|. \quad (7.17)$$

Our procedure to compute d will be different for $d/L = \mathcal{O}(1)$, when fibers are far apart and minimizing over a discrete set of nodes is sufficiently accurate, than for $d/L = \mathcal{O}(\epsilon)$, where it is more efficient to actually solve the continuous minimization problem.

3. Use a special quadrature scheme to compute the integral for $d = \mathcal{O}(\epsilon)$. The scheme we use here is taken directly from [140] and is based on extending the ideas used for the self integral in Section 6.2 to near-singular quadrature.

To begin, we define an acceptable tolerance for the integrals. Since slender body theory itself is only accurate to $\mathcal{O}(\epsilon)$, it does not make sense to set a tolerance less than ϵ . For actin filaments $\epsilon \approx 10^{-3}$, so we define the tolerance as 10^{-3} and set $\epsilon = 10^{-3}$ in our accuracy tests. That is, our goal is to compute the interaction velocities $\mathbf{v}(\mathbf{x})$ to three digits of accuracy regardless of the distance between a target \mathbf{x} and fiber $\mathbb{X}(s)$. While this method cannot guarantee 3 digits of accuracy, it does so for most target-fiber pairs of interest to us; see Fig. 7.3 for numerical results.

7.2.2.1 Distance where direct quadrature breaks down

Our first goal is to determine when direct quadrature breaks down. To do this, we simply measure the accuracy of direct quadrature on (7.16) with $N = 16$ and $N = 32$ Chebyshev nodes for randomly-generated pairs of fibers and targets. Specifically, we generate 100 inextensible fibers with 16 nonzero Chebyshev coefficients decaying exponentially (in expectation) by four orders of magnitude. We place 100 targets around each fiber a distance d in the normal direction and compute the integrals (7.16) using direct quadrature. Measuring error with respect to a refined direct quadrature, we show in Fig. 7.2 that $N = 16$ gives 3 digits of accuracy for all test cases when the non-dimensional distance $d/L \geq 0.15$. Likewise, direct quadrature with $N = 32$ points gives 3 digits of accuracy when $d/L \geq 0.06$.

7.2.2.2 Estimating d for $d/L = \mathcal{O}(1)$

Since direct quadrature breaks down for $d/L < 0.15$ when $N = 16$, we need to determine whether a target \mathbf{x} is indeed a distance less than $d/L = 0.15$ from the fiber centerline $\mathbf{X}(s)$ (the analogous statement holds for $d/L = 0.06$, $N = 32$). To do this quickly, we resample

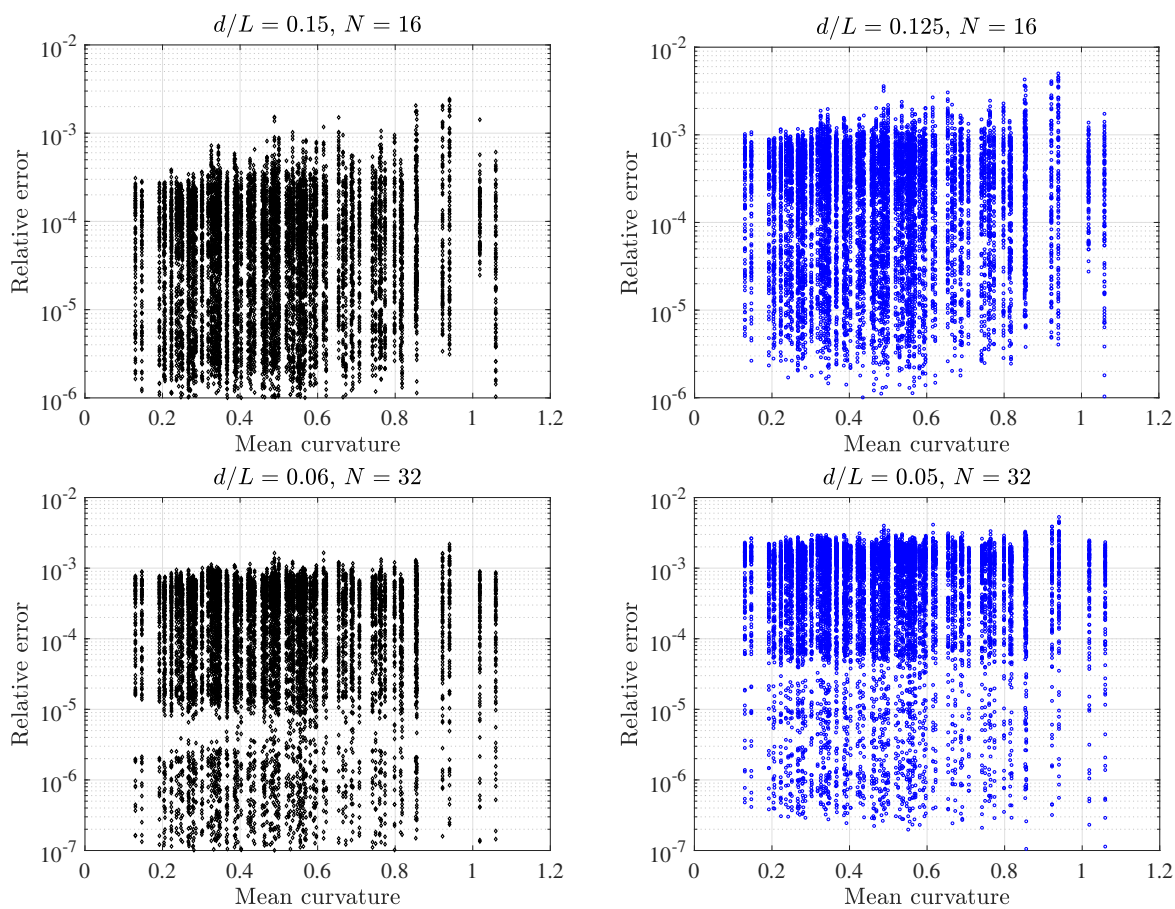


Figure 7.2: Relative errors using direct quadrature for the RPY kernel of (7.16). We consider (top) $N = 16$ and (bottom) $N = 32$ points and vary the distance d/L of the targets from the fiber. We see that we can get 3 digits of accuracy using $N = 16$ when $d/L \geq 0.15$. For $N = 32$, the 3 digit threshold occurs around $d/L = 0.06$.

the fiber centerline at $N_u = 16$ uniformly spaced points (in the arclength coordinate s) and perform a discrete minimization over the uniform fiber points to estimate the distance d . We denote this approximation by \tilde{d} . Using the same random set of fibers and targets as in Section 7.2.2.1, we found the estimate of d by \tilde{d} has a relative error in the distance $|\tilde{d} - d|/d$ of at most 5% for $d/L = 0.15$ and 20% for $d/L = 0.06$. Since direct quadrature breaks down at $d/L = 0.15$ for $N = 16$, after accounting for the error in estimating d , we have that if $\tilde{d}/L \geq 0.15 \times 1.05$, direct quadrature with $N = 16$ gives the integral to 3 digits. If $\tilde{d}/L < 0.15 \times 1.05$, we either need to use more direct quadrature points or special quadrature, which is described next.

7.2.2.3 Special quadrature.

The special quadrature routine is taken directly from [140], and is used to correct integrals where direct quadrature is insufficient. As in Section 6.2, the underlying idea is to find the near singularity in the integrand, factor it out, expand what remains in a monomial basis, and compute integrals with the singularity and monomials analytically. This time, however, the integrand is not actually singular on the fiber centerline. By expanding the fiber representation to the complex plane, the nearby singularity can be found in the complex plane and the entire procedure of Section 6.2 can be repeated.

In more detail, the interaction velocity integral (7.16) can be written so that the numerator is smooth as \mathbf{x} approaches the centerline of $\mathbb{X}(s)$. Starting from

$$\mathbf{v}(\mathbf{x}) = \frac{1}{8\pi\mu} \int_0^L \left(\frac{\mathbf{f}(s)}{\|\mathbf{R}\|} + \frac{\left((\mathbf{R}\mathbf{R}) + \frac{2\hat{a}^2}{3}\mathbf{I} \right) \mathbf{f}(s)}{\|\mathbf{R}\|^3} - 2\hat{a}^2 \frac{(\mathbf{R}\mathbf{R}) \mathbf{f}(s)}{\|\mathbf{R}\|^5} \right) ds, \quad (7.18)$$

where $\mathbf{R} = \mathbf{x} - \mathbb{X}(s)$, we rescale s by $\eta = -1 + 2s/L$. Then each of the terms in the integral

(7.18) can be written in the form

$$\int_{-1}^1 \frac{\mathbf{h}_m(\mathbf{x}, \eta)}{\|\mathbf{x} - \mathbf{X}(\eta)\|^m} d\eta, \quad (7.19)$$

where $m = 1, 3, 5$. For each m , $\mathbf{h}_m(\mathbf{x}, \eta)$ is a density that depends on the target point \mathbf{x} and varies smoothly along the fiber j arclength coordinate η .

Now, the idea of [140] is to extend the representation of $\mathbb{X}(\eta)$ from $\eta \in [-1, 1] \subset \mathbb{R}$ to the *complex* plane \mathbb{C} and compute the complex root of $\|\mathbf{x} - \mathbb{X}(\eta)\| = 0$. Because the centerline representation $\mathbb{X}(\eta)$ is available as a Chebyshev series, it is simple to solve for the root via Newton iteration. We denote this root by η^* , i.e., $\|\mathbf{x} - \mathbb{X}(\eta^*)\| = 0$ with $\eta^* \in \mathbb{C}$.

Once the root is known, the algorithm of [140] finds the radius of the Bernstein ellipse associated with η^* . This radius then bounds the direct quadrature error for the integral (7.18). If the upper bound on the direct quadrature error with the number of oversampled points used is less than 10^{-3} , we proceed with direct quadrature. Otherwise, we use the special quadrature scheme of [140], which is the same as in the finite part integration in Section 6.2. The complex singularity is factored out of the integrand, leaving a “smooth” function which can be expanded in a monomial basis. Integrals involving monomials multiplied by the singularity are computed analytically, and an inner product of the monomial coefficients with the analytical integrands yields the approximation to $\mathbf{v}(\mathbf{x})$.

The only difference from Section 6.2 is that the location of the singularity (complex root η^*) now depends on the fiber position \mathbf{X} . This means that the roots and monomial coefficients must be computed at every time step using an *LU* factorization of the Vandermonde matrix. Since the Vandermonde matrix is a function of the nodes s_p on the fiber, its *LU* factorization is the same for all fibers and can be precomputed once at the start of a dynamic simulation. We refer the reader to [140, Section 3] for more details on this quadrature scheme.

While error bounds exist for direct quadrature, the special quadrature scheme of [140]

does not provide error bounds or a method for selecting the fiber discretization (number of panels, points per panel, etc.). Because of this, we performed an empirical study on the same set of 100 fibers and targets as in Section 7.2.2.1. The randomized testing described in Section 7.2.2.5 showed that, for most fiber configurations of interest to us, 1 panel of 32 points is sufficient to compute the integral to 3 digits using special quadrature as long as the non-dimensional distance between the target and fiber is $d/(\epsilon L) > 8$. Otherwise, 2 panels of 32 points are required. Fibers with high curvature typically give the largest errors for a given discretization.

7.2.2.4 Estimating d for $d/L = \mathcal{O}(\epsilon)$

We still require a robust numerical procedure to determine when the target point is too close to the cross section of the fiber. Our idea is to use the real part of the root η^* as the closest arclength coordinate on the fiber to the target. We know the root η^* solves $\|\mathbf{x} - \mathbb{X}(\eta^*)\| = 0$. It seems sensible, therefore, for the real part of the root to approximately minimize (over real η) $\|\mathbf{x} - \mathbb{X}(\eta)\|$ when the root is close to the real line. We therefore define s^* , the closest point on the fiber to the target, from the complex root η^* by removing the imaginary part of the root and rescaling,

$$s^* = \begin{cases} \frac{L}{2}(\operatorname{Re}(\eta^*) + 1) & -1 \leq \operatorname{Re}(\eta^*) \leq 1 \\ 0 & \operatorname{Re}(\eta^*) < -1 \\ L & \operatorname{Re}(\eta^*) > 1 \end{cases} . \quad (7.20)$$

The shortest distance from \mathbf{x} to the fiber can then be estimated as

$$\hat{d} := \|\mathbf{x} - \mathbb{X}(s^*)\| , \quad (7.21)$$

where the position $\mathbb{X}(s^*)$ is computed by evaluating the Chebyshev interpolant at s^* . Our randomized tests showed that this estimate gives an error of at most 10% for $d/(\epsilon L) \leq 8$, which is sufficiently accurate for our purpose. Because we use a point s^* on the fiber centerline to estimate \hat{d} , this 10% error is always an overestimation. For this reason we use 2 panels for special quadrature when $\hat{d}/(\epsilon L) \leq 8.8$.

7.2.2.5 Near fiber accuracy

We now test the accuracy of our algorithm for computing the slender body interaction integrals (7.16). We generate 100 smooth inextensible fibers by initializing an unnormalized tangent vector that is an exact Chebyshev series with 15 exponentially decaying terms. More precisely, the k th coefficient of the series is a Gaussian random variable with mean 0 and standard deviation $e^{-10k/N}$, where $k = 0, \dots, 15$. We then normalize this tangent vector to obtain $\boldsymbol{\tau}$ and integrate to obtain the fiber positions $\mathbb{X}(s)$. To make sure the resulting fiber is smooth after tangent vector normalization, we compute the Chebyshev series of the fiber position $\mathbb{X}(s)$. Denoting the coefficients of the position Chebyshev series by \hat{a}_k , we only accept fibers with Chebyshev series coefficients $|\hat{a}_k| \leq e^{-0.61k}$ for $k = 2, \dots, 15$ (the constant and linear modes play no role in the fiber smoothness). This means that the last coefficient \hat{a}_{15} has value at most 10^{-4} .

We consider fibers with $L = 2$ and $\epsilon = 10^{-3}$ in a fluid of viscosity $\mu = 1/8\pi$. Our goal is to evaluate the velocity (7.16) due to a fiber $\mathbb{X}(s)$ at a target \boldsymbol{x} . To measure the accuracy of our quadrature in doing so, we place 100 targets a distance d away in a random normal direction from each fiber's centerline. To get a reference answer, we compute the integral (7.16) directly by upsampling the fiber to 6000 type 1 Chebyshev points. We then compute the integral using our algorithm. We show the maximum relative error, $\mathbf{E}_i / \|\mathbf{v}(\boldsymbol{x})\|_\infty$, where the maximum is over the direction $i = 1, 2, 3$ and \mathbf{E} is the absolute difference between the approximate and reference values of the velocity (7.16).




Figure 7.3: Errors in algorithm for near singular quadratures. We show a histogram of the number of digits obtained in the velocity $\mathbf{v}(\mathbf{x})$, given in (7.16), for 100 different fibers and 100 targets per fiber. Here d is the distance from the target point to the fiber centerline and we show histograms of the number of digits obtained in the integral $\mathbf{v}(\mathbf{x})$. The number of digits is computed as $-\log_{10}(\max_i \mathbf{E}_i / \|\mathbf{v}(\mathbf{x})\|_\infty)$, where the maximum is over the direction $i = 1, 2, 3$ and \mathbf{E} is the absolute difference between the approximate and reference values of the velocity (7.16). (a) Errors from short distances $2 \leq d/(\epsilon L) \leq 10$. We see that we are over-working in most cases, since most of the time we obtain many more than 3 digits of accuracy, but there are some cases when we only obtain 3 digits. (b) Long distances $0.04 < d/L < 0.20$, where we obtain 4-7 digits most of the time.

We separate our results into short distances, $2\epsilon L < d < 10\epsilon L$, and long distances, $0.01L < d < 0.2L$ (there is overlap between the two regions since $10\epsilon L = 0.01L = 0.02$ with our parameters). Fig. 7.3(a) shows the errors for short distances, where we see that we obtain many more digits than necessary in most cases. There are, however, a few cases where we obtain 3 digits. Since d will rarely be $\mathcal{O}(\hat{a})$, it is acceptable to expend extra computational effort to guarantee accuracy. In Fig. 7.3(b), we show the errors for long distances. In particular, we see that we obtain 4-5 digits most of the time, and that $\approx 5\%$ of the time we obtain exactly 3 digits of accuracy.

7.2.3 Fast direct summation

The oversampled mobility definition (7.10) shows that we need an algorithm to sum the RPY kernel over many point force pairs. In (7.10), the pairs are obtained by upsampling

the fiber position function to N_u points and then multiplying by the Chebyshev weight on the upsampled grid to obtain the force there. The focus of this section is on how to apply the matrix $\widetilde{\mathbf{M}}_{\text{RPY},u}$ fast on the upsampled grid. We will restrict ourselves to triply periodic boundary conditions, for which we can use fast Ewald summation [103, 104, 141, 142] to split the RPY kernel into a near field part (nonsmooth in Fourier space, but rapidly decaying in real space) and a far field part (smooth in Fourier space). When we perform rheology experiments, the directions of periodicity change with time; thus we begin this section by introducing the domain we want to solve the equations on, and how this affects the resulting Ewald splitting algorithm.

7.2.3.1 Sheared coordinate system

In order to implement a shear flow in periodic boundary conditions, a strained coordinate system is necessary. The spectral Ewald method was extended to sheared cells in [75], but here we give a more detailed description for completeness. We assume (without loss of generality) that x is the flow direction, y is the gradient direction, and z is the vorticity direction. Let the total nondimensional strain be $g(t)$. In Fig. 7.4 we define a strained coordinate system with axes

$$\mathbf{e}_{x'} = \mathbf{e}_x, \quad \mathbf{e}_{y'} = \mathbf{e}_y + g(t)\mathbf{e}_x, \quad \mathbf{e}_{z'} = \mathbf{e}_z, \quad (7.22)$$

and strained wave numbers

$$k'_x = k_x, \quad k'_y = k_y + g(t)k_x, \quad k'_z = k_z. \quad (7.23)$$

Here k_x , k_y , and k_z are the wave numbers when the periodicity is over the x , y , and z directions, while k'_x , k'_y , and k'_z are the wave numbers when the periodicity is over the x' , y' , and z' directions.




Figure 7.4: The sheared/parallelepiped periodic simulation cell. We denote the dimensionless shear by g . The green area shows the periodic cell (2D projection) when $g = 0$, and the blue area is the periodic cell for nonzero g .

The transformation between the two coordinate systems is given by

$$\mathbf{x}' := \begin{pmatrix} x' \\ y' \\ z' \end{pmatrix} = \begin{pmatrix} 1 & -g(t) & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} := \mathbf{L}\mathbf{x}. \quad (7.24)$$

In the unsheared to sheared transformation (7.24), the sheared coordinates x' , y' , and z' are all periodic on $[0, L]^3$ (the blue simulation cell in Fig. 7.4).

Now we use the transformation (7.24) to transform the derivative operators in the un-sheared coordinate system to the sheared one,

$$\frac{\partial}{\partial x} = \frac{\partial}{\partial x'} \quad \frac{\partial}{\partial y} = \frac{\partial}{\partial y'} - g(t) \frac{\partial}{\partial x'} \quad \frac{\partial}{\partial z} = \frac{\partial}{\partial z'}. \quad (7.25)$$

We therefore have the Laplacian in the transformed space as [143]

$$\Delta = \left(\frac{\partial^2}{\partial x'^2} + \left(\frac{\partial}{\partial y'} - g(t) \frac{\partial}{\partial x'} \right)^2 + \frac{\partial^2}{\partial z'^2} \right). \quad (7.26)$$

In Fourier space, $\widehat{\Delta} = \mathbf{k}' \cdot \mathbf{k}'$, where

$$\mathbf{k}' = (k'_x, k'_y - g(t)k'_x, k'_z), \quad (7.27)$$

Using the sheared to unsheared transformation in (7.23), it is easy to see that $k' := \|\mathbf{k}'\| = \|(k_x, k_y, k_z)\| := k$. It follows that we can simply replace k in any isotropic Fourier calculations by k' to use a Fourier method in the sheared coordinate system [143, 144].

7.2.3.2 Ewald splitting for direct quadratures

We use the Ewald splitting of [75, 103] to accelerate the computation of the many-body sum in $\widetilde{\mathbf{M}}_{\text{RPY},u}$ in (7.10) on a periodic domain. The idea of Ewald splitting or Ewald summation is to split the RPY kernel into a smooth long-ranged part and a remaining short-ranged part. The smooth “far field” part has an exponential decay in Fourier space and can be done by standard Fourier methods (namely the non-uniform FFT), and the “near field” part decays exponentially in real space and can be truncated so that it is nonzero for $\mathcal{O}(1)$ neighboring points (sources) per target.

Let \mathbf{x}' and \mathbf{y}' be the coordinates of the points \mathbf{x} and \mathbf{y} in the sheared domain using the coordinate transformation (7.24). The periodic RPY tensor for a sphere with radius \hat{a} can then be written on the sheared domain as

$$\mathbf{S}_{\text{RPY}}^{(P)}(\mathbf{x}', \mathbf{y}'; \hat{a}) = \frac{1}{V\mu} \sum_{\mathbf{k}' \neq \mathbf{0}} e^{i\mathbf{k}' \cdot (\mathbf{x}' - \mathbf{y}')} \frac{1}{k'^2} \left(\mathbf{I} - \hat{\mathbf{k}}' \hat{\mathbf{k}}'^T \right) \text{sinc}^2(k' \hat{a}), \quad (7.28)$$

where V is the domain volume and $\mathbf{k}' = 2\pi\mathbf{m}/L_d$, where \mathbf{m} is a vector of three integers and L_d is the periodic domain length. Using the screening function of Hasimoto [145],

$$H(k', \xi) = \left(1 + \frac{k'^2}{4\xi^2} \right) e^{-k'^2/4\xi^2}, \quad (7.29)$$

we split the periodic kernel $\mathbf{S}_{\text{RPY}}^{(P)}$ into a far field and near field component, $\mathbf{S}_{\text{RPY}}^{(P)} = \mathbf{S}_{\text{RPY}}^{(\text{FF})} + \mathbf{S}_{\text{RPY}}^{(\text{NF})}$, where the far field is given in Fourier space by

$$\mathbf{S}_{\text{RPY}}^{(\text{FF})}(\mathbf{x}', \mathbf{y}'; \hat{a}) = \frac{1}{V\mu} \sum_{\mathbf{k}' \neq \mathbf{0}} e^{i\mathbf{k}' \cdot (\mathbf{x}' - \mathbf{y}')} \frac{1}{k'^2} \left(\mathbf{I} - \hat{\mathbf{k}}' \hat{\mathbf{k}}'^T \right) \text{sinc}^2(k' \hat{a}) H(k', \xi). \quad (7.30)$$

Here ξ is a splitting parameter that controls the decay of the far field kernel in Fourier space and of the near field kernel in real space, and is chosen to optimize performance. The total far field sum is obtained by summing the far field kernel (7.30) over all points \mathbf{y}' , with the $k' = 0$ mode set to zero since in continuum the total force on the system is zero. We use an NUFFT to compute these sums at all points \mathbf{x}' in log-linear time in the number of points (see [75, 103] for more details).

Assuming that the near field decays rapidly enough that Fourier series can be replaced by Fourier integrals, the near field mobility can be computed in real space by inverse transforming its Fourier space representation,

$$\mathbf{S}_{\text{RPY}}^{(\text{NF})}(\mathbf{x}, \mathbf{y}; \hat{a}) = F(r, \xi, \hat{a}) (\mathbf{I} - \hat{\mathbf{r}}\hat{\mathbf{r}}^T) + G(r, \xi, \hat{a})\hat{\mathbf{r}}\hat{\mathbf{r}}^T, \quad (7.31)$$

where $\mathbf{r} = (\mathbf{x} - \mathbf{y})^*$, $r = \|\mathbf{r}\|$, and the $*$ denotes the nearest periodic image in the sheared domain (blue in Fig. 7.4). The exact forms of F and G are given in [103, Appendix A]. The total near field sum is computed at \mathbf{x} by summing the near field kernel (7.31) over neighboring points whose minimum image distance from \mathbf{x} is less than a precomputed value r^* . We choose r^* so that the velocities are computed to a relative tolerance of 10^{-3} , and set ξ such that r^* is small enough that only the nearest periodic image contributes to the near field sum for each pair of points.

While the nearest image for near field calculations is over the sheared domain, \mathbf{r} and r are computed using the Euclidean metric. We search for pairs of points closer than r^* apart in log-linear time using a multithreaded linked-list algorithm for a rectangular periodic cell. To adjust for the fact that the \mathbf{x}' coordinates are given on a non-orthogonal coordinate system, we bound the Euclidean distance between points in the unsheared coordinates by

their “distance” in sheared coordinates,

$$\|\mathbf{r}\| = \sqrt{\mathbf{r}^T \mathbf{r}} = \sqrt{(\mathbf{r}')^T \mathbf{L}^{-T} \mathbf{L}^{-1} \mathbf{r}'} \geq \left(1 + \frac{1}{2} \left(g^2 + \sqrt{g^2(g^2 + 4)}\right)\right)^{-1} \|\mathbf{r}'\| := \frac{1}{\psi} \|\mathbf{r}'\|, \quad (7.32)$$

where we have used the maximum eigenvalue of $\mathbf{L}^{-T} \mathbf{L}^{-1}$ to bound the norm [75]. The factor ψ defined in (7.32) can be thought of as a “safety factor” in the sense that points that are ψr^* apart using the Euclidean metric in primed coordinates are at least r^* apart in physical space.

7.2.3.3 Verification for sheared unit cell

To test our implementation of sheared periodic boundary conditions, we consider a packing of points that is hexagonal in the xy plane. As shown in Fig. 7.5, the points are positioned on a (green) periodic slanted cell at $(0, 0, 0)$ (red blob), $(1, 0, 0)$ (black), $(0.5, 1, 0)$ (orange), and $(1.5, 1, 0)$ (sky blue). To form a periodic hexagonal packing in the xy plane, we set $g = 0.5$ with periodic domain length $L_x = L_y = L_z = 2$. Using a coloring scheme (see Fig. 7.5), it is easy to see that this arrangement is equivalent to the same set of points on a (gray) rectangular unit cell, with additional points at $(1, 2, 0)$, $(0, 2, 0)$, $(1.5, 3, 0)$, $(0.5, 3, 0)$, with the ordering of forces in the second set of points being the same as the first, and periodic length $L_y = 4$. We place a force of strength $+1$ in each direction (including z) on the first (red) pair of points, -1 in each direction on the second (black) pair, $+2$ on the third (orange), and -2 on the fourth (sky blue). Note that the z direction is also periodic in all cases with length $L_z = 2$, so that we are actually considering a set of stacked copies of Fig. 7.5.

We solve for the RPY velocities induced by the forces at each point using the Ewald splitting technique described in Section 7.2.3. We set $\xi = 5$, sphere radius $\hat{a} = 10^{-2}$, and fluid viscosity $\mu = 3$. The maximum relative 2-norm error in the velocity of the four points is less than 10^{-5} for all values of the NUFFT tolerance less than 10^{-2} , with decay to 10^{-11} when the tolerance is 10^{-8} . We conclude that our modified Ewald splitting scheme of Section




Figure 7.5: The lattice for the sheared unit cell test. The points are positioned on a lattice with a 1-particle (primitive) unit cell (shaded green) at $(0, 0, 0)$ (red), $(1, 0, 0)$ (black), $(0.5, 1, 0)$ (orange), and $(1.5, 1, 0)$ (blue). The lattice can also be viewed as periodic on the larger 2-particle rectangular unit cell shaded in gray. Colors indicate the magnitude of the force placed on each set of points: $+1$ in all three directions on the red points, -1 in each direction on the black points, $+2$ on the orange points, and -2 on the blue points.

7.2.3 properly treats the strain in the periodic coordinate system.

7.2.4 The choice of mobility

A trade-off now ensues between the upsampling factor in (7.10) and the resulting amount of effort that has to be put into the special quadrature in Section 7.2.2. Indeed, if we consider a small number of upsampled points, then we will need to correct a lot of nearby fibers for which the direct quadrature fails. Conversely, if we take the limit of infinitely many upsampled points, we won't need the special quadrature scheme at all! Our choice is for the latter, based on the following considerations:

1. In triply periodic systems, the mobility (7.10) with large N_u can be applied fast on a GPU using the UAMMD library [105]. As such, for the actin filaments we consider with $\epsilon \geq 10^{-3}$ it turns out that it is *always* faster to use oversampled quadrature on the GPU.
2. The special quadrature scheme in Section 7.2.2 does not account for the change in the RPY kernel when $R < 2\hat{a}$. Indeed, because the RPY kernel changes on $R < 2\hat{a}$,

we would need a separate procedure for fiber-target pairs which are separated by a distance less than $2\hat{a}$. So far, we have not developed such a procedure because it is faster to use oversampling.

3. Computing the mobility via (7.10) without modifications guarantees an SPD mobility matrix, and gives us a formula for $\widetilde{\mathbf{M}}^{1/2}$. This will be necessary when we introduce thermal fluctuations in Chapter 8.

For the reasons listed here, we will *always* use the mobility matrix given by (7.10), allowing for the possibility that the self term be treated differently as in (7.15). In practice, we will choose a number of oversampling points, then make sure that the quantities we are interested in don't change when we double the number of points.

Summarizing then, the nonlocal mobility is always computed by the off-diagonal parts of (7.10). We allow for the *local* mobility to be computed with special quadrature (Section 6.2). When we describe a mobility calculation, then, there are two questions which completely define the calculation

1. What kind of mobility was used for the local integrals? Special quadrature or over-sampled direct quadrature?
2. How many points were used in the upsampling for the RPY integrals in the mobility (7.10)?

7.3 Temporal integration

We now discuss how we split up the mobility in the temporal integrator to produce a minimum number of inter-fiber hydrodynamic evaluations per time step. We will introduce a first order and second order temporal integrator for this purpose. In the second order case, we avoid nonlinear solves by using extrapolated values as arguments to nonlinear functions.

For example, we define $\widetilde{\mathbf{M}}^{n+1/2,*} = \widetilde{\mathbf{M}}(\mathbf{X}^{n+1/2,*})$, with

$$\mathbf{X}^{n+1/2,*} = \frac{3}{2}\mathbf{X}^n - \frac{1}{2}\mathbf{X}^{n-1}, \quad (7.33)$$

where we have used the notation \mathbf{X}^n to denote the fiber positions \mathbf{X} at the n th time step. We will also consider a backward Euler temporal integrator, for which no extrapolations are necessary. This discretization is only first-order accurate, but is more robust and useful in biological applications.

To discretize the saddle point system (7.3) in time, we split the mobility matrix a local and nonlocal part,

$$\widetilde{\mathbf{M}} = \widetilde{\mathbf{M}}_{\text{L}} + \widetilde{\mathbf{M}}_{\text{NL}} \quad (7.34)$$

$$= \text{BDiag}\{\widetilde{\mathbf{M}}\} + \left(\widetilde{\mathbf{M}} - \text{BDiag}\{\widetilde{\mathbf{M}}\}\right). \quad (7.35)$$

While other options are possible, including making $\widetilde{\mathbf{M}}_{\text{L}}$ itself diagonal, here we will restrict to the case of (7.35), where the local mobility on fiber i is just the i th diagonal block of the matrix $\widetilde{\mathbf{M}}$. We will analyze the differences between the case where the mobility is given by (7.10) (oversampled RPY for all fibers) and the case where special quadrature is used for the self terms (7.15).

7.3.1 Second-order discretization

The second-order temporal discretization of the evolution equation (7.3) begins by solving the saddle-point system

$$\begin{aligned} & \widetilde{\mathbf{M}}_{\text{L}}^{n+1/2,*} \left(\boldsymbol{\Lambda}^{n+1/2} - \frac{1}{2} \mathbf{L} (\mathbf{X}^n + \mathbf{X}^{n+1,*}) + \mathbf{F}_{\text{ext}}^{n+1/2,*} \right) \\ & + \widetilde{\mathbf{M}}_{\text{NL}}^{n+1/2,*} \left(\boldsymbol{\Lambda}^{n+1/2,*} - \mathbf{L} \mathbf{X}^{n+1/2,*} + \mathbf{F}_{\text{ext}}^{n+1/2,*} \right) + \mathbf{u}_0 \left(\mathbf{X}^{n+1/2,*} \right) = \mathbf{K}^{n+1/2,*} \boldsymbol{\alpha}^{n+1/2}, \\ & \left(\mathbf{K}^{n+1/2,*} \right)^T \boldsymbol{\lambda}^{n+1/2} = \mathbf{0} \end{aligned} \quad (7.36)$$

for $\boldsymbol{\lambda}^{n+1/2}$ and $\boldsymbol{\alpha}^{n+1/2}$, where $\mathbf{F}_{\text{ext}}(\mathbf{X})$ is an external force. To obtain a second-order block-diagonal system, we extrapolate previous $\boldsymbol{\Lambda}$ values to the midpoint of the next time step,

$$\boldsymbol{\Lambda}^{n+1/2,*} = 2\boldsymbol{\Lambda}^{n-1/2} - \boldsymbol{\Lambda}^{n-3/2}. \quad (7.37)$$

We also introduce the approximation

$$\mathbf{X}^{n+1,*} = \mathbf{X}^n + \Delta t \mathbf{K}^{n+1/2,*} \boldsymbol{\alpha}^{n+1/2} \quad (7.38)$$

to make (7.36) a linear system in $\boldsymbol{\alpha}^{n+1/2}$ and $\boldsymbol{\lambda}^{n+1/2}$.

By substituting the approximation $\mathbf{X}^{n+1,*}$ in (7.38) into saddle-point system (7.36), we obtain the following saddle-point linear system to solve at every time step,

$$\begin{aligned} & \begin{pmatrix} -\widetilde{\mathbf{M}}_{\text{L}} & \mathbf{K} + \frac{\Delta t}{2} \widetilde{\mathbf{M}}_{\text{L}} \mathbf{L} \mathbf{K} \\ \mathbf{K}^T & \mathbf{0} \end{pmatrix}^{n+1/2,*} \begin{pmatrix} \boldsymbol{\Lambda}^{n+1/2} \\ \boldsymbol{\alpha}^{n+1/2} \end{pmatrix} = \\ & \begin{pmatrix} \widetilde{\mathbf{M}}_{\text{L}}^{n+1/2,*} \left(-\mathbf{L} \mathbf{X}^n + \mathbf{F}_{\text{ext}}^{n+1/2,*} \right) + \widetilde{\mathbf{M}}_{\text{NL}}^{n+1/2,*} \left(\boldsymbol{\Lambda}^{n+1/2,*} - \mathbf{L} \mathbf{X}^{n+1/2,*} + \mathbf{F}_{\text{ext}}^{n+1/2,*} \right) + \mathbf{u}_0 \left(\mathbf{X}^{n+1/2,*} \right) \\ \mathbf{0} \end{pmatrix}. \end{aligned} \quad (7.39)$$

This system can be solved fiber by fiber, since the matrix $\widetilde{\mathbf{M}}_{\text{L}}$ on the left hand side of

(7.39) is block diagonal (the matrix \mathbf{K} here is a shorthand for a block diagonal matrix whose diagonal entries are the \mathbf{K} for each filament).

In the case when the fibers are packed densely enough to make the temporal discretization (7.39) unstable, we treat the bending force in the nonlocal hydrodynamics implicitly and use GMRES to solve for $\Lambda^{n+1/2}$ and $\alpha^{n+1/2}$. The new linear system of equations is

$$\begin{aligned} & \left(\widetilde{\mathbf{M}}_{\text{L}}^{n+1/2,*} + \widetilde{\mathbf{M}}_{\text{NL}}^{n+1/2,*} \right) \left(\Lambda^{n+1/2} - \frac{1}{2} \mathbf{L} (\mathbf{X}^n + \mathbf{X}^{n+1,*}) + \mathbf{F}_{\text{ext}}^{n+1/2,*} \right) \\ & = \mathbf{K}^{n+1/2,*} \alpha^{n+1/2} - \mathbf{u}_0 (\mathbf{X}^{n+1/2,*}), \\ & \left(\mathbf{K}^{n+1/2,*} \right)^T \Lambda^{n+1/2} = \mathbf{0}. \end{aligned} \quad (7.40)$$

Now, let us denote the solutions of the block diagonal system (7.39) by $\widetilde{\Lambda}^{n+1/2}$ and $\widetilde{\alpha}^{n+1/2}$. By subtracting the fully implicit system (7.40) from the locally implicit system (7.36), we obtain the residual form of the saddle-point system

$$\begin{aligned} & \left(\begin{array}{cc} - \left(\widetilde{\mathbf{M}}_{\text{L}} + \widetilde{\mathbf{M}}_{\text{NL}} \right) & \mathbf{K} + \frac{\Delta t}{2} \left(\widetilde{\mathbf{M}}_{\text{L}} + \widetilde{\mathbf{M}}_{\text{NL}} \right) \mathbf{L} \mathbf{K} \\ \mathbf{K}^T & \mathbf{0} \end{array} \right)_{n+1/2,*} \begin{pmatrix} \Delta \Lambda^{n+1/2} \\ \Delta \alpha^{n+1/2} \end{pmatrix} = \\ & \left(\begin{array}{c} \widetilde{\mathbf{M}}_{\text{NL}}^{n+1/2,*} \left(-\mathbf{L} \left(\mathbf{X}^n + \frac{\Delta t}{2} \mathbf{K}^{n+1/2,*} \widetilde{\alpha}^{n+1/2} - \mathbf{X}^{n+1/2,*} \right) + \widetilde{\Lambda}^{n+1/2} - \Lambda^{n+1/2,*} \right) \\ \mathbf{0} \end{array} \right) \end{aligned} \quad (7.41)$$

to be solved using GMRES for the perturbations

$$\begin{aligned} \Delta \Lambda^{n+1/2} &= \Lambda^{n+1/2} - \widetilde{\Lambda}^{n+1/2} \quad \text{and} \\ \Delta \alpha^{n+1/2} &= \alpha^{n+1/2} - \widetilde{\alpha}^{n+1/2}. \end{aligned} \quad (7.42)$$

At $t = 0$ and $t = \Delta t$, we solve the system (7.40) with a tolerance of 10^{-6} to initialize Λ^1 , α^1 , Λ^2 , and α^2 .

7.3.2 First-order discretization

The equations for the first order temporal discretization are simpler than those of the previous section. The initial saddle-point system to be solved is

$$\begin{aligned} \widetilde{\mathbf{M}}_{\text{L}}^n (\boldsymbol{\Lambda}^{n+1} - \mathbf{L}\mathbf{X}^{n+1,*} + \mathbf{F}_{\text{ext}}^n) + \widetilde{\mathbf{M}}_{\text{NL}}^n (\boldsymbol{\Lambda}^n - \mathbf{L}\mathbf{X}^n + \mathbf{F}_{\text{ext}}^n) + \mathbf{u}_0(\mathbf{X}^n) = \mathbf{K}^n \boldsymbol{\alpha}^{n+1}, \quad (7.43) \\ (\mathbf{K}^n)^T \boldsymbol{\Lambda}^{n+1} = \mathbf{0} \end{aligned}$$

for $\boldsymbol{\Lambda}^{n+1}$ and $\boldsymbol{\alpha}^{n+1}$. We again introduce the approximation

$$\mathbf{X}^{n+1,*} = \mathbf{X}^n + \Delta t \mathbf{K}^n \boldsymbol{\alpha}^{n+1} \quad (7.44)$$

to make (7.43) a linear system in $\boldsymbol{\alpha}^{n+1}$ and $\boldsymbol{\Lambda}^{n+1}$,

$$\begin{aligned} \begin{pmatrix} -\widetilde{\mathbf{M}}_{\text{L}} & \mathbf{K} + \Delta t \widetilde{\mathbf{M}}_{\text{L}} \mathbf{L} \mathbf{K} \\ \mathbf{K}^T & \mathbf{0} \end{pmatrix}_n \begin{pmatrix} \boldsymbol{\Lambda}^{n+1} \\ \boldsymbol{\alpha}^{n+1} \end{pmatrix} = \\ \begin{pmatrix} \widetilde{\mathbf{M}}_{\text{L}}^n (-\mathbf{L}\mathbf{X}^n + \mathbf{F}_{\text{ext}}^n) + \widetilde{\mathbf{M}}_{\text{NL}}^n (\boldsymbol{\Lambda}^n - \mathbf{L}\mathbf{X}^n + \mathbf{F}_{\text{ext}}^n) + \mathbf{u}_0(\mathbf{X}^n) \\ \mathbf{0} \end{pmatrix}. \end{aligned} \quad (7.45)$$

In the case when the fibers are packed densely enough to make the temporal discretization (7.45) unstable, we treat the bending force in the nonlocal hydrodynamics implicitly and use GMRES to solve for $\boldsymbol{\Lambda}^{n+1}$ and $\boldsymbol{\alpha}^{n+1}$. The new linear system of equations is

$$\begin{aligned} (\widetilde{\mathbf{M}}_{\text{L}}^n + \widetilde{\mathbf{M}}_{\text{NL}}^n) (\boldsymbol{\Lambda}^{n+1} - \mathbf{L}\mathbf{X}^{n+1,*} + \mathbf{F}_{\text{ext}}^n) + \mathbf{u}_0(\mathbf{X}^n) = \mathbf{K}^n \boldsymbol{\alpha}^{n+1}, \\ (\mathbf{K}^n)^T \boldsymbol{\Lambda}^{n+1} = \mathbf{0}. \end{aligned} \quad (7.46)$$

Now, let us denote the solutions of the block diagonal system (7.45) by $\widetilde{\boldsymbol{\Lambda}}^{n+1}$ and $\widetilde{\boldsymbol{\alpha}}^{n+1}$. By subtracting the fully implicit system (7.46) from the locally implicit system (7.43), we

obtain the residual form of the saddle-point system

$$\begin{pmatrix} -\left(\widetilde{\mathbf{M}}_{\text{L}} + \widetilde{\mathbf{M}}_{\text{NL}}\right) & \mathbf{K} + \Delta t \left(\widetilde{\mathbf{M}}_{\text{L}} + \widetilde{\mathbf{M}}_{\text{NL}}\right) \mathbf{L}\mathbf{K} \\ \mathbf{K}^T & \mathbf{0} \end{pmatrix}_n \begin{pmatrix} \Delta \boldsymbol{\Lambda}^{n+1} \\ \Delta \boldsymbol{\alpha}^{n+1} \end{pmatrix} = \begin{pmatrix} \widetilde{\mathbf{M}}_{\text{NL}}^n \left(-\mathbf{L} (\Delta t \mathbf{K}^n \widetilde{\boldsymbol{\alpha}}^{n+1}) + \widetilde{\boldsymbol{\Lambda}}^{n+1} - \boldsymbol{\Lambda}^n\right) \\ \mathbf{0} \end{pmatrix} \quad (7.47)$$

to be solved using GMRES for the perturbations

$$\begin{aligned} \Delta \boldsymbol{\Lambda}^{n+1} &= \boldsymbol{\Lambda}^{n+1} - \widetilde{\boldsymbol{\Lambda}}^{n+1} \quad \text{and} \\ \Delta \boldsymbol{\alpha}^{n+1} &= \boldsymbol{\alpha}^{n+1} - \widetilde{\boldsymbol{\alpha}}^{n+1}. \end{aligned} \quad (7.48)$$

At $t = 0$, we solve the system (7.46) with a tolerance of 10^{-6} to initialize $\boldsymbol{\lambda}^1$ and $\boldsymbol{\alpha}^1$.

7.3.3 Preconditioning

To solve systems (7.41) and (7.47) rapidly with GMRES, we replace the full mobility with just the local part to get a “block diagonal” preconditioner

$$\mathbf{P} = \begin{pmatrix} -\widetilde{\mathbf{M}}_{\text{L}} & \mathbf{K} + \Delta t c \widetilde{\mathbf{M}}_{\text{L}} \mathbf{L}\mathbf{K} \\ \mathbf{K}^T & \mathbf{0} \end{pmatrix}, \quad (7.49)$$

where $c = 1$ for backward Euler and $c = 1/2$ for Crank Nicolson. This amounts to removing the $\widetilde{\mathbf{M}}_{\text{NL}}$ terms from the left hand sides of (7.41) and (7.47). This preconditioner ought to be effective since $\widetilde{\mathbf{M}}_{\text{L}}$ typically dominates $\widetilde{\mathbf{M}}_{\text{NL}}$ for slender filaments.

7.4 Suspensions of bent filaments

In this section, we examine the accuracy and stability of our method for deterministic suspensions of flexible filaments. We study spatio-temporal accuracy in Sections 7.4.1 and 7.4.2 with simple examples of four falling fibers in free space and three fibers in periodic shear flow. Section 7.4.3 gives our most important result for computational complexity: the number of hydrodynamic evaluations per time step required to stably evolve the dynamics of a fiber suspension is at most five. By varying the fiber number density and bending modulus, we show that one block diagonal solve (7.45) combined with at most three iterations of GMRES per time step are needed to maintain stability (the extra hydrodynamic evaluation to give a total of five comes in the conversion to residual form (7.41)).

In general, we will use an L^2 function norm to compute the differences between configurations throughout this section. Given two fiber configurations, we evaluate the Chebyshev interpolant of each on a 1000 point type 2 Chebyshev grid and calculate the discrete L^2 error using Clenshaw-Curtis quadrature. Whenever there are multiple fibers, we compute the error on the first fiber $\mathbf{X}^{(1)}$ unless otherwise specified.

Since fibers in shear flow are our primary interest, some of our examples will use shear flows. The general form of a time-oscillatory shear flow is given by

$$\mathbf{u}_0(\mathbf{x}, t) = \dot{\gamma}_0 \cos(\omega t)(y, 0, 0). \quad (7.50)$$

The corresponding strain is given by $g(t) = (\dot{\gamma}_0/\omega) \sin(\omega t)$ for $\omega > 0$ and $g(t) = \dot{\gamma}_0 t$ for $\omega = 0$.

7.4.1 Spatio-temporal accuracy: comparing to strong formulation

We first verify our method for fibers in gravity by comparing the results to those obtained using the line-tension-based method prescribed in [15] and in the process show improved

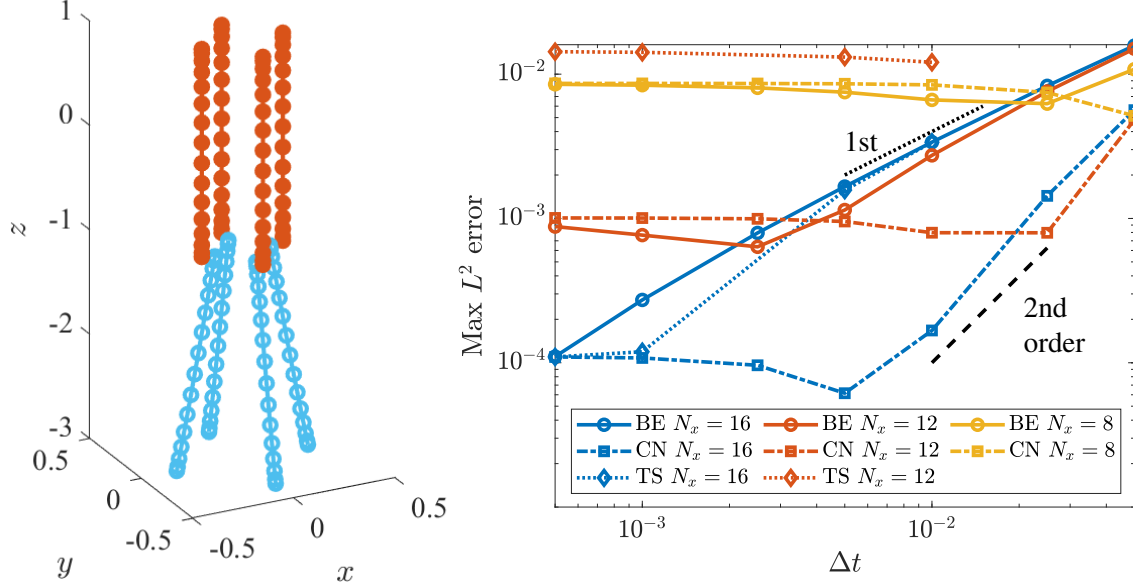


Figure 7.6: Four fibers in gravity. (a) Initial (filled orange) and final (light blue) configurations of the fibers. There are $N_x = 16$ points on each fiber. (b) Spatio-temporal convergence of our weak formulation (circles for backward Euler, squares for Crank-Nicolson) compared to the strong formulation of [15] (squares). In both cases, the exact solution is a trajectory with $N_x = 24$, $\Delta t = 5 \times 10^{-4}$. For small Δt , the spatial error dominates and the spatio-temporal error saturates.

robustness and temporal accuracy. We consider four fibers centered around a circle of radius $d = 0.2$. The fibers have initial tangent vector $\boldsymbol{\tau}(t = 0) = (0, 0, 1)$ (they are aligned in the z direction) and positions $\mathbf{X}^{(1)}(t = 0) = (d, 0, s - 1)$, $\mathbf{X}^{(2)}(t = 0) = (0, d, s - 1)$, $\mathbf{X}^{(3)}(t = 0) = (-d, 0, s - 1)$, and $\mathbf{X}^{(4)}(t = 0) = (0, -d, s - 1)$, where $0 \leq s \leq L = 2$. For simplicity, we set $\mu = \kappa = 1$, and $\epsilon = 10^{-3}$. For this test only, we use fibers with ellipsoidal cross sections and set the local drag coefficient $c(s) = -\ln(\epsilon^2)$ in (1.24). We simulate until $t = 0.25$.

Each fiber has a uniform gravitational force density $\mathbf{f}^g = (0, 0, -5)$ placed on it. In the absence of nonlocal interactions, the fibers fall straight downward. When the interactions between fibers are included, however, the fibers influence each other and have an x and y direction to their motion. Figure 7.6 (left panel) shows the initial and final ($t = 0.25$) configurations of the fibers in this test.

Our goal here is to verify our results by comparing to results obtained using the method

of [15]. Because the method of [15] uses regularization for the finite part integral (and we do not), we drop the finite part integral in this calculation and only include local drag and cross fiber interactions in the mobility. The fibers are sufficiently far from each other that the dipole term in the RPY kernel, which has coefficient of 0 in [15], has a very small effect on the result (two orders of magnitude less than our smallest spatio-temporal error). We use free space boundary conditions (no periodicity) and compute all nonlocal integrals by direct quadrature, without any upsampling. We use the block diagonal solves (7.39) for temporal integration and do *not* perform any GMRES iterations.

We first verify that our results match those of the strong formulation [15] when the spatial and temporal discretizations are well-refined. Considering $N_x = 24$ and $\Delta t = 5 \times 10^{-4}$ in both algorithms, we obtain a maximal L^2 difference of 1.8×10^{-4} in the position of the first fiber, which, as we show in Fig. 7.6, is on the order of magnitude of the discretization error.

It is instructive to compare the spatio-temporal error between the two algorithms. We define the “exact” solution for both algorithms to be a trajectory with $N = 24$ and $\Delta t = 5 \times 10^{-4}$. Figure 7.6 (right panel) shows the maximum L^2 errors over the time interval $[0, 0.25]$ for both algorithms with different spatial and temporal discretizations. For small Δt , the spatial error dominates and the spatio-temporal error saturates. We observe that our weak formulation outperforms the strong formulation of [15] in two ways. First, for coarse discretizations (e.g. $N = 12$, red lines in Fig. 7.6), our saturated spatial error is more than an order of magnitude lower than the saturated spatial error of [15]. This is likely because the line tension equation of [15] has larger aliasing errors for coarser spatial discretizations, and because of our improved treatment of the free fiber boundary conditions. Secondly, when we use the second-order Crank Nicolson discretization (dashed-dotted lines) our errors saturate at a much larger time step size than those of [15]. For example, when $N = 16$ our error saturates for $\Delta t = 5 \times 10^{-3}$, whereas the error from [15] does not saturate until $\Delta t = 1 \times 10^{-3}$. This occurs because the Crank-Nicolson is second-order accurate; indeed, we

see that the corresponding errors for Backward Euler (solid lines) saturate at the same value as those from [15], which is sensible since both use a backward Euler discretization. While there are limitations to this simple example, it does show the improved spatial accuracy of our new formulation over the tension-based one, and the improved accuracy of our second-order temporal discretization, even in the absence of GMRES iterations. In the next section, we will see that the applicability of this temporal integrator is limited, however, and so we should ultimately be comparing the backward Euler discretization to that developed in [15].

7.4.2 Spatio-temporal convergence: sheared fibers

We further study the spatio-temporal convergence of our algorithm for periodic sheared domains by choosing an example where interactions between the fibers contribute significantly to the dynamics. We consider three fibers in shear flow with $L = 2$ and $\mathbf{X}^{(1)}(s) = (s - 1, -0.6, -0.04)$, $\mathbf{X}^{(2)}(s) = (0, s - 1, 0)$, and $\mathbf{X}^{(3)}(s) = (s - 1, 0.6, 0.06)$. As shown in Figure 7.7, this corresponds to an “I” shaped initial configuration of the fibers, with the fibers stacked in the z direction.

We set the periodic domain length $L_d = 2.4$, the Ewald parameter $\xi = 4$, and set $\mu = 1$ and $\kappa = 0.01$. We use a constant shear flow (7.50) with $\gamma_0 = 1$ and $\omega = 0$. Because of the small bending rigidity, the fibers deform from their straight configurations in a shear flow. In this example, two of the fibers are initially aligned with the x direction. Without nonlocal hydrodynamics, they would stay aligned with the x direction and simply translate. When nonlocal interactions are included, the flows generated by the middle fiber $\mathbf{X}^{(2)}$ induce deformations of the top and bottom fibers. This is evident in Fig. 7.7, which shows the final fiber positions at $t = 2.4$ for $\epsilon = 10^{-3}$.




Figure 7.7: Fiber configurations at $t = 0$ and $t = 2.4$ for three fibers in shear flow with $\epsilon = 10^{-3}$.

7.4.2.1 Spatio-temporal convergence under self-refinement

We begin in Fig. 7.8 by looking at the spatio-temporal convergence of the fiber trajectories when the diagonal parts of the mobility $\widetilde{\mathbf{M}}_L$ are the RPY *special quadrature* matrices, and the off-diagonal parts are given by (7.10) with $N_u = 40$ upsampling points. We consider two different temporal integrators: backward Euler and Crank Nicolson, and show the errors relative to a reference solution with $N = 32$ and $\Delta t = 0.00125$ (and 9 additional GMRES iterations to remove any extrapolation errors). For $\epsilon = 10^{-3}$, the behavior is excellent: the Crank-Nicolson trajectories rapidly converge at a second order rate until the error is saturated by spatial error, which occurs around $\Delta t = 0.05$. Backward Euler has a slower (first order) rate of convergence, but we still obtain good accuracy (2 digits) with $\Delta t = 0.05$. When $\epsilon = 10^{-3}$, adding GMRES iterations makes little difference in the spatio-temporal errors, so it is not prudent to do any of these.

It is when we switch to $\epsilon = 10^{-2}$ that we start to see problems with our algorithm and interesting dynamics. With Crank-Nicolson, we cannot generate stable trajectories without GMRES iterations, but we can rapidly saturate the spatial errors by applying 2 additional GMRES iterations per time step (again at $\Delta t = 0.05$). The alternative is to use backward Euler, where we can obtain stable, convergent dynamics by solving the block diagonal system (7.45). Now, these time steps are four times cheaper than their Crank-Nicolson counterparts

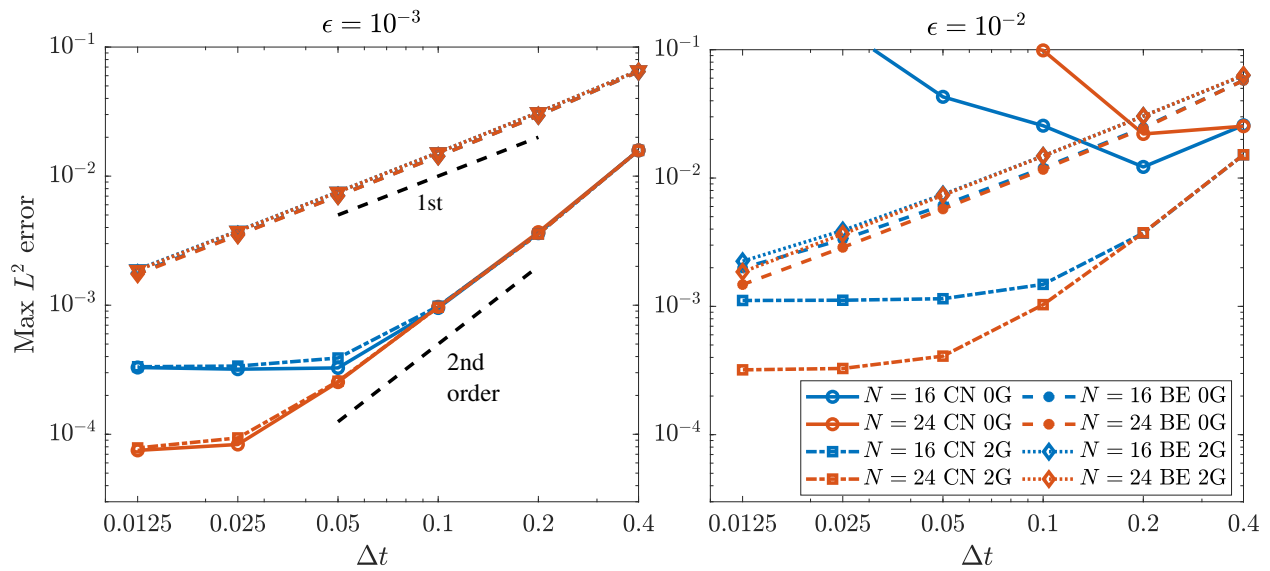


Figure 7.8: Spatio-temporal convergence for three sheared fibers with RPY special quadrature. The local mobility $\widetilde{\mathbf{M}}_L$ is given by the RPY special quadrature matrix. The nonlocal mobility $\widetilde{\mathbf{M}}_{NL}$ is given by upsampling the integrals via (7.10) with $N_u = 40$ points. In each case, a reference solution is obtained using $N = 32$, $\Delta t = 0.00125$, and 10 GMRES iterations to remove any extrapolation error (this time step is small enough for the Crank-Nicolson and backward Euler integrators to generate the same reference solution with a difference smaller than those observed here). Legend entries “2G” denote two additional GMRES iterations for the residual system (7.41)/ (7.47).




Figure 7.9: Error in the fiber trajectories using the special quadrature mobility (7.15), relative to a reference result obtained using oversampled RPY quadrature (7.10) with $1/\epsilon$ points. For a given N , we converge the trajectories in time and plot the spatial error using (7.15) as a solid blue line. We compare these results to the errors made in using oversampled RPY quadrature with a smaller number of points (red lines show $0.5/\epsilon$ points, yellow lines show $0.25/\epsilon$). We also consider the special quadrature mobility (7.15), but with twice the eigenvalue truncation level for $\widetilde{\mathcal{M}}_{\text{SQS}}$ in (7.14) (purple diamonds). We use $N_u = 40$ points for the nonlocal parts of the mobility in (7.15).

(2 GMRES iterations + 1 BD solve + 1 conversion to residual form), and so we need to compare backward Euler with time step size Δt to the Crank-Nicolson with GMRES iterations with time step size $4\Delta t$. Doing this, we observe that the errors are roughly the same. And so it makes no real difference what integrator we use for these suspensions. Going forward, we will use the more robust backward Euler scheme for testing.

7.4.2.2 Spatio-temporal error relative to oversampling

Now that we have studied the error with respect to self-refinement, the spatial error we really want to study is the error relative to oversampled RPY quadrature on *all* fibers. To obtain a reference solution, we use the oversampled RPY mobility (7.10) with $1/\epsilon$ points and converge the trajectory in time for $N = 8, 16$, and 24 . We then repeat the same runs using the oversampled mobility (7.10) with $0.5/\epsilon$ points and $0.25/\epsilon$ points. The idea is that sampling at fewer points provides a potential alternative to the special quadrature mobility

(7.15), which we also simulate with the same values of N .

Figure 7.9 shows the errors relative to oversampled RPY quadrature with $1/\epsilon$ points for these three mobility alternatives. We consider both $\epsilon = 10^{-3}$ and $\epsilon = 10^{-2}$, finding a similar result in both cases: the oversampled mobility (7.10) with $0.25/\epsilon$ points has a spatial error that is roughly 5 times as large as the oversampled mobility with $0.5/\epsilon$ points. The latter provides competition for special quadrature, giving similar errors for the values of N we consider. However, unlike oversampling with fewer points, the special quadrature approximation has an error which *decreases* with increasing N ; as N increases special quadrature approaches oversampled RPY with $1/\epsilon$ points. The same is not true with oversampling at fewer points, where the errors are roughly constant when N changes. Given that special quadrature is considerably cheaper than oversampling the RPY kernel at $0.5/\epsilon$ points (at least for $\epsilon = 10^{-3}$), we can clearly say that the mobility (7.15) has significant advantages in deterministic simulations of slender filaments.

It will also behoove us, as we show in the next section, to consider bumping up the level at which we truncate the eigenvalues of the special quadrature matrix $\widetilde{\mathbf{M}}_{\text{SQS}}$. We recall the procedure for constructing the special quadrature matrix, where we form $\widetilde{\mathbf{M}}_{\text{SQS}}$, then symmetrize in (7.14), and then truncate the eigenvalues to a threshold σ which is equal to the smallest eigenvalue of the reference mobility (7.10) for a straight fiber (see Section 7.2.1). In Fig. 7.9, we also consider trajectories where we truncate the eigenvalues at 2σ instead. The errors, shown in purple, behave similarly to the special quadrature errors in blue, in that they decrease with increasing N . For small N , the error is significantly higher (roughly doubled) when we truncate the eigenvalues at the higher level, which means that the truncated eigenvalues are contributing to the dynamics. For large N , the errors are smaller, and for $\epsilon = 10^{-2}$ the error is the same regardless of if we truncate at σ or 2σ . Thus the main takeaway is that bumping up the eigenvalue threshold has less of an effect on accuracy for larger N and ϵ . This will help us in Section 7.4.3.1, where we will need to bump up the

eigenvalue threshold to get robust convergence of GMRES for the nonlocal problem (7.47).

7.4.3 Stability

To examine stability in more detail, we consider a suspension of F fibers of length $L = 2$ and vary the density of fibers by changing F and the periodic domain length L_d . If $f = F/L_d^3$ is the number density of fibers and L is the length of a fiber, a dimensionless density $fL^3 < 1$ is considered a dilute fiber suspension, while a semi-dilute suspension is one with $afL^2 = \epsilon fL^3 \ll 1$, and a semi-concentrated one has $\epsilon fL^3 = \mathcal{O}(1)$ [10]. Here we explore the semi-dilute and semi-concentrated regimes and derive empirical bounds on how many GMRES iterations are required to maintain stability for a variety of bending moduli. Our conclusion is that at most five nonlocal hydrodynamic evaluations are sufficient to maintain stability, even for semi-concentrated suspensions.

We simulate F initially straight fibers. In the oscillatory shear flow (7.50), we set $\omega = 2\pi$ and $\dot{\gamma}_0 = \omega/10$, so that the maximum strain is $g = 0.1$ and the time for one cycle is 1. We expect to need at least 20 time steps per cycle to obtain reasonable accuracy, so we set $\Delta t = 0.05$, although in reality we find that smaller time step sizes are needed to accurately resolve the dynamics of dense suspensions. To obtain interesting results, we want a suspension in the semi-dilute or semi-concentrated regime where at least a few GMRES iterations are required for stability. Consequently, we use a domain size of $L_d = 2.4$ when $\epsilon = 10^{-3}$, for which we consider $F = 200$ fibers and $F = 1000$ fibers. This translates to $fL^3 \approx 120 \rightarrow \epsilon fL^3 \approx 0.12$ and $fL^3 \approx 600 \rightarrow \epsilon fL^3 \approx 0.6$, respectively, placing these suspensions in the semi-dilute and semi-concentrated regimes. (The corresponding mesh sizes $\ell_m = \sqrt{L_d^3/(FL)}$ are 0.19 and 0.08, respectively). For the more difficult case of $\epsilon = 10^{-2}$, we consider one suspension with $F = 200$ and $L_d = 5$, which gives $fL^3 = 12 \rightarrow \epsilon fL^3 = 0.12$ (semi-dilute, mesh size 0.56) and another with $F = 1000$ and $L_d = 3$, which gives $fL^3 \approx 300 \rightarrow \epsilon fL^3 = 3$ (semi-concentrated, mesh size 0.12).

7.4.3.1 Convergence of GMRES at $t = 0$

We first look at the convergence of GMRES at $t = 0$ in Fig. 7.10, where here the fibers have stiffness $\kappa = 0.01$ (the number for larger κ is slightly larger, but the overall behavior is the same). We consider two options for the mobility: oversampled quadrature with $1/\epsilon$ points (solid lines), and special quadrature for the intra-fiber (self interaction) term with oversampling at $N_u = 128$ points for the nonlocal integrals (dashed-dotted lines). While the number of oversampled points (128) is slightly large for $\epsilon = 10^{-2}$, the results are similar when we drop to $N_u = 40$ points for the nonlocal integrals, so we stick with $N_u = 128$ in Fig. 7.10. In both cases, the local mobility $\widetilde{\mathbf{M}}_L$ is the block diagonal part of the overall mobility $\widetilde{\mathbf{M}}$, and the preconditioner is given by (7.49).

Our first takeaway from Fig. 7.10 is the robustness of oversampled RPY quadrature: it always converges to 6 digits in some 30 iterations, and the convergence does not depend too strongly on N or ϵ (the number of iterations is a strong function of density, however, which is expected). Given this, for $\epsilon = 10^{-2}$ we can say that oversampled quadrature is a viable option for the mobility. But for $\epsilon = 10^{-3}$, it is simply too expensive to be a worthwhile choice.

Fortunately our special quadrature scheme can reduce the cost while not significantly affecting the convergence of GMRES. In the semi-dilute regime, which is the top row of Fig. 7.10, we observe that special quadrature (dashed-dotted lines) has the exact same convergence properties as oversampled quadrature (solid lines), with the exception of $N = 32$ and $\epsilon = 10^{-2}$, where we see some hiccups in the convergence. These problems likely have to do with changes in the smallest eigenvalues for this large N and ϵ (where our slender body quadrature is expected to lead to more spurious negative eigenvalues). The issues are magnified when we consider semi-concentrated suspensions, shown in the bottom row of Fig. 7.10. For slender fibers ($\epsilon = 10^{-3}$) and especially for small N , we see that our slender body quadrature scheme can still yield convergence, but for larger values of N and larger ϵ the




Figure 7.10: Convergence of GMRES at $t = 0$ for the two different mobility options (7.10) (over-sampled RPY, solid lines) and (7.15) (special quadrature for block diagonal terms, dashed-dotted lines). The top row shows semi-dilute suspensions ($\epsilon f L^3 \ll 1$), while the bottom row shows semi-concentrated suspensions ($\epsilon f L^3 \sim 1$). When we use special quadrature, the off-diagonal parts of the mobility are given by (7.10) with $N_u = 128$ points. For both mobilities, the preconditioner is given by (7.49) with $\widetilde{\mathbf{M}}_L = \text{BDiag}\{\widetilde{\mathbf{M}}\}$. The dotted lines show the special quadrature mobility with a new eigenvalue truncation level; see Section 7.4.3.2 for details on this.




Figure 7.11: Convergence of GMRES at $t = 0$ for the semi-dilute suspension with $N = 32$ points and $\epsilon = 10^{-2}$ (red lines in top right of Fig. 7.10). We show again the oversampled mobility as a solid line, and the special quadrature mobility as a dashed-dotted line (these are the same as in Fig. 7.10). We then look at mixing and matching preconditioners: the yellow line shows convergence when the mobility is the oversampled mobility (7.10), but $\tilde{\mathbf{M}}_L$ in the preconditioner (7.49) is the special quadrature mobility. The purple line shows the opposite case (special quadrature mobility (7.15), oversampled RPY for the preconditioner). The dotted red line shows the special quadrature mobility (7.15), but with twice the eigenvalue truncation threshold.

convergence stalls. Thus the semi-concentrated suspensions are at the limit of the validity of the numerical approximation (7.15), where we used special quadrature in the diagonal blocks and oversampled RPY in the off-diagonal ones.

7.4.3.2 Modifying the special quadrature matrix for robust convergence

Let us focus more closely on the problems with the special quadrature mobility for large N and ϵ . We consider in particular the top right of Fig. 7.10, where we run a semi-dilute suspension with $N = 32$ points per fiber and $\epsilon = 10^{-2}$. (While this is pushing the limits of where special quadrature is necessary, it is still important to understand what the problems are with special quadrature and how to fix them. In practice for $\epsilon = 10^{-2}$ it seems pretty clear that oversampled quadrature for all fibers is the best choice of mobility). To diagnose the




Figure 7.12: Eigenvalues of the full nonlocal mobility for various levels of special quadrature truncation, and the resulting GMRES convergence. We consider $F = 25$ fibers with $\epsilon = 10^{-2}$ and $N = 32$ points, so that we can form the full matrix $\tilde{\mathbf{M}}$ (size $3N_x F = 3(33)(25) = 2475$) explicitly and compute its eigenvalues. We show the oversampled mobility as a solid blue line, and the special quadrature mobility with various eigenvalue truncation levels using lines of other colors. The left panel shows the eigenvalues, with the inset showing the full spectrum and the main plot showing only the last 10%. The right panel shows the GMRES convergence that results (the preconditioner is the block diagonal part of the mobility in all cases). The results show that GMRES convergence is clearly tied to the eigenvalues of the matrix $\tilde{\mathbf{M}}$.

convergence problems, in Fig. 7.11 we vary the mobilities and preconditioners and see how the convergence changes. The solid blue and dashed-dotted red lines are the same as in Fig. 7.10, and show the oversampled and special quadrature mobilities, respectively, with preconditioners equal to the block diagonal parts of the mobility. We then consider two alternative cases: the solid yellow line shows the case when the mobility is computed using oversampled quadrature, but with the special quadrature preconditioner, and the dashed-dotted purple line shows the special quadrature mobility, but with an oversampled quadrature preconditioner. It is clear that, regardless of the preconditioner, taking the *mobility* to be the special quadrature matrix leads to stalls in the convergence of GMRES. Thus the issue is with the special quadrature mobility (7.14) itself, not with the preconditioner.

To understand the problem more, we reduce the number of fibers in the suspension

to $F = 25$, keeping $\epsilon = 10^{-2}$ and $N = 32$. We can now form the full matrix $\widetilde{\mathbf{M}}$ (size $3N_x F = 3(33)(25) = 2475$) explicitly and compute its eigenvalues, as we do in Fig. 7.12. At first glance (inset), the full spectra of the two matrices are indistinguishable, but a zoom in on the last 10% of the eigenvalues shows that there are indeed discrepancies. We see that the oversampled quadrature matrix (7.10) has step-like eigenvalues, which cutoff smoothly at 0.2, while the special quadrature matrices (7.15) have eigenvalues which seem to plateau at the correct value (here $\sigma = 0.18$, so the red line ought to plateau at 0.18), but then have a dip at the end. When the truncation threshold is too small (e.g., 1σ in this case), the eigenvalues get close to or drop below zero, which results in poor convergence of GMRES.

To work around this problem, we propose a simple increase in the eigenvalue truncation level to make the matrix more SPD. In our previous example of three sheared fibers (Fig. 7.9), we found that increasing the truncation level to 2σ has a negligible impact on the dynamics for larger N . By contrast, Figures 7.11 and 7.12 show that its impact on the convergence of GMRES is quite significant, as simply doubling the eigenvalue threshold yields roughly the same convergence as the oversampled RPY matrix, and eigenvalues that are bumped up away from zero. Using this fact, we return to Fig. 7.10, where in each case that the convergence stalls we demonstrate that using the truncation level of 2σ restores rapid GMRES convergence (the notable exception is the semi-concentrated suspension with $N = 32$ and $\epsilon = 10^{-2}$, which requires a slightly higher level of 5σ). The new truncation levels are needed primarily for large N and ϵ , which is precisely the regime where Fig. 7.10 shows that the eigenvalue truncation level has little effect on accuracy. Thus, when GMRES stalls, we can simply increase the truncation level without affecting the accuracy of our solution.

7.4.3.3 Stability in dynamic simulations

We now consider how many iterations are required to obtain *stability* in dynamic simulations. Since we find that changing the frequency ω has a negligible impact on the results, we non-

ϵ	N	SD/SC	$\Delta t_\kappa = 10^{-5}$	$\Delta t_\kappa = 10^{-4}$	$\Delta t_\kappa = 10^{-3}$	$\Delta t_\kappa = 10^{-2}$	$\Delta t_\kappa = 10^{-1}$
10^{-3}	16	SD	1	1	3	3	3
10^{-3}	16	SC	3	3	3	4	4
10^{-3}	16	SC*	3	3	3	4	4
10^{-3}	32	SD	1	1	3	3	3
10^{-3}	32	SC	3	3	3	4	4
10^{-3}	32	SC*	3	3	3	4	4
10^{-2}	16	SD	1	1	1	1	1
10^{-2}	16	SC	3	3	3	4	4
10^{-2}	16	SC*	3	3	3	4	4
10^{-2}	32	SD	1	1	1	1	1
10^{-2}	32	SC	5	5	4	4	4
10^{-3}	32	SC*	3	3	3	4	4

Table 7.1: Number of hydrodynamic evaluations needed for stability in semi-dilute (SD) and semi-concentrated (SC) deterministic fiber suspensions. The number of GMRES iterations for (7.47) is two less than the numbers in this table (except when the number is 1, for which no GMRES iterations are required). We report Δt in units of $\Delta t_\kappa = \Delta t / (1.72\tau)$, where the elastic timescale $\tau = 8\pi\mu L^4 / (\ln(\epsilon^{-2})\kappa)$ is varied by changing κ and $\Delta t = 0.05$ is fixed. Simulations marked with a * use 2σ as the eigenvalue truncation cutoff for $\widetilde{\mathbf{M}}_{\text{SQS}}$ in (7.15).

dimensionalize Δt by the bending timescale $\tau = 8\pi\mu L^4 / (\ln(\epsilon^{-2})\kappa)$, where here we use $\mu = 1$. We simulate 5 cycles of motion, until $t = 5$, and check if the simulation has gone “unstable” by examining the maximum position (absolute value of the Chebyshev point coordinates \mathbf{X}). For $L_d \geq 4$, we stop the simulation when the maximum position exceeds $2L_d$. For $L_d = 3$, we consider unstable to be a position of 7 or higher, and for $L_d = 2.4$, unstable means a position of 6 or higher.

In Table 7.1, we look at the number of hydrodynamic evaluations (one for the block diagonal system (7.45), one for the conversion to residual form (7.47), and one for each subsequent GMRES iteration) required for stability as a function of the number of fiber points N and slenderness ϵ . In this table, the mobility is always given by (7.15), with special quadrature for the self term (with truncation at the default value of 1σ), plus oversampling for the off-diagonal terms with $N_u = 128$ points (the preconditioner is always the block diagonal components of the mobility). We consider five different vales of $\Delta t_\kappa = \Delta t / (1.72\tau)$, where the elastic timescale $\tau = 8\pi\mu L^4 / (\ln(\epsilon^{-2}))$ is varied by changing κ and $\Delta t = 0.05$ is

fixed.

In the semi-dilute suspensions (top row of Fig. 7.10), the block diagonal solves (7.45) are sufficient to give stability except for $\epsilon = 10^{-3}$ and $\Delta t_\kappa \geq 10^{-3}$, for which one additional GMRES iteration is required for both $N = 16$ and $N = 32$. For semi-concentrated suspensions, the number of additional hydrodynamic evaluations is generally small (at most 4 additional evaluations) and independent of the number of collocation points N and slenderness ϵ . The one exception is $N = 32$ and $\epsilon = 10^{-2}$, which we were previously (Fig. 7.10) unable to converge at $t = 0$ without modifying the eigenvalue truncation cutoff for $\widetilde{\mathbf{M}}_{\text{SQS}}$ in (7.15). Thus, in Table 7.1, we obtain the iteration count for semi-concentrated suspensions in the case when we truncate the eigenvalues at level 2σ (starred rows). Increasing the truncation level reduces the iteration count only when $\epsilon = 10^{-2}$ and $N = 32$, which was most problematic to converge at $t = 0$ when the truncation level was 1σ .¹ Even when the truncation level is 1σ , however, we still get stability for all semi-concentrated suspensions, despite the fact that we never were able to converge the full implicit system at $t = 0$ (Fig. 7.10). This is unsurprising because stability is a weaker restriction than accuracy.

7.4.4 Bundled systems

We now look at how the number of required iterations for the two different mobility options changes when we introduce bundling into the system. The model system we consider is 200 filaments of length $L = 1$ in a box of size $L_d = 2$, with $N = 12$, $\epsilon = 4 \times 10^{-3}$, and all of the other parameters that we use in Part III (see there and Table 11.1 for more details on the bundling dynamics). Here we set the persistence length $\ell_p = 17 \mu\text{m}$ to be equal to that of an actin filament. The precise definitions of bundling are in Chapter 12; here it suffices for the reader to know that more bundles means fibers are closer together, so that the mobility is dominated less and less by intra-fiber hydrodynamics as the bundle density increases. See

¹Increasing the eigenvalue truncation level to 2σ is a substantially better strategy than dropping the time step, as decreasing to $\Delta t = 0.01$ has no effect on the numbers in Table 7.1 for $\epsilon = 10^{-2}$ and $N = 32$.

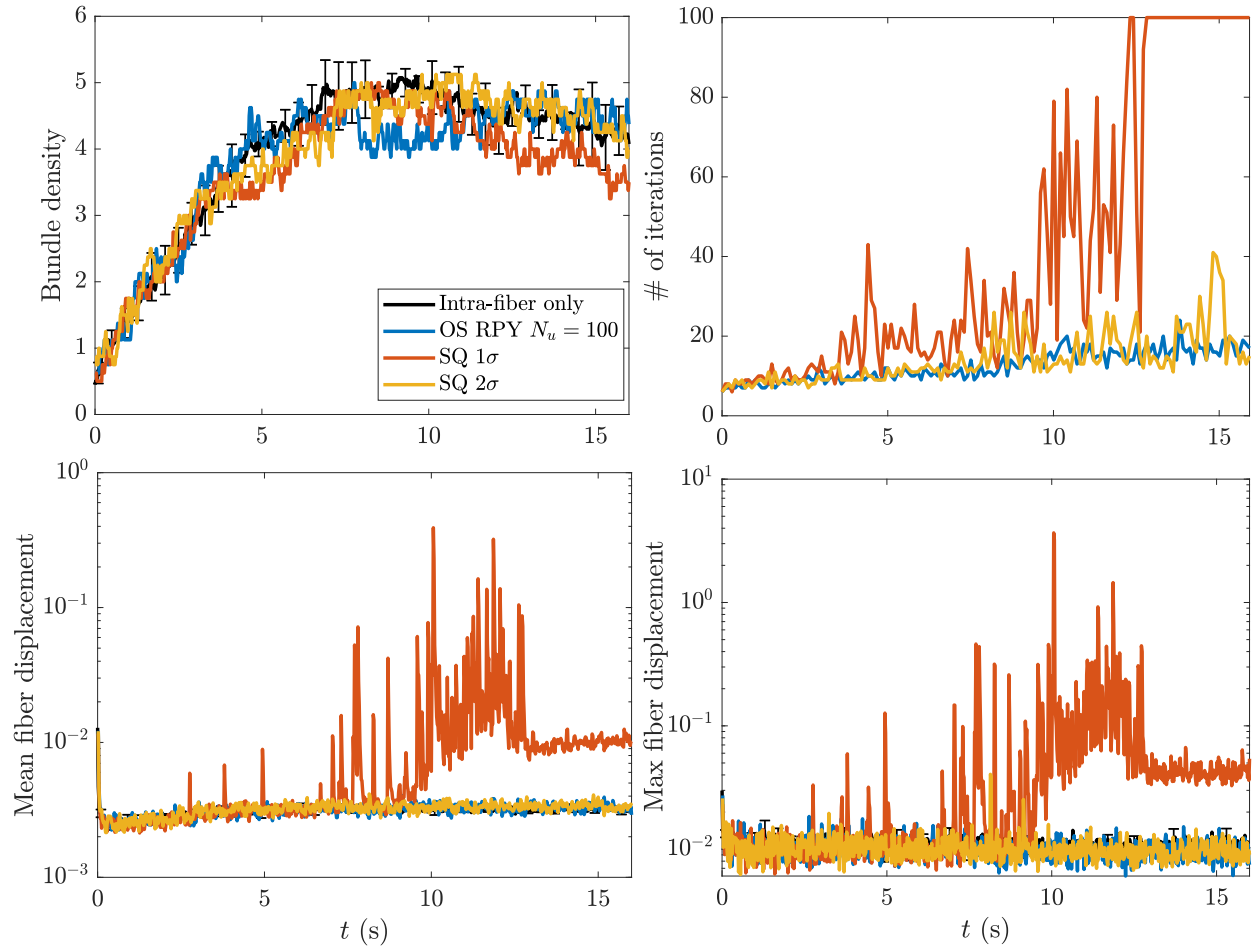


Figure 7.13: Stability and accuracy of nonlocal hydrodynamic mobility in bundled systems. We show over time the bundle density, the number of required iterations to converge (7.47) to a tolerance 10^{-3} , and the mean and maximum fiber displacements (over a time span of 0.02 s), considering three different mobilities: the oversampled RPY quadrature (7.10) with $N_u = 100$ oversampling points (blue), and the special quadrature mobility (for the diagonal parts) (7.15), with cutoffs 1σ (red) and 2σ (yellow). Black lines are the intra-fiber mobility only (no inter-fiber interactions), and are shown for reference.

Figs. 12.1 and 12.2 for how the system looks at a certain “bundle density.”

In order to look at the behavior of the mobility in time, at each time step we converge the residual system (7.47) to a tolerance of 10^{-3} . We consider three different choices for the mobility: the oversampled RPY quadrature (7.10) with $N_u = 100$ oversampling points, and the special quadrature mobility (for the diagonal parts) (7.15), with cutoffs 1σ and 2σ . Figure 7.13 shows how the trajectory of the bundling process plays out in these cases: we see that stable dynamics are only obtained at late times when we use a truncation level of 2σ in the special quadrature matrix (7.14). This tells us that the behavior of the overall mobility (7.15) (with special quadrature on the diagonal) can be sensitive to the amount of bundling in the system.

We examine the convergence of GMRES further in Fig. 7.14, where we converge GMRES at times 0, 5, 10, and 15 s to a tolerance of 10^{-6} . At $t = 0$, when the fibers are isotropically and uniformly distributed in the domain, the convergence is essentially the same regardless of the mobility. As bundling increases, however, the convergence of the special quadrature mobility with truncation level 1σ stalls, and eventually we see instabilities and no convergence at all. This same pattern might be starting to happen for truncation level 2σ as well, as at $t = 15$ we do see some strange convergence behavior.




Figure 7.14: GMRES convergence at various times in deterministic bundled systems. We show the convergence of GMRES to tolerance 10^{-6} at times $t = 0, 5, 10,$ and 15 (see Fig. 7.13 for system statistics at these times), considering three different mobilities: the oversampled RPY quadrature (7.10) with $N_u = 100$ oversampling points (blue), and the special quadrature mobility (for the diagonal parts) (7.15), with cutoffs 1σ (red) and 2σ (yellow).

Chapter 8

Spectral discretization for Brownian fluctuations

In Chapter 7, we developed a discretization for free-ended filaments. Our focus in particular in Sections 7.1 and 7.2 was on reformulating the deterministic dynamics so the velocity $\partial_t \mathbf{X}$ takes the form $\partial_t \mathbf{X} = -\widetilde{\mathbf{M}} \mathbf{L} \mathbf{X}$, where $\widetilde{\mathbf{M}}$ is an SPD matrix and \mathbf{L} is the matrix discretizing the bending energy. This included making the special quadrature matrix $\widetilde{\mathbf{M}}_{\text{SQS}}$ in (7.14) SPD by symmetrizing and truncating the eigenvalues, and defining a reference mobility (7.10) that applies to multiple fibers. Thus we have a discretization where $\widetilde{\mathbf{M}}$ is guaranteed to be SPD, which allows us to develop a temporal integrator which takes into account the Brownian fluctuations of the filament. This integrator satisfies fluctuation-dissipation balance with respect to the constrained inextensible Gibbs-Boltzmann distribution, which is the first thing we introduce here.

Seeing that this is, to our knowledge, the first work to consider a spectral discretization of constrained Brownian hydrodynamics, our program is very much experimental, in part because our premise is counterintuitive. The conventional wisdom, as taught in a first semester of numerical analysis, is that spectral methods perform well for *smooth* problems (with exponentially-decaying spectra), while for nonsmooth ones (power law spectra) a sparse

finite difference method is a better choice. While it is certainly true that Brownian motion induces a power law spectrum in the fiber positions, our hope is that restricting to *semiflexible* fibers (with $\ell_p/L \gtrsim 1$) will make a spectral discretization tractable by reducing the number of modes required for accurate simulation. In other words, as long as the Chebyshev grid resolves the fiber *persistence length* (as opposed to the fiber radius), many quantities should be approximated well since lengthscales smaller than ℓ_p will not contribute. We demonstrate that this is indeed the case through examples, but it should be noted that this is just a first step, with Chapter 10 laying out much of the work that remains.

8.1 Langevin equation for semiflexible filaments

The previous chapters (in particular Sections 6.1.1 and 7.1) established our discrete fiber model. From this, we can write the Gibbs-Boltzmann probability distribution for the fiber energy. The overdamped Langevin equation then follows from this and the deterministic dynamics (7.3).

8.1.1 Gibbs-Boltzmann distribution

Once we define a discretization, we can postulate the constrained Gibbs-Boltzmann distribution in terms of the tangent vectors $\boldsymbol{\tau}$. This distribution depends on the fiber elastic energy $\mathcal{E}_{\text{bend}}$, which we discretize using the method of Section 6.1.3.1. Following the notation there, we let $\mathcal{E}_{\text{bend}}(\bar{\boldsymbol{\tau}})$ denote the discrete bending energy of the fiber, and take the Gibbs-Boltzmann distribution to be

$$dP_{\text{eq}}(\bar{\boldsymbol{\tau}}) = Z^{-1} \exp(-\mathcal{E}_{\text{bend}}(\bar{\boldsymbol{\tau}})/k_B T) d\mu_0(\bar{\boldsymbol{\tau}}) \quad (8.1)$$

$$P_{\text{eq}}(\bar{\boldsymbol{\tau}}) = Z^{-1} \exp(-\mathcal{E}_{\text{bend}}(\bar{\boldsymbol{\tau}})/k_B T) \prod_{p=1}^N \delta(\boldsymbol{\tau}_{\{p\}}^T \boldsymbol{\tau}_{\{p\}} - 1).$$

What we mean by the product of δ functions is really that the base measure $d\mu_0(\bar{\boldsymbol{\tau}})$ in the Gibbs-Boltzmann distribution (8.1) corresponds to the tangent vectors $\boldsymbol{\tau}_{\{p\}}$ being independently uniformly distributed on the unit sphere for $p = 1, \dots, N$. For blob-link chains, the tangent vectors $\boldsymbol{\tau}_{\{p\}}$ are uniformly spaced, as shown in Fig. 6.1, and this distribution follows naturally from their inextensibility and statistical independence in a freely-jointed chain, in which case (8.1) defines the standard worm-like chain model.

In our spectral discretization, the tangent vectors $\boldsymbol{\tau}_{\{p\}}$ are defined on a Chebyshev grid, and this model lacks the physical motivation that characterizes the blob-link model (see Fig. 6.1). In fact, the base measure $d\mu_0$ on a uniform (blob-link) grid has no continuum limit, since a freely-jointed continuum chain would require the tangent vector $\boldsymbol{\tau}(s)$ to perform infinitely fast Brownian motion in s on the unit sphere; thus only (8.1) with the Gibbs-Boltzmann weight makes sense in continuum. Furthermore, a Chebyshev *polynomial* cannot represent an inextensible curve everywhere (since the number of zeros in $\boldsymbol{\tau} \cdot \boldsymbol{\tau} - 1$ is limited by the polynomial degree), and therefore it is not obvious how to uniquely define a discrete equilibrium distribution for a spectral grid that is a good approximation to the continuum worm-like chain (when the Chebyshev grid resolves the persistence length)¹. In the first step toward a more mathematically-justified approach, we postulate (8.1) as a reasonable guess, but do not claim any precise sense of convergence of (8.1) for a spectral grid to the equilibrium distribution of a continuum worm-like chain. We do show (in Sections 8.4.1 and 8.3) that for a sufficient number of Chebyshev nodes over the persistence length, the equilibrium distribution for a spectral chain approximates well the equilibrium distribution for the blob-link chain, in the sense that samples from the two distributions give the same large-scale statistical properties of the chain, e.g., its end-to-end distance.

¹Of course, we could also take $d\mu_0$ as the base measure and put additional “entropic” or “metric” factors in $\mathcal{E}_{\text{bend}}$. Here, however, we take $\mathcal{E}_{\text{bend}}$ to be the standard elastic bending energy of a semiflexible chain.

8.1.2 Overdamped Langevin dynamics

To formulate the overdamped Langevin equation, we convert the saddle point system (7.3) into an equation for \mathbf{X} and $\bar{\boldsymbol{\tau}}$. This can be done by eliminating the Lagrange multipliers $\boldsymbol{\Lambda}$ from (7.3) to yield

$$\partial_t \mathbf{X} = \mathbf{K} \boldsymbol{\alpha} = -(\mathbf{K} \mathbf{N} \mathbf{K}^T) \mathbf{L} \mathbf{X}, \quad (8.2)$$

$$\text{where } \mathbf{N} = \left(\mathbf{K}^T \widetilde{\mathbf{M}}^{-1} \mathbf{K} \right)^\dagger \quad (8.3)$$

is the mobility matrix projected onto the space of inextensible motions (the discrete form of the continuum operator \mathcal{N} defined in (4.3)). We can transform this equation to obtain the evolution of $\bar{\boldsymbol{\tau}}$ by substituting $\mathbf{X} = \boldsymbol{\chi} \bar{\boldsymbol{\tau}}$ and $\mathbf{K} = \boldsymbol{\chi} \bar{\mathbf{C}}$ into (8.2) to obtain

$$\partial_t \bar{\boldsymbol{\tau}} = \boldsymbol{\chi}^{-1} (\mathbf{K} \mathbf{N} \mathbf{K}^T) \mathbf{L} \boldsymbol{\chi} \bar{\boldsymbol{\tau}} = \left(\bar{\mathbf{C}} \mathbf{N} \bar{\mathbf{C}}^T \right) \boldsymbol{\chi}^T \mathbf{L} \boldsymbol{\chi} \bar{\boldsymbol{\tau}} = \left(\bar{\mathbf{C}} \mathbf{N} \bar{\mathbf{C}}^T \right) \mathbf{L}_\tau \bar{\boldsymbol{\tau}}, \quad (8.4)$$

where \mathbf{L}_τ is defined in (6.12). We note that these equations apply for continuous time. In discrete time, we will solve the saddle point system (7.3) for $\boldsymbol{\alpha} = (\boldsymbol{\Omega}, \mathbf{U}_{\text{MP}})$, then update the tangent vectors on the unit sphere by $\boldsymbol{\Omega} \Delta t$ using the nonlinear rotation (6.71).

8.1.2.1 Overdamped Langevin equation in $\bar{\boldsymbol{\tau}}$

The deterministic equation (8.4) takes the same form as that for rigid bodies, and in the blob-link picture it can be seen as describing the motion of a series of N connected rigid rods. Following the same process used to derive the overdamped Ito Langevin equation [116, Eq. (12)] in the rigid body case, we have the overdamped Ito Langevin equation

$$\partial_t \bar{\boldsymbol{\tau}} = - \left(\bar{\mathbf{C}} \mathbf{N} \bar{\mathbf{C}}^T \right) \mathbf{L}_\tau \bar{\boldsymbol{\tau}} + k_B T \partial_{\bar{\boldsymbol{\tau}}} \cdot \left(\bar{\mathbf{C}} \mathbf{N} \bar{\mathbf{C}}^T \right) + \sqrt{2k_B T} \bar{\mathbf{C}} \mathbf{N}^{1/2} \boldsymbol{\mathcal{W}} \quad (8.5)$$

describing the evolution of $\bar{\tau}$. Here $\mathbf{W}(t)$ is a collection of white noise processes (the formal derivative of Brownian motions), the divergence with respect to $\bar{\tau}$ is defined in (6.27), and $\mathbf{N}^{1/2}$ satisfies the fluctuation-dissipation relation

$$\mathbf{N}^{1/2} \left(\mathbf{N}^{1/2} \right)^T = \mathbf{N}. \quad (8.6)$$

The arguments in [116, Sec. II.B.2] can be used to show that (8.5) is time-reversible with respect to the Gibbs-Boltzmann equilibrium distribution (8.1). As described there, because we formulate (8.5) with respect to τ and not \mathbf{X} , there are no additional drift terms arising from metric/entropic forces.

An important point in practice is that $\mathbf{N}^{1/2}$, which is not unique and only needs to satisfy (8.6), can be applied by adding Brownian noise with covariance $\widetilde{\mathbf{M}}$ to the right-hand side of saddle point system (7.3). In an Euler-Maruyama discretization, this corresponds to solving the saddle point system

$$\begin{aligned} \mathbf{K}\boldsymbol{\alpha} &= \widetilde{\mathbf{M}}(-\mathbf{L}\mathbf{X} + \boldsymbol{\Lambda}) + \sqrt{\frac{2k_B T}{\Delta t}} \widetilde{\mathbf{M}}^{1/2} \boldsymbol{\eta} \\ \mathbf{K}^T \boldsymbol{\Lambda} &= \mathbf{0}, \end{aligned} \quad (8.7)$$

where $\boldsymbol{\eta}$ is an i.i.d. vector of standard normal random variables. As with the deterministic saddle point system (7.3), a variational technique can be used to derive the saddle point system (8.7) by adding the work done (entropy dissipated) by the random force $\mathbf{F}_{\text{stoch}}$

$$\left(\frac{\mathbf{X} - \mathbf{X}^n}{\Delta t} \right)^T \mathbf{F}_{\text{stoch}} \Delta t = \left(\frac{\mathbf{X} - \mathbf{X}^n}{\Delta t} \right)^T \sqrt{\frac{2k_B T}{\Delta t}} \left(\widetilde{\mathbf{M}}^n \right)^{-1/2} \boldsymbol{\eta}^n \Delta t, \quad (8.8)$$

to the Lagrangian (7.5). This does not account for additional stochastic drift terms, which are required to ensure time-reversibility (detailed balance). So far, we have not been able to produce those terms via a variational argument.

As discussed at length in [146, Sec. II(B)], solving the saddle point system (8.7) gives

$$\begin{aligned}\partial_t \mathbf{X} = \mathbf{K} \boldsymbol{\alpha} &= - (\mathbf{K} \mathbf{N} \mathbf{K}^T) \mathbf{L} \mathbf{X} + \sqrt{\frac{2k_B T}{\Delta t}} (\mathbf{K} \mathbf{N} \mathbf{K}^T) \widetilde{\mathbf{M}}^{-1/2} \boldsymbol{\eta} \\ &= - (\mathbf{K} \mathbf{N} \mathbf{K}^T) \mathbf{L} \mathbf{X} + \sqrt{\frac{2k_B T}{\Delta t}} \mathbf{K} \mathbf{N}^{1/2} \boldsymbol{\eta},\end{aligned}\tag{8.9}$$

giving $\mathbf{N}^{1/2} = \mathbf{N} \mathbf{K}^T \widetilde{\mathbf{M}}^{-1/2}$. Thus, generating noise of the form $\mathbf{N}^{1/2} \boldsymbol{\eta}$ reduces to the simpler process of solving a saddle point system with right hand side $\widetilde{\mathbf{M}}^{-1/2} \boldsymbol{\eta}$. For hydrodynamics which is localized to each fiber, we do this using the eigenvalue decomposition of $\widetilde{\mathbf{M}}$, which already must be computed for the purposes of eigenvalue truncation (see Section 7.2.1). Given that $\mathbf{N}^{1/2}$ can also be computed via dense linear algebra, the real savings in the saddle point solve come when we need to generate $\mathbf{N}^{1/2}$ with *nonlocal* hydrodynamics (between the many fibers), where dense linear algebra is infeasible, but the action of $\widetilde{\mathbf{M}}^{1/2}$ can be computed via the Lanczos algorithm [147] or the positively-split Ewald method [103] (which uses a closed-form Cholesky factor for the square root of the far field part and Lanczos iterations for the square root of the near field). The latter is how we form $\widetilde{\mathbf{M}}^{1/2}$ when we add nonlocal hydrodynamics in Section 8.6.

8.1.2.2 Overdamped Langevin equation in \mathbf{X}

We now use (8.5) to derive an overdamped Langevin equation in terms of \mathbf{X} . If we multiply (8.5) on the left by $\boldsymbol{\mathcal{X}}$ and expand $\mathbf{L}_\tau = \boldsymbol{\mathcal{X}}^T \mathbf{L} \boldsymbol{\mathcal{X}}$, we obtain

$$\partial_t \mathbf{X} = - (\mathbf{K} \mathbf{N} \mathbf{K}^T) \mathbf{L} \mathbf{X} + k_B T \partial_{\bar{\tau}} \cdot (\mathbf{K} \mathbf{N} \bar{\mathbf{C}}^T) + \sqrt{2k_B T} \mathbf{K} \mathbf{N}^{1/2} \boldsymbol{\mathcal{W}}.\tag{8.10}$$

Using the chain rule to write differentiation with respect to \mathbf{X} as

$$\frac{\partial f}{\partial \mathbf{X}_k} = \frac{\partial f}{\partial \bar{\tau}_p} \frac{\partial \bar{\tau}_p}{\partial \mathbf{X}_k} = \frac{\partial f}{\partial \bar{\tau}_p} \boldsymbol{\mathcal{X}}_{pk}^{-1}$$

We now rewrite the divergence in (8.10) as

$$\frac{\partial}{\partial \bar{\tau}_j} \left(\mathbf{K} \mathbf{N} \bar{\mathbf{C}}^T \right)_{ij} = \frac{\partial}{\partial \mathbf{X}_k} \left(\mathbf{K} \mathbf{N} \bar{\mathbf{C}}^T \right)_{ij} \frac{\partial \mathbf{X}_k}{\partial \bar{\tau}_j} = \frac{\partial}{\partial \mathbf{X}_k} \left(\mathbf{K} \mathbf{N} \bar{\mathbf{C}}^T \right)_{ij} \boldsymbol{\chi}_{kj} = \frac{\partial}{\partial \mathbf{X}_k} \left(\mathbf{K} \mathbf{N} \mathbf{K}^T \right)_{ik} \quad (8.11)$$

so that the Ito equation (8.5) could equivalently be formulated in terms of \mathbf{X} as might be expected from the deterministic equation (8.2),

$$\partial_t \mathbf{X} = - \left(\mathbf{K} \mathbf{N} \mathbf{K}^T \right) \mathbf{L} \mathbf{X} + k_B T \partial_{\mathbf{X}} \cdot \left(\mathbf{K} \mathbf{N} \mathbf{K}^T \right) + \sqrt{2k_B T} \mathbf{K} \mathbf{N}^{1/2} \boldsymbol{\mathcal{W}} \quad (8.12)$$

$$\stackrel{d}{=} -\widehat{\mathbf{N}} \mathbf{L} \mathbf{X} + \sqrt{2k_B T} \widehat{\mathbf{N}} \circ \widehat{\mathbf{N}}^{-1/2} \boldsymbol{\mathcal{W}}, \quad (8.13)$$

where $\widehat{\mathbf{N}} = \mathbf{K} \mathbf{N} \mathbf{K}^T$ and the second equality denotes that paths of (8.12) and (8.13) have the same probability distribution. Equation (8.13) is the Langevin equation written in a split Stratonovich-Ito [116] or kinetic [148] form, where the terms before the \circ are evaluated at the midpoint of a given time step, while the terms after are evaluated at the beginning of the time step (c.f. [116, Eq. (26)]). When we develop our numerical methods, we will do so with the equation for \mathbf{X} in mind, since ultimately we will evolve and track the fiber positions. That said, it will be simpler when analyzing the Langevin equation to work with the equation (8.5) for $\bar{\boldsymbol{\tau}}$, knowing that the one in \mathbf{X} can be obtained by this simple transformation. Note that (8.12) is much simpler than the overdamped equations derived previously for bead-link models [114, 115].

8.2 Temporal integration

In this section, we discuss our temporal integrator for the overdamped Langevin equation (8.5). The scheme, which is in the spirit of the Fixman method [149] and similar to that of Westwood et. al for rigid bodies [150], is able to integrate the overdamped Langevin equation using one saddle-point solve per time step. The key idea is to first move to the midpoint

to compute the mobility, then solve a saddle point system using the midpoint values, which generates the required drift term in expectation.

Before introducing our numerical scheme, it is helpful to simplify the drift term in (8.5).

We first separate it into three terms,

$$\frac{\partial}{\partial \bar{\boldsymbol{\tau}}_j} \left(\bar{\mathbf{C}} \mathbf{N} \bar{\mathbf{C}}^T \right)_{ij} = (\partial_j \bar{\mathbf{C}}_{ik}) \mathbf{N}_{kp} \bar{\mathbf{C}}_{pj}^T + \bar{\mathbf{C}}_{ik} (\partial_j \mathbf{N}_{kp}) \bar{\mathbf{C}}_{pj}^T + \bar{\mathbf{C}}_{ik} \mathbf{N}_{kp} (\partial_j \bar{\mathbf{C}}_{pj}^T), \quad (8.14)$$

where ∂_j is shorthand for $\partial/\partial \bar{\boldsymbol{\tau}}_j$. As shown in [116, Sec. III(A)], rotating the tangent vectors at every time step n using the Euler-Maruyama method

$$\boldsymbol{\tau}^{n+1} = \text{rotate} \left(\boldsymbol{\tau}^n, \sqrt{2k_B T \Delta t} (\mathbf{N}^n)^{1/2} \boldsymbol{\eta}^n + \mathcal{O}(\Delta t) \right) \quad (8.15)$$

where $\boldsymbol{\eta}^n$ is a vector of i.i.d. standard normal random variables, is sufficient to capture the first drift term. The third term in (8.14) is zero by (6.27). Therefore, our schemes simply need to generate the additional drift term

$$(k_B T) \bar{\mathbf{C}}_{ik} (\partial_j \mathbf{N}_{kp}) \bar{\mathbf{C}}_{pj}^T \Delta t \quad \text{i.e.,} \quad (k_B T) \bar{\mathbf{C}} \left(\nabla_{\bar{\boldsymbol{\tau}}} \mathbf{N} : \bar{\mathbf{C}}^T \right) \Delta t \quad (8.16)$$

This term corresponds to $\mathcal{O}(\Delta t^2)$ to a rotate procedure by an angle $\Delta t (k_B T) \nabla_{\bar{\boldsymbol{\tau}}} \mathbf{N} : \bar{\mathbf{C}}^T$. Thus our task will be to design numerical methods to produce the stochastic drift term in $\boldsymbol{\Omega}$

$$\text{Drift} = \nabla_{\bar{\boldsymbol{\tau}}} \mathbf{N} : \bar{\mathbf{C}}^T \quad \text{i.e.,} \quad \text{Drift}_k = (\partial_j \mathbf{N}_{kp}) \bar{\mathbf{C}}_{jp}^T \quad (8.17)$$

in expectation, which will give (8.16) after rotation.

8.2.1 Implicit methods

We first motivate our temporal discretization of (8.5) by considering the discretization of the unconstrained *linearized* SDE

$$\partial_t \mathbf{X} = -\widetilde{\mathbf{M}}[\mathbf{X}_0] \mathbf{L} \Delta \mathbf{X} + \mathbf{U}_B, \quad (8.18)$$

where \mathbf{X}_0 is an equilibrium position, $\Delta \mathbf{X} = \mathbf{X} - \mathbf{X}_0$, the matrix \mathbf{L} discretizes the energy $\mathcal{E}_{\text{bend}} = \frac{1}{2} \Delta \mathbf{X}^T \mathbf{L} \Delta \mathbf{X}$, and \mathbf{U}_B is the Brownian velocity given by fluctuation-dissipation balance as $\sqrt{2k_B T} \left(\widetilde{\mathbf{M}}[\mathbf{X}_0] \right)^{1/2} \boldsymbol{\mathcal{W}}(t)$. Because this SDE is unconstrained, the equilibrium covariance of $\Delta \mathbf{X}$ is known from statistical mechanics,

$$\mathbb{E} [\Delta \mathbf{X} \Delta \mathbf{X}^T] = k_B T \mathbf{L}^{-1} \quad (8.19)$$

In our case, the bending force resulting from the matrix \mathbf{L} is very stiff (fourth derivative), and so we need to discretize it implicitly. Our goal is to design numerical methods that preserve the covariance (8.19) for arbitrary Δt when applied to (8.18). To do this, we follow the analysis of [151, Sec. III(B)] to derive the steady state covariance for a given temporal integrator.

We consider an implicit-explicit method for (8.18) of the form

$$\mathbf{X}^{n+1} = \mathbf{X}^n + \Delta t \left(-\widetilde{\mathbf{M}} \mathbf{L} (c \mathbf{X}^{n+1} + (1-c) \mathbf{X}^n) + \mathbf{U}_B^n \right),$$

which can be rearranged to yield

$$\left(\mathbf{I} + c \Delta t \widetilde{\mathbf{M}} \mathbf{L} \right) \mathbf{X}^{n+1} = \left[\left(\mathbf{I} - (1-c) \Delta t \widetilde{\mathbf{M}} \mathbf{L} \right) \mathbf{X}^n + \Delta t \mathbf{U}_B^n \right].$$

We now take an outer product of the two sides of the equation, then substitute the desired

covariance from (8.19), which at steady state is independent of the time step n , to obtain the matrix equation

$$2k_B T \Delta t \widetilde{\mathbf{M}} + k_B T (2c - 1) \Delta t^2 \widetilde{\mathbf{M}} \mathbf{L} \widetilde{\mathbf{M}} = \Delta t^2 \mathbb{E} [\mathbf{U}_B \mathbf{U}_B^T].$$

To obtain the exact covariance for $c = 1$ (backward Euler), we therefore set

$$\mathbf{U}_B = \sqrt{\frac{2k_B T}{\Delta t}} \left(\widetilde{\mathbf{M}}^{1/2} \boldsymbol{\eta} + \sqrt{\frac{\Delta t}{2}} \widetilde{\mathbf{M}} \mathbf{L}^{1/2} \widetilde{\boldsymbol{\eta}} \right), \quad (8.20)$$

where $\widetilde{\boldsymbol{\eta}}$ is another standard normal random vector. Another option is to use Crank-Nicolson ($c = 1/2$), which gives the exact covariance for arbitrary Δt with the usual Brownian velocity $\sqrt{2k_B T / \Delta t} \widetilde{\mathbf{M}}^{1/2} \boldsymbol{\eta}$. While this choice has been preferred for other applications [151], we find it to be less accurate than our “modified” backward Euler scheme for the higher order modes, which take too long to equilibrate using $c = 1/2$. As such, we will use $c = 1$ and the Brownian velocity (8.20) throughout this paper, where we can precompute $\mathbf{L}^{1/2}$ via Cholesky or eigenvalue (dense matrix) decomposition (since \mathbf{L} is a block diagonal matrix for a suspension of fibers). When we have constraints and nonlinear updates, the velocity does not generate the exact covariance for arbitrary Δt , but it gives a covariance which converges more rapidly to the correct answer.

8.2.2 RFD scheme

We first present an RFD scheme which explicitly tells us what drift terms are being included, so that we can exclude them to demonstrate that (8.5) *without* the correct drift terms does not give the correct equilibrium statistical mechanics. This scheme begins by solving the

system of equations

$$\widetilde{\mathbf{M}}^n (\mathbf{L} (c\mathbf{X}^{n+1,*} + (1-c)\mathbf{X}^n) + \boldsymbol{\Lambda}) + \mathbf{U}_B^n = \mathbf{K}^n \boldsymbol{\alpha} \quad (8.21)$$

$$(\mathbf{K}^n)^T \boldsymbol{\Lambda} = \mathbf{0},$$

$$\text{where } \mathbf{X}^{n+1,*} = \mathbf{X}^n + \Delta t \mathbf{K}^n \boldsymbol{\alpha} \quad (8.22)$$

is the linearized position at the next time step that we use in the implicit method determined by the coefficient c . Substituting (8.22) into (8.21), we begin by solving the saddle point system

$$\begin{pmatrix} -\widetilde{\mathbf{M}} & (\mathbf{I} + c\Delta t \widetilde{\mathbf{M}}\mathbf{L}) \mathbf{K} \\ \mathbf{K}^T & \mathbf{0} \end{pmatrix}^n \begin{pmatrix} \boldsymbol{\Lambda} \\ \boldsymbol{\alpha} \end{pmatrix}^n = \begin{pmatrix} -\widetilde{\mathbf{M}}\mathbf{L}\mathbf{X}^n + \mathbf{U}_B^n \\ \mathbf{0} \end{pmatrix}, \quad (8.23)$$

where all time-dependent quantities are evaluated at time n , the Brownian velocity \mathbf{U}_B is defined in (8.20), and $\boldsymbol{\alpha} = (\boldsymbol{\Omega}, \mathbf{U}_{\text{MP}})$. Using the Schur complement to eliminate $\boldsymbol{\Lambda}$, we obtain the solution

$$\begin{aligned} \boldsymbol{\alpha}^n &= \left(-\mathbf{N}\mathbf{K}^T \mathbf{L}\mathbf{X} + \sqrt{\frac{2k_B T}{\Delta t}} \mathbf{N}^{1/2} \boldsymbol{\eta} \right)^n + \mathcal{O}(\Delta t) \\ &= \left(-\mathbf{N}\bar{\mathbf{C}}^T \mathbf{L}_\tau \bar{\boldsymbol{\tau}} + \sqrt{\frac{2k_B T}{\Delta t}} \mathbf{N}^{1/2} \boldsymbol{\eta} \right)^n + \mathcal{O}(\Delta t) \end{aligned} \quad (8.24)$$

where the terms of order Δt come from using an implicit method. Comparing (8.24) with the SDE (8.5), we see that we still need to capture the drift term (8.17). To do this, we add

a random finite difference term [116, Sec. IIIC] by setting

$$\bar{\boldsymbol{\tau}}^{n,\text{RFD}} = \begin{pmatrix} \text{rotate} \left(\boldsymbol{\tau}^n, \delta \boldsymbol{\eta}_{\{1:3N\}}^{n,\text{RFD}} \right) \\ \mathbf{X}_{\text{MP}} + \delta L \boldsymbol{\eta}_{\{3N+1:3N+3\}}^{n,\text{RFD}} \end{pmatrix} \quad (8.25)$$

$$\boldsymbol{\alpha}^{n,\text{RFD}} = \frac{k_B T}{\delta} \left(\mathbf{N}(\bar{\boldsymbol{\tau}}_{\text{RFD}}^n) - \mathbf{N}(\bar{\boldsymbol{\tau}}^n) \right) \begin{pmatrix} \boldsymbol{\eta}_{\{1:3N\}}^{n,\text{RFD}} \\ L^{-1} \boldsymbol{\eta}_{\{3N+1:3N+3\}}^{n,\text{RFD}} \end{pmatrix} \quad (8.26)$$

where $\delta \ll 1$ is a dimensionless small parameter and $\boldsymbol{\eta}^{n,\text{RFD}}$ is generated independent of $\boldsymbol{\eta}^n$. The procedure (8.25) corresponds to first rotating the tangent vectors around a random axis by a random small angle, which we follow in (8.26) by solving two saddle point systems with the same random numbers on the right hand side. The Taylor expansion of the average $\boldsymbol{\alpha}^{(\text{RFD})}$ is (dropping the time superscript)

$$\begin{aligned} \mathbb{E} [\boldsymbol{\alpha}^{\text{RFD}}] &= (k_B T) \partial_j (\mathbf{N}_{ik}) \mathbb{E} \left[\boldsymbol{\eta}_k^{(\text{RFD})} (\bar{\boldsymbol{\tau}}^{(\text{RFD})} - \bar{\boldsymbol{\tau}})_j \right] + \mathcal{O}(\delta) \\ &= (k_B T) \partial_j (\mathbf{N}_{ik}) \mathbb{E} \left[\boldsymbol{\eta}_k^{(\text{RFD})} \boldsymbol{\eta}_m^{(\text{RFD})} \right] \bar{\mathbf{C}}_{jm} + \mathcal{O}(\delta) \\ &= (k_B T) \partial_j (\mathbf{N}_{ik}) \bar{\mathbf{C}}_{jk}. \end{aligned}$$

which is exactly the drift term in (8.17). In the second equality, we use the fact that the update (8.25) is $\bar{\mathbf{C}}\boldsymbol{\eta}$ to first order in δ . Thus the RFD term $\boldsymbol{\alpha}^{\text{RFD}}$ captures (8.17) in expectation, and the fiber update

$$\begin{aligned} \boldsymbol{\tau}^{n+1} &= \text{rotate} \left(\boldsymbol{\tau}^n, \Delta t (\boldsymbol{\Omega}^n + \boldsymbol{\Omega}^{n,\text{RFD}}), \right) \\ \mathbf{X}_{\text{MP}}^{n+1} &= \mathbf{X}_{\text{MP}}^n + \Delta t \left(\mathbf{U}_{\text{MP}}^n + \mathbf{U}_{\text{MP}}^{n,\text{RFD}} \right) \end{aligned} \quad (8.27)$$

is a weakly first-order accurate temporal integrator for (8.5).

8.2.3 Midpoint scheme

We can now present our “midpoint” method which can produce the drift term (8.17) in expectation with only one saddle-point solve per time step. Given that an inextensible chain can be viewed as a collection of interacting rigid rods (tangent vectors), our method is similar in spirit to that of Westwood et al. [150] for rigid body suspensions, but differs in the ways we detail in Section 8.2.3.3.

At each time step n , we perform the following steps

1. Compute a rotation rate for the tangent vectors based on the Brownian velocity

$$\widetilde{\mathbf{M}}^{1/2} \boldsymbol{\eta}^n$$

$$\boldsymbol{\alpha}^{n,*} = \sqrt{\frac{2k_B T}{\Delta t}} (\mathbf{K}^n)^{-1} \left(\widetilde{\mathbf{M}}^n \right)^{1/2} \boldsymbol{\eta}^n = \begin{pmatrix} \boldsymbol{\Omega}^{n,*} \\ \mathbf{U}_{\text{MP}}^{n,*} \end{pmatrix} \quad (8.28)$$

where \mathbf{K}^{-1} is defined in (6.29).

2. Rotate the tangent vectors by $\boldsymbol{\Omega}^{n,*} \Delta t/2$ to generate a new configuration

$$\bar{\boldsymbol{\tau}}^{n+1/2,*} = \begin{pmatrix} \text{rotate}(\boldsymbol{\tau}^n, (\Delta t/2) \boldsymbol{\Omega}^{n,*}) \\ \mathbf{X}_{\text{MP}}^n + (\Delta t/2) \mathbf{U}_{\text{MP}}^{n,*} \end{pmatrix} \quad (8.29)$$

$$= \sqrt{\frac{k_B T \Delta t}{2}} \bar{\mathbf{C}}^n (\mathbf{K}^n)^{-1} \left(\widetilde{\mathbf{M}}^n \right)^{1/2} \boldsymbol{\eta}^n + \mathcal{O}(\Delta t^{3/2}) \quad (8.30)$$

3. Evaluate the mobility $\widetilde{\mathbf{M}}^{n+1/2,*}$ and use it to compute the additional drift velocity using the random finite difference (RFD) [152] with $\delta \sim \sqrt{\Delta t}$,

$$\mathbf{U}_{\text{MD}}^n = \sqrt{\frac{2k_B T}{\Delta t}} \left(\widetilde{\mathbf{M}}^{n+1/2,*} - \widetilde{\mathbf{M}}^n \right) \left(\left(\widetilde{\mathbf{M}}^n \right)^{-1/2} \right)^T \boldsymbol{\eta}^n. \quad (8.31)$$

This term might be impractical for large systems because it is based on solving a resistance problem to obtain $\left(\widetilde{\mathbf{M}}^{-1/2} \right)^T$. An alternative way of obtaining the same drift

term (8.31), in which $\widetilde{\mathbf{M}}$ only has to be applied rather than inverted, is via an RFD, which avoids potentially expensive resistance problems. The alternative expression for \mathbf{U}_{MD} in this case is

$$\begin{aligned}\boldsymbol{\mu}^{(\text{RFD})} &= \mathbf{K}^{-1}\boldsymbol{\eta}^{(\text{RFD})} & \boldsymbol{\tau}^{(\text{RFD})} &= \text{rotate}\left(\boldsymbol{\tau}, \delta L \boldsymbol{\mu}^{(\text{RFD})}\right) \approx \boldsymbol{\tau}_p - \delta L \mathbf{C} \mathbf{K}^{-1} \boldsymbol{\eta}^{(\text{RFD})} \\ \mathbf{U}_{\text{MD}} &= \frac{k_B T}{\delta L} \left(\widetilde{\mathbf{M}}\left(\boldsymbol{\tau}^{(\text{RFD})}\right) - \widetilde{\mathbf{M}}\left(\boldsymbol{\tau}\right) \right) \boldsymbol{\eta}^{(\text{RFD})},\end{aligned}\quad (8.32)$$

Numerical tests show that using the expression (8.31), rather than (8.32), gives a closer approximation to the equilibrium distribution for larger time step sizes, whereas for smaller time step sizes the two expressions perform similarly. Thus when the mobility is localized to a single fiber, we use (8.31), but when we treat nonlocal hydrodynamics we use (8.32).

4. To obtain the tangent vector rotation rates, solve the saddle point system

$$\begin{aligned}\begin{pmatrix} -\widetilde{\mathbf{M}} & \left(\mathbf{I} + c\Delta t \widetilde{\mathbf{M}} \mathbf{L}\right) \mathbf{K} \\ \mathbf{K}^T & \mathbf{0} \end{pmatrix}^{n+1/2,*} \begin{pmatrix} \boldsymbol{\Lambda} \\ \boldsymbol{\alpha} \end{pmatrix}^{n+1/2} \\ = \begin{pmatrix} -\widetilde{\mathbf{M}}^{n+1/2,*} \mathbf{L} \mathbf{X}^n + \mathbf{U}_B^n + \mathbf{U}_{\text{MD}}^n \\ \mathbf{0} \end{pmatrix}\end{aligned}\quad (8.33)$$

for $\boldsymbol{\Lambda}^{n+1/2}$ and $\boldsymbol{\alpha}^{n+1/2} = \left(\boldsymbol{\Omega}^{n+1/2}, \mathbf{U}_{\text{MP}}^{n+1/2}\right)$. The Brownian velocity \mathbf{U}_B is defined in (8.20), and the first part of it is used in (8.28) to generate the midpoint configuration, i.e., the same $\boldsymbol{\eta}$ is used in steps 1, 3, and 4.

5. Update the fiber via (6.71),

$$\bar{\boldsymbol{\tau}}^{n+1} = \begin{pmatrix} \text{rotate}\left(\boldsymbol{\tau}^n, \Delta t \boldsymbol{\Omega}^{n+1/2}\right) \\ \mathbf{X}_{\text{MP}}^n + \Delta t \mathbf{U}_{\text{MP}}^{n+1/2} \end{pmatrix}.\quad (8.34)$$

Solving (8.33) yields (to leading order in Δt)

$$\boldsymbol{\alpha}^{n+1/2} = \left(\mathbf{N} \mathbf{K}^T \widetilde{\mathbf{M}}^{-1} \right)^{n+1/2,*} \left(-\widetilde{\mathbf{M}}^{n+1/2,*} \mathbf{L} \mathbf{X} + \mathbf{U}_B^n + \mathbf{U}_{\text{MD}}^n \right) \quad (8.35)$$

In Section 8.2.3.1, we show that using this value of $\boldsymbol{\alpha}$ in the nonlinear update (8.35) generates the drift term (8.17) in expectation. Thus, after the rotation, we obtain dynamics consistent with (8.5).

In this chapter, we consider both single filament hydrodynamics and the dynamics of multiple filaments. When we only consider one filament, we use direct solvers for the saddle point system (8.33), and in this case we can vary the mobility matrix because we can always symmetrize and truncate the eigenvalues (see Section 7.2.1). When multiple filaments are involved, it is not possible to form the matrix $\widetilde{\mathbf{M}}$ explicitly, and we instead solve (8.33) iteratively. We recall that the mobility matrix for nonlocal hydrodynamics is the oversampled RPY kernel (7.10). In Chapter 7, we spoke at length about replacing the diagonal parts of this matrix with the RPY special quadrature method in order to reduce the number of oversampled points on the fiber centerline. When we do this, it removes the guarantee that the overall matrix $\widetilde{\mathbf{M}}$ is SPD, and it makes it difficult to apply $\widetilde{\mathbf{M}}^{1/2}$. Because of this, when we consider nonlocal dynamics we will **only** use the mobility (7.10), with a preconditioner $\widetilde{\mathbf{M}}_L$ that is just the oversampled mobility on each fiber. While this method does not scale well in the limit $\epsilon \rightarrow 0$, in Section 8.5.1.1 we show that roughly $0.4/\hat{\epsilon}$ points are required to achieve convergence. Thus when $\hat{\epsilon} = 10^{-3}$, we need 400 oversampling points, which is feasible.

8.2.3.1 Drift terms in midpoint scheme

We now show that the tangent vector update (8.35) generates the drift term (8.17) in expectation. To simplify the notation, in this appendix we will assume that all time/position-dependent matrices and vectors are evaluated at time n , unless otherwise indicated. Substi-

tuting for \mathbf{U}_B and \mathbf{U}_{MD} , we expand (8.35) to read

$$\begin{aligned}
\boldsymbol{\alpha}^{n+1/2,*} &= -(\mathbf{N}\mathbf{K}^T)^{n+1/2,*} \mathbf{L}\mathbf{X} \\
&+ \sqrt{\frac{2k_B T}{\Delta t}} (\mathbf{N}\mathbf{K}^T \widetilde{\mathbf{M}}^{-1})^{n+1/2,*} \widetilde{\mathbf{M}}^{1/2} \boldsymbol{\eta} \\
&+ \sqrt{\frac{2k_B T}{\Delta t}} (\mathbf{N}\mathbf{K}^T \widetilde{\mathbf{M}}^{-1})^{n+1/2,*} (\widetilde{\mathbf{M}}^{n+1/2,*} - \widetilde{\mathbf{M}}) (\widetilde{\mathbf{M}}^{-1/2})^T \boldsymbol{\eta}.
\end{aligned} \tag{8.36}$$

Here we have used the dense matrix expression (8.31) for \mathbf{U}_{MD} , but it is not hard to see the RFD (8.32) generates the same drift term. We will Taylor expand the second and third line in succession. Beginning with the second line, we have

$$\begin{aligned}
\sqrt{\frac{2k_B T}{\Delta t}} \left[(\mathbf{N}\mathbf{K}^T \widetilde{\mathbf{M}}^{-1})^{n+1/2,*} \widetilde{\mathbf{M}}^{1/2} \boldsymbol{\eta} \right]_i &= \sqrt{\frac{2k_B T}{\Delta t}} (\mathbf{N}\mathbf{K}^T \widetilde{\mathbf{M}}^{-1/2} \boldsymbol{\eta})_i \\
&+ \sqrt{\frac{2k_B T}{\Delta t}} \partial_j (\mathbf{N}\mathbf{K}^T \widetilde{\mathbf{M}}^{-1})_{ih} (\bar{\boldsymbol{\tau}}_j^{n+1/2,*} - \bar{\boldsymbol{\tau}}_j) \widetilde{\mathbf{M}}_{hp}^{1/2} \boldsymbol{\eta}_p + \mathcal{O}(\Delta t).
\end{aligned}$$

We recognize the first term from (8.9) as $\mathbf{N}^{1/2} \boldsymbol{\eta}$, which is the typical term associated with the Brownian noise. Substituting the expansion of $\boldsymbol{\tau}^{n+1/2,*}$ in (8.30) into the second line, we get the average drift

$$\begin{aligned}
\text{DriftUB}_i &= \sqrt{\frac{2k_B T}{\Delta t}} \partial_j (\mathbf{N}\mathbf{K}^T \widetilde{\mathbf{M}}^{-1})_{ih} \sqrt{\frac{k_B T \Delta t}{2}} \bar{\mathbf{C}}_{ja} \mathbf{K}_{ab}^{-1} \widetilde{\mathbf{M}}_{bc}^{1/2} \mathbb{E} [\boldsymbol{\eta}_c \boldsymbol{\eta}_p] \widetilde{\mathbf{M}}_{hp}^{1/2} \\
&= (k_B T) \partial_j (\mathbf{N}\mathbf{K}^T \widetilde{\mathbf{M}}^{-1})_{ih} (\widetilde{\mathbf{M}}_{hb} \mathbf{K}_{ba}^{-T} \bar{\mathbf{C}}_{aj}^T)
\end{aligned}$$

where in the last equality we used $\widetilde{\mathbf{M}} = \widetilde{\mathbf{M}}^{1/2} \left(\widetilde{\mathbf{M}}^{1/2} \right)^T$. Rearranging this gives $\widetilde{\mathbf{M}} \left(\widetilde{\mathbf{M}}^{-1/2} \right)^T = \widetilde{\mathbf{M}}^{1/2}$, which allows us to expand the last line in (8.36) as

$$\begin{aligned}
\text{DriftUMD}_i &= \mathbb{E} \left[\sqrt{\frac{2k_B T}{\Delta t}} \left(\mathbf{N} \mathbf{K}^T \widetilde{\mathbf{M}}^{-1} \right)_{ih} \partial_j \left(\widetilde{\mathbf{M}}_{hp} \right) \left(\bar{\tau}_j^{n+1/2,*} - \bar{\tau}_j \right) \widetilde{\mathbf{M}}_{qp}^{-1/2} \boldsymbol{\eta}_q \right] + \mathcal{O}(\Delta t) \\
&= (k_B T) \left(\mathbf{N} \mathbf{K}^T \widetilde{\mathbf{M}}^{-1} \right)_{ih} \partial_j \left(\widetilde{\mathbf{M}}_{hp} \right) \bar{\mathbf{C}}_{ja} \mathbf{K}_{ab}^{-1} \widetilde{\mathbf{M}}_{bc}^{1/2} \mathbb{E} [\boldsymbol{\eta}_c \boldsymbol{\eta}_q] \widetilde{\mathbf{M}}_{qp}^{-1/2} \\
&= (k_B T) \left(\mathbf{N} \mathbf{K}^T \widetilde{\mathbf{M}}^{-1} \right)_{ih} \partial_j \left(\widetilde{\mathbf{M}}_{hp} \right) \bar{\mathbf{C}}_{ja} \mathbf{K}_{ap}^{-1} \\
&= (k_B T) \left(\mathbf{N} \mathbf{K}^T \widetilde{\mathbf{M}}^{-1} \right)_{ih} \partial_j \left(\widetilde{\mathbf{M}}_{hp} \bar{\mathbf{C}}_{ja} \mathbf{K}_{ap}^{-1} \right). \tag{8.37}
\end{aligned}$$

The last equality makes use of the identity

$$\mathbf{N}_{ik} \mathbf{K}_{kp}^T \partial_j \left(\bar{\mathbf{C}}_{jr} \mathbf{K}_{rp}^{-1} \right) = 0, \tag{8.38}$$

which is proven later.

We can now add the two drift terms to obtain the total drift

$$\begin{aligned}
\text{Drift}_i &= \text{DriftUB}_i + \text{DriftUMD}_i \\
&= k_B T \partial_j \left(\left(\mathbf{N} \mathbf{K}^T \widetilde{\mathbf{M}}^{-1} \right) \left(\widetilde{\mathbf{M}} \mathbf{K}^{-T} \bar{\mathbf{C}}^T \right) \right)_{ij} \\
&= k_B T \partial_j \left(\mathbf{N} \left(\bar{\mathbf{C}} \mathbf{K}^{-1} \mathbf{K} \right)^T \right)_{ij} = k_B T \partial_j \left(\mathbf{N} \bar{\mathbf{C}}^T \right)_{ij} = k_B T \partial_j \left(\mathbf{N}_{ik} \right) \bar{\mathbf{C}}_{kj}^T.
\end{aligned}$$

In the last line above, we used (6.30) in the second equality, thereby establishing that the one-solve scheme produces the drift (8.17), as desired.

8.2.3.2 Proof of (8.38)

We now prove that (8.38) is indeed zero by substituting the definitions of \mathbf{K} and \mathbf{K}^{-1} to get

$$\begin{aligned} \mathbf{N}_{ik} \mathbf{K}_{kp}^T \partial_j (\bar{\mathbf{C}}_{jr} \mathbf{K}_{rp}^{-1}) &= \mathbf{N}_{ik} \bar{\mathbf{C}}_{ky}^T \boldsymbol{\chi}_{yp}^T \partial_j \left(\bar{\mathbf{C}}_{jr} \bar{\mathbf{C}}_{rb}^T \right) \boldsymbol{\chi}_{bp}^{-1} \\ &= \mathbf{N}_{ik} \bar{\mathbf{C}}_{ky}^T \partial_j \left(\bar{\mathbf{C}}_{jr} \bar{\mathbf{C}}_{ry}^T \right). \end{aligned}$$

Now, the matrix $\left(\bar{\mathbf{C}}_{jr} \bar{\mathbf{C}}_{rb}^T \right)$ is a $3(N+1) \times 3(N+1)$ matrix composed of $N+1$ diagonal blocks. The first N diagonal blocks are $\mathbf{I}_3 - \boldsymbol{\tau}_{\{p\}} \boldsymbol{\tau}_{\{p\}}$, while the last block is the 3×3 identity.

It follows that the derivative

$$\frac{\partial}{\partial \bar{\boldsymbol{\tau}}_j} \left(\bar{\mathbf{C}}_{jr} \bar{\mathbf{C}}_{ry}^T \right) = -4\bar{\boldsymbol{\tau}}_y, \quad (8.39)$$

for $b \leq 3N$, and zero otherwise (for the parts of $\bar{\boldsymbol{\tau}}$ associated with constant motions). Multiplying by $\bar{\mathbf{C}}^T$ and using (6.26) gives (8.38).

8.2.3.3 Comparison to scheme of Westwood et al. [150]

The midpoint method proposed here applies for any kinematic matrix $\mathbf{K}(\mathbf{X})$ and mobility $\mathbf{M}(\mathbf{X})$ and can therefore be used for suspensions of rigid bodies as well. In fact, our scheme is related to the generalized drift-correcting (gDC) scheme proposed by Westwood et al. [150, Sec. 4.1] for rigid body suspensions.

For a single rigid body modeled as a rigid-multiblob with blob positions \mathbf{X}_i (generalization to many rigid bodies is straightforward) the authors of [150] define the kinematic matrix $\mathbf{K}(\mathbf{X})$ through its action on a 6×1 rigid body velocity $\mathbf{U} = [\mathbf{u}, \boldsymbol{\omega}]$, so that

$$\partial_t \mathbf{X}_i = \mathbf{u} + \boldsymbol{\omega} \times (\mathbf{X}_i - \mathbf{q}) = [\mathbf{K}(\mathbf{X}) \mathbf{U}]_i, \quad (8.40)$$

where \mathbf{q} is an arbitrary tracking point for the rigid structure. Equation (E.5) in [150] gives

the drift term for rigid bodies as

$$\text{Drift}_i = k_B T \partial_j \mathbf{N}_{ij}. \quad (8.41)$$

Similar to the midpoint scheme presented here, the gDC scheme captures the thermal drift term by splitting it into two parts

$$\begin{aligned} \text{Drift}_i &= (k_B T) \partial_j \mathbf{N}_{ij} \\ &= (k_B T) \partial_j \left(\mathbf{N} \mathbf{K}^T \mathbf{M}^{-1/2} \mathbf{M}^{1/2} \mathbf{K}^{-T} \right)_{ij} \\ &= (k_B T) \partial_j \left(\mathbf{N} \mathbf{K}^T \mathbf{M}^{-1/2} \right)_{ik} \left(\mathbf{M}^{1/2} \mathbf{K}^{-T} \right)_{kj} \\ &\quad + k_B T \left(\mathbf{N} \mathbf{K}^T \mathbf{M}^{-1/2} \right)_{ik} \partial_j \left(\mathbf{M}^{1/2} \mathbf{K}^{-T} \right)_{kj}, \end{aligned} \quad (8.42)$$

where $\mathbf{K}^{-T} = \mathbf{K} (\mathbf{K}^T \mathbf{K})^{-1}$ is the left pseudoinverse of \mathbf{K}^T and can be computed very efficiently. By contrast, the midpoint scheme presented in this work splits the drift term according to

$$\begin{aligned} \text{Drift}_i &= (k_B T) \partial_j \mathbf{N}_{ij} \\ &= (k_B T) \partial_j \left(\mathbf{N} \mathbf{K}^T \mathbf{M}^{-1} \mathbf{M}^{1/2} \mathbf{M}^{1/2} \mathbf{K}^{-T} \right)_{ij} \\ &= (k_B T) \partial_j \left(\mathbf{N} \mathbf{K}^T \mathbf{M}^{-1} \right)_{ik} \left(\mathbf{M} \mathbf{K}^{-T} \right)_{kj} + k_B T \left(\mathbf{N} \mathbf{K}^T \mathbf{M}^{-1} \right)_{ik} \partial_j \left(\mathbf{M} \mathbf{K}^{-T} \right)_{kj}. \end{aligned} \quad (8.43)$$

While equations (8.42) and (8.43) appear very similar, the approach used in our midpoint scheme only requires one application of $\mathbf{M}^{1/2}$, while the gDC scheme requires four. In some cases, saving on applications of $\mathbf{M}^{1/2}$ can significantly reduce the computational cost of the whole scheme.

To elaborate on why the gDC scheme requires extra applications of $\mathbf{M}^{1/2}$, we will briefly summarize how the scheme captures each term in (8.42). The second term is calculated

efficiently by using an RFD for

$$\boldsymbol{\nu}_k = \partial_j \left(\mathbf{M}^{1/2} \mathbf{K}^{-T} \right)_{kj} \quad (8.44)$$

since it does not involve \mathbf{N} , while the first term is captured using a midpoint time integration scheme of the form

$$\mathbf{X}^{n+1/2,*} - \mathbf{X}^n = \sqrt{\frac{k_B T \Delta t}{2}} \mathbf{K}^{-1} \mathbf{M}^{1/2} \boldsymbol{\eta}^n \quad (8.45)$$

$$\tilde{\mathbf{V}}^{n+1/2,*} = \sqrt{\frac{2k_B T}{\Delta t}} \left(\mathbf{N} \mathbf{K}^T \mathbf{M}^{-1/2} \right)^{n+1/2,*} \boldsymbol{\eta}^n, \quad (8.46)$$

so that

$$\begin{aligned} \mathbb{E} \left[\tilde{\mathbf{V}}^{n+1/2,*} \right] &= (k_B T) \partial_j \left(\mathbf{N} \mathbf{K}^T \mathbf{M}^{-1/2} \right)_{ik} \mathbb{E} [\boldsymbol{\eta}_k^n \boldsymbol{\eta}_r^n] \left(\mathbf{K}^{-1} \mathbf{M}^{1/2} \right)_{jr} \\ &= (k_B T) \partial_j \left(\mathbf{N} \mathbf{K}^T \mathbf{M}^{-1/2} \right)_{ik} \left(\mathbf{M}^{1/2} \mathbf{K}^{-T} \right)_{kj}. \end{aligned}$$

The gDC scheme combines these parts to compute

$$\text{Drift}_i = k_B T \left(\mathbf{N} \mathbf{K}^T \mathbf{M}^{-1/2} \right)_{ik} \boldsymbol{\nu}_k + \mathbb{E} \left[\tilde{\mathbf{V}}^{n+1/2,*} \right] = (k_B T) \partial_j \mathbf{N}_{ij}. \quad (8.47)$$

Note that calculating $\tilde{\mathbf{V}}^{n+1/2,*}$ according to (8.46) requires two applications of $\mathbf{M}^{1/2}$ (one at each time level), and calculating $\boldsymbol{\nu}_k$ using an RFD requires two additional applications of $\mathbf{M}^{1/2}$. Hence the gDC scheme requires four total applications of $\mathbf{M}^{1/2}$ while the midpoint scheme presented in this work only requires one.

8.3 Small equilibrium fluctuations of a curved fiber

To start our numerical analysis of our spectral discretization and temporal integrators, it is logical to study the performance in the context of small fluctuations around an equilibrium state. The motivation for studying this problem is both physical and numerical: physically, it is useful to study because we are eventually interested in fibers that are constrained by cross linkers, which roughly hold them in place while they fluctuate. Numerically, starting with small fluctuations [83], and a curved, rather than straight, fiber allows us to confirm that our method correctly handles nonlinearities. We discuss the curved configuration we choose, and how we keep the dynamics from drifting away from it, in Section 8.3.1.

Because the dynamics are linearized around a particular state, we can compute a theoretical covariance matrix and break the dynamics into a set of modes (from the linearized covariance). In Section 8.3.2, we compute the equilibrium variance of each of these modes via Markov Chain Monte Carlo (MCMC) calculations using the equilibrium probability distribution (8.1). By comparing our spectral discretization to a blob-link one, we verify that our equilibrium distribution (8.1) gives behavior similar to that of the more physical one with uniformly-spaced links. Following this, we take advantage of the small fluctuations to linearize the SDE (8.12) around the base state \mathbf{X}_0 , which we then diagonalize in Section 8.3.3.1 to compute relaxation timescales of the fundamental modes. This aids us in Section 8.3.4, where we study the time step size required for our midpoint temporal integrator to accurately sample from the equilibrium probability distribution (8.1).

8.3.1 Set-up for small fluctuations

For the base state \mathbf{X}_0 , we introduce a curved configuration with $L = 2 \mu\text{m}$ and

$$\boldsymbol{\tau}_0(s) = \frac{1}{\sqrt{2}} \begin{pmatrix} \cos(s^3(s-L)^3), & \sin(s^3(s-L)^3), & 1 \end{pmatrix} \quad (8.48)$$

with \mathbf{X}_0 defined as the integral of this on the $N + 1$ point Chebyshev grid with $\mathbf{X}_{\text{MP}} = \mathbf{0}$. To keep the dynamics near this base state, we introduce an additional energy which penalizes the discrete squared L^2 norm of $\Delta\mathbf{X} = \mathbf{X} - \mathbf{X}_0$,

$$\mathcal{E}_P = \frac{P}{2} \Delta\mathbf{X}^T \widetilde{\mathbf{W}} \Delta\mathbf{X}, \quad P = (1.6 \times 10^4) \frac{k_B T}{L^3}. \quad (8.49)$$

This choice of P ensures that the dimensionless discrete L^2 norm $\Delta\mathbf{X}^T \widetilde{\mathbf{W}} \Delta\mathbf{X} / L^3$ remains constant when $k_B T$ and L change. While this is the required scaling of P to keep the relative magnitude of $\Delta\mathbf{X}$ roughly constant (and small), throughout this section we will fix $k_B T = 4.1 \times 10^{-3}$ pN $\cdot\mu\text{m}$ and $L = 2 \mu\text{m}$, thus setting $P = 8.2$ pN/ μm^2 . To modify the filament's persistence length, we will modify the bending stiffness κ in the modified bending energy that keeps the curved \mathbf{X}_0 as an equilibrium configuration

$$\mathcal{E}_{\text{bend}} = \frac{1}{2} \kappa \Delta\mathbf{X}^T (\mathbf{D}^2)^T \widetilde{\mathbf{W}} \mathbf{D}^2 \Delta\mathbf{X}. \quad (8.50)$$

In this case, the Ito Langevin equation (8.12) can be viewed as an equation for $\Delta\mathbf{X}$ rather than \mathbf{X} , and the matrix \mathbf{L} becomes

$$\mathbf{L} = P \widetilde{\mathbf{W}} + \kappa (\mathbf{D}^2)^T \widetilde{\mathbf{W}} \mathbf{D}^2, \quad (8.51)$$

which is the original definition (6.10) modified to account for the penalty force. Notice that the first term in the energy (the penalty term) is independent of κ , while the bending term is proportional to κ . Figure 8.1 shows \mathbf{X}_0 along with some samples from the Gibbs-Boltzmann distribution (8.1) with this energy.

We will use the covariance in $\Delta\mathbf{X}$ to determine the accuracy of our spatial and temporal discretizations. Because the energy is nonlinearly constrained, we cannot write down the exact covariance, but it is informative to project the modes onto the covariance we would

obtain if we replace the nonlinear inextensibility constraint with the *linearized* version $\Delta\mathbf{X} = \mathbf{K}[\mathbf{X}_0]\boldsymbol{\alpha} := \mathbf{K}_0\boldsymbol{\alpha}$. The resulting energy $\frac{1}{2}\Delta\mathbf{X}^T\mathbf{L}\Delta\mathbf{X} = \frac{1}{2}\boldsymbol{\alpha}^T\mathbf{K}_0^T\mathbf{L}\mathbf{K}_0\boldsymbol{\alpha}$ dictates that $\boldsymbol{\alpha}$ has covariance

$$\mathbf{C}_\alpha = \mathbb{E}[\boldsymbol{\alpha}\boldsymbol{\alpha}^T] = (k_B T) (\mathbf{K}_0^T \mathbf{L} \mathbf{K}_0)^\dagger. \quad (8.52)$$

Now pre- and post-multiplying by \mathbf{K}_0 , we get the expected covariance for $\Delta\mathbf{X}$ as

$$\mathbf{C} = \mathbb{E}[\Delta\mathbf{X}\Delta\mathbf{X}^T] = (k_B T) \mathbf{K}_0 (\mathbf{K}_0^T \mathbf{L} \mathbf{K}_0)^\dagger \mathbf{K}_0^T, \quad (8.53)$$

which is valid in the limit of small fluctuations.

To check how close the true covariance is to \mathbf{C} , we will generate samples of $\Delta\mathbf{X}$ through both MCMC sampling and Langevin dynamics, then use the procedure outlined in Section 8.3.1.1 to project the resulting covariance onto the eigenmodes of \mathbf{C} . This amounts to computing the Rayleigh quotient of the eigenmodes using the true covariance and normalizing with respect to the eigenvalues of \mathbf{C} . We focus on the variance in each of the modes (diagonal entries of the covariance matrix projected onto the eigenmodes), as the off-diagonal entries for the smallest persistence length we consider ($\ell_p/L = 1$) are zero within the statistical error of our MCMC calculations.

8.3.1.1 Covariance calculations

To compare the covariance from MCMC to (8.53), we first use a blob-link discretization to compute the matrix \mathbf{C} exactly as defined in (8.53). In order for the eigenmodes of \mathbf{C} to be the same across all discretizations, we need to ensure that they are orthonormal in L^2 and not \mathbb{R}^n . This can be accomplished by computing the eigenvalues of $\widetilde{\mathbf{W}}^{1/2} \mathbf{C} \widetilde{\mathbf{W}}^{1/2}$, which is a Hermitian matrix and therefore has eigendecomposition

$$\widetilde{\mathbf{W}}^{1/2} \mathbf{C} \widetilde{\mathbf{W}}^{1/2} = \widetilde{\mathbf{V}} \boldsymbol{\Lambda} \widetilde{\mathbf{V}}^T \quad \widetilde{\mathbf{V}}^T \widetilde{\mathbf{V}} = \mathbf{I} \quad (8.54)$$




Figure 8.1: Three samples of penalty-bound blob-link chains (with 100 links) for various ℓ_p/L , sampled using MCMC. We show the “random” filament shapes $\mathbf{X} = \mathbf{X}_0 + \Delta\mathbf{X}$, where $\Delta\mathbf{X}$ is sampled from the equilibrium distribution (8.57). The blue line shows \mathbf{X}_0 .

The matrix $\widetilde{\mathbf{W}}$ is used to compute the L^2 inner product, $\|\mathbf{X}\|_{L^2}^2 \approx \mathbf{X}^T \widetilde{\mathbf{W}} \mathbf{X}$. In the case of the blob-link method, $\widetilde{\mathbf{W}} = \Delta s \text{diag}(1/2, 1, \dots, 1, 1/2)$. Now, if we use (8.54), we observe that

$$\mathbf{C} = \widetilde{\mathbf{W}}^{-1/2} \widetilde{\mathbf{V}} \Lambda \left(\widetilde{\mathbf{W}}^{-1/2} \widetilde{\mathbf{V}} \right)^T$$

Thus, if we define $\mathbf{V} = \widetilde{\mathbf{W}}^{-1/2} \widetilde{\mathbf{V}}$, we have (using the orthonormality of $\widetilde{\mathbf{V}}$ in \mathbb{R}^n) the desirable property that

$$\mathbf{C} = \mathbf{V} \Lambda \mathbf{V}^T \quad \mathbf{V}^T \widetilde{\mathbf{W}} \mathbf{V} = \mathbf{I}, \quad (8.55)$$

which means the columns of \mathbf{V} are orthonormal in L^2 . Then, to compute the covariance, we can project the modes of $\Delta\mathbf{X}$ onto the columns of \mathbf{V} in L^2 as $\mathbf{V}^T \widetilde{\mathbf{W}} \Delta\mathbf{X}$, since

$$\begin{aligned} \mathbb{E} [\Delta\mathbf{X} \Delta\mathbf{X}^T] &= \mathbf{C} \\ \mathbb{E} \left[\mathbf{V}^T \widetilde{\mathbf{W}} \Delta\mathbf{X} \left(\mathbf{V}^T \widetilde{\mathbf{W}} \Delta\mathbf{X} \right)^T \right] &= \mathbf{V}^T \widetilde{\mathbf{W}} \mathbf{C} \widetilde{\mathbf{W}} \mathbf{V} = \Lambda, \end{aligned}$$

using (8.55) in the last equality.

8.3.2 MCMC estimation of reference covariance

To establish a reference result for the variance of each of the modes of \mathbf{C} , we use MCMC to sample from the Gibbs-Boltzmann measure (8.57), but with $\Delta\mathbf{X} = \mathbf{X} - \mathbf{X}_0$ taking the place of \mathbf{X} in (8.57) and (8.58).

8.3.2.1 Blob-link MCMC

To establish a benchmark for the spectral code, we perform MCMC on a blob-link chain with $L = 2$ and $N = 100$ links. To do this, we initialize the chain with shape $\mathbf{X} = \mathbf{X}_0$. Then, at each Monte Carlo step, we propose a new chain $\widetilde{\mathbf{X}}$ that is generated by rotating the tangent vectors by an oriented angle

$$\widehat{\boldsymbol{\Omega}}^\perp = 0.1\sqrt{k_B T}\mathbf{L}^{-1/2}\boldsymbol{\eta}, \quad (8.56)$$

where $\boldsymbol{\eta}$ is a $3N + 3$ vector of i.i.d standard normal random variables, and the constant of 0.1 is added to make the acceptance ratio roughly 40%. Note that this proposal is informed by the bending energy (unconstrained covariance) matrix because there are many links in the blob-link chain, and so the multiplication of the Gaussian by $\mathbf{L}^{-1/2}$ effectively projects the proposal onto the expected (unconstrained) covariance, which reduces the number of samples required to equilibrate the higher-order modes. When $P = 0$ (for a free fiber), we compute the proposal (8.56) using the pseudo-inverse of \mathbf{L} , since constant and linear modes have zero eigenvalues. Note that in this special case it is possible to construct direct (exact) independent samplers to obtain the chain configurations, though these were not used in this work.

8.3.2.2 Acceptance probabilities

Once we generate the proposal, we use the Metropolis-Hastings algorithm to accept/reject it. Since we are sampling from the energy

$$d\pi(\mathbf{X}) \propto \exp\left(-\frac{\mathcal{E}(\mathbf{X})}{k_B T}\right) d\mu_0(\mathbf{X}) := \bar{\pi}(\mathbf{X}) d\mu_0(\mathbf{X}), \quad (8.57)$$

where $\mathcal{E}(\mathbf{X}) = \frac{1}{2}\Delta\mathbf{X}^T \mathbf{L}\Delta\mathbf{X}$ and the base measure $d\mu_0(\mathbf{X})$ is defined in (8.1), the probability of acceptance is simply the Metropolis factor

$$p_{\text{acc}} = \max\left(\frac{\bar{\pi}(\Delta\widetilde{\mathbf{X}})}{\bar{\pi}(\Delta\mathbf{X})}, 1\right). \quad (8.58)$$

8.3.2.3 Results

Figure 8.1 shows some sample 100-link chains with various values of κ , which we quantify by the dimensionless persistence length $\ell_p/L = \kappa/(Lk_B T)$. While the relative mean-square deviation of the chains from the blue base state is approximately the same (1%) in all cases, we see more and more “wiggles” in the chain as we decrease ℓ_p/L .

To compute the variance of each mode and compare it to the expected result (8.53), we use the MCMC sampler to generate 10^6 samples of \mathbf{X} for $\ell_p/L = 100, 10$, and 1. We then throw out the first 20% of the samples and compute the covariance of the rest of the samples. We repeat this ten times to generate error bars, and show the variance in each eigenmode in Fig. 8.2(a), where we normalize by the corresponding eigenvalue and do the projection in L^2 as discussed in Section 8.3.1.1. In addition to running simulations with 100 links, we also run with 200 links to confirm that the results are converged in space. We see that, for $\ell_p/L \gtrsim 10$, the variance of each mode is exactly that predicted by the theory (8.53). When we decrease κ so that $\ell_p/L = 1$, some of the intermediate modes (from about 5 to 20) are damped relative to the theoretical prediction. In this case the bending fluctuations are

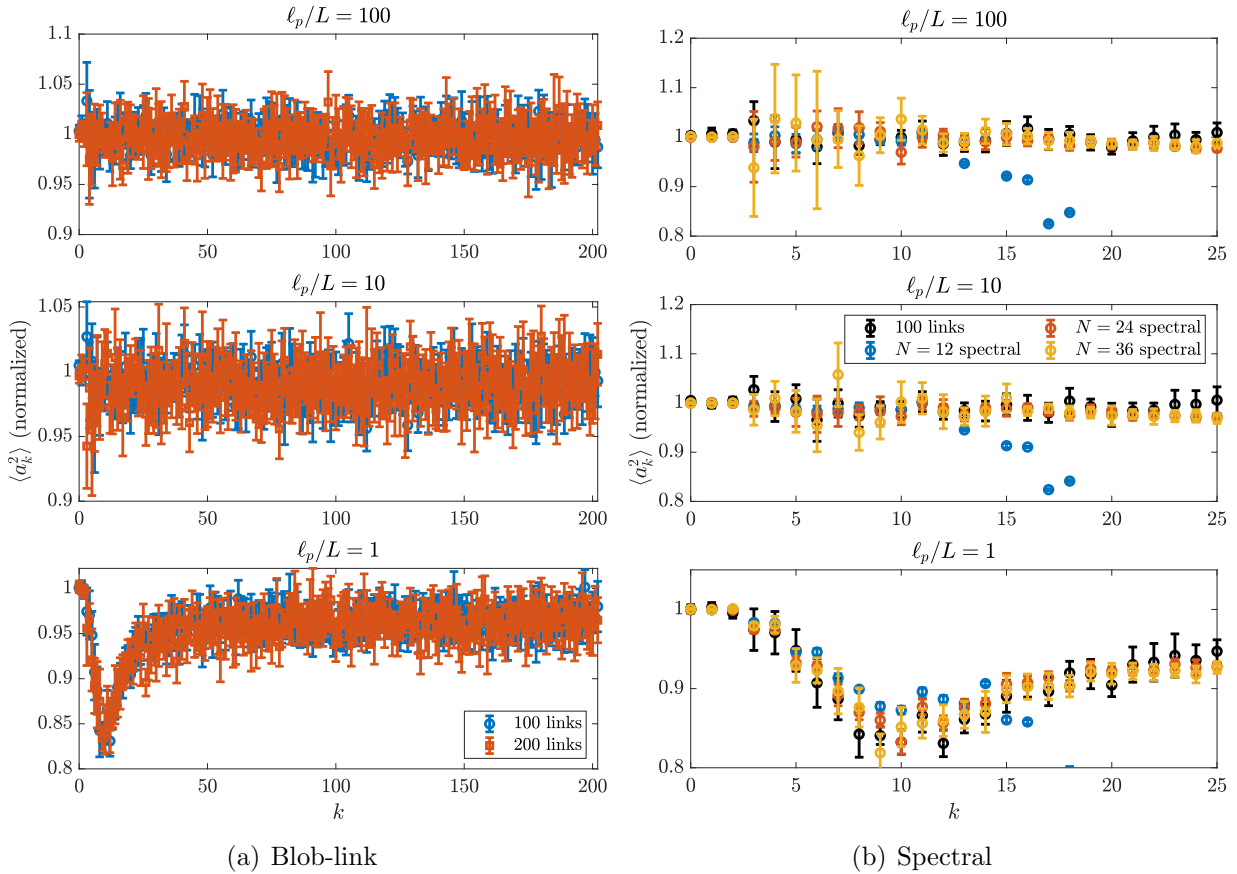


Figure 8.2: Covariance of linearized modes for various ℓ_p/L using blob-link and spectral discretizations, estimated using MCMC sampling. We always project the variance onto the same set of L^2 orthonormal modes given by the eigenvectors of the linearized covariance (8.53), as discussed in Section 8.3.1.1. (a) The covariance for a blob-link discretization with 100 links and 101 blobs (203 modes). Results are shown for 100 links (in blue), and 200 links (in red) to verify convergence. (b) Comparison of the variance of the first 25 modes using a spectral grid with those of the blob-link discretization. The black points show the results of MCMC calculations with 100 links, while the blue, red, and yellow symbols show results with $N = 12, 24,$ and 36 Chebyshev nodes, respectively. When $N = 12$ and $k \geq 17$, the variance is reduced significantly to about 0.2 and is not shown.

sufficiently large for nonlinear effects to matter.

8.3.2.4 MCMC for spectral discretization

Let us now perform MCMC on a chain discretized using Chebyshev collocation points, rather than uniformly-spaced nodes. To do this, we choose a number of Chebyshev collocation points N and set $\mathbf{X} = \mathbf{X}_0$. Then, we propose a rotation of the tangent vectors at each of the Chebyshev nodes²

$$\tilde{\Omega}_{\{p\}} = \frac{0.48}{N} \sqrt{\frac{L}{\ell_p}} \boldsymbol{\eta}_{\{p\}} \quad (8.59)$$

where $\boldsymbol{\eta}_{\{p\}}$ is a vector of three i.i.d. standard normal numbers for each node p , and the constant in front is chosen so that roughly 50% of the samples are accepted if the tangent vectors are rotated by $\tilde{\Omega}_p$ and the middle of the fiber is held in place.

The second part of the proposal is to update the fiber midpoint. To do this, we propose a new midpoint via

$$\tilde{\mathbf{X}}_{\text{MP}} = \mathbf{X}_{\text{MP}} + (7.5 \times 10^{-3}) L \boldsymbol{\eta}, \quad (8.60)$$

where $\boldsymbol{\eta}$ is a vector of three i.i.d. standard normal numbers. Now, notice that the scaling of $\Delta \mathbf{X}_{\text{MP}}$ does not depend on N or κ , since the penalty energy is independent of these two quantities. As such, the constant 7.5×10^{-3} is chosen so that roughly 50% of the samples are accepted if the tangent vectors are not updated and the fiber only translates. We combine the two proposals $\tilde{\Omega}_p$ and $\tilde{\mathbf{X}}_{\text{MP}}$ into a single proposal

$$\tilde{\mathbf{X}} = \boldsymbol{\mathcal{X}} \left(\begin{array}{c} \text{rotate}(\boldsymbol{\tau}, \tilde{\Omega}) \\ \tilde{\mathbf{X}}_{\text{MP}} \end{array} \right), \quad (8.61)$$

and compute the energy via $\frac{1}{2} \Delta \tilde{\mathbf{X}}^T \mathbf{L} \Delta \tilde{\mathbf{X}}$. With this proposal, the Metropolis algorithm in

²Note that multiplication of $\boldsymbol{\eta}$ by $\mathbf{L}^{-1/2}$ (c.f. (8.56)) is not necessary since in the spectral discretization the number of Chebyshev nodes is intended to be relatively small.

Section 8.3.2.2 gives a total acceptance ratio between 20 and 30%.

Repeating our MCMC procedure on the spectral discretization, we look at the variance of each of the eigenmodes of the linearized covariance (8.53) in Fig. 8.2(b). We show three different values of N , and observe that the spectral equilibrium distribution for $N = 12$ has variance which diverges slightly from the blob-link variance for small mode numbers when $\ell_p/L = 1$, and for larger mode numbers ($k \geq 17$) for all ℓ_p . Increasing to $N = 24$ and $N = 36$, we see overlap with the variance from the blob-link discretization with 100 links for all values of ℓ_p , which shows that the spectral method with a sufficiently large N can properly reproduce the magnitude of the small equilibrium fluctuations of the blob-link chain (at least for this example).

8.3.3 Dynamics

Our task now is to see if our temporal integrators can be used to sample from the Gibbs-Boltzmann distribution (8.1) via Langevin dynamics. Before doing that, we look at the fundamental timescales governing system dynamics. The largest of these is τ_{fund} , which we use in the rest of this chapter to non-dimensionalize the time step sizes of free filaments.

8.3.3.1 Linearized timescales

One of the benefits of considering dynamics that are only slightly perturbed from the equilibrium configuration \mathbf{X}_0 is that we can linearize the SDE (8.12) and diagonalize the result to compute the fundamental timescales in the system. Because the linearized mobility matrix is constant (evaluated at \mathbf{X}_0), the drift term disappears and the linearized form of the SDE (8.12) around $\mathbf{X} = \mathbf{X}_0$ is

$$\frac{d\Delta\mathbf{X}}{dt} = -\widehat{\mathbf{N}}_0 \mathbf{L} \Delta\mathbf{X} + \sqrt{2k_B T} \widehat{\mathbf{N}}_0^{1/2} \boldsymbol{\mathcal{W}}, \quad (8.62)$$

where $\widehat{\mathbf{N}}_0 = \mathbf{K}_0 \mathbf{N}_0 \mathbf{K}_0^T$, with \mathbf{N} defined in (8.3). Since the null space of $\widehat{\mathbf{N}}_0$ is the N *extensible* motions around $\mathbf{X} = \mathbf{X}_0$, there are $2N + 3$ remaining directions, and the SDE (8.62) diagonalizes when we consider the first $2N + 3$ generalized eigenvectors of $\widehat{\mathbf{N}}_0^\dagger$ and \mathbf{L} , which satisfy

$$\widehat{\mathbf{N}}_0^\dagger \mathbf{V} = \mathbf{L} \mathbf{V} \mathbf{\Lambda}, \quad (8.63)$$

where the $2N + 3$ eigenvectors, which make up the $(3N + 3) \times (2N + 3)$ matrix \mathbf{V} , are normalized such that

$$\mathbf{V}^T \mathbf{L} \mathbf{V} = \mathbf{I} \rightarrow \mathbf{V}^T \widehat{\mathbf{N}}_0^\dagger \mathbf{V} = \mathbf{\Lambda}. \quad (8.64)$$

These equations imply that

$$\widehat{\mathbf{N}}_0^\dagger \mathbf{V} = \mathbf{L} \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T \mathbf{L} \mathbf{V} \rightarrow \widehat{\mathbf{N}}_0^\dagger = \mathbf{L} \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T \mathbf{L} \rightarrow \widehat{\mathbf{N}}_0^{-1/2} = \mathbf{L} \mathbf{V} \mathbf{\Lambda}^{1/2}, \quad (8.65)$$

which gives a square root of $\widehat{\mathbf{N}}_0$ as $\widehat{\mathbf{N}}_0^{1/2} = \widehat{\mathbf{N}}_0 \widehat{\mathbf{N}}_0^{-1/2}$. Substituting this into the linearized Langevin equation (8.62) and using the definition of the eigenvalues (8.63), we obtain the SDE governing the evolution of the eigenmodes $\widehat{\mathbf{X}} = \mathbf{V}^{-1} \Delta \mathbf{X}$,

$$\begin{aligned} \mathbf{V} \frac{d\widehat{\mathbf{X}}}{dt} &= -\widehat{\mathbf{N}}_0 \mathbf{L} \mathbf{V} \widehat{\mathbf{X}} + \sqrt{2k_B T} \widehat{\mathbf{N}}_0 \mathbf{L} \mathbf{V} \mathbf{\Lambda}^{1/2} \mathcal{W} \\ \frac{d\widehat{\mathbf{X}}}{dt} &= -\mathbf{\Lambda}^{-1} \widehat{\mathbf{X}} + \sqrt{2k_B T} \mathbf{\Lambda}^{-1/2} \mathcal{W}. \end{aligned} \quad (8.66)$$

Thus, each of the $2N + 3$ inextensible eigenmodes relaxes with characteristic timescale equal to its eigenvalue, and the modes satisfy the equipartition principle $\mathbb{E} [\widehat{\mathbf{X}} \widehat{\mathbf{X}}^T] = (k_B T) \mathbf{I}$. The average elastic energy is then $\frac{1}{2} \mathbb{E} [\Delta \mathbf{X}^T \mathbf{L} \Delta \mathbf{X}] = \frac{1}{2} \mathbb{E} [\widehat{\mathbf{X}}^T \mathbf{V}^T \mathbf{L} \mathbf{V} \widehat{\mathbf{X}}] = (N + 3/2) k_B T$.

We first examine the dependence of the timescales on N in the left panel of Fig. 8.3, beginning by plotting the timescales for $N = 12$ (blue) and $N = 24$ (red) with the penalty force included in \mathbf{L} . Including the penalty force makes \mathbf{L} invertible, which gives finite timescales for the first six modes (three translation and three rotation) that have zero eigenvalues when

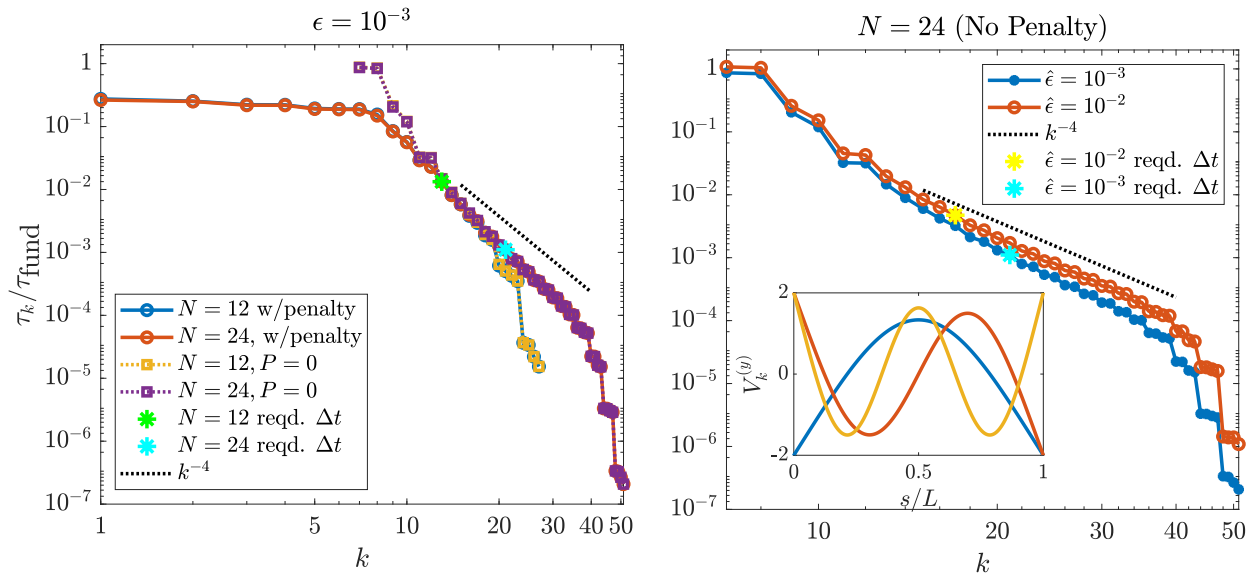


Figure 8.3: Timescales in the linearized filament problem. In each plot, we give the timescales (the $2N + 3$ nonzero diagonal entries of $\mathbf{\Lambda}$ in (8.63)) made dimensionless by τ_{fund} in (8.69) (the largest timescale for free fibers), then mark with a star the time step size required for the midpoint temporal integrator to have an error less than 5% in the equilibrium variance when $\ell_p/L = 10$ (see Section 8.3.4). (Left) Dependence of timescales on N . We show the timescales for two different N with (blue and red) and without (yellow and purple) the penalty force. (Right) Dependence on the mobility (aspect ratio $\hat{\epsilon}$) with $P = 0$. We show $\hat{\epsilon} = 10^{-3}$ in blue and $\hat{\epsilon} = 10^{-2}$ in red. Inset: the first three unique modes for a straight filament (compare to [83, Fig. 2(a)]). The eigenvalue of the blue mode is the timescale τ_{fund} .

$P = 0$. After the first few modes, the timescales begin to decay with the expected k^{-4} scaling [153] until the spatial discretization error causes shorter timescales than we would obtain in continuum. We see in particular that for a given N we have about 6 incorrect timescales from spatial discretization error. Excepting these modes, we see that increasing N only adds additional small timescales into the problem, since $N = 12$ is sufficient to correctly give the timescales of the first 15 or so modes. Since the penalty force contributes a fixed multiple of the identity to \mathbf{L} , the largest eigenvalues are changed very little, and so the effect of the penalty force on the relaxation timescales is negligible for modes beyond the first 10 or so (see the left panel of Fig. 8.3).

A slight complication in Fig. 8.3 is that the linearized timescales for two different $\hat{\epsilon}$ do *not* collapse onto the same curve when rescaled by the estimate of (8.69). In fact, the expected log scaling, which comes from slender body theory [19, 20], approximately holds only for the smoothest modes ($k \lesssim 10$), with the timescales of the high-frequency modes scaling at a much sharper rate. Indeed, at the shortest scales we expect to see $\hat{\epsilon}^{-1}$ scaling, corresponding to the timescales on which individual blobs relax (Stokes drag law).

8.3.4 Temporal accuracy

We now use the timescale analysis of the previous section to understand the performance of our temporal integrator. For each set of parameters, we first compute the relevant timescales using the method of Section 8.3.3.1. Then, we choose a subset of the modes and run the Langevin dynamics with $\Delta t = \tau_k$, for each k in the set of modes chosen. To obtain statistics, we initialize the fiber in its equilibrium state and run until $10\tau_1$, removing the first τ_1 as a burn-in period. We repeat this a total of 20 times to generate a mean, then perform the 20 trials a total of 5 times to generate error bars. The resulting variance of each mode is shown in Figs. 8.4 and 8.5. Here we follow the pattern of Fig. 8.3 to systematically vary the parameters N (top of Fig. 8.4), $\hat{\epsilon}$ (bottom), and κ (Fig. 8.5), while keeping the others

constant.

8.3.4.1 Effect of changing N

Beginning with the top of Fig. 8.4, we see that the number of modes whose dynamics we need to resolve with Δt depends on N . For $N = 12$, where there are a total of 27 modes, we can use a time step corresponding to the $k = 13$ mode and obtain an accurate variance for most modes, while for $N = 24$, where there are 51 modes, we need a time step which matches the timescale for the $k = 21$ mode. Thus the number of modes we need to resolve increases with N . Since the timescale of the modes scales as k^{-4} , doubling N requires a time step refinement of at least 10. This is obviously sub-optimal, and tells us that the optimal way to simulate the dynamics of the first k modes is to choose the minimum N that resolves those k modes.

8.3.4.2 Effect of changing $\hat{\epsilon}$

We now move on to the variation in Δt with changing $\hat{\epsilon}$ for $N = 24$. The bottom row of Fig. 8.4 shows that, for a fixed relative time step size, $\hat{\epsilon} = 10^{-2}$ gives results closer to the equilibrium distribution. The required time step size for a 5% error in the variance is about $\Delta t = \tau_{17}$ for $\hat{\epsilon} = 10^{-2}$ and $\Delta t = \tau_{21}$ for $\hat{\epsilon} = 10^{-3}$. Because the scaling of the timescales for intermediate modes with $\hat{\epsilon}$ is somewhere between log scaling (factor of 1.5) and $\hat{\epsilon}^{-1}$ scaling (factor of 10), the net result of this is that the absolute time step size we need decreases by a factor of six (see the stars in Fig. 8.3(b)). That said, changing aspect ratio cannot be accommodated by a simple rescaling of time since it scales each mode differently depending on its smoothness.

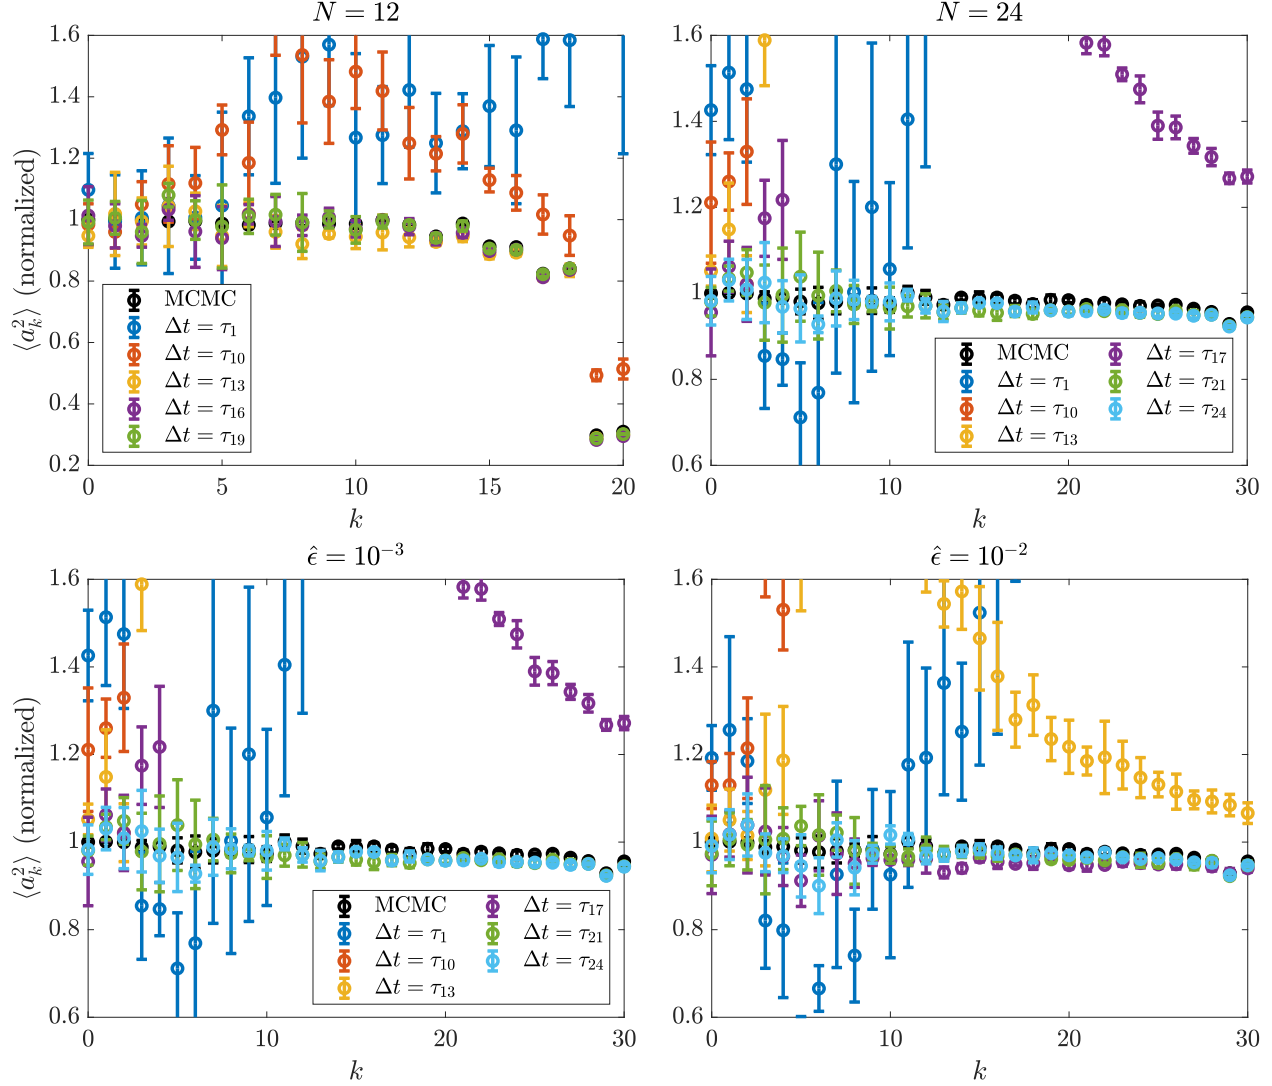


Figure 8.4: Convergence of the equilibrium covariance in our midpoint temporal integrator, with Δt expressed in terms of relaxation timescales. In each case, we plot convergence to the MCMC results (see Fig. 8.2) as a function of the time step size, which is chosen from the modes in Fig. 8.3. (Top) Dependence of timescales on N with fixed $\ell_p/L = 10$ and $\hat{\epsilon} = 10^{-3}$. (Bottom) Dependence on the mobility (aspect ratio $\hat{\epsilon}$) with fixed $N = 24$ and $\ell_p/L = 10$.




Figure 8.5: Convergence of the equilibrium covariance in our midpoint temporal integrator for changing κ , with Δt expressed in terms of relaxation timescales. We fix $N = 12$ and $\hat{\epsilon} = 10^{-3}$ and plot the results for changing κ in terms of ℓ_p/L .

8.3.4.3 Effect of changing κ

Finally, we consider the variation with ℓ_p/L in Fig. 8.5, where we observe the expected behavior: as we increase ℓ_p/L , we need to resolve the dynamics of fewer modes to get the correct variance. However, we find that the number of modes required does *not* scale as $1/\kappa$; rather, we need ≈ 3 less modes for every factor of 10 increase in κ since the chain becomes smoother. Because increasing κ causes the timescale of each mode to decrease, the net effect of this is that the actual time step size we need is roughly constant as κ increases (for the range of parameters we consider here).

8.3.5 Importance of drift terms

To conclude this section, we look at what happens to the variance when we exclude the drift terms from the overdamped Langevin equation (8.12). We do this with the RFD scheme of Section 8.2.2, which allows us to explicitly exclude the drift terms to see what kind of covariance we obtain in \mathbf{X} . To demonstrate that the drift terms are important even when

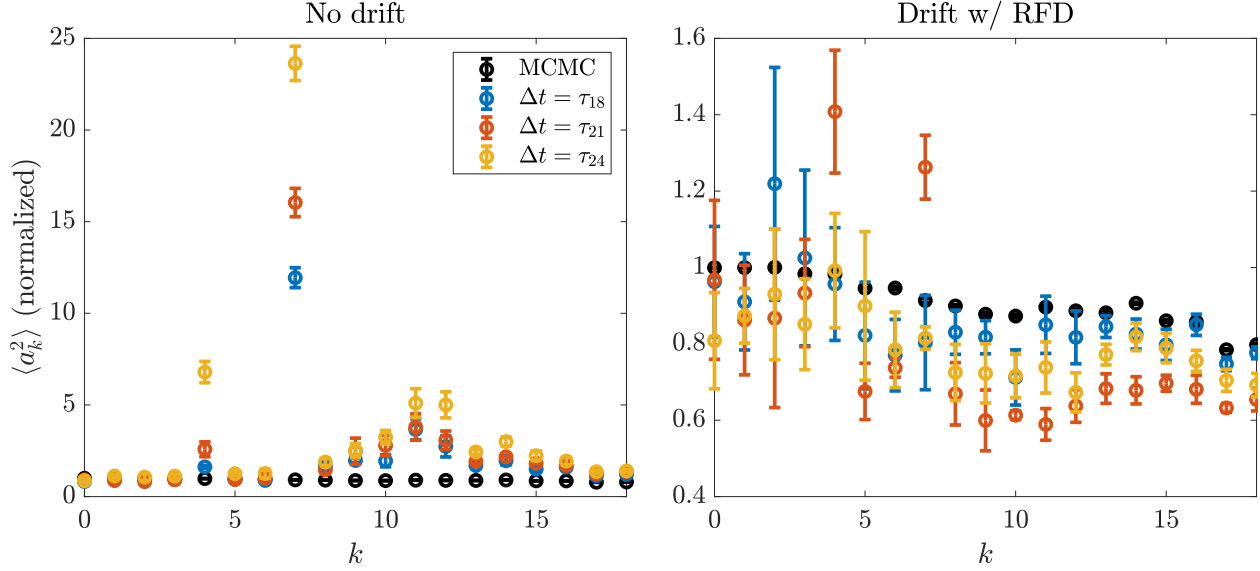


Figure 8.6: Results for linearized fluctuations with constant mobility $\mathbf{M} = \mathbf{I}/(8\pi\mu)$. We fix $N = 12$ and $\ell_p/L = 1$ and show the variance of each of the L^2 -orthonormal modes of the linearized covariance matrix (8.53) (as described in Section 8.3.1.1) using the simple RFD scheme in Section 8.2.2. We show the results without (left) and with (right) the additional RFD drift term in (8.26) with $\delta = 10^{-5}$. In all cases, the MCMC results in black are those obtained in Fig. 8.2(b).

the mobility is position-independent, we use the constant mobility $\mathbf{M} = \mathbf{I}/(8\pi\mu) \rightarrow \widetilde{\mathbf{M}} = \widetilde{\mathbf{W}}/(8\pi\mu)$, which we apply to a fiber with $N = 12$ tangent vectors and $\ell_p/L = 1$. In Fig. 8.6, we show the variance of each of the L^2 orthonormal modes with several different values of Δt , compared to the results from MCMC, which we obtained on the *spectral* grid in Section 8.3.2.4. In the left panel, we show the results when we simulate the dynamics *without* adding the RFD term (8.26). As $\Delta t \rightarrow 0$, many of the modes (but especially 4 and 7) increase in variance without a visible bound. When we add the RFD term back in, the right panel shows that the dynamics converge to the MCMC dynamics as $\Delta t \rightarrow 0$. This shows that, even if the mobility is constant, proper handling of the stochastic drift terms that arise from the inextensibility constraints is necessary to obtain the correct dynamics. We note also that the RFD scheme appears to require *all* of the modes (there are 27 for $N = 12$) to be resolved to obtain accurate results, which is in contrast to the midpoint scheme which requires less than half (c.f. Figs. 8.4 and 8.5).

8.4 Equilibrium statistical mechanics of free fibers

Now that we verified our that our temporal integrator can effectively sample from the equilibrium distribution (8.1) for small fluctuations, we are ready to consider the full nonlinear problem of a fluctuating free filament. As for small fluctuations, we focus first on comparing the Gibbs-Boltzmann distribution (8.1) for blob-link and spectral chains through MCMC calculations, then transition to showing that our midpoint temporal integrator can also sample from the Gibbs-Boltzmann distribution if the time step size is sufficiently small.

For free fibers, we will use end-to-end distance

$$r(t) = \frac{1}{L} \|\mathbb{X}(s=0, t) - \mathbb{X}(s=L, t)\|. \quad (8.67)$$

as a metric to compare statistics across different discretizations. When configurations are sampled from the equilibrium distribution of inextensible filaments, the distribution of r is approximately [154]

$$G(r) = \frac{1}{Z} \sum_{\ell=1}^{\infty} \frac{1}{(\ell_p^* (1-r))^{3/2}} \exp\left(-\frac{(\ell-1/2)^2}{\ell_p^* (1-r)}\right) H_2\left(\frac{\ell-1/2}{\sqrt{\ell_p^* (1-r)}}\right) r^2, \quad (8.68)$$

where $\ell_p^* = \ell_p/L$ is the dimensionless persistence length, $H_2(x) = 4x^2 - 2$ is the second Hermite polynomial, and we have multiplied by the Jacobian factor r^2 to effectively compare (8.68) to a one-dimensional histogram of distances that we will generate from our data. We will also look at other distance metrics, including the distance from the fiber end to its middle, the end to the quarter point, and the distance between the two interior quarter points (the middle half of the fiber).

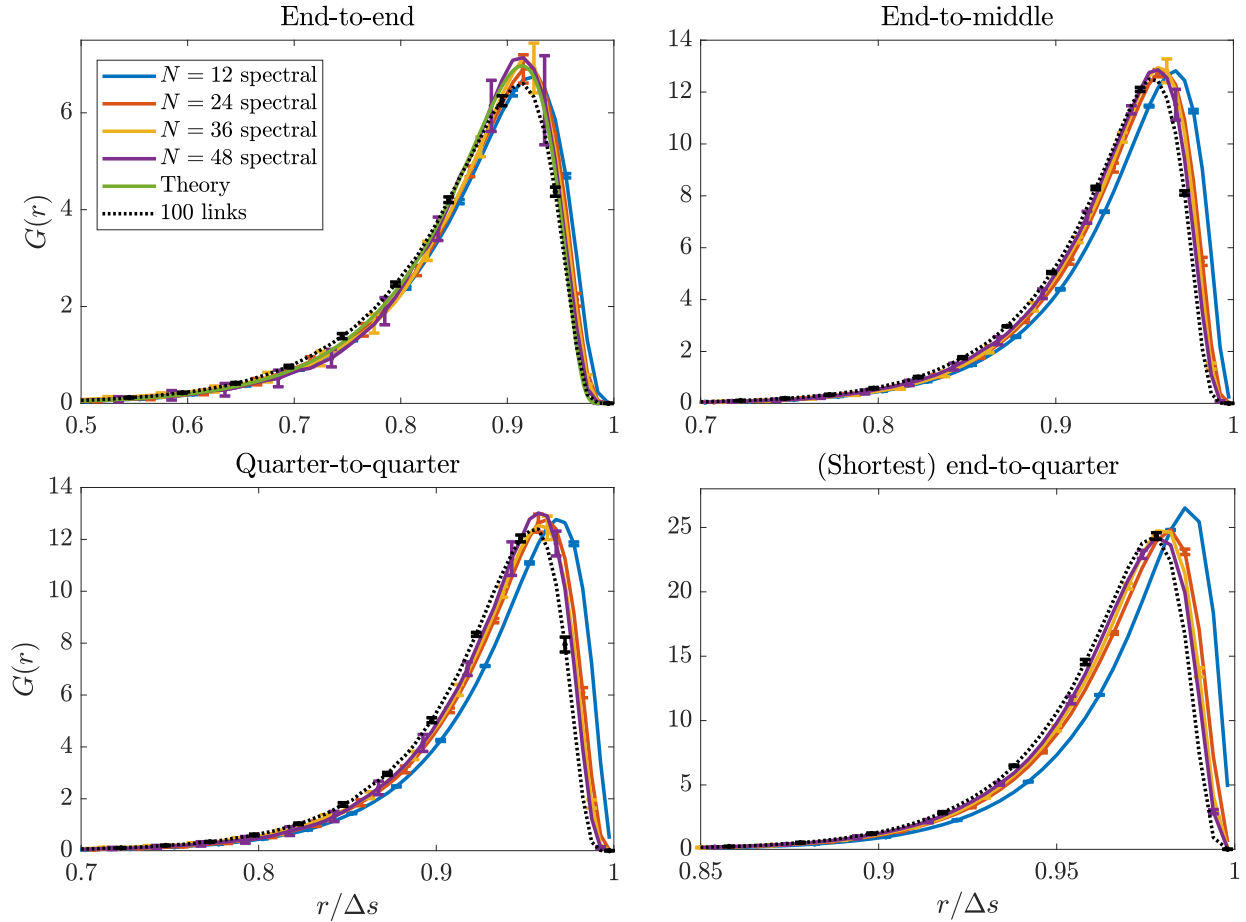


Figure 8.7: MCMC results for a freely fluctuating spectral fiber with $\ell_p/L = 1$. We compute distances along the fiber including end-to-end ($\Delta s = L$, top left), end-to-middle ($\Delta s = L/2$, top right), quarter-to-quarter ($\Delta s = L/2$, bottom left) and the nearest end-to-quarter ($\Delta s = L/4$, bottom right). Note that there are twice as many observations in the right column as the left, since we measure from both the left and right endpoint. We compare the blob-link data for 100 links (which is the same with 200 links; the mean end-to-end distance is $(0.851 \pm 0.001)L$ in both cases) with several spectral discretizations and the theory (8.68). The spectral discretization results are approaching the 100 link results as N increases (to within error bars), with a more rapid approach for larger Δs .

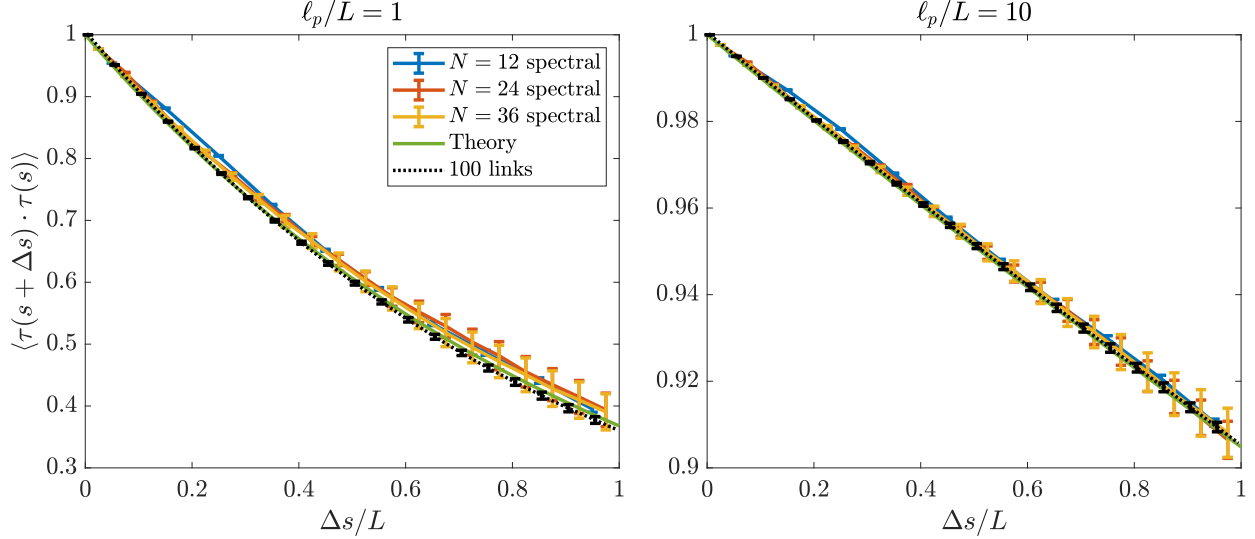


Figure 8.8: Tangent vector correlation function $\langle \tau(s + \Delta s) \cdot \tau(s) \rangle$, computed from MCMC sampling. We show the data for a blob-link discretization with 100 links in black, and $N = 12, 24,$ and 36 spectral nodes in blue, red, and yellow. The theory $e^{-\Delta s/\ell_p}$ is shown in green.

8.4.1 Quantifying the Gibbs-Boltzmann distribution with MCMC

We first examine if the equilibrium distribution (8.1) for a freely fluctuating *spectral* filament approximates well that of a blob-link chain. We use the MCMC procedure described in Section 8.3.2 to generate 10^6 – 10^7 sample chains, removing the first 20% of the chains as a burn-in period. Repeating this ten times to generate error bars, we report the distribution of end-to-end, end-to-middle, quarter-to-quarter, and end-to-quarter distances for both the spectral and blob-link chains in Fig. 8.7. We show only $\ell_p/L = 1$, as the relative errors for $\ell_p/L = 10$ are the same as $\ell_p/L = 1$. The spectral discretization has a relatively small error even when $N = 12$, and the distributions it generates move towards the blob-link ones as N increases. This occurs at a faster rate for larger scales (end-to-end) than for smaller scales (end-to-quarter), as expected.

An additional metric we can use to study the equilibrium distribution of a freely fluctuating chain is the correlation in the tangent vectors $\langle \tau(s + \Delta s) \cdot \tau(s) \rangle$ for $\Delta s \in [0, L]$. According to the definition of persistence length, this correlation should decay exponentially as $e^{-\Delta s/\ell_p}$. To measure the correlation function in the spectral discretization, we compute

the correlation for all Δs on the type 1 Chebyshev grid on which τ is defined (see Fig. 6.1). We then assign these measurements into bins corresponding to 10 (for $N = 12$) or 20 (for $N \geq 24$) uniformly-spaced values of Δs on $[0, L]$. Figure 8.8 shows how our spectral results compare to the theory (and 100 link discretization). While the blob-link chain has a correlation function which matches the theory exactly, there is an apparent small bias in our spectral chain at both small and large distances, with the correlation being larger than expected for small N . This bias is especially noticeable for $\ell_p/L = 1$, but larger error bars for larger N make it harder to make a definitive statement. Combining this with our observations in Section 8.3.2 for small filament fluctuations, we can conclude that the Gibbs-Boltzmann distribution (8.1) for *spectral* filaments is a good approximation of the more physical one for blob-link chains.

8.4.2 Sampling with the midpoint temporal integrator

We now discuss how the midpoint integrator of Section 8.2.3 can also give samples from the Gibbs-Boltzmann distribution (8.1). To make our analysis in this regard universal, we need to understand how a certain time step size generalizes to a set of arbitrary parameters. We are once again aided by our example of a filament with small fluctuations, where we can linearize the SDE (8.12) around a certain state and compute a set of eigenmodes and associated timescales for the dynamics. The analysis in Section 8.3.3.1 shows that the largest timescale for free filaments is associated with the first “fundamental” bending mode [83], which we show in the inset of Fig. 8.3. The timescale associated with this mode is roughly

$$\tau_{\text{fund}} = 0.003 \frac{4\pi\mu L^4}{\kappa \ln(\hat{\epsilon}^{-1})}, \quad (8.69)$$

and so we will report time in units of τ_{fund} .

8.4.2.1 Required time step for midpoint integrator

To examine the accuracy of the midpoint integrator relative to our MCMC calculations, we run Langevin dynamics from $t = 0$ to $t = 10\tau_{\text{fund}}$ on an initially straight filament using the RPY mobility (2.63) with slenderness $\hat{\epsilon} = 10^{-3}$. We record a histogram of the end-to-end distance after the first τ_{fund} (burn-in), ignoring the other distance metrics which we have already seen behave similarly.

Figure 8.9 shows that for sufficiently small Δt the end-to-end distributions from Langevin dynamics converge to those of MCMC, validating our temporal integrator. It also gives an indication of how the required time step sizes change with N and κ (reported in terms of ℓ_p/L). Focusing on N first, we see that for a fixed ℓ_p/L the required time step size decreases by a factor of about 15 as N doubles from 12 to 24. This is in accordance with our finding in Section 8.3.4 that we need to resolve roughly twice as many modes when we double N . The scaling of the timescale of each mode $\tau_k \sim k^{-4}$ implies a decrease in the time step size of 16, which points to the limitations of our temporal integrator for larger N . It seems from Fig. 8.9 that $N = 16$ might be a happy medium between $N = 12$ and $N = 24$, as the required time step is roughly a factor of $3 \approx (16/12)^4$ less than that for $N = 12$.

Switching our focus to κ , Fig. 8.9 shows that the *relative* time step size required as we increase from $\ell_p/L = 1$ to $\ell_p/L = 10$ increases by a factor of roughly 10. As we saw for the linearized modes (see Section 8.3.4), the relative increase in Δt_f is offset by a factor of 10 decrease in τ_{fund} for 10 times larger κ , so that the net effect is no change in the required time step size.

8.5 Dynamics of relaxation to equilibrium

So far, we have only examined equilibrium statistical mechanics, finding that samples from the spectral and blob-link Gibbs-Boltzmann distributions generate similar statistics for a

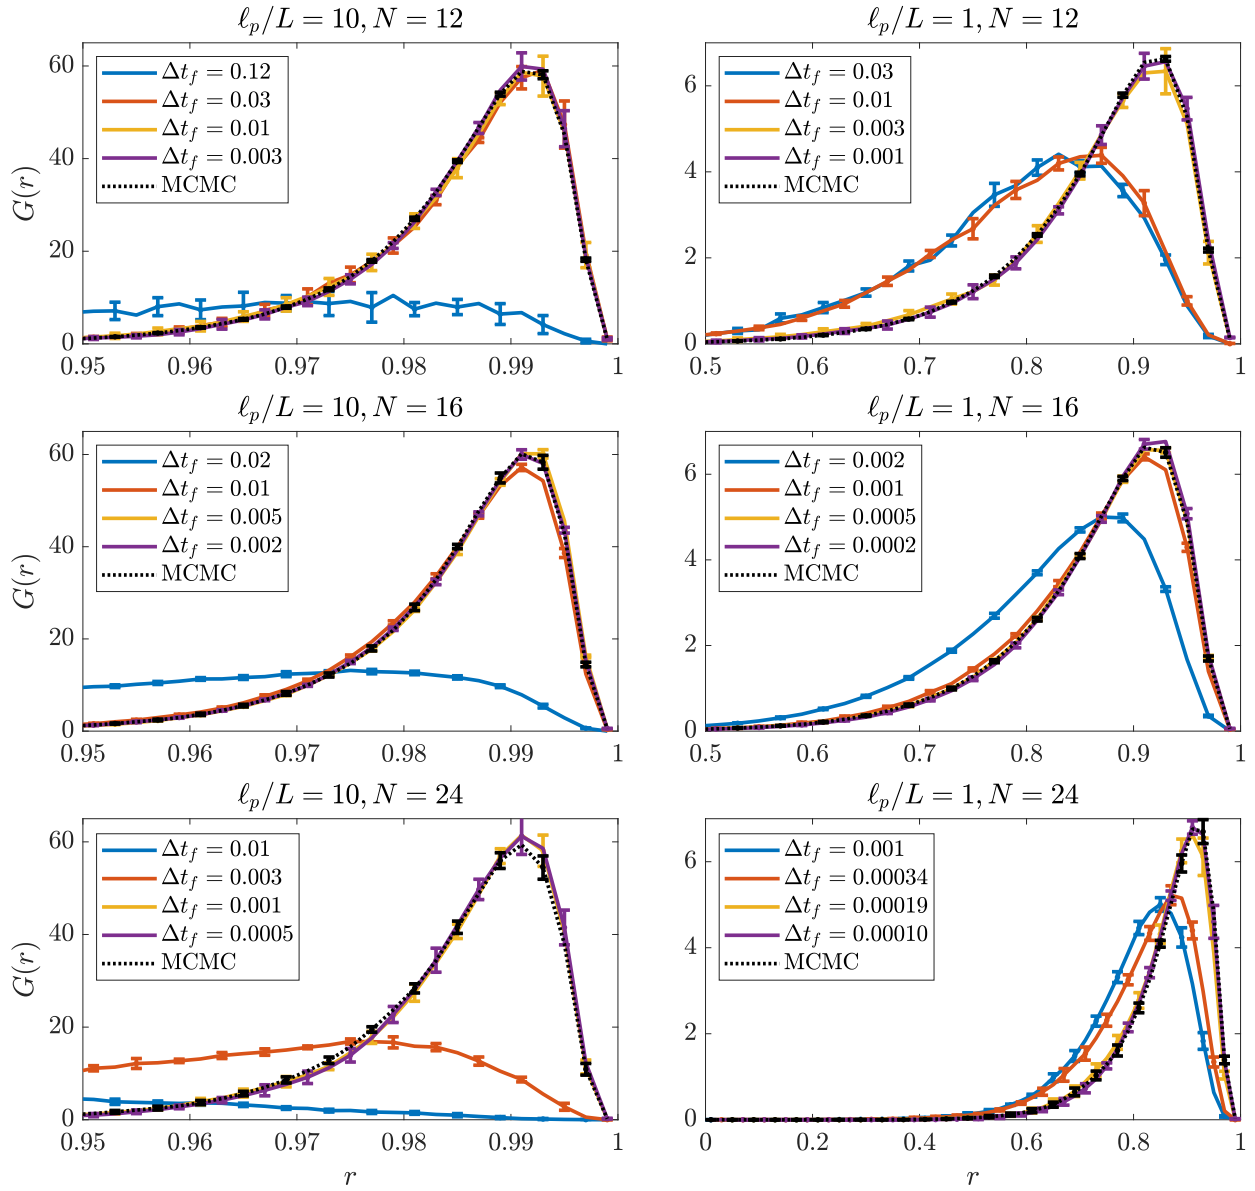


Figure 8.9: Distribution of end-to-end distance for the midpoint temporal integrator discussed in Section 8.2.3 with the RPY mobility and $\hat{\epsilon} = 10^{-3}$. Time step sizes are reported as $\Delta t_f = \Delta t / \tau_{\text{fund}}$, where τ_{fund} is the slowest system relaxation time, defined in (8.69). For each ℓ_p/L and N , we show the MCMC results in black, together with the distributions obtained from the midpoint integrator with various time step sizes.

given set of parameters. But what about dynamics, and in particular, resolving the hydrodynamic interactions in slender filaments?

The temporal integrator we developed here performs similarly regardless of the spatial discretization, in the sense that the number of modes we need to resolve scales with N . Thus, if we want to simulate slender filaments without having to take unreasonably small time steps, our only hope is to resolve the hydrodynamic interactions with a small number of collocation points, or a number of collocation points that is independent of $\hat{\epsilon}$. We already showed (Section 6.3) that we can achieve this goal for *deterministic* hydrodynamics. In this section (and Section 8.5.1), we verify that this is also the case for *Brownian* hydrodynamics.

To do this, we consider the dynamic problem of an initially straight semiflexible chain relaxing to its equilibrium fluctuations [83, 153, 155, 156]. Our focus here is on the relaxation of the mean end-to-end distance to its mean value, i.e., to the mean of the distributions shown in Fig. 8.9. As discussed in [155], the scenario that we simulate is not really physical, since it is not possible for a fluctuating chain to ever reach an exactly straight configuration. As such, the more physically-relevant timescales are those that correspond to long-wavelength modes, where the shorter wavelength modes (which affect the end-to-end distance relatively little) have already reached their equilibrium state. Thus we will accept errors in the end-to-end distance on short timescales and concentrate on long-time behavior [156].

On long timescales, numerical results verify that the data for various μ , L , κ , $\hat{\epsilon}$, and $k_B T$ can be (roughly) collapsed onto a single master curve with rescaled time and end-to-end variables

$$t^* = \frac{t}{\bar{t}}, \quad \bar{t} = 0.0008 \frac{4\pi\mu L^4}{\kappa \ln(\hat{\epsilon}^{-1})} \approx 0.27\tau_{\text{fund}}, \quad r^* = \frac{r - \bar{r}}{1 - \bar{r}}, \quad (8.70)$$

where \bar{r} is the mean end-to-end distance computed in Section 8.4.1 and \bar{t} is the long-time decay rate, i.e., $r^*(t^* = 1) \approx e^{-1} \approx 0.3$. In terms of our modal analysis, the timescale \bar{t} is between the longest and second-longest timescales in the system (see Fig. 8.3), meaning that




Figure 8.10: Two samples of free fibers with $\ell_p/L = 1$ relaxing to their equilibrium fluctuations. In both cases $\hat{\epsilon} = 10^{-2}$, but at left we use a blob-link discretization, while at right we use a spectral discretization with $N = 12$ nodes (we do not try to match the random forcing between the two discretizations, so the comparison is only in the qualitative look of the chain). The fibers are initialized straight (magenta $t^* = 0$ lines), and then assume their equilibrium end-to-end length on the timescale \bar{t} (defined in (8.70)). In both cases, realizations of the fibers at different times are artificially staggered along the direction of the arrows for visual clarity.

all of the modes except the first should be relaxed by \bar{t} . Simulating until \bar{t} thus provides a set of intermediate times at which we can measure non-equilibrium statistics (see Fig. 8.10 for pictures of the relaxation process).

To compare the spectral method to a blob-link method, we fix $\ell_p/L = 1$ and compare four different discretizations of the chain: the spectral discretization with $N = 12$ (which requires a time step size $\Delta t_f \approx 0.003$ for accurate dynamics; this time step is the same as that needed for equilibrium statistical mechanics), $N = 24$ (time step size $\Delta t_f \approx 2 \times 10^{-4}$), and $N = 36$ ($\Delta t_f \approx 7 \times 10^{-5}$) and the blob-link discretization with 100 blobs (required time step size $\Delta t^* \approx 10^{-4}$). The blob-link discretization is considerably more expensive to simulate for small $\hat{\epsilon}$ (even with a GPU-accelerated implementation), which requires us to limit our comparison to $\hat{\epsilon} = 10^{-2}$.

Figure 8.11 shows how the two discretizations compare with each other for three different statistics: the average end-to-end distance $r^* = \langle \|\mathbb{X}(0, t) - \mathbb{X}(L, t)\| \rangle$, the mean-square displacement of the center-of-mass $r_{\text{COM}}^2 = \langle \|\mathbb{X}(L/2, t) - \mathbb{X}(L/2, 0)\|^2 \rangle$, and the average square

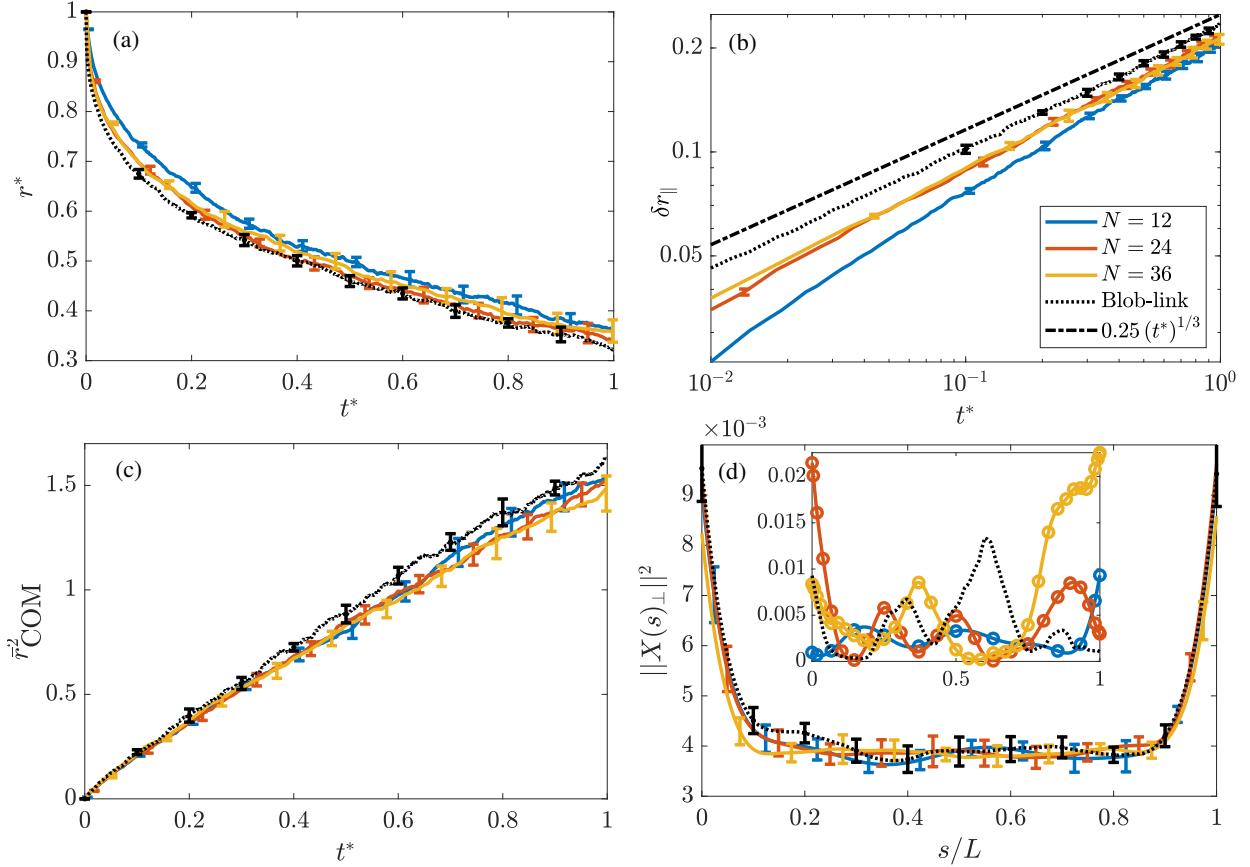


Figure 8.11: Comparing the spectral and blob-link trajectories for relaxation of a straight filament to its equilibrium fluctuations. We set $\ell_p/L = 1$ and $\hat{\epsilon} = 10^{-2}$ and simulate until $t^* = 1$, as defined in (8.70), comparing spectral discretizations with $N = 12$ (blue), $N = 24$ (red), and $N = 36$ (yellow) to the blob-link discretization (black). (a and b) The end-to-end distance over time. In (a), we isolate the *dynamic* error by normalizing each curve by its average \bar{r} obtained from MCMC in Section 8.4.1 (for $N = 12$ this is $\bar{r} = 0.865 \pm 0.001$, for $N = 24$ it is $\bar{r} = 0.862 \pm 0.005$, for $N = 36$ it is $\bar{r} = 0.859 \pm 0.006$, and for the blob-link it is $\bar{r} = 0.851 \pm 0.001$). In (b), we plot δr_{\parallel} , which is the change in end-to-end distance in the parallel direction, for the four different discretizations, observing a $1/3$ power law scaling at long times. (c) The squared-norm of the displacement of the center of mass, normalized by the theoretical value for a rigid fiber (see text for details). (d) The squared perpendicular displacement at $t^* = 0.1$. The inset shows a single sample, where we observe some high-frequency behavior in the blob-link results that is smoothed by the spectral method.

perpendicular displacement $\langle \|\mathbb{X}_\perp(s, t)\|^2 \rangle$ at $t^* = 0.1$. We normalize the center-of-mass MSD by the value at $t = \bar{t}$ for a rigid fiber, which is $2k_B T \bar{t} \text{trace}(\mathbf{N}_{tt})$, where \mathbf{N}_{tt} is the 3×3 matrix relating forces on a rigid fiber to its translational velocity. The normalized displacement is denoted \bar{r}_{COM}^2 .

To separate the error in the *dynamics* from that of equilibrium statistical mechanics, in Fig. 8.11(a), we normalize r^* by the average \bar{r} obtained from MCMC in Section 8.4.1. With this normalization, there is little difference between the spectral method with $N \geq 24$ and the blob-link method at later times, and the difference between the spectral method with $N = 12$ and the blob-link method is small. Combining this with Fig. 8.11(c), which shows that the diffusion of the center of mass is the same (within error bars) across the different discretizations, we can conclude that a small number of spectral nodes can indeed resolve the dynamics at later times, as desired. This is an important statement because the spectral method with $N = 12$ (resp. $N = 24$) uses a time step that is two (resp. one) orders of magnitude larger than that of the blob-link chains, as well as a number collocation points that is an order of magnitude fewer.

To compare our results to theory, in Fig. 8.11(b), we plot the shortening of the end-to-end distance projected onto the initial tangential direction, $\delta r_\parallel(t) = \langle L - (\mathbb{X}(L, t) - \mathbb{X}(0, t)) \cdot \boldsymbol{\tau}(t = 0) \rangle$. Because we no longer normalize by the equilibrium end-to-end distance, we see a larger difference between the spectral and blob-link codes. Focusing on long times, we see that the data approach a $1/3$ power law for $t^* \gtrsim 0.1$, with a faster growth for short times, matching what is observed in [155, Fig. 5(d)]. This $1/3$ exponent is predicted to be universal independent of the way the initial state is prepared; see the second column in [155, Table I].

At early times, we see a more significant difference between the blob-link and spectral code, which can be explained by the fast relaxation of high-frequency modes. In Fig. 8.11(d) we examine the perpendicular displacement along the curve at an early time of $t^* = 0.1$. In the inset, we show a single sample of the (squared) perpendicular displacement, and observe




Figure 8.12: Fiber relaxation results when $\hat{\epsilon} = 10^{-3}$. In this case the blob-link method is too expensive to simulate, so we show the results in the end-to-end distance for the spectral method only with $N = 12$ (blue), $N = 24$ (red), and $N = 36$ (yellow). These plots are analogous to Fig. 8.11(a) and (b), and the difference between the two discretizations is similar to the $\hat{\epsilon} = 10^{-2}$ case, as desired.

the small-length fluctuations in the blob link code which appear at short times. These fluctuations, which control the early-time relaxation, are smoothed out by the spectral code, and therefore treated incorrectly. At late times, their contribution is sufficiently small for the spectral and blob-link codes to match.

While we cannot obtain statistics for the blob-link code when $\hat{\epsilon} = 10^{-3}$ because of the expense in resolving hydrodynamics, we can still repeat our fiber relaxation test using $\hat{\epsilon} = 10^{-3}$ in the spectral method. Figure 8.12 shows the results of this test in end-to-end distance (compare to Fig. 8.11(a) and (b)). We see that the error between the three different values of N is roughly the same across the two aspect ratios, which indicates that the number of points required for a given accuracy is *not* sensitive to the fiber aspect ratio. In addition, we continue to observe the $1/3$ universal power law scaling for δr_{\parallel} at long times [155]. Without blob-link data, it is difficult to say this for certain how the error in the spectral code scales with $\hat{\epsilon}$. Still, the data strongly suggest that we can effectively resolve hydrodynamics of slender fibers with the same number of points, independent of the fiber aspect ratio.

8.5.1 Comparing special quadrature to other mobilities

To flush out this claim even further, in this part of the section we consider different options for the matrix $\widetilde{\mathbf{M}}$, which maps force to velocity in fluctuating fibers. We consider four possibilities:

1. RPY special quadrature as discussed in Section 7.2.1, in which we first compute the matrix \mathbf{M} , then symmetrize $\mathbf{M}\widetilde{\mathbf{W}}^{-1}$ and truncate its negative eigenvalues.
2. RPY oversampled quadrature on a grid of size N_u , which is the SPD matrix (7.10). The main part of this computation is the action of the matrix $\widetilde{\mathbf{M}}_{\text{RPY},u}$, which maps forces to velocities on the upsampled grid.
3. Direct evaluation of the RPY kernel on the spectral grid, i.e., $\widetilde{\mathbf{M}} = \widetilde{\mathbf{M}}_{\text{RPY}}$, which is automatically SPD. This makes the action of the mobility the same as that of the blob-link method, but with uneven spacing of the grid points. While this calculation does not formally see the Chebyshev weights if we work on forces directly, it can still be viewed as Clenshaw-Curtis quadrature on force *density*, with the conversion given in (6.14).
4. The local drag approximation to the matrix \mathbf{M} , which is a block-diagonal matrix with 3×3 blocks given by

$$8\pi\mu \left(\widetilde{\mathbf{M}}_{\text{L}} \right)_{\{i,i\}} = \left(\frac{a_S(s_i) + a_{\text{CLI}}(s_i)}{\|\partial_s \mathbb{X}(s_i)\|} + \frac{2\hat{a}^2}{3} \frac{a_D(s_i)}{\|\partial_s \mathbb{X}(s_i)\|^3} \right) \mathbf{I} \quad (8.71)$$

$$+ \left(\frac{a_S(s_i) + a_{\text{CLT}}(s_i)}{\|\partial_s \mathbb{X}(s_i)\|} - 2\hat{a}^2 \frac{a_D(s_i)}{\|\partial_s \mathbb{X}(s_i)\|^3} \right) \boldsymbol{\tau}_{\{i\}} \boldsymbol{\tau}_{\{i\}}.$$

As in step 1, we symmetrize the matrix $\widetilde{\mathbf{M}} = \mathbf{M}\widetilde{\mathbf{W}}^{-1}$ and truncate its negative eigenvalues (the matrix is not automatically symmetric) to obtain the force-velocity mobility matrix. We emphasize that the local drag matrix here is the actual mobility, and not a preconditioner or approximation of it (as it was in Section 7.4).

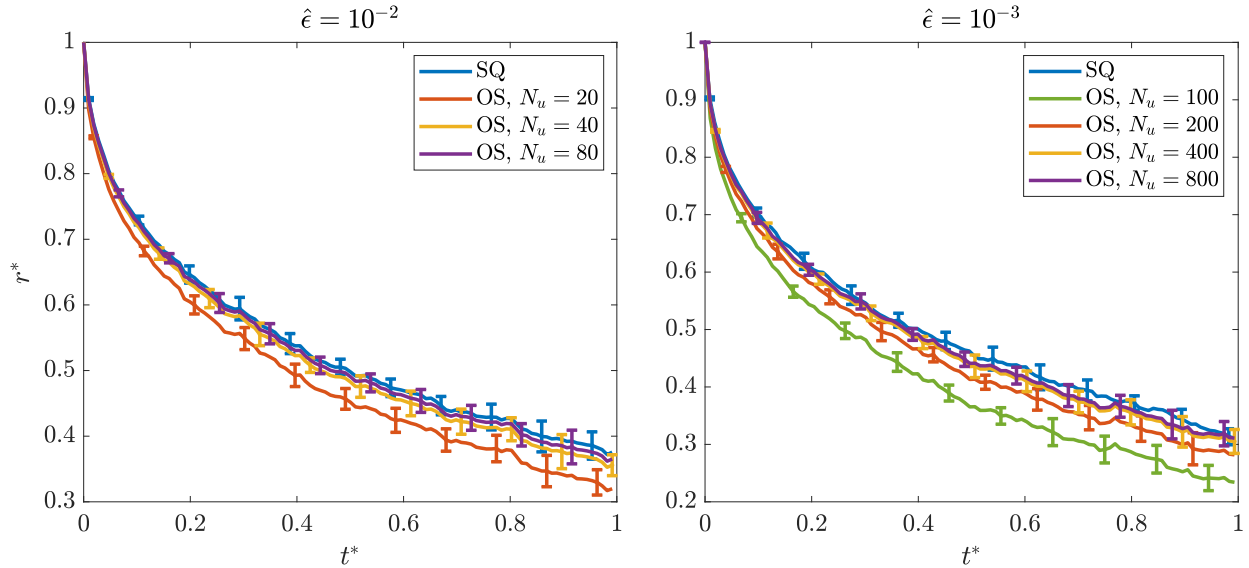


Figure 8.13: End-to-end distance trajectory comparing special quadrature (SQ) with oversampled quadrature (OS) with various numbers of oversampling points N_u , and a fixed number of $N = 12$ collocation points. The left plot shows $\hat{\epsilon} = 10^{-2}$, for which we can use a small number of oversampled points. The more slender $\hat{\epsilon} = 10^{-3}$ at right requires many more oversampled points, but eventually converges to the special quadrature results.

Our strategy here is to first consider only trajectories that are converged in time, thereby isolating the difference between spatial discretizations. We will then zero in on potential temporal convergence problems for the mobilities that appear advantageous at first. The overall picture that emerges is that the special quadrature and local drag mobilities are the only ones which can simulate the dynamics with a cost independent of $\hat{\epsilon}$. The special quadrature mobility converges (in time) at a much larger time step size than the local drag one, which makes it overall the best mobility for fluctuating fibers.

8.5.1.1 Oversampling converges to special quadrature

We can eliminate the ambiguity between special and oversampled quadrature by showing that trajectories using oversampled quadrature converge to those of symmetrized special quadrature as N_u increases. Results for this are shown in Fig. 8.13, where we fix $N = 12$ and show that the number of oversampled points N_u required for the special quadrature

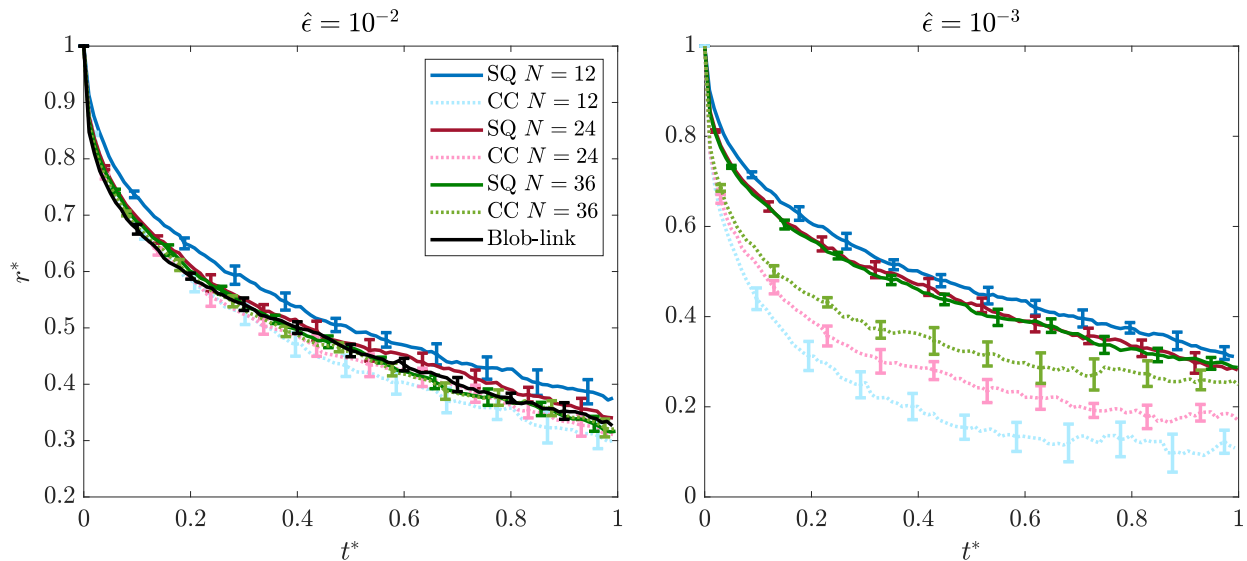


Figure 8.14: End-to-end distance trajectory comparing special quadrature (SQ) with direct quadrature (CC) with various numbers of collocation points. The left plot shows $\hat{\epsilon} = 10^{-2}$, for which we can use a small number of collocation points. The more slender $\hat{\epsilon} = 10^{-3}$ at right shows the weaknesses of direct quadrature (similar to an under-resolved bead link model).

to match oversampled quadrature depends on $\hat{\epsilon}$. We require roughly $0.4/\hat{\epsilon}$ points to obtain equivalence of the two trajectories: thus 40 points are acceptable for $\hat{\epsilon} = 10^{-2}$, which may make the special quadrature worthless for this value of $\hat{\epsilon}$. Decreasing to $\hat{\epsilon} = 10^{-3}$, we need about 400 points for the trajectory to match special quadrature. Considering that the cost of the quadrature is constant with respect to $\hat{\epsilon}$, this tells us that, for sufficiently slender fibers, special quadrature will be advantageous, and going forward we will only compare special quadrature to direct quadrature and local drag.

8.5.1.2 Special vs. direct quadrature

Let's now move on to comparing special quadrature to direct quadrature on the spectral grid. With direct quadrature, the spectral mobility becomes analogous to a blob-link mobility, and so we would expect fewer collocation points to be acceptable for large $\hat{\epsilon}$. Figure 8.14 confirms this result; there we see that when $\hat{\epsilon} = 10^{-2}$, the results for direct quadrature and special quadrature for $N \geq 24$ are all essentially the same. When $N = 12$, we see that the direct

quadrature is actually more accurate than special quadrature (relative to the blob-link or self-refined solutions), though this may be accidental (for this specific example).

As in the previous section, the special quadrature scheme shines only when we drop to $\hat{\epsilon} = 10^{-3}$. This time (right panel of Fig. 8.14), we see that special quadrature converges in space (to within statistical error) for $N \geq 24$, while direct quadrature fails to give an accurate trajectory even for $N = 36$. This is not surprising; for $\hat{\epsilon} = 10^{-3}$ we would expect at least several hundred points (not computationally feasible due to the small required time step size) would be needed to accurately resolve the hydrodynamics via direct quadrature. Still, it is reassuring that those trajectories move toward the converged special quadrature ones.

8.5.1.3 Local drag

So far, we have established that special quadrature is only effective in the slender limit, which is also the regime where we expect local drag to give a correct answer. So what about the difference between special quadrature and local drag? This is what we study in this section. We will compare both local drag and special quadrature to special quadrature with $N = 36$, which matches the blob-link calculation that is available to us, and gives a reference answer for smaller $\hat{\epsilon}$.

Figure 8.15 shows how our local drag theory (2.20) compares with special quadrature for $N = 12, 24$, and 36 (using special quadrature with $N = 36$ as a reference). The key takeaway is that local drag performs better as $\hat{\epsilon}$ decreases, which is not surprising. When $\hat{\epsilon} = 10^{-2}$ and $N = 12$, we see that local drag and special quadrature have roughly the same error at intermediate times, with local drag being more accurate at short times ($t^* < 0.1$) and special quadrature winning out on long times. For $N = 24$ and $\hat{\epsilon} = 10^{-2}$, special quadrature is a clear winner, as the trajectory for local drag, which is unchanged when we increase to $N = 36$, gives faster relaxation than the reference solution. Decreasing to $\hat{\epsilon} = 10^{-3}$, we see




Figure 8.15: End-to-end distance trajectory comparing special quadrature (SQ) with the local drag theory (2.20) with various numbers of collocation points. We use special quadrature with $N = 36$, which for $\hat{\epsilon} = 10^{-2}$ gives the same trajectory as the blob-link code, as a reference result in both cases. The left plot shows $\hat{\epsilon} = 10^{-2}$, while the right shows $\hat{\epsilon} = 10^{-3}$.

that local drag performs better: with $N = 12$ it gives an error identical to special quadrature with $N = 24$. This is likely just a coincidence, since with $N = 24$ and $N = 36$ the accuracy of local drag is clearly inferior to that of special quadrature. Still, it certainly provides a good approximation for $\hat{\epsilon} = 10^{-3}$.

From a spatial perspective then, we see that local drag can be just as effective as special quadrature when $\hat{\epsilon}$ is small. We still have not however, considered the temporal convergence of local drag, which we show in Fig. 8.16 for $\hat{\epsilon} = 10^{-3}$. In the left panel, we show how well behaved the special quadrature is: the data are basically converged when $\Delta t_f = 0.001$ for $N = 12$ (and we could probably use an even larger Δt) and $\Delta t_f = 10^{-4}$ for $N = 24$. The time step sizes for local drag are much smaller; the right panel shows that we need a time step size of $\Delta t_f = 5 \times 10^{-4}$ for $N = 12$ and $\Delta t_f = 2 \times 10^{-5}$ for $N = 24$. Thus, while the special quadrature might be more expensive in space, it is ultimately the most efficient, since local drag requires smaller time step sizes to achieve convergence.

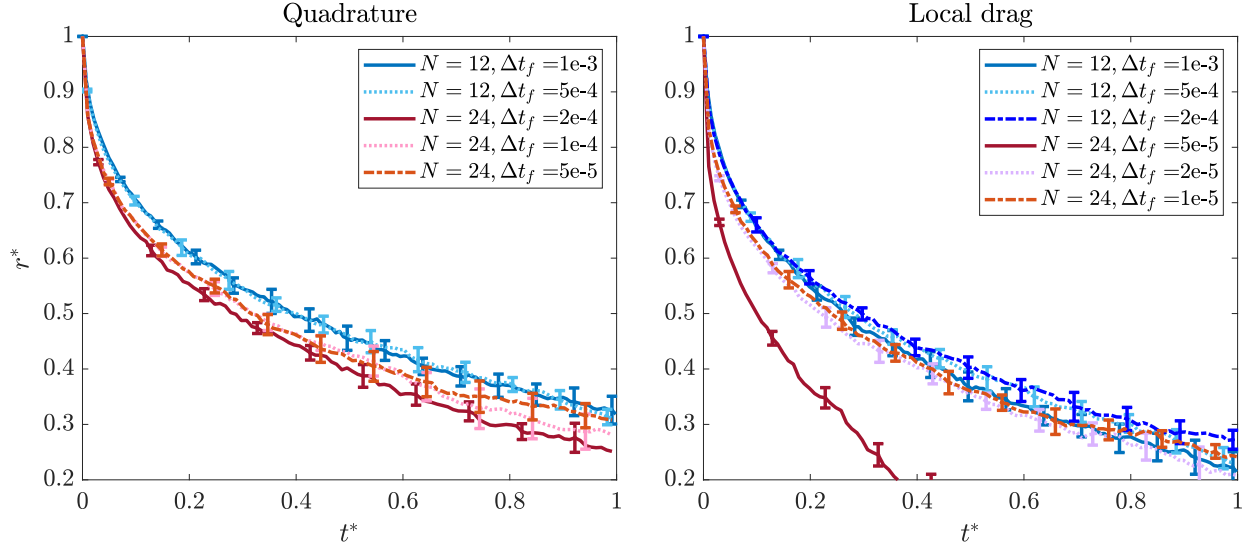


Figure 8.16: Studying temporal convergence of special quadrature (left) and local drag (right). Blue colors show $N = 12$, while red colors show $N = 24$. This is for $\hat{\epsilon} = 10^{-3}$, where local drag is most effective. Time step sizes are reported in terms of $\Delta t_f = \Delta t / \tau_{\text{fund}}$, where τ_{fund} is defined in (8.69).

8.6 Nonlocal hydrodynamics

We now perform a set of corresponding tests when there are multiple fibers in a suspension, all interacting hydrodynamically through the fluid. In this case, the mobility is given by (7.10), since this ensures a symmetric positive definite $\widetilde{\mathbf{M}}$ by construction, and we solve the saddle point system (8.33) using the block diagonal preconditioner (7.49) with $\widetilde{\mathbf{M}}_L$ = the diagonal parts of $\widetilde{\mathbf{M}}_{\text{ref}}$. Unlike in the deterministic case, here the solution of (8.33) is needed for fluctuation-dissipation balance (accuracy) as well as stability. Therefore we need to solve (8.33) to a fixed tolerance (see Section 8.6.1.2), rather than the fixed number of iterations we used in the deterministic case (Section 7.4.3.3). Because we use $\widetilde{\mathbf{M}}_{\text{ref}}$ in (7.10) to do this, the mobility does *not* have a cost independent of $\hat{\epsilon}$. An avenue of future work is therefore to do what we did in deterministic suspensions: replace the diagonal parts of (7.10) with the special quadrature matrix to achieve a cost independent of $\hat{\epsilon}$. The difficult part of this is to keep the mobility $\widetilde{\mathbf{M}}$ SPD and the computation of $\widetilde{\mathbf{M}}^{1/2}$ efficient.

8.6.1 Equilibrium fluctuations

We first test whether the temporal integrator can successfully sample from the end-to-end distance distribution of free fibers at equilibrium. For this we initialize 100 fibers with aspect ratio $\hat{\epsilon} = 10^{-2}$ and $N = 12$ collocation tangent vectors in a periodic box with length L_d . We run the overdamped Langevin dynamics over a time $10\tau_{\text{fund}}$, and save the end-to-end distance after $t = \tau_{\text{fund}}$. Since $\hat{\epsilon} = 10^{-2}$, we can obtain accurate solutions with a small number of oversampling points; in this case we use $N_u = 64$.

8.6.1.1 Comparing to MCMC

In Fig. 8.17, we plot the observed distributions of end-to-end distance for two different values of L_d (columns; $L_d = 8$ corresponds to the limit when the inter-fiber hydrodynamics makes little difference) and ℓ_p/L (rows). We observe convergence to the expected equilibrium distribution. The dimensionless time step required is approximately independent of L_d and is similar to the case without inter-fiber hydrodynamics. Specifically, for $\ell_p/L = 10$ (resp. 1), we see that $\Delta t_f = 0.01$ (0.001) gives convergence, while in Section 8.4 we reported $\Delta t_f = 0.03$ (0.003). The differences for dilute suspensions can be attributed to the new way of estimating the drift term via (8.32), which is required for linear scaling but gives somewhat poorer accuracy than (8.31).

8.6.1.2 Optimizing GMRES and Ewald tolerance

To generate Fig. 8.17, we used a GMRES and Ewald tolerance of 10^{-6} in the system (8.33). This makes the simulations more expensive than necessary, as presumably we can still get good accuracy in the equilibrium distribution with a larger tolerance. With that in mind, let us fix $\Delta t_f = 0.01$ and $\ell_p/L = 10$, and look at how the results change as a function of the GMRES tolerance. We will fix the positively split Ewald (and Lanczos) tolerance at 10^{-3} throughout, as this is the base accuracy we want in deterministic simulations.

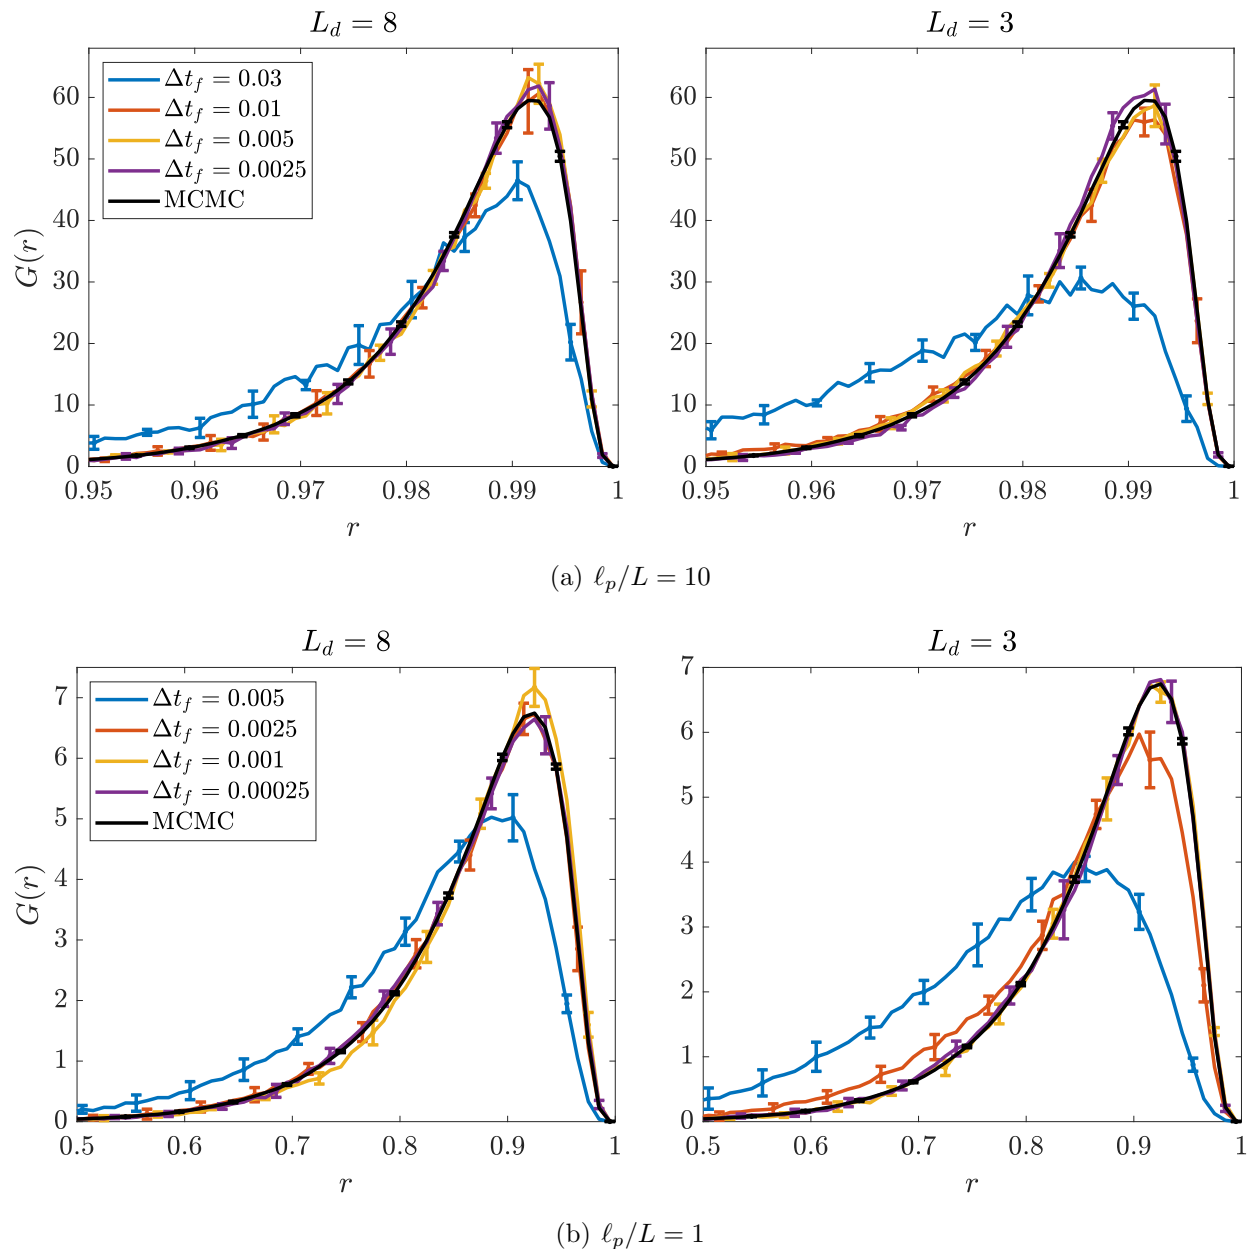


Figure 8.17: End-to-end distance distribution with nonlocal hydrodynamics. (a) Top row: $\ell_p/L = 10$; (b) bottom row: $\ell_p/L = 1$. We consider 100 fibers in a domain of size $L_d = 8$ (left) and $L_d = 3$ (right). In all cases, we see convergence to the distribution from MCMC, but at a slightly smaller time step size than $\Delta t = 0.003\ell_p/L$, which is what we obtained in Section 8.4. The difference can be attributed to the different way of estimating the drift term via (8.32).




Figure 8.18: Experimenting with the GMRES tolerance when $\ell_p/L = 10$ and $L_d = 3$. The blue line is the reference result from Fig. 8.17.

Figure 8.18 shows what happens when we experiment with the GMRES tolerance used to solve (8.33). We compare runs with various values of the tolerance with the reference result from Fig. 8.17, observing that a tolerance of 10^{-3} is sufficient to give the same end-to-end distribution. In this plot, we use $L_d = 3$, since a larger $L_d = 8$ gives the same distribution regardless of the GMRES tolerance (it is completely dominated by local hydrodynamics). This begs the question: is the tolerance a function of the concentration, or is there a maximum tolerance at which the results don't change regardless of concentration?

To see if the value of 10^{-3} for the tolerance is indeed universal, we consider a different set of parameters: $\ell_p/L = 1$ with $L_d = 3$ and $L_d = 2.4$. From Fig. 8.17, we expect that $\Delta t_f = 0.001$ will be sufficient to obtain the correct end-to-end distance distribution. Figure 8.19 confirms this, where we have again run with a tolerance of 10^{-3} . Thus the required tolerance appears to be universal. The number of iterations, however, is not, as for $L_d = 3$ we typically need 6–7 iterations while for $L_d = 2.4$ we need 7–8 iterations. Considering that such a minor change in the domain size produces an extra iteration, it is probably best to fix the GMRES tolerance and vary the number of iterations to get the same accuracy. This is in




Figure 8.19: End-to-end distance distribution (left) and number of required GMRES iterations (right) for fibers interacting hydrodynamically when $\ell_p/L = 1$. We consider domain sizes of $L_d = 3$ (blue) and $L_d = 2.4$ (red). The GMRES tolerance is 10^{-3} , and for both values of L_d we use $\Delta t_f = 0.001$ (in accordance with the required value from Fig. 8.17).

contrast to what we did for deterministic dynamics in Chapter 7, where we fixed the number of iterations. The reason for the difference is that in Chapter 7 we were concerned primarily with stability, while here our main concern is accuracy (specifically, discrete fluctuation-dissipation balance to 3 digits).

8.6.2 Out of equilibrium dynamics

We now look at how the convergence of the algorithm is affected by out of equilibrium dynamics, in particular what happens when we introduce bundling into the system. The model system we consider is 200 filaments of length $L = 1$ in a box of size $L_d = 2$, with all of the same parameters that we use in Part III (see there for more details on the bundling dynamics). Here we set the persistence length $\ell_p = 17 \mu\text{m}$ to be equal to that of an actin filament. In this case, the characteristic bundling time, or the peak in the bundle density, occurs around $t = 3$ to 4 seconds. The precise definitions of bundling are in Chapter 12; here it suffices for the reader to know that more bundles means fibers are closer together,

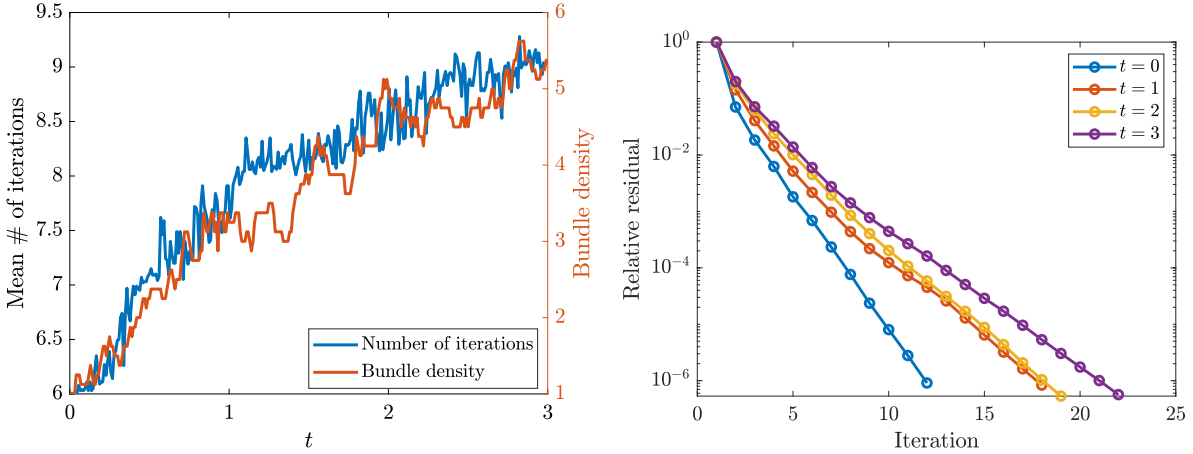


Figure 8.20: GMRES convergence in bundled systems. We consider a system with bundling time $\tau_c \approx 3$ seconds, and look at how the number of iterations depends on the morphology in the regime $0 \leq t \leq \tau_c$. Left panel: the number of iterations (averaged over 0.1 second intervals) as a function of time, with tolerance 10^{-3} . We also plot the bundle density on the right y axis. Right panel: GMRES convergence to 10^{-6} every second (see the plot at left for how bundle density changes over time).

so that the mobility is dominated less and less by intra-fiber hydrodynamics as the bundle density increases.

8.6.2.1 Number of iterations in bundled systems

To avoid super agglomerated networks, we run the system with nonlocal hydrodynamics and filament fluctuations until $t = 3$ seconds. At each time step, we converge GMRES in the midpoint integrator until a tolerance of 10^{-3} . The number of iterations needed over time is shown in the left panel of Fig. 8.20, where we observe a moderate increase as the bundle density increases. Indeed, while the bundle density goes from 1 to 5, the number of iterations increases by only 3, which is a sign that our algorithm and preconditioner is robust in this regime. The full convergence to tolerance 10^{-6} is shown in the right panel of Fig. 8.20, where we see a moderate slow-down as time (and bundle density) increases.

Finally, the preconditioner can be made sparse by only including a finite number of bands around the diagonal. However, in solving (8.33), we precompute and store the preconditioner

and its factorization at the first application. Thus subsequent applications are essentially free. This means the cost of the algorithm is dominated by applying the nonlocal mobility $\widetilde{\mathbf{M}}_{\text{NL}}$, and so it is a waste of time to optimize the preconditioner.

Chapter 9

Other nonlocal interactions: cross linking and sterics

The purpose of this chapter is to outline how we compute the external forces \mathbf{F}_{ext} in the saddle point system. In our framework, these forces are treated *explicitly* (so that each solve is fiber-by-fiber), and so all we need to do is evaluate $\mathbf{F}_{\text{ext}}(\mathbf{X})$ at every time step. We will deal with two types of external forcing in this chapter: cross linking (which is a dynamic process where CLs can bind and unbind) and steric interactions, which are intended to keep the fibers from crossing.

9.1 Cross linking model

This section deals with how we compute the forces due to cross linking. This is a two step process: first, when the cross linkers are transient, we must update the list of connections between filaments. Then, given the connectivity list, we can compute forces.




Figure 9.1: Our model of dynamic cross-linking with fiber turnover. We coarse-grain the dynamics of individual CLs into a rate for each end to bind to a fiber. The first end (purple) can bind to one fiber at any binding site. Once bound, we account for thermal fluctuations of the CL length by allowing the second end (red) to bind to any other fiber within a distance $\ell \pm \delta\ell$ of the first end, where ℓ is the CL rest length. To model actin turnover, we allow each fiber to disassemble with rate $1/\tau_f$ and for a new (straight) fiber to assemble in a random place (nascent fiber is shown in green) at the same time.

9.1.1 Markov chain for transient linkers

To model transient cross linking, we assume that the density of cross-links is sufficiently high that CLs do not get locally depleted by binding, and that the links diffuse rapidly. By “rapid” diffusion, we mean that the Brownian motion of the links is fast relative to the residence time of a single link, which is on the order of one second (see Table 11.1). To justify this, we consider the translational diffusion of a rod of length $\ell = 50$ nm and radius $a = 4$ nm. The translational diffusion coefficient of the rigid rod can be approximated by Broersma’s relations as [157]

$$D_t = \frac{kT}{3\pi\mu\ell} \left(\zeta - \frac{1}{2} (\gamma_{\parallel} + \gamma_{\perp}) \right) \quad (9.1)$$

$$\text{where } \zeta = \ln \left(\frac{\ell}{a} \right), \quad \gamma_{\parallel} = 1.27 - 7.4 (\zeta^{-1} - 0.34)^2, \quad \gamma_{\perp} = 0.19 - 4.2 (\zeta^{-1} - 0.39)^2$$

For our parameters in Table 11.1, we obtain $\zeta = 2.53$, $\gamma_{\parallel} = 1.25$, $\gamma_{\perp} = 0.19$, and $D_t = 0.15 \mu\text{m}^2/\text{s}$. This implies a center of mass displacement of $\langle r^2 \rangle = 6D_t t \approx 0.9t$. Thus the time for a CL to travel a mesh size of $r = 0.2 \mu\text{m}$ is about 0.04 seconds, which is much faster than the one second lifetime of each link. These assumptions considerably simplify our simulations,

as we do *not* have to keep track of explicit CLs, but only the number of links bound to each fiber, which can be tuned by setting the ratio of the binding and unbinding rates. A more theoretical study of this limit of infinitely large diffusion (perfect mixing) has recently been performed in [158], with limiting equations consistent with the model we summarize next.

There are four reactions in the model (parameters are summarized in Table 11.1):

1. The attachment of one end of a free CL to a fiber, which occurs with rate k_{on} (units $1/(\text{length} \times \text{time})$). Here k_{on} accounts for the rate of diffusion of the CL until one end gets sufficiently close to a fiber, and then the reaction for the CL to bind to the site.
2. The detachment of a singly-bound CL end to become a free-floating CL, which occurs with rate k_{off} (1/time). This is the reverse of reaction 1.
3. The attachment of the second end of a singly-bound CL to *another* (distinct) fiber, which occurs with rate $k_{\text{on,s}}$ (units $1/(\text{length} \times \text{time})$). This necessarily follows after reaction 1. The rate $k_{\text{on,s}}$ accounts for the diffusion of the second end to find a fiber within a sufficient distance, and then the binding reaction for the end to latch on to the fiber. We model the thermal fluctuations in the CL length by allowing a singly-bound CL to bind its free end to another attachment site on a different fiber that is within a distance interval

$$(\ell - \delta l, \ell + \delta l), \quad \text{where} \quad \delta l = \sqrt{\frac{kT}{K_c}}. \quad (9.2)$$

Note that this relationship (using a constant $k_{\text{on,s}}$) does *not* guarantee that the links are passive. When we consider Brownian fluctuations, it will be important to maintain detailed balance in the link dynamics, so that they neither consume nor create energy. The modifications we make in this case are in Section 9.1.2.

4. The detachment of a single end of the doubly-bound CL, so that it becomes singly-bound, which occurs with rate $k_{\text{off,s}}$ (1/time). This is the reverse of reaction 3.

A schematic of the four reactions is shown in Fig. 9.1.

To simulate these reactions stochastically, we employ a variant of the Stochastic simulation/ Gillespie algorithm [159, 160]. To advance the system by a time step Δt , we calculate the rate of each reaction and use the rate to sample an exponentially-distributed time for each reaction to occur. We choose the minimum of these times, advance the system to that time, and then recalculate the rates for the reactions that were affected by the previous step. We organize the event queue efficiently using a heap data structure [161], so that we can compute the next reaction in $\mathcal{O}(\log n)$ time.

We break each filament into N_s possible uniformly-spaced binding sites, where each of the sites represents a binding region of size $L/(N_s - 1) = \Delta s_u$. To properly resolve the possible connections between fibers, we require $\Delta s \lesssim \delta \ell$. We then compute the rates of each reaction in the following way:

1. Single end binding: the rate of a single end binding to one of the N_s attachment points on one of the F fibers is $k_{\text{on}}\Delta s_u$. We implement this as a single event with rate $N_s F k_{\text{on}}\Delta s_u$, where the specific site is chosen uniformly at random once the event is selected. Multiple CLs are allowed to bind to a site, since each site represents a fiber segment of length Δs_u that is typically longer than the minimal possible distance between two CLs.
2. Single end unbinding: we keep track of the number of bound ends on each of the $N_s F$ sites. If a site has b ends bound to it, we schedule a single end unbinding event from that site with rate $k_{\text{off}}b$. We do this for each site separately, so that this reaction contributes $N_s F$ events to the event queue.
3. Second end binding: we first construct a neighbor list of all possible pairs of sites (i, j) that are on distinct fibers and separated by a distance $\ell + \delta \ell$ or less, using a linked list cell data structure [162]. If site i has b free ends bound to it, we schedule a binding

event for that pair of sites with rate $bk_{\text{on},s}\Delta s_u$. Notice that this naturally implies a zero binding rate if there are no free ends bound to site i . We treat the (j, i) connection separate from the (i, j) one for simplicity at the cost of increasing the number of events in the queue.

4. Second end unbinding: if a link exists between sites i and j , it can unbind from one of the sites with rate $k_{\text{off},s}$. Letting β be the number of links in the system, the total rate for which bound links unbind from one end is $2\beta k_{\text{off},s}$. We therefore schedule a single event with rate $2\beta k_{\text{off},s}$, and, once it is chosen, select the link and end to unbind randomly and uniformly.

9.1.2 Keeping the CL dynamics in detailed balance

When we account for thermal translation and rotation of the fibers, we also want to be sure that the binding and unbinding of the links is consistent with detailed balance, which is not the case for the constant rates we introduced in the previous section. Let \mathbf{C} denote a configuration of C links (list of fiber pair connections), and \mathbf{x} denote the configuration of fibers (binding site positions). The desired Gibbs-Boltzmann equilibrium distribution is

$$P_{\text{eq}}(\mathbf{C}, \mathbf{x}) = \zeta(\mathbf{C}) \prod_{k \in \mathbf{C}} \exp\left(-\frac{K_c}{2} \frac{(\ell_k - \ell)^2}{k_B T}\right), \quad (9.3)$$

where ℓ_k is the length of link k , and $\zeta(\mathbf{C})$ determines the probability to observe the cross-linking configuration \mathbf{C} . Now consider a transition to/from a state \mathbf{C}' with one added link C' , which has length ℓ'_k . Then at equilibrium the transition between the two states must obey

$$P_{\text{eq}}(\mathbf{C}, \mathbf{x}) k_{\text{on},s}(\ell'_k) = P_{\text{eq}}(\mathbf{C}', \mathbf{x}) k_{\text{off},s}(\ell'_k). \quad (9.4)$$

Substituting P_{eq} from (9.3) into (9.4), we obtain the constraint of detailed balance

$$\frac{k_{\text{on},s}(\ell'_k)}{k_{\text{off},s}(\ell'_k)} = \frac{\zeta(\mathbf{C}')}{\zeta(\mathbf{C})} \exp\left(-\frac{K_c(\ell'_k - \ell_c)^2}{2k_B T}\right). \quad (9.5)$$

To satisfy (9.6) for every choice of \mathbf{C} and \mathbf{C}' with binding and unbinding rates that only depend on spring length and not \mathbf{C} and \mathbf{C}' , we must have [158, Sec. II(G).1]

$$\frac{k_{\text{on},s}(\ell'_k)}{k_{\text{off},s}(\ell'_k)} = \frac{k_{\text{on},s}^0}{k_{\text{off},s}^0} \exp\left(-\frac{K_c(\ell'_k - \ell_c)^2}{2k_B T}\right), \quad (9.6)$$

where $k_{\text{on},s}^0$ and $k_{\text{off},s}^0$ are the transition rates for a link at rest length.

To satisfy (9.6), we maintain a constant $k_{\text{off},s}(\ell'_k) = k_{\text{off},s}^0$ and set

$$k_{\text{on},s}(\ell'_k) = k_{\text{on},s}^0 \exp\left(-\frac{K_c(\ell'_k - \ell_c)^2}{2k_B T}\right). \quad (9.7)$$

Other choices are possible; for example, the rate of unbinding can depend on the stretch [163, 164]. To efficiently search for possible binding pairs, we approximate the set of all binding combinations by setting the maximum link stretch in (9.2) to be two standard deviations of the Gaussian (9.7), i.e., $\delta\ell = 2\sqrt{k_B T/K_c}$. This change effectively increases the amount a CL can stretch when we move from “non-Brownian” dynamics ($\delta\ell = \sqrt{k_B T/K_c}$) to “Brownian” dynamics ($\delta\ell = 2\sqrt{k_B T/K_c}$). This leads to misleading results when comparing Brownian and non-Brownian filaments (see Section 12.2 for more discussion).

9.1.3 Computing forces from fiber connectivity

Letting $\overline{\mathbf{X}}$ denote the set of upsampled uniform points (binding sites in this case), the stochastic algorithm gives pairs of two uniform points $\overline{\mathbf{X}}_{\{p\}}^{(i)}$ and $\overline{\mathbf{X}}_{\{q\}}^{(j)}$ between which a CL force is exerted. We offer two ways of computing a force on the fiber centerline from this connectivity list. The first way is energy-based, but dangerous because it can lead to nonsmooth forces

on the Chebyshev grid. The second way is more computationally expensive, but leads to smoother forces.

In both cases, our model assumes that the CL acts as a linear spring between binding sites on two distinct fibers, and is an approximation to the true behavior of biological cross linkers; as shown in [165], the stress-strain relationship for α -actinin, which is the model CL we consider here, is nonlinear and exhibits hysteresis for loading and unloading. We follow a number of other modeling studies [166, 167] in approximating α -actinin as a linear spring, although we do not include an angular stiffness, since recent experimental results [168] indicate that the dynamics of the α -actinin-actin bond are insensitive to rotation.

9.1.3.1 Energy-based approach

When we resample the fiber $\mathbb{X}(s)$ at N_u uniform points to form the vector $\bar{\mathbf{X}}$ of CL binding site locations, the Chebyshev points are related by the simple relationship

$$\bar{\mathbf{X}} = \mathbf{R}^{(u)} \mathbf{X} \leftrightarrow \bar{\mathbf{X}}_{\{p\}} = \mathbf{R}_{\{p,:\}}^{(u)} \mathbf{X}, \quad (9.8)$$

where $\mathbf{R}^{(u)}$ is the resampling matrix and $\mathbf{R}_{\{p,:\}}^{(u)}$ is its p th row.

Denoting the uniform points connected by the CL with $\bar{\mathbf{X}}_{\{p\}}^{(i)}$ and $\bar{\mathbf{X}}_{\{q\}}^{(j)}$, we have the displacement

$$\mathbf{r} = \bar{\mathbf{X}}_{\{p\}}^{(i)} - \bar{\mathbf{X}}_{\{q\}}^{(j)} = \mathbf{R}_{\{p,:\}}^{(u)} \mathbf{X}^{(i)} - \mathbf{R}_{\{q,:\}}^{(u)} \mathbf{X}^{(j)} \quad (9.9)$$

with $r = \|\mathbf{r}\|$. Then, let us postulate the CL energy

$$\mathcal{E}(r) = \frac{K_c}{2} (r - \ell)^2 \quad (9.10)$$

$$\frac{\partial \mathcal{E}}{\partial r} = K_c (r - \ell), \quad (9.11)$$

between the two uniform points, where K_c is the spring constant for the CL (units force/length)

and ℓ is the rest length. The corresponding force on fibers i and j can then be obtained by differentiating the energy

$$\begin{aligned} \mathbf{F}_{\{a\}}^{(i)} &= -\frac{\partial \mathcal{E}}{\partial r} \frac{\partial r}{\partial \mathbf{X}_{\{a\}}^{(i)}} & \mathbf{F}_{\{b\}}^{(j)} &= -\frac{\partial \mathcal{E}}{\partial r} \frac{\partial r}{\partial \mathbf{X}_{\{b\}}^{(j)}} \\ \mathbf{F}_{\{a\}}^{(i)} &= -\frac{\partial \mathcal{E}}{\partial r} \hat{\mathbf{r}} \mathbf{R}_{\{p,a\}}^{(u)} & \mathbf{F}_{\{b\}}^{(j)} &= \frac{\partial \mathcal{E}}{\partial r} \hat{\mathbf{r}} \mathbf{R}_{\{q,b\}}^{(u)} \end{aligned} \quad (9.12)$$

In the special case when the two uniform points are Chebyshev points, we get $\mathbf{R}_{\{p,a\}}^{(u)} = \delta_{ap}$ and $\mathbf{R}_{\{q,b\}}^{(u)} = \delta_{qb}$, and the resulting force is a delta function between the two points on the Chebyshev grid, which demonstrates that it can be nonsmooth in nature. By superposition, an equivalent (though slower) formulation is to compute all forces on the uniform points, then multiply by the matrix \mathbf{R}^T to obtain the forces on the Chebyshev points.

9.1.3.2 Smooth forces

It may be desirable to preserve the smoothness of the filaments in time and therefore the spectral accuracy of our numerical method. In this case, another option is to define the cross-linking force *density* between two fibers $\mathbf{X}^{(i)}$ and $\mathbf{X}^{(j)}$, then convert to force via (6.14). Suppose that a CL connects two fibers by attaching to arclength coordinate s_i^* on fiber i and s_j^* on fiber j . Then a smooth force density on fiber i due to the CL is

$$\mathbf{f}_{\text{ext}}^{(i)}(s) = -K_c \left(1 - \frac{\ell_c}{\left\| \overline{\mathbf{X}}_{\{p\}}^{(i)} - \overline{\mathbf{X}}_{\{q\}}^{(j)} \right\|} \right) \delta_h(s - s_i^*) \int_0^L \left(\mathbf{X}^{(i)}(s) - \mathbf{X}^{(j)}(s') \right) \delta_h(s' - s_j^*) ds' \quad (9.13)$$

where s_i^* and s_j^* are the anchor points on fibers i and j , and δ_h is a Gaussian density with standard deviation σ . For the force to be smooth, the Gaussian width σ must be a function of N . In particular, for $N = 16$, the width $\sigma/L = 0.1$ is sufficient for smoothness. In the limit $\sigma \rightarrow 0$, the Gaussian density becomes a Dirac delta function, and the force density (9.13) becomes the expected expression for a linear spring connecting the points $\overline{\mathbf{X}}_{\{p\}}^{(i)}$ and

$\overline{\mathbf{X}}_{\{q\}}^{(j)}$.

9.1.4 Fiber turnover

We account for the turnover of actin filaments by removing a filament (and all of its attached links) from the system and rebirthing a new, straight filament with a random location and orientation elsewhere in the system (see schematic in Fig. 9.1). Our reasoning for this is twofold: first, it is simple to implement, and second, our simplifying assumptions are supported by experimental data. Experiments have shown that pointed-end depolymerization from one end of an actin filament is too slow to explain the observed rates of filament turnover *in vivo* [169], and that depolymerization actually occurs in bursts as filaments abruptly sever into smaller pieces [170]. In this sense, our model of depolymerization as an instantaneous event is supported by the data. For polymerization, the process is normally linear, with monomers being added at the plus end until it is capped. It has been shown experimentally [171], however, that the combined time of growth, capping, and disassembly is still significantly smaller than the time in which the filament is at a finite length, so we assume the former to be instantaneous relative to the latter. We refer to [15, Sec. 2.5] for a mathematical treatment of linear polymerization.

Having chosen this model, fiber turnover can be added as another reaction in our stochastic simulation algorithm. Defining the mean turnover time of each fiber as τ_f , we add a fiber turnover reaction to our reaction list with rate F/τ_f . If this reaction is chosen, we select the fiber to turnover randomly and uniformly. We emphasize that our time steps are at most 10^{-3} seconds while the turnover times are of the order seconds, so we almost always see zero or one fiber turnover events per time step.

9.1.5 Temporal integration

We have three different events to process at each time step: fiber turnover, link turnover, and fiber movement. We will treat them sequentially using a first order Lie splitting scheme [172]. For each time step n of duration ∂_t , from time $n\Delta t$ to $(n+1)\Delta t$, we

1. Turnover filaments to update configuration $\mathbf{X}^n \rightarrow \mathbf{X}^{n,*}$.
2. Process binding/unbinding events to evolve $\mathbf{C}^n \rightarrow \mathbf{C}^{n+1}$, where \mathbf{C} represents the state of the link Markov chain.
3. Evolve the fiber dynamics $\mathbf{X}^n \rightarrow \mathbf{X}^{n+1}$ using the external forcing $\mathbf{F}_{\text{ext}}(\mathbf{X}^n; \mathbf{C}^{n+1})$

9.2 Steric interactions

In addition to cross linking forces, when appropriate the external force \mathbf{F}_{ext} in the previous section can also include steric interactions, designed to keep the fibers from contacting each other. In this section, we present two models of steric interactions: one based on resampling the fiber to $\sim 1/\epsilon$ uniform points to find nearly touching regions, and another based on a combination of straight segment approximations and nonlinear optimization. The former method is simpler, but performs poorly for small ϵ , while the latter method is complex, but shines when ϵ is small. Both methods give the same steric forces if sufficiently refined.

The steric interaction energy between fibers i and j can be written as

$$\begin{aligned} \mathcal{E}^{(ij)} &= \int_0^L \int_0^L \hat{\mathcal{E}}(r(s^{(i)}, s^{(j)})) ds^{(i)} ds^{(j)}, \\ r(s^{(i)}, s^{(j)}) &= \|\mathbb{X}^{(i)}(s^{(i)}) - \mathbb{X}^{(j)}(s^{(j)})\|, \end{aligned} \tag{9.14}$$

where $\hat{\mathcal{E}}$ is the potential density function (units energy/area), which we assume is compactly supported on $0 \leq r \leq r_{\text{max}}$. For the potential density function, we use an error function so

that the force will ultimately be a Gaussian with standard deviation δ

$$\begin{aligned}\hat{\mathcal{E}}(r) &= \frac{\mathcal{E}_0}{a^2} \operatorname{erf}\left(r/(\delta\sqrt{2})\right) \\ \frac{d\hat{\mathcal{E}}}{dr} &= \frac{\mathcal{E}_0}{a^2\delta} \sqrt{\frac{2}{\pi}} \exp\left(-r^2/(2\delta^2)\right)\end{aligned}\tag{9.15}$$

In this potential density, we have two free parameters: δ , which is the scale on which the force decays, and \mathcal{E}_0 , which controls the magnitude of the forces (and therefore the temporal stiffness) has units of energy. We set $\mathcal{E}_0 = 4k_B T$ [173], which leaves δ as the parameter that controls the decay. If we truncate the Gaussian at four standard deviations, the truncation distance $r_{\max} = 4\delta$. We set $\delta = a$ to make the steric repulsion short ranged.

To determine the forces, we will compute the double integral (9.14) via upsampling the position to a new Chebyshev grid. If the Chebyshev points used to discretize $\mathbb{X}(s)$ are given by \mathbf{X} , then the upsampled points will be denoted by $\mathbf{X}^{(u)} = \mathbf{E}\mathbf{X}$. The double integral can then be evaluated and differentiated via

$$\begin{aligned}\mathcal{E} &= \sum_k \sum_j \hat{\mathcal{E}}\left(\left\|\mathbf{X}_{\{k\}}^{(u)} - \mathbf{X}_{\{j\}}^{(u)}\right\|\right) w_k w_j \\ &= \sum_k \sum_j \hat{\mathcal{E}}\left(\left\|\mathbf{E}_{kp}\mathbf{X}_{\{p\}} - \mathbf{E}_{jq}\mathbf{X}_{\{q\}}\right\|\right) w_k w_j\end{aligned}\tag{9.16}$$

$$\mathbf{F}_{\{a\}}^{(i)} = -\frac{\partial \mathcal{E}}{\partial \mathbf{X}_{\{a\}}^{(i)}} = -\sum_k \sum_j \frac{\partial \hat{\mathcal{E}}}{\partial r}(r_{kj}) \hat{\mathbf{r}}_{kj} \mathbf{E}_{ka} w_k w_j.\tag{9.17}$$

The last equation gives the force at Chebyshev node a on fiber i , and is a function of the integration weights w_k and w_j of points k and j on the upsampled grid.

In our algorithm, we will *always* use (9.17) to compute forces. The freedom we have is how to choose the oversampled points $\mathbf{X}^{(u)}$ (and therefore the matrix entry \mathbf{E}_{ka} in (9.17)). We present two possibilities for doing this: global upsampling (guaranteed to work, but inefficient in the limit $\epsilon \rightarrow 0$), as well as a second algorithm based on selective upsampling

for pieces of the fibers that are nearly in contact. This second algorithm achieves our goal of a cost independent of ϵ .

9.2.1 Global uniform point resampling

Let us suppose that we resample the fiber $\mathbb{X}(s)$ at N_u uniform points to form the vector $\bar{\mathbf{X}}$. We will typically use $N_u = 1/\epsilon = 1/\delta$, but more upsampling might be necessary to achieve higher accuracy. Thus the points we choose are at arclength coordinates $s = 0, \Delta s_u, \dots, L$, where $\Delta s_u = L/(N_u - 1)$ is the spacing. The corresponding weights are $w = \Delta s_u/2, \Delta s_u, \dots, \Delta s_u, \Delta s_u/2$, so that the first and last point have a weight of $1/2$, in accordance with the trapezoid rule. Then a simple algorithm to evaluate (9.17) is as follows:

1. Resample every fiber at N_u points.
2. Perform a (linear-cost) neighbor search to determine pairs of points for which the potential function $\partial \hat{\mathcal{E}}/\partial r$ is nonzero to within a certain tolerance, i.e., find pairs of points a distance r_{\max} or less apart.
3. For each pair of points with $\left\| \mathbf{X}_{\{k\}}^{(u)} - \mathbf{X}_{\{j\}}^{(u)} \right\| < r_{\max}$ (determined in step 2), compute the corresponding term in the sum (9.17) and add it to obtain the forces on the Chebyshev nodes.

This is obviously an extremely simple algorithm, but it can become costly as the fibers get more slender, since resolving all of the potential contacts requires a large number uniform points (and the quadratures are only second-order accurate, so there could be accuracy limitations as well). In the next section, we aim for an algorithm which is independent of ϵ .

9.2.2 Segment-based algorithm

The following algorithm theoretically achieves a cost that is linear in the number of fibers (at constant density) and independent of $\delta \sim a$. It is based on dividing the fiber into N_{seg}

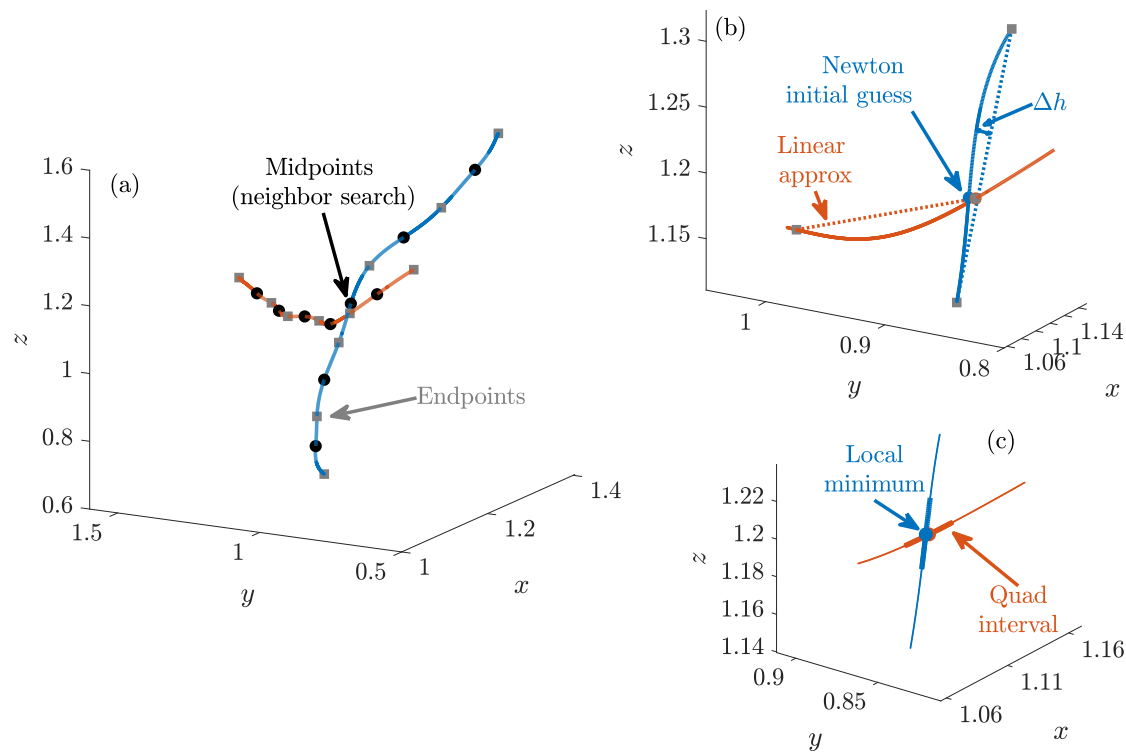


Figure 9.2: Schematic of segment-based steric interaction algorithm. (a) We divide the fiber into pieces and use the midpoint of each piece (black circles) to identify possible close segments. (b) We then approximate each piece using a straight segment (dashed lines), find the minimum distance between the two segments, and use that as an initial guess for Newton's method. (c) Newton's method finds the local minimizer of the distance on the two fibers, around which we draw quadrature intervals (thick lines).

(curved) fiber pieces of length L_{seg} , then approximating those pieces as straight segments for the purposes of a cheap distance estimation. If the segments are close enough, we perform a nonlinear Newton solve to obtain the closest points on the corresponding fibers. The number of fiber pieces in this algorithm is chosen to minimize computational cost, and should be independent of ϵ . A schematic of the algorithm is shown in Fig. 9.2, and a more detailed explanation of the steps is as follows.

1. We first sample each fiber at N_{seg} points, corresponding to the midpoints of the fiber pieces (circles in Fig. 9.2(a)). Then, we perform a neighbor search over all the midpoints using the cutoff distance $r_{\text{cut}} = L_{\text{seg}} + r_{\text{max}}$. This gives a list of potentially interacting segments.
2. For each pair of fiber pieces, we approximate curved fiber pieces as straight segments (dashed lines in Fig. 9.2(b)) and solve a quadratic equation to determine the closest point of interaction between the straight segments. Let us denote this distance $d_{\text{segs}}^{(ij)}$, with the corresponding closest points on the fibers given by $s_{\text{seg}}^{(i)}$ and $s_{\text{seg}}^{(j)}$.
3. We next determine if $d_{\text{segs}}^{(ij)} < r_{\text{max}} + \Delta h^{(i)} + \Delta h^{(j)}$, where $\Delta h^{(i)}$ represents the distance the fiber piece $\mathbb{X}^{(i)}(s)$ curves from the straight segment we use to approximate it (see Section 9.2.2.1). If $d_{\text{segs}} \geq r_{\text{max}} + \Delta h^{(i)} + \Delta h^{(j)}$, then exit, as the corresponding fiber pieces could never be close enough to interact.
4. We now proceed to the case when the fiber pieces could be close enough to interact. In this case, we solve the minimization problem

$$\begin{aligned}
 (s_*^{(i)}, s_*^{(j)}) = & \arg \min \quad \|\mathbb{X}^{(i)}(s^{(i)}) - \mathbb{X}^{(j)}(s^{(j)})\|^2, & (9.18) \\
 & 0 \leq s^{(i)} \leq L \\
 & 0 \leq s^{(j)} \leq L
 \end{aligned}$$

with Newton's method, with initial guess $s^{(i)} = s_{\text{seg}}^{(i)}$ and $s^{(j)} = s_{\text{seg}}^{(j)}$. This is technically

a constrained calculation, so we need to give a few details about how we actually calculate (9.18). Our approach is to first try to find a point where the gradient is zero

$$\begin{pmatrix} 2\mathbb{X}^{(i)}(s^{(i)}) \cdot \partial_s \mathbb{X}^{(i)}(s^{(i)}) - 2\partial_s \mathbb{X}^{(i)}(s^{(i)}) \cdot \mathbb{X}^{(j)}(s^{(j)}) \\ 2\mathbb{X}^{(j)}(s^{(j)}) \cdot \partial_s \mathbb{X}^{(j)}(s^{(j)}) - 2\partial_s \mathbb{X}^{(j)}(s^{(j)}) \cdot \mathbb{X}^{(i)}(s^{(i)}) \end{pmatrix} = \mathbf{0} \quad (9.19)$$

by running Newton's method with initial guess $s^{(i)} = s_{\text{seg}}^{(i)}$ and $s^{(j)} = s_{\text{seg}}^{(j)}$. If this converges to a tolerance 10^{-3} in 10 or less iterations, we return the interior points $s_*^{(i)}$ and $s_*^{(j)}$ (assuming the point passes the second derivative test for a local minimum). Otherwise (if Newton's method does not converge or if one of the minimizers is outside of $[0, L]$), we fix $s_*^{(i)} = 0$ and perform a Newton solve on the second term in (9.19). If this converges to a point $0 \leq s_*^{(j)} \leq L$, we add this point to a list of potential minima (if it passes the second derivative test). Changing the fixed point to $s_*^{(i)} = L$, then switching the roles of fibers i and j gives a total of four possible pairs. This list is then augmented with four additional combinations of the distances between the fiber endpoints, making for a total of eight possible pairs of points. We return the minimum of the eight distances, and the corresponding points $s_*^{(i)}$ and $s_*^{(j)}$.

5. We identify $\mathbf{d}_* = \mathbb{X}_*^{(i)} - \mathbb{X}_*^{(j)}$, where the notation $\mathbb{X}_*^{(i)}$ means $\mathbb{X}^{(i)}(s_*^{(i)})$. If the minimum distance $\|\mathbf{d}_*\|$ is less than r_{max} , we identify the domains on the fibers over which they could be interacting within distance r_{max} . We do this by approximating the fibers via a quadratic interpolant, $\mathbb{X}^{(i)}(s) \approx \mathbb{X}^{(i)}(s_*^{(i)}) + \partial_s \mathbb{X}^{(i)}(s_*^{(i)}) \Delta s^{(i)} + \partial_s^2 \mathbb{X}^{(i)}(s_*^{(i)}) (\Delta s^{(i)})^2 / 2$, so that the distance between the fibers is roughly (dropping

terms of order Δs^3 or larger),

$$\begin{aligned}\Delta d^2 (s^{(i)}, s^{(j)}) &\approx a (\Delta s^{(i)})^2 + b (\Delta s^{(j)})^2 + 2c\Delta s^{(i)} + 2d\Delta s^{(j)} + 2e\Delta s^{(i)}\Delta s^{(j)} + f \\ a &= (\partial_s \mathbb{X}_*^{(i)} \cdot \partial_s \mathbb{X}_*^{(i)}) + \mathbf{d}_* \cdot \partial_s^2 \mathbb{X}_*^{(i)}, \quad b = (\partial_s \mathbb{X}_*^{(j)} \cdot \partial_s \mathbb{X}_*^{(j)}) - \mathbf{d}_* \cdot \partial_s^2 \mathbb{X}_*^{(j)} \\ c &= \mathbf{d}_* \cdot \partial_s \mathbb{X}_*^{(i)}, \quad d = -\mathbf{d}_* \cdot \partial_s \mathbb{X}_*^{(j)}, \quad e = -\partial_s \mathbb{X}_*^{(i)} \cdot \partial_s \mathbb{X}_*^{(j)}, \quad f = \mathbf{d}_* \cdot \mathbf{d}_* - r_{\max}^2 < 0.\end{aligned}\tag{9.20}$$

The solution of $\Delta d^2 = 0$ for $\Delta s^{(i)}$ and $\Delta s^{(j)}$ is an ellipse. The curves that bound the ellipse are given by

$$\begin{aligned}\Delta s_{\pm}^{(j)} (\Delta s^{(i)}) &= \frac{1}{2b} \left(-2d - 2e\Delta s^{(i)} \pm \sqrt{(2d + 2e\Delta s^{(i)})^2 - 4(a(\Delta s^{(i)})^2 + 2c\Delta s^{(i)} + f)b} \right) \\ \Delta s_{\pm}^{(i)} (\Delta s^{(j)}) &= \frac{1}{2a} \left(-2c - 2e\Delta s^{(j)} \pm \sqrt{(2c + 2e\Delta s^{(j)})^2 - 4(b(\Delta s^{(j)})^2 + 2d\Delta s^{(j)} + f)a} \right)\end{aligned}$$

We are interested in the bounding box that we can draw around the ellipse in $(\Delta s^{(i)}, \Delta s^{(j)})$ space. These points are where the discriminants are zero, and are given by

$$\begin{aligned}(2d + 2e\Delta s^{(i)})^2 - 4(a(\Delta s^{(i)})^2 + 2c\Delta s^{(i)} + f)b &= 0 \\ \Delta s_{\pm}^{(i)} &= \frac{8cb - 8de \pm \sqrt{(8de - 8cb)^2 - 4(4e^2 - 4ab)(4d^2 - 4bf)}}{8e^2 - 8ab} \\ (2c + 2e\Delta s^{(j)})^2 - 4(b(\Delta s^{(j)})^2 + 2d\Delta s^{(j)} + f)a &= 0 \\ \Delta s_{\pm}^{(j)} &= \frac{8ad - 8ce \pm \sqrt{(8ce - 8ad)^2 - 4(4e^2 - 4ab)(4c^2 - 4af)}}{8e^2 - 8ab}\end{aligned}\tag{9.21}$$

We then set $\Delta s^{(i)} = \max(|\Delta s_+^{(i)}|, |\Delta s_-^{(i)}|)$, and likewise for $\Delta s^{(j)}$, then form an interval $S^{(i)} = \left(\max(0, s_*^{(i)} - \Delta s^{(i)}), \min(L, s_*^{(i)} + \Delta s^{(i)}) \right)$, and likewise for $S^{(j)}$. The pieces of the fiber $\mathbb{X}^{(i)}(S^{(i)})$ and $\mathbb{X}^{(j)}(S^{(j)})$ could interact sterically (see Fig. 9.2(c)).

6. Because there are multiple segments on each fiber, it is possible that the interval associated with a particular minimum might overlap with another interval for the same fiber pair. Because of that, we then assemble a list of fiber pairs i and j and intervals $S^{(i)}$ and $S^{(j)}$ over which we do the integral (9.17) *by taking a union of all intervals obtained from step 5.*
7. Finally, we put a grid of Gauss-Legendre points over each interval. Letting L_{int} be the length of an interval, the number of points we use is

$$N_{\text{GL}} = \left\lfloor N_{\delta} \frac{L_{\text{int}}}{\delta} \right\rfloor + 1, \quad (9.22)$$

where δ is the standard deviation of the Gaussian potential (9.15) and N_{δ} is a constant which gives the number of Gauss-Legendre points we use per standard deviation. We then apply the formula (9.17) over the two Gauss-Legendre grids. To implement this step efficiently, we precompute a maximum number of Gauss-Legendre points by setting $L_{\text{int}} = L_{\text{seg}}$ in (9.22). We then pretabulate all possible points and weights on $[-1, 1]$ for $N_{\text{GL}} = 1, \dots, N_{\text{GL}}^{(\text{max})}$. Then, for a given L_{int} , we compute N_{GL} using (9.22), and look up the points and weights in the table, scaling appropriately by L_{int} .

To make the cost of the algorithm formally independent of ϵ , the neighbor search in step 1 must be on a number of segments that is independent of ϵ . The cost of step 2 is then independent of ϵ because we are minimizing segment distances. Steps 3–5 are the most difficult, but have cost independent of ϵ because we use nonlinear solves to obtain the closest points. And because $r_{\text{max}} \sim a$, the grid in steps 5–7 will have a thickness that scales with ϵ , so the cost of the quadrature in step 7 is independent of ϵ .

Ultimately, the efficiency of the algorithm (and its favorability over the one presented in Section 9.2.1) comes down to the expense of the neighbor search vs. root finding and quadrature. If the root finding is expensive, then it makes sense to just use $1/\epsilon$ fiber pieces

(which become spheres), and then the algorithm just becomes equivalent to Section 9.2.1. If the root finding and quadrature are cheaper, then we can use less points in the neighbor search and this segment-based algorithm might be more efficient than the uniform-point algorithm.

9.2.2.1 Errors from straight segments

One difficult part of the segment-based algorithm is step 3, where we need to estimate the curvature of the fibers and use Newton's method to obtain the closest points on the two segments. This step is actually strictly necessary to preserve independence from ϵ . If we use segments (dashed lines in Fig. 9.2(b)), we will have some error in estimating the distance between curved fiber pieces. Since $r_{\max} \sim a$, this error will matter more when ϵ is smaller, and so there will be more erroneous pairs of fiber pieces as ϵ decreases (assuming the curvature of the fibers does not change with ϵ). Thus we need to eliminate errors by performing the nonlinear optimization (9.18).

This section deals with how we estimate the curvature of the fiber segments relative to their straight backbones. Figure 9.2(b) shows what we mean by this: the pieces of the fibers $\mathbb{X}(s)$ that we are concerned with are shown as solid lines, while the straight segment approximations are shown as dotted lines. We define a distance $\Delta h(s)$ that is simply the distance between the straight segment and the curved fiber. Our goal is to get an estimate for the maximum Δh along the segment. We expect this distance to depend on the number of segments we put in each persistence length, which is the ratio ℓ_p/L_{seg} .

There are two possible ways to do this. First, we can pre-tabulate, for each ℓ_p/L_{seg} , the expected deviation from straightness, and use that as Δh . This is a good idea for equilibrium simulations, but in reality other external forces could bend the fibers, so we should really do something with the actual fiber shape. The simplest thing we can do is sample the fiber at the segment midpoint, and then use this to estimate the deviation from straight. For fibers

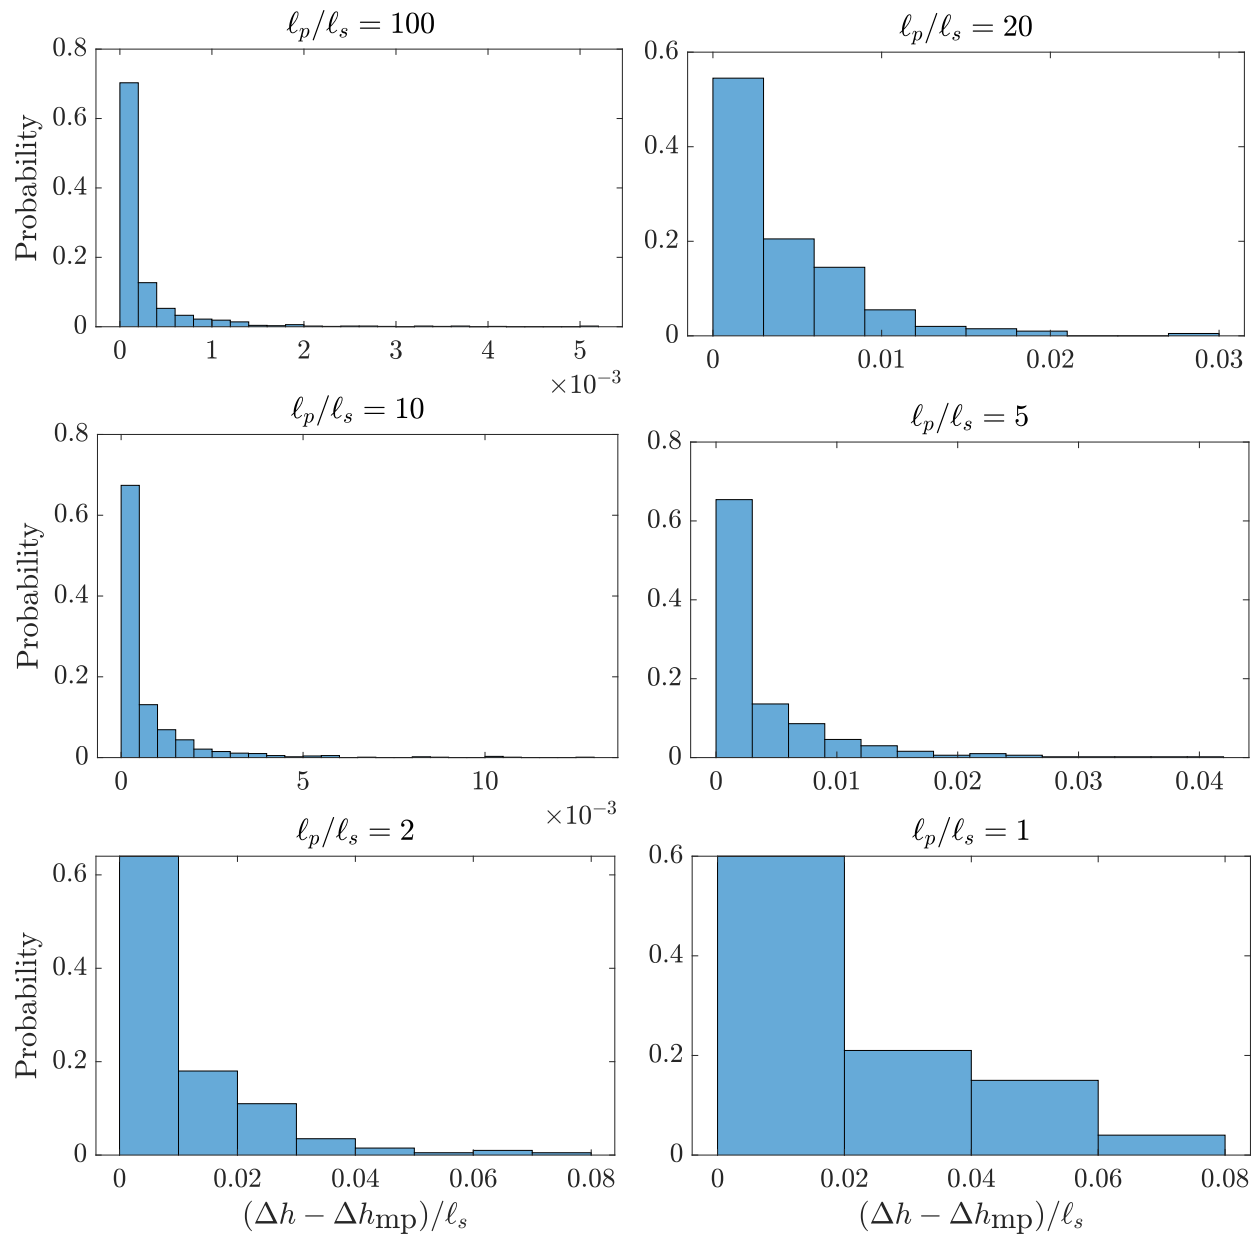


Figure 9.3: Error in estimating Δh using its value at the midpoint of the segment. When there are at least 2 panels per persistence length the error is typically 5% of the segment length or less.

sampled from the equilibrium distribution, Fig. 9.3 shows that we can use segments of size ℓ_p and still make an error of no more than 10% in doing this. Therefore, we just sample the segment at the midpoint, compute Δh_{mp} (at the fiber midpoint), and set $\Delta h = \Delta h_{\text{mp}}$ in step 3.

9.2.2.2 Examples of failure

The uniform point implementation and segment-based implementation should identify the same intervals on pairs of fibers, and consequently evaluate the same integrals. Thus, as the number of uniform points gets larger and N_δ in (9.22) gets larger, we should converge to the same forces on the Chebyshev points. While we cannot guarantee a certain accuracy in our segment-based algorithm, we can set a goal for *the segment-based algorithm to give the same accuracy as uniform points spaced $1/\epsilon$ apart* for “most” configurations.

To test this assumption, we consider configurations generated from simulations of fiber bundling with steric interactions (Section 12.6), in particular the bottom row of Fig. 12.13. For each configuration, we compute steric forces using $8/\epsilon$ uniform points (to obtain a reference force), then look at the errors with various N_δ and numbers of segments. The L^∞ errors in the steric forces are shown in Fig. 9.4.

Figure 9.4 shows that using $1/\epsilon$ uniform points gives 1–2 digits of accuracy in the forces. Typically, using $N_\delta = 1$ for the Gauss-Legendre intervals gives the same accuracy as $1/\epsilon$ uniform points, independent of how many segments we use. This is expected because the average point spacing ($\epsilon L = a = \delta$) is the same in both cases. Increasing to $N_\delta = 2$ gives an extra digit of accuracy, but at the cost of more expensive quadratures. There are, however, two cases where we see large errors in the segment-based implementation, which are indicated by arrows in Fig. 9.4. At $t \approx 6$ (black arrow), we have a case where the *quadrature* errors in the segment-based implementation are quite large (since the error is independent of the number of segments and reduced as we increase N_δ), while at $t \approx 7$ we have a case where




Figure 9.4: L^∞ errors in the steric forces, computed using the segment-based algorithm (colored lines) and the uniform point algorithm with $1/\epsilon$ points (black lines). In all cases, the reference solution uses $8/\epsilon$ uniform points. For the segments, we consider 10 (blue) and 20 segments (red), and the dotted lines show $N_\delta = 1$ in (9.22), while the solid lines show $N_\delta = 2$. There are two cases with high errors (indicated by the arrows).

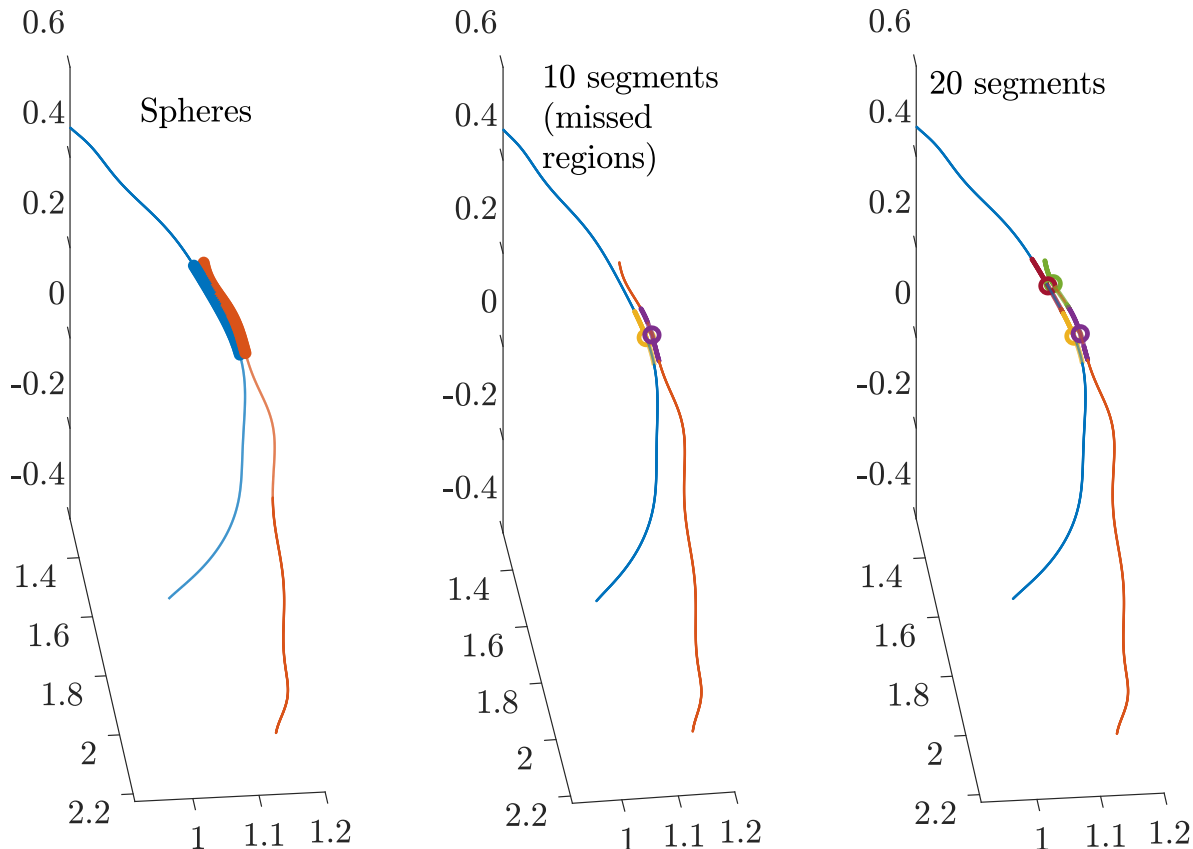


Figure 9.5: Analyzing the “problem case” where 10 segments are insufficient to give the correct steric force. The blue and red lines show the fiber backbones. Circles show local minima obtained from Newton solves, and the thicker lines show the corresponding intervals of interaction. In the left panel, we show the interval from uniform points (spheres) using thicker lines of the same color.

using 10 segments causes us to miss some contacts, and the error using 10 segments is much larger than when we use 20.

We first analyze the case when 10 segments are insufficient to give the correct force (red arrow in Fig. 9.4). Figure 9.5 shows the pair of fibers that causes the problem here. Over a small region, the fibers are nearly parallel and curve together. Because each fiber has its own distinct curvature, there are multiple local minima, and so we need a variety of initial guesses in the Newton solve to obtain the full interval where the fibers interact. For this reason, using 10 segments causes some of the interval to be missed (in particular, the top part of the interval). Using 20 segments gives more initial guesses for Newton, and consequently allows




Figure 9.6: Analyzing the “problem case” where $N_\delta = 1$ is insufficient. In this case, we have fibers which overlap at a 90 degree angle. This means the integrand is sharply peaked and requires more Gauss-Legendre points to resolve. Here the regions where the integrand is nonzero are shown using the thick lines.

us to find the entire interval of interaction. Note that the requirement for more segments here happens because of the fiber configuration, and not because of the fiber slenderness.

The last case we deal with is where the *quadrature* errors are large. This is the case at $t \approx 6$ (black arrow) in Fig. 9.4, where we see that the error is a strong function of N_δ and independent of the number of segments. The specific fibers that are causing the problem are isolated in Fig. 9.6, where the regions of overlap ($r \leq r_{\max}$) are shown using thick lines. We see that the fibers intersect each other at a 90 degree angle (distance between the lines $< 10^{-3}$ for both intersections), which means that the integrand is sharply peaked and decays rapidly. This is what causes the inaccuracies in the quadrature, and why we need larger N_δ to get a more accurate result.

Chapter 10

Conclusions

This part of the dissertation developed numerical methods to discretize the equations we derived in the previous part. Based on its length, the reader would be correct in concluding that our contributions in this area surpass those in the theoretical and biological parts. These contributions can be centered around one central idea: using spectral methods to reduce the number of degrees of freedom required in simulations.

For deterministic suspensions of filaments (Chapters 6 and 7), spectral methods reduce the number of points needed to obtain a given accuracy in the filament trajectories. While spectral methods have previously been used with SBT to model deterministic cytoskeletal filaments [15], we made a number of advances in the deterministic context that reduce the overall cost of simulations. For example, we formulated the inextensibility constraint on a discrete set of tangent vectors $\boldsymbol{\tau}$, which evolve via rotations on the unit sphere. We introduced constraint forces $\boldsymbol{\Lambda}$ and closed the system by requiring that the constraint forces perform no work. This formulation allowed us to preserve inextensibility discretely without penalty parameters, while at the same time derive a closed form evolution equation for the nodal points \boldsymbol{X} . The latter was vital when we formulated the overdamped Langevin equation for \boldsymbol{X} in Chapter 8.

Using the formulation in Chapter 2, we defined the hydrodynamic mobility operators

for translation-translation, translation-rotation, rotation-translation, and rotation-rotation coupling in terms of integrals of the corresponding Rotne-Prager-Yamakawa (RPY) kernels, then evaluated these integrals numerically using the “slender-body” quadrature schemes presented in Section 6.2. All of these schemes follow the same pattern: separate the integral into the part on $R > 2\hat{a}$ and $R \leq 2\hat{a}$, do the first integral via singularity subtraction (adding and subtracting the leading order *asymptotic* representation, which comes from Chapter 2), and the second via direct Gauss-Legendre quadrature (which is possible by resampling the global interpolant of the fiber centerline $\mathbb{X}(s)$ at the Gauss-Legendre points). For the trans-trans mobility, we showed that our quadrature schemes remove the spurious negative eigenvalues that result when the trans-trans integral is evaluated asymptotically as in (2.24). The only remaining uncontrolled approximation in our mobility evaluation is to use local drag, and not the full RPY integral, for the rot-rot mobility. We did this because the $\mathcal{O}(\epsilon^2)$ accuracy of this approximation was sufficient for our purposes, and because it makes the calculation of the rot-rot mobility \mathbf{M}_{rr} trivial. Nevertheless, we also developed a slender body quadrature for the rot-rot mobility that we compared to the asymptotic one in Fig. 6.7.

Our new quadrature schemes allowed us to take a detailed look into how the force and torque from twisting affect the overall dynamics of slender filaments. The most interesting example in this regard was a clamped spinning filament, which requires finite twist elasticity to drive the filament into an unstable trajectory for spinning frequencies $\omega > \omega_c$. We showed that our numerical method reproduces exactly the semi-numerical prediction of [99] for ω_c when using ellipsoidally-tapered filaments with local drag, and that the shape of the filament and the addition of nonlocal trans-trans hydrodynamics modify the critical frequency by at most 5% [54]. We showed that rot-trans coupling is more significant, lowering the frequency by another 15% to about 80% of the value predicted for ellipsoidally-tapered local drag, which explains about half of the deviation observed experimentally [102]. Above the critical

frequency ω_c , the fiber transitions to a steady crankshafting rotation in the shape of a shepherd's crook (overwhirling) [100]. While we found this state to be stable up to the highest frequencies we investigated, getting to that state from an initially straight filament requires complex dynamics that can lead to self-intersections, suggesting that in practice plectonemes may be observed instead [102].

While the overwhirling example is certainly more interesting than the relaxing bent filament we first considered, the result that twist elasticity has a negligible $\mathcal{O}(\epsilon^2)$ effect on the dynamics of free fibers allowed us to neglect it, and in Chapter 7 we developed specialized spatial and temporal schemes for deterministic fibers that resist bending and interact hydrodynamically through the viscous fluid. We reformulated the equations in terms of *force*, thereby deriving the symmetric positive definite matrix $\widetilde{\mathbf{M}}$ that relates fiber velocity to fiber forces in terms of the matrix \mathbf{M} which relates velocity to force density. The requirement that $\widetilde{\mathbf{M}}$ be SPD is important for thermal fluctuations, and as such we filtered the matrix we obtain from our slender body quadrature to symmetrize it and remove any negative eigenvalues. We use the special quadrature matrix for the self-mobility, while the nonlocal parts of the mobility (interactions between distinct fibers) are done by upsampling to a grid of size $N_u \sim 1$ (the SPD matrix (7.10)). This leads to an overall mobility which has cost independent of ϵ . Time-lagging the nonlocal parts of the mobility relative to the local parts then reduces the number of nonlocal hydrodynamic evaluations per time step to something that can be counted with the fingers on one hand.

In Chapter 8, we applied our spectral methods to fluctuating inextensible filaments in Stokes flow. To include fluctuations in the numerical method we already proposed, we postulated the Gibbs-Boltzmann distribution (8.1) on the spectral grid of N nodes, then constructed spatially-discrete overdamped Langevin dynamics that is in detailed balance with respect to that distribution. Unlike deterministic problems, it was not totally clear to us what to expect, since the fiber shapes are inherently nonsmooth when Brownian fluctuations

are included. We showed through a series of equilibrium and non-equilibrium tests that spectral methods can be advantageous in the regime where the fibers are semiflexible $\ell_p \gtrsim L$ and slender $\epsilon = a/L \ll 1$, which is therefore the regime where the persistence length (smallest lengthscale on which fluctuations are visible) is much larger than the fiber radius. Since our primary interest is in relatively stiff (and quite slender) fibers like actin, there is some promise for spectral methods to give good approximations of blob-link chains by faithfully modeling the hydrodynamics with fewer degrees of freedom. In this case, the fluctuations on lengthscales a , which are smoothed out by our spectral method, do not impact the fiber dynamics (other than on very short timescales) or equilibrium statistics, and so we are able to approximate both well using a small number of Chebyshev nodes. Our tests on dynamics showed that the error in the rate of relaxation of a straight chain for $N \geq 24$ is small and dominated by the error from *equilibrium* statistical mechanics, which means that the main driver of the difference between the blob-link and spectral methods comes from the difference between our coarse-grained Gibbs-Boltzmann distribution (8.1) and the GB distribution of a refined blob-link chain. Nevertheless, for $\ell_p/L = 1$ and $\epsilon \approx 10^{-2}$, we showed that even $N = 12$ provides a reasonable approximation to the relaxation dynamics of a stretched chain, with a time step size two orders of magnitude larger than the corresponding blob-link discretization with 100 blobs. This sort of spectral coarse-grained (fully discrete rather than continuum) representation of semiflexible fibers can allow for large-scale simulations of networks of fibers over physically relevant timescales, unlike fully or finely resolved blob-link models.

In Chapter 9, we designed a novel algorithm to treat steric (excluded volume) interactions in filaments discretized with Chebyshev collocation points. The key was to formulate the energy as a double integral, similar to what was done in [174] for straight cylinders. However, in [174], a singular energy density is chosen, which allows for analytic integration over straight segments. Because our simulations include Brownian fluctuations, fibers might

overlap regardless of the strength of the steric potential, and so we chose an energy density that is Gaussian (nonsingular and smooth) in nature, and consequently not possible to integrate analytically. As such, we needed to design efficient quadrature schemes to integrate the density in the limit $\epsilon \rightarrow 0$.

We considered two different quadrature schemes to integrate the steric interaction potential and obtain the steric forces at the Chebyshev collocation points. In the first scheme, we simply resample the fiber at a set of uniform points, then perform a neighbor search to find pairs of interactions that contribute to the sum. This scheme, which is similar to how we treat the cross-linking forces, does not scale well in the limit $\epsilon \rightarrow 0$, so we developed a second scheme based on breaking the fibers into $\mathcal{O}(1)$ pieces, then finding and progressively zeroing in on pieces that could be interacting.

10.1 Future directions

While this part of the dissertation was certainly exhaustive, there are still a number of future improvements awaiting attention for both single and multiple filament hydrodynamics. For single filaments, for instance, we chose *not* to regularize the shape of the filament, keeping the radius constant all the way up to the endpoints. Because we use integrals of the RPY tensor, the mobility for constant regularized radius \hat{a} is well-defined at the fiber endpoints, although its lack of smoothness causes slow convergence in our spectral method. This is in contrast to SBT, in which the local drag coefficient is not well-defined for cylindrical fibers, and some kind of ellipsoidally-tapered radius function has to be assumed. Future work could develop a hybrid approach that uses the RPY kernel for *unequally-sized spheres* [78] to simulate fibers with radius functions that decay smoothly to zero at the endpoints. This should smooth out the mobility, keep the eigenvalues of the trans-trans mobility away from zero, and make the local drag part of the single fiber mobility a better preconditioner in nonlocal simulations. The integrals involved are more complex than those for equally sized spheres

because the integrand includes the radius function $\hat{a}(s)$. Still, because the singularities (Stokeslet, doublet, rotlet, etc.) in the polydisperse integral kernels are unchanged, it should be possible to account for non-constant $\hat{a}(s)$ using the slender body quadrature schemes developed here.

Another choice we made for the single filament was to stick with a first kind integral equation for the force-velocity relationship, i.e., the trans-trans mobility. While this yielded sufficient accuracy for our purposes, the eigenvalues of the first kind mobility cluster around zero, leading to ill-defined solutions near the fiber endpoints. Of course, one argument around this is that both the ill-conditioning and negative eigenvalues of the SBT mobility occur when $N \gtrsim 1/\hat{\epsilon}$, which is too large for slender body approaches to be efficient. Still, a better-conditioned formulation (for trans-trans mobility) would be to use a second-kind integral equation of the form given in [95, Eq. (4)]. To enable rapid dynamic simulation of the model in [95], it is necessary to develop slender body quadratures to evaluate integrals of regularized Stokeslets and doublets along the centerline, as we did for the first kind RPY mobility in Section 6.2. The nearly-singular quadrature scheme of [140], which is based on finding the *complex* value of s' that makes the regularized Stokeslet $\mathbb{S}_\epsilon(\mathbb{X}(s) - \mathbb{X}(s'))$ singular, is perfectly tailored to this application. In fact, if the fiber is inextensible, the displacement $\|\mathbb{X}(s) - \mathbb{X}(s')\|$ is roughly constant in time when $s' \approx s$, so the complex roots can be precomputed for each s . If implemented, this would make the scheme of [95] comparable to the ones we use.

We had to take extreme care in designing the spectral discretization because of aliasing errors, which manifest themselves most especially in simulations with Brownian fluctuations. An open question is whether this level of scrupulousness, which is also necessary because Chebyshev polynomials cannot define everywhere inextensible curves, could have been avoided by using a different (rapidly converging) series representation for the fiber centerline. Other, non-Chebyshev, representations for (everywhere) inextensible filaments

like curvature-torsion require higher order derivatives of $\boldsymbol{\tau} = \partial_s \mathbf{X}$ which have non-decaying spectra, since for Brownian filaments $\boldsymbol{\tau}$ is almost everywhere differentiable but its derivative is white noise. An optimal representation would be in terms of some basis functions whose contributions to the fiber shape decay fast, especially for large persistence lengths. Such a representation would ideally nearly diagonalize the bending elasticity operator and be everywhere inextensible, the latter being vital to remove ambiguity in how to represent the fiber (tangent vectors vs. positions).

One of the reasons we were most impressed with the computational savings in our spectral method for fiber fluctuations is the correlation between the number of nodes in the spatial discretization and the required time step size. In particular, our midpoint temporal integrator, which is based on taking a Fixman-like predictor step to the midpoint to capture the drift term in expectation, can only generate accurate equilibrium statistics and dynamics when the time step is sufficiently small. The required time step size scales roughly as N^{-4} , which means we must hold the number of modes down if we want to simulate with a reasonable time step size. An alternative approach would be to try to increase the required time step size by using an exponential integrator. For rotations on the unit sphere, this requires Lie integrators like that used by Schoeller et al. in [47]. Stochastic exponential Lie integrators have not yet been developed (to our knowledge), and so they represent an interesting avenue of exploration for future work.

For suspensions of multiple filaments, there are two important future directions, one for deterministic filaments and one for fluctuating ones. For deterministic filaments, a future extension of this work is to consider the dynamics of multiple filaments in a suspension with twist. The mobility of Chapter 2 directly generalizes to multiple filaments, since the same RPY kernel can be used between distinct filaments and a filament and itself. The challenge lies in efficiently numerically computing the solution. Focusing on the static case first, we need to develop an efficient numerical method to sum the rot-trans and trans-rot

hydrodynamic interactions over many bodies, since Chapter 7 already treated the trans-trans kernel with a combination of the spectral Ewald method [103] (for well-separated fibers)¹ and corrections for nearly-touching fibers [140]. Extending our many-body-summation approach to twist is straightforward, but arduous. In particular, spectral Ewald methods have been developed for the rot-trans and trans-rot singularities of Stokes flow [141, 178], and the nearly-singular quadratures of [140] readily generalize to the other singularities. For rot-rot coupling, the fast decay of the doublet singularity makes it a good approximation to only include the flows generated by nearly contacting fibers.

Then there is the question of computing the dynamics of many twisting filaments. For trans-trans dynamics, we treated terms involving nonlocal hydrodynamic interactions explicitly, and only treated the local drag terms implicitly. This allowed us to solve essentially diagonal linear systems, and to evaluate long-ranged hydrodynamics only once or at most a few times per time step. A challenge for future exploration is how to efficiently handle pairs of filaments (or pieces of a single filament) in near contact in a semi-implicit temporal integrator, without expensive GMRES iterations at each time step.

Additionally, the mobility that we ultimately use for multiple *fluctuating* filaments (Section 8.6) does not achieve our goal of being independent of ϵ . Indeed, in this case we were unable to repeat our trick of separating the intra-fiber mobility from the inter-fiber mobility, and doing the former by special quadrature and the latter with $\mathcal{O}(1)$ points. The reason for this is because it is unclear how to take the square root of the sum of the two pieces. Future work should come up with a method for this, so that we can simulate multiple Brownian filaments interacting hydrodynamically with cost independent of ϵ .

We chose to treat steric forces explicitly in our temporal update. As shown in Section 12.6, this leads to a reduction of a factor of roughly 5 to 10 in the allowable time step

¹We chose an Ewald approach because we are interested in triply-periodic systems; in free space a fast multipole method [175] is likely the best choice, and systems with mixed periodicity are an avenue of active research [176, 177].

size, with our original time step leading to unstable simulations when sterics is included. It is tempting in this case to consider an implicit approach to achieve stability at larger time step sizes, such as the complementarity formulation used in [87, 179–181], whereby an optimization problem is solved to obtain a configuration that is contact-free while minimizing the contact forces between particles. While these methods do indeed increase the time step size, the method used to solve the optimization problem requires multiple applications of the hydrodynamic mobility, and as such makes each time step at least 10 times more expensive than when contacts are not resolved. Thus employing a more complex contact algorithm would make our simulations just as expensive as treating the steric force explicitly.

Part of the reason we can say this is that our potential is relatively “soft,” in the sense that we allow some particle overlaps to keep the time step relatively large. If we were to stiffen up the potential to prevent more overlaps, then we likely would arrive at a point where contact-based algorithms *would* be more efficient. It remains unclear, however, how to extend these algorithms from rigid particles [179–181] to flexible fibers, as the most recent works of Yan et al. [87, 180] treat flexible fibers via a series of rigid rods, which is not a model that scales well in the slender limit (in particular, it is not clear how to resolve hydrodynamics between the rods in the slender limit). Thus a numerical method to prevent any contacts between flexible Brownian filaments is still an open research direction.

Part III

Macroscopic behavior of cross-linked actin gels

Chapter 11

Introduction

In most cells, as much as 10% of all protein is actin [182]. The majority of actin is the F-actin cytoskeleton – a gel made of rapidly-turning-over (assembling and disassembling) actin fibers, which are dynamically cross linked by a vast host of actin binding proteins (cross linkers or CLs). The actin cytoskeleton is largely responsible for the cell’s shape, movements, division, and mechanical response to its external environment [11]. Thus, understanding the mechanical and rheological properties of dynamically cross-linked cytoskeletal networks is the first step to understanding the mechanical properties of the cell at large.

In this part of the dissertation, we will apply our theoretical and numerical methods from the previous two parts to examine the role of hydrodynamics and morphology on the macroscopic mechanical properties of cross-linked actin networks. We will divide our study into two parts: the dynamics of morphology changes (Chapter 12), and dynamic steady state rheology (Chapter 13). This introduction reviews the literature in these categories (including *in vitro* experiments and simulations), and discusses our program for what follows.

11.1 Rheology

Several *in vitro* experimental techniques, including active poking, parallel plate shearing, and embedded-microbead tracking, have characterized the mechano-rheological behavior of cross-linked actin gels [182]. Depending on the experimental conditions, these studies report viscoelastic moduli in a wide range, from 0.1 to hundreds of Pascals [118, 183–185]. The mixed viscoelastic mechanical response of the densely cross-linked gel is expected: simply speaking, at short time scales the network of elastic fibers and CLs can be considered permanently interconnected and deform elastically, while at long time scales dynamic CLs are expected to connect fibers only transiently, enabling the brief storage of elastic energy of deformations. This elastic energy is dissipated after the CLs detach, which causes effective viscous behavior on long timescales. Often, the measured elastic modulus for actin gels is about an order of magnitude greater than the viscous one [118, 183, 185]; however, some studies measure elastic and viscous moduli of similar magnitude [186], and in some cases the elastic modulus is smaller than the viscous one [187]. Both moduli are increasing functions of actin and CL concentrations [188, 189]; see especially [189] for master curves over a wide range of CL and actin concentrations.

To understand the intrinsic timescales in the gel, many experimental studies examine the dependence of the loss and storage moduli on the *frequency* of oscillatory shear deformation ω . All experimental studies agree that the behavior of transiently cross-linked networks on long timescales (low frequencies) is qualitatively viscous both *in vitro* [190, 191] and *in vivo* [186], with many studies reporting a decay rate of $\omega^{1/2}$ for both moduli at low frequencies in the absence of actin turnover [190, 192, 193]. As the frequency increases, some studies report a monotonically-increasing viscous modulus [189], while others display a curious local maximum and minimum in the viscous modulus data [192, 194, 195]. Proceeding to the high-frequency limit, there are again contradictions, as some studies report a viscous modulus that scales as $\omega^{0.75}$ [189] (consistent with the thermal fluctuations of a single semiflexible

filament [185, 196]), while others observe a viscous-fluid-like scaling of ω^1 [197, 198], and still others yield a scaling of $\omega^{1/2}$ [190, 191, 199], which is characteristic of the Rouse model of polymer physics (fluctuating beads joined by harmonic springs) [200]. Thus the short and intermediate timescale behavior, and its dependence on microscopic parameters, is still an open question, as is the exact meaning of “long” and “short” timescales (i.e., when these observed scalings begin to dominate).

Part of the reason for the variations in experimental data is that each actin binding protein produces a unique change in network morphology [201], which in turn uniquely affects the viscoelastic moduli. For compliant CLs such as filamin, Kasza et al. observe a strong frequency-dependence of $\omega^{0.7}$ for the viscous modulus and a weak power law scaling of the elastic modulus, with both taking values in the range 0.1 – 1 Pa [202]. As the CLs get shorter and stiffer, the frequency-dependence in the elastic [12, 187, 189, 202] and sometimes even viscous modulus disappears [202], which suggests a change in the network morphology with the type of cross-linker. It is shown in [190] that changes in the moduli with chemical cross linking are only relevant when the cross-linker length is shorter than the mesh size, because these kind of CLs group the fibers into bundles, thereby changing the macroscopic structure of the network. Indeed, changes in cell mechanical properties, from more viscous to more elastic, relate to corresponding cytoskeletal morphology changes, from a weakly cross-linked network of actin filaments to a network of clustered bundles [203].

Wachsstock et al. [184] study the relationship between mechanical response and morphology using actin filaments with α -actinin cross linkers. Their experiments and modeling show a transition from elastic, solid-like behavior to viscous, fluid-like behavior as the α -actinin and actin concentrations increase, which corresponds to the formation of bundles within the network. These bundles are no longer entangled in a complex network and are free to slide past each other, which leads to viscous behavior. This observation is compatible with the work of Tseng et al. [204], which shows that more homogeneous networks are more

elastic in nature, and suggests that the cell makes the strongest elastic structures by combining bundling and cross-linking proteins to form a cross-linked network of bundles [205]. It is therefore pertinent to study the dynamic formation of bundles in cross-linked actin networks.

11.2 Morphology

The formation of bundles has previously been observed in actin suspensions with CLs such as filamin [206, 207], scruin [188], and α -actinin [184, 208]. In all of these systems, increasing the concentration of the cross-linking protein progressively transitions the steady state from a homogeneous meshwork, where filaments are distributed isotropically, through a composite bundle state, where bundles are composed of only a few filaments, to the clustered bundle state, where bundles can be separated by distances as large as 100 μm [206, 208].

While it was observed over 30 years ago [207] that Brownian motion drives bundle formation, its precise mechanism for doing so remains unclear. For instance Hou et al. [207], speculated that rotational diffusion aids in bundling, as filaments that are linked at one location rotate until other locations can be linked together, resulting in a bundle. More recently, it was shown that bundling is most efficient in a fluid-like environment, where actin filaments can diffuse more readily [209, 210]. At minimum, these studies imply that bundling is more difficult without Brownian motion, but could actin filaments still arrange into bundles without it?

The importance of Brownian motion in bundling can be seen in experiments where filament length varies, or when polymerization and bundling are initiated simultaneously. In this case, shorter filaments form a more stable clustered bundle state [202, 210], with the shortest filaments organizing into spindle-type structures [211, 212]. In systems where polymerization and bundling happen simultaneously, it has been shown that the formation of the clustered bundle state can be prevented via an increase in the actin polymerization rate

[209]. Mean field theory and simulations show that the slow down in bundling at high polymerization rates could be driven by a combination of steric interactions and the Brownian motion of the fibers being constrained by cross linkers [210]. It remains unclear, however, to what extent the attenuation of bundling is driven by sterics vs. cross linking, and even if a composite bundle state can form if the length of the filaments is larger than the initial mesh size.

An underlying assumption in conceptual explanations of bundling is that sufficient CL is available to cross link filaments once they move closer together. The literature is conflicted, however, on exactly how much CL is sufficient. For instance, in the same experimental system of filamin and actin, some authors report a constant ratio of CL to actin necessary for bundling [184, 207], while others report that the relative amount of CL necessary for bundling decreases as actin concentration increases [206]. There is also a nontrivial effect of temperature on the amount of CL required for bundling; with higher CL-to-actin ratios bundling can occur at lower temperatures [213]. Experimental investigation of the precise amount of CL necessary for the clustered bundle state to form is difficult since the observation of bundles is a qualitative phenomenon with a subjective definition, and therefore varies based on the tools used.

11.3 Other modeling studies

In view of this experimental complexity, many modeling studies examine the changes in actin network morphology and mechanics as a function of the underlying kinetic and microscopic mechanical parameters. Some of the rheological behavior of pure actin networks can be explained with semiflexible polymer theory [214, 215]. For instance, strain-stiffening behavior in filament networks can be accounted for by an entropic model, in which each thermally fluctuating polymer resists affine (stretching) deformations, or an enthalpic model, in which the filaments bend before they stretch, allowing them to reorganize along the direction of

shear [216]. In a similar approach, Head et al. [13] solve an energy minimization problem for straight filaments in two dimensions and use their results to characterize the transition from affine (stretching-dominated) to non-affine (bending-dominated) deformation in cross-linked networks. Semiflexible filament theories of this nature have recently been extended to transiently cross-linked networks [195, 217, 218], and it has been shown that a broad spectrum of relaxation times appears for a cross linker with a single unbinding rate, because larger and larger wavelength bending modes are able to relax as cross linkers unbind [193, 217]. These theories may explain the soft glassy rheology observed in living cells [193] and the low-frequency regime where the moduli scale with the square root of frequency [195, 217].

Similarly, the bundling of actin filaments has been studied using equilibrium thermodynamics to find conditions at which the free energy, consisting of translational and rotational entropy of rod-like filaments and enthalpy and entropy of the CL distribution, is lower in the bundled state than in the unbundled mesh [219, 220]. The important results of these theories were that a critical CL concentration is needed for the bundling phase transition and that ultimately one giant bundle has to form, but transiently the filaments could be kinetically trapped in multiple bundles [219, 221]. However, actin bundling is not taking place in thermodynamic equilibrium, and a Brownian dynamics approach is therefore a useful complement.

In bundled networks, an early Brownian dynamics study [222] looked at the roles of translational and rotational diffusion in bundling of uniformly laterally attracting filaments. In a more detailed approach, Kim et al. [107] use Brownian dynamics simulations on a bead-spring model of actin to show that bundling is reduced when CLs are forced to bind in the perpendicular direction (as they do when fascin is the cross-linker). They propose that the ability of a network to flow on long timescales is due to the breaking of CLs with shear, followed by their reformation on a timescale determined by network reorganization and not individual CL binding. Last, but not least, a combination of scaling estimates and Brownian

dynamics simulations with simplified CL properties revealed multi-scale transitions from the isotropic to bundled phase [210]. Most of these previous studies focused on the actin network structure, rather than on the temporal evolution of the bundled state, which will be our focus here.

On the rheology side, two recent studies [166, 218] combine computational rheology with microscopic Brownian dynamics simulations to measure the rheological properties of transiently cross-linked actin networks *without* filament turnover and without hydrodynamic interactions between the filaments. In these studies, a local maximum in the viscous modulus is observed [194], which was shown to correspond to a maximum in CL binding and unbinding events [166]. However, the networks considered in these studies are either homogeneous [166, Fig. 2] or highly bundled [218, Fig. 1], and we know of no detailed simulations that have reported the rheological properties over a range of structures and microscopic parameters. Furthermore, due to the absence of turnover, in both of these studies the structure being simulated is not in equilibrium and evolves significantly (by bundling) with the imposed flow, and so the measured low-frequency viscous and elastic moduli depend on history and how the measurement is performed, and are consequently not well-defined.

Because of the complexity of microscopic simulations, there has been a considerable amount of literature in the last few decades on *continuum* models of the cell cytoskeleton and cross-linked actin networks. For example, Karcher et al. employ a continuum finite element model to measure the stress induced by magnetic-bead-forcing [223] (a similar model was also used in [224]). They model the cytoskeleton as either a Maxwell (spring and dashpot in series) or Kelvin-Voigt (spring and dashpot in parallel) material and report a large sensitivity of the mechanical behavior on the choice of model. Other studies have approximated the cell cytoskeleton as a poroelastic [225], porous viscoelastic [226], or Brinkmann [227] fluid. In either case, as discussed in [182, 223], continuum models are only as good as their fits to the data, and they in general have difficulty relating macroscopic parameters (such as

porosities, discrete timescales, and stiffnesses) to microscopic parameters. In fact, it has been suggested that the cell possesses a continuum of relaxation timescales [193, 228], which would invalidate any continuum model with a discrete number of elements. To us, a happy medium is to use a microscopic model to inform a continuum fluid model which can be used in larger-scale simulations, as we attempt to do in Chapter 13.

11.4 Program

To interpret the experimental data from these previous studies, it is vital to understand the dynamics of the bundling process and how the resulting morphology feeds back onto the mechanical properties of the gel as a whole. We do this in two steps. In Chapter 12, we study the bundling process and how each of the microscopic parameters influences the overall bundling timescale. We conclude the chapter by showing that the dynamic steady state comes from a competition between bundling and fiber turnover. Then, in Chapter 13, we perform an analysis of the rheology of two representative steady states, showing how morphology and hydrodynamics contribute to the overall viscoelastic behavior.

11.4.1 Network statistics

In order to describe the evolution of the actin network, we need to introduce some statistics to quantify the “look” of the gel (how clustered it is). We examine the connectivity of the fibers in two ways. First, given the total number of CLs in the system C , we compute an average link density per fiber via the formula “Link density” = $2C/(LF)$. Second, we map the network to a connected graph to study how the structure evolves in time [211]. We define a “bundle” as a connected group of at least $F_B = 2$ filaments, where a connection between two fibers is a pair of links with anchoring locations at least $d_{\text{bund}} = L/4$ apart on each fiber, so that the links limit the fibers’ rotational degrees of freedom. We then define two measures

of the degree of bundling in the system. The first measure is the bundle density, which is the number of bundles per unit volume B/L_d^3 , where B is the number of bundles and L_d is the length of the simulation cell. The second measure is the percentage of fibers in bundles, defined as the percentage of filaments connected to at least one other filament by two links at least $d_{\text{bund}} = L/4$ apart. The bundle density statistic preferentially weights smaller bundles, since a bundle of two filaments is counted the same as a bundle of five filaments, while the percentage of fibers in bundles is independent of F_B . In fact, we can think of a bundle of F_B filaments as contributing a weight $\sim 1/F_B$ to the bundle density, but a weight $\sim F_B$ to the percentage of fibers in bundles.

For a bundle of b filaments, we define an orientation parameter as the maximum eigenvalue of the matrix [211]

$$\mathbf{Q} = \frac{1}{bL} \sum_{i=1}^b \int \boldsymbol{\tau}^{(i)}(s) (\boldsymbol{\tau}^{(i)}(s))^T ds. \quad (11.1)$$

The orientation parameter takes values in $[1/3, 1]$, with 1 being the value for a group of straight fibers with the same tangent vector. Given information about the bundles, we compute an average bundle orientation parameter by taking an average over bundles with at least two filaments, weighted by the number of filaments in each bundle.

Throughout this paper, we will quantify the concentration of fibers in terms of the initial mesh size [196] of the suspension, $\ell_m = \sqrt{L_d^3/(FL)}$ (parameters are defined in Table 11.1). Note that this estimate for ℓ_m applies to non-bundled (disordered) suspensions of fibers, so really when we use ℓ_m we mean the *initial* mesh size, prior to the bundling process beginning. We will operate in the regime where the fluctuations in the CL rest length as defined in (9.2), which are of magnitude $\delta\ell = 20$ nm (see parameters in Table 11.1), are several times smaller than the typical filament spacing, which is at most the initial mesh size $\ell_m = \mathcal{O}(100)$ nm and at least the cross linker length of 50 nm.

Parameter	Definition	Value	Unit	Notes
a	Fiber radius	4	nm	[117]
L	Fiber length	1	μm	[118, 119]
F	Number of fibers	200 – 1600		$\ell_m = 0.2$ initially
L_d	Simulated volume's extent	2 – 4	μm	Cubic unit cell
μ	Cytoplasm viscosity	0.1	Pa·s	Cytoplasm [229]
κ	Fiber bending stiffness	0.07	$\text{pN}\cdot\mu\text{m}^2$	$\ell_p = 17 \mu\text{m}$ [18]
K_c	CL spring stiffness	10	$\text{pN}/\mu\text{m}$	[165]
ℓ	CL rest length	50	nm	[230]
k_{on}	CL first end binding rate	5	$1/(\mu\text{m}\cdot\text{s})$	
$k_{\text{on},s}$	CL second end binding rate	50	$1/(\mu\text{m}\cdot\text{s})$	$k_{\text{on},s} \gg k_{\text{on}}$,
k_{off}	CL (one end) unbinding rate	1	1/s	[231, 232]
$k_{\text{off},s}$	CL (second end) unbinding rate	1	1/s	$k_{\text{off},s} = k_{\text{off}}$
c_w	Actin binding site width	20	nm	[230]
$k_B T$	Thermal energy	4.1×10^{-3}	$\text{pN}\cdot\mu\text{m}$	25°C
N_x	Number of Chebyshev points	12–16		
Δs_u	Binding site spacing	0.026	μm	
Δt	Time step size	10^{-4}	s	Limited by K_c

Table 11.1: Simulation parameters for transiently cross-linked actin networks.

11.4.2 Simulation parameters

In the case when we want to study the effect of the parameter, we pick values higher or lower than the “base” parameter value. The base values for every parameter we consider are given in Table 11.1. Most of the parameters here can be measured experimentally, with the notable exception of k_{on} , the rate (per length) at which a CL binds to a filament. In experimental studies, what is measured for k_{on} is a *macroscopic* binding rate of the order $1 \mu\text{M}^{-1} \text{s}^{-1}$, which cannot easily be converted to the effective microscopic rate (units $1/(\text{length}\times\text{time})$) that we need. To find a suitable on rate, we make the observation from [184] that 10 – 50% of a $1 \mu\text{m}$ long filament is decorated with α -actinin [184], which corresponds to about 1 – 10 CLs per filament of length $1 \mu\text{m}$; we find that $k_{\text{on}} = 5/(\mu\text{m}\times\text{s})$ is of the correct order of magnitude to give about 5 CL ends attached to each filament. Likewise for $k_{\text{on},s}$, we simply assume that $k_{\text{on},s} \gg k_{\text{on}}$.

When we simulate bundling, we limit the number of CLs that can be bound to one site.

We introduce a CL width $c_w = 20$ nm [230] and set the maximum number of bound CLs at each site to $\lceil \Delta s_u / c_w \rceil$. We implement this in the stochastic simulation algorithm using rejection: if a binding event is selected and the binding site is full, we simply move on to the next possible event.

Our simulations use a background fluid with viscosity $\mu = 0.1$ Pa·s, mimicking the larger viscosity of the cytoplasm [229]. We will, however, assume the background fluid is Newtonian, in contrast to actual cytoplasm, so that we can isolate the contribution of the fibers to the viscoelastic moduli.

11.4.3 Calculating stress and SAOS rheology

When we perform simulations to measure the mechanical properties of cross-linked actin networks, we need to measure the stress in the simulation cell. This section describes how to do this for *deterministic* fiber networks. We do not make stress-strain measurements in Brownian fiber networks, as more theoretical work is necessary to determine the correct formula for the Brownian stress, as we discuss in the conclusions.

Recalling the time-oscillatory shear flow (7.50),

$$\mathbf{u}_0 = \dot{\gamma}_0 \cos(\omega t) (y, 0, 0), \quad (11.2)$$

the only nonzero component of the rate of strain tensor is constant in space and is given by

$$\dot{\gamma}_{21}(t) = \frac{\partial u_0^x}{\partial y} = \dot{\gamma}_0 \cos(\omega t), \quad (11.3)$$

and the relevant component of the *strain* tensor is therefore

$$\gamma_{21}(t) = \int_0^t \dot{\gamma}_{21}(t') dt' = \frac{\dot{\gamma}_0}{\omega} \sin(\omega t) := \gamma_0 \sin(\omega t), \quad (11.4)$$

where we have defined γ_0 as the maximum strain in the system. We define the bulk elastic (G') and viscous modulus (G'') from γ_0 by [233]

$$\frac{\sigma_{21}}{\gamma_0} = G' \sin(\omega t) + G'' \cos(\omega t). \quad (11.5)$$

Notice that the elastic modulus G' gives the part of the stress that is in phase with the strain, and the viscous modulus G'' gives the part of the stress that is in phase with the rate of strain. This analysis applies when the stress response is linear, so that higher frequency modes do not enter, hence we must keep the strain amplitude γ_0 small (Section 13.2.4 has a quantitative analysis of how small this needs to be). This gives rise to the name small amplitude oscillatory shear (SAOS) rheology, by which this type of experiment is known.

The stress tensor itself can be decomposed into a part coming from the background fluid and a part coming from the internal fiber stresses,

$$\sigma_{21} = \sigma_{21}^{(\mu)} + \sigma_{21}^{(f)} = \mu \frac{\partial u_0^x}{\partial y} + \sigma_{21}^{(f)}. \quad (11.6)$$

For pure viscous fluid, the fiber contribution to the stress is zero and the stress is given entirely by the viscous stress tensor,

$$\frac{\sigma_{21}^{(\mu)}}{\gamma_0} = \frac{\omega\mu}{\dot{\gamma}_0} (\dot{\gamma}_0 \cos(\omega t)) = \omega\mu \cos(\omega t). \quad (11.7)$$

Thus the viscous modulus due to the fluid is $G'' = \omega\mu$, and the elastic modulus is zero.

The stress due to the fibers depends on the force the fibers exert on the fluid. Because the CLs exert no torque or force on the fluid, we use Batchelor's formula [23] for the volume-averaged stress due to the fibers and CLs. In the slender limit (i.e., the case when the surface area force density is constant on fiber cross sections), the bulk stress due to the

fibers and CLs in a volume V is given by $\boldsymbol{\sigma}^{(f)} = \boldsymbol{\sigma}^{(i)} + \boldsymbol{\sigma}^{(\text{CL})}$, where

$$\boldsymbol{\sigma}^{(i)} = -\frac{1}{V} \left(\sum_{i=1}^F \sum_{p=1}^N \left(\mathbf{X}_{\{p\}}^{(i)} \left(\boldsymbol{\Lambda}_{\{p\}}^{(i)} + \left(\mathbf{F}^{(\kappa)} \right)_{\{p\}}^{(i)} \right) \right) \right), \quad (11.8)$$

$$\boldsymbol{\sigma}^{(\text{CL})} = -\frac{1}{V} \left(\sum_{c=1}^C \sum_{p=1}^N \left(\mathbf{X}_{\{p\}}^{(c_1)} \left(\mathbf{F}_{\text{ext}} \right)_{\{p\}}^{(c_1)} + \mathbf{X}_{\{p\}}^{(c_2)} \left(\mathbf{F}_{\text{ext}} \right)_{\{p\}}^{(c_2)} \right) \right), \quad (11.9)$$

where the double sum for $\boldsymbol{\sigma}^{(i)}$ is over points on fibers and the double sum for $\boldsymbol{\sigma}^{(\text{CL})}$ is over points on pairs of cross-linked fibers. In the cross-linker stress, C is the number of CLs and the notation c_1 and c_2 means that cross-linker c links fibers c_1 and c_2 . The sums must be separated since different periodic images of $\mathbf{X}^{(i)}$ could be involved for different cross linkers. In $\boldsymbol{\sigma}^{(\text{CL})}$, the positions $\mathbf{X}^{(c_1)}$ and $\mathbf{X}^{(c_2)}$ must be the periodic images of the two fibers that are connected by the CL. In this case, the total torque exerted by the cross-linker on the pair of fibers it connects is zero, and we expect the stress tensor (11.8) to be symmetric, since the constraint force $\boldsymbol{\Lambda}$ and elastic forces $\mathbf{F}\mathbf{X}$ exert exactly zero torque discretely [23].¹ Note that the formulas (11.8) and (11.9) are traditionally represented as integrals of the positions and force densities. Here we have replaced those by sums of forces, which gives the same result since $\mathbf{F} = \widetilde{\mathbf{W}}\mathbf{f}$, and $\widetilde{\mathbf{W}}$ is the matrix that appears when computing the integrals. Because we are not considering Brownian fluctuations in these rheology simulations, these are the only forces we need to consider (no additional random terms).

We discretize the stress tensor to first order accuracy in time. Specifically, we substitute \mathbf{X}^n for \mathbf{X} and $\boldsymbol{\Lambda}^n$ for $\boldsymbol{\Lambda}$ in the internal fiber stress $\boldsymbol{\sigma}^{(i)}$, and evaluate the cross-linking forces at \mathbf{X}^n in $\boldsymbol{\sigma}^{(\text{CL})}$. Assuming that the final time T is an integer multiple of the period $2\pi/\omega$, we then compute the bulk moduli by discretizing with the first-order quadrature the integrals

$$G' = \frac{2}{\gamma_0 T} \int_0^T \sigma_{21} \sin(\omega t) dt \quad G'' = \frac{2}{\gamma_0 T} \int_0^T \sigma_{21} \cos(\omega t) dt. \quad (11.10)$$

¹This is only true if we use the energy-based formulation of elastic force in Section 6.1.3.1, and if the positions \mathbf{X} in (11.8) are the same as those used to form the matrix $\mathbf{K}[\mathbf{X}]$ that enforces $\mathbf{K}^T \boldsymbol{\Lambda}$, a condition which ensures that $\boldsymbol{\Lambda}$ exerts no torque on the positions \mathbf{X} .

Chapter 12

Bundling dynamics

In this chapter, we use our agent-based framework to quantify how a homogeneous meshwork of filaments can evolve into a state of composite, and then clustered, bundles through Brownian motion. We introduce a timescale, τ_c , that quantifies the time to reach the composite bundle state, and show that the dynamics on shorter, but not so much on longer, timescales are accelerated by Brownian motion of the filaments. We also show that the bundling timescale is limited by filament diffusion for smaller CL concentrations, while for larger CL concentrations this diffusion has a minor effect. We find that the relative CL-to-actin ratio required to achieve the same bundling time decreases with increasing actin network mesh size. We also show that the diffusion effect explains faster bundling for shorter filaments, and that including steric interactions slows down the bundling process, but does not change its fundamental trajectory.

Throughout this chapter, we will only consider hydrodynamic interactions on one filament at a time. That is, there are *no* hydrodynamic interactions between different fibers in these simulations. We will consider hydrodynamic interactions between fibers in Chapter 13.

12.1 Non-Brownian filaments

We begin this section by discussing the kinetics of bundling for *non-Brownian, semiflexible* fibers, and establish that semiflexible fibers with a persistence length similar to that of actin [18] can be well-approximated by rigid filaments. This simulations are done without thermal fluctuations; see Section 12.5 for the corresponding results with fluctuations.

We begin with simulations in a system of initial mesh size $\ell_m = 0.2 \mu\text{m}$, which translates to $F = 200$ filaments in an $L_d = 2 \mu\text{m}$ domain, $F = 675$ filaments in an $L_d = 3 \mu\text{m}$ domain, and $F = 1600$ filaments in an $L_d = 4 \mu\text{m}$ domain. The mesh size we use is of the same order of magnitude as that in cell cortex in vivo [234] and corresponds to 10 to 15 μM G-actin concentration often used in in vitro experiments [184, 208]. In Fig. 12.2, we show that the statistics of the bundling process are insensitive to the domain size up to the point where there is mass coalescence of almost all the fibers in the simulation cell. For this reason, we will consider results from only one set of simulations, the one with $F = 675$ filaments and $L_d = 3$.

We initialize the set of F filaments with random locations and orientations, then during each time step we evolve the fibers by updating the dynamic CLs and then updating the fiber positions in sequential order. Figure 12.1 shows how the bundling process evolves in small and large systems. On the microscopic scale, filaments that are initially not parallel are linked by CLs, which pull them closer together and allow more links to bind. The binding of additional links leads to the alignment of filaments. Note that the key to the bundling process is the flexibility of the CL, in particular rapid thermal fluctuation of the CL length, which is present implicitly in our model from (9.2). Because the CLs are small, fluctuations in their length occur on a time scale which is much faster than other characteristic time scales, and so we do not model the fluctuations explicitly. The combination of the CLs' elasticity and length fluctuations is crucial, as the length fluctuations effectively allow the CLs to “find” the neighboring fibers and bind them, whereupon the elasticity of the CL

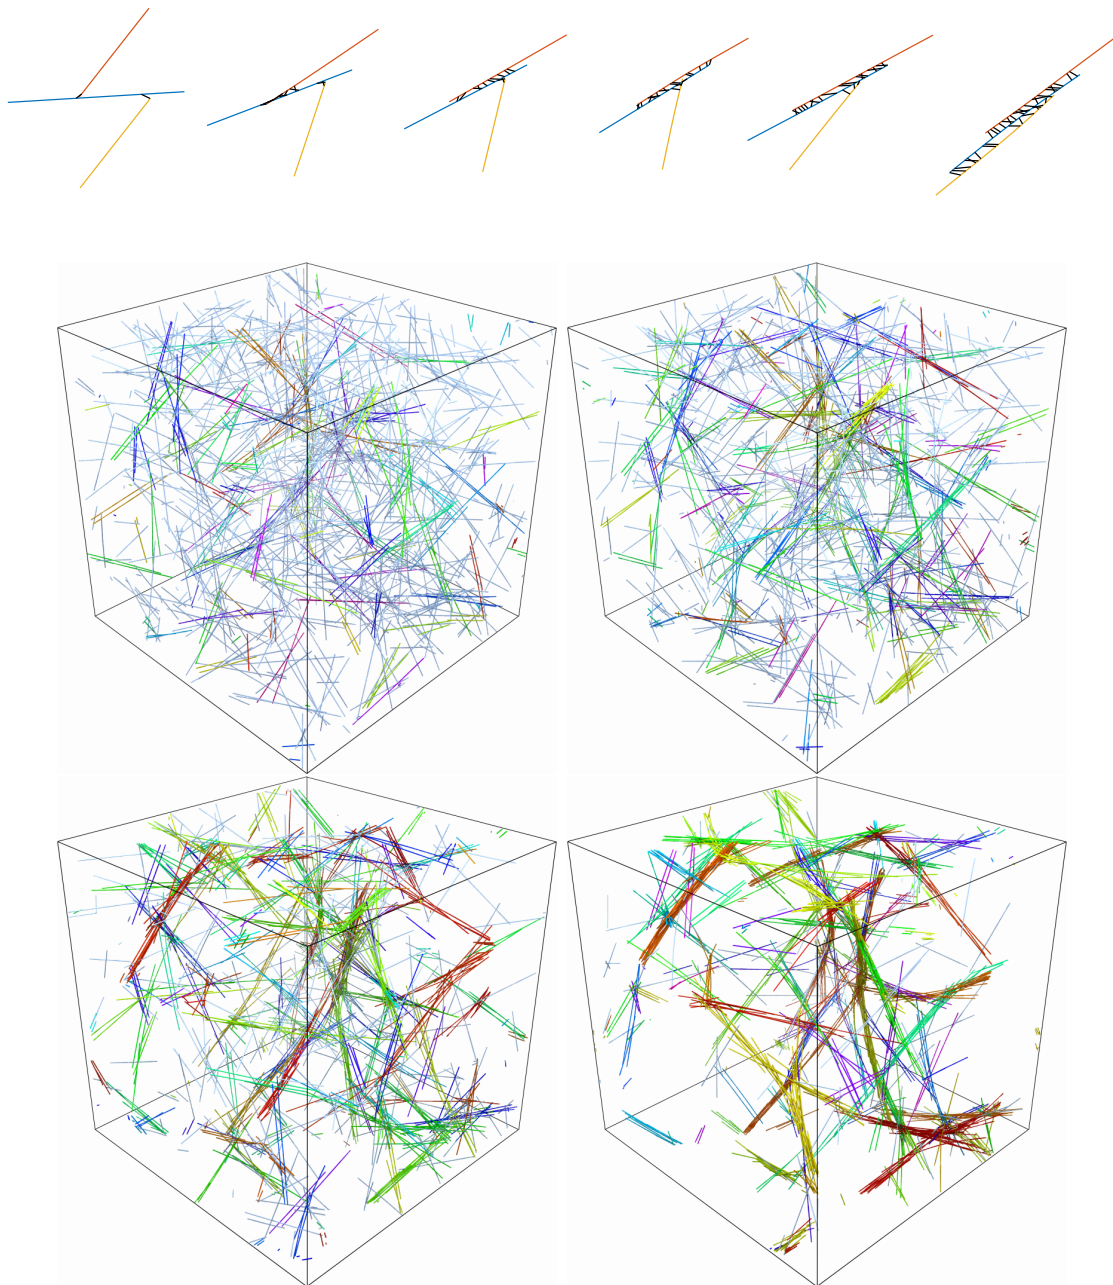


Figure 12.1: Bundling dynamics on small and large scales. Top: a small-scale bundling process with three filaments and snapshots taken at times $t = 0, 2, 4, 6, 8,$ and 10 s. Bottom: Snapshots of the bundling process taken (from left to right and top down) at $t = 5, 10, 20,$ and 40 seconds for semiflexible fibers with stiffness $\kappa = 0.07 \text{ pN} \cdot \mu\text{m}^2$. Fibers in the same bundle are colored with the same color. The two networks at the middle are before the coalescence transition time $\tau_c \approx 16$ s, while the two networks at the bottom are after the coalescence time.

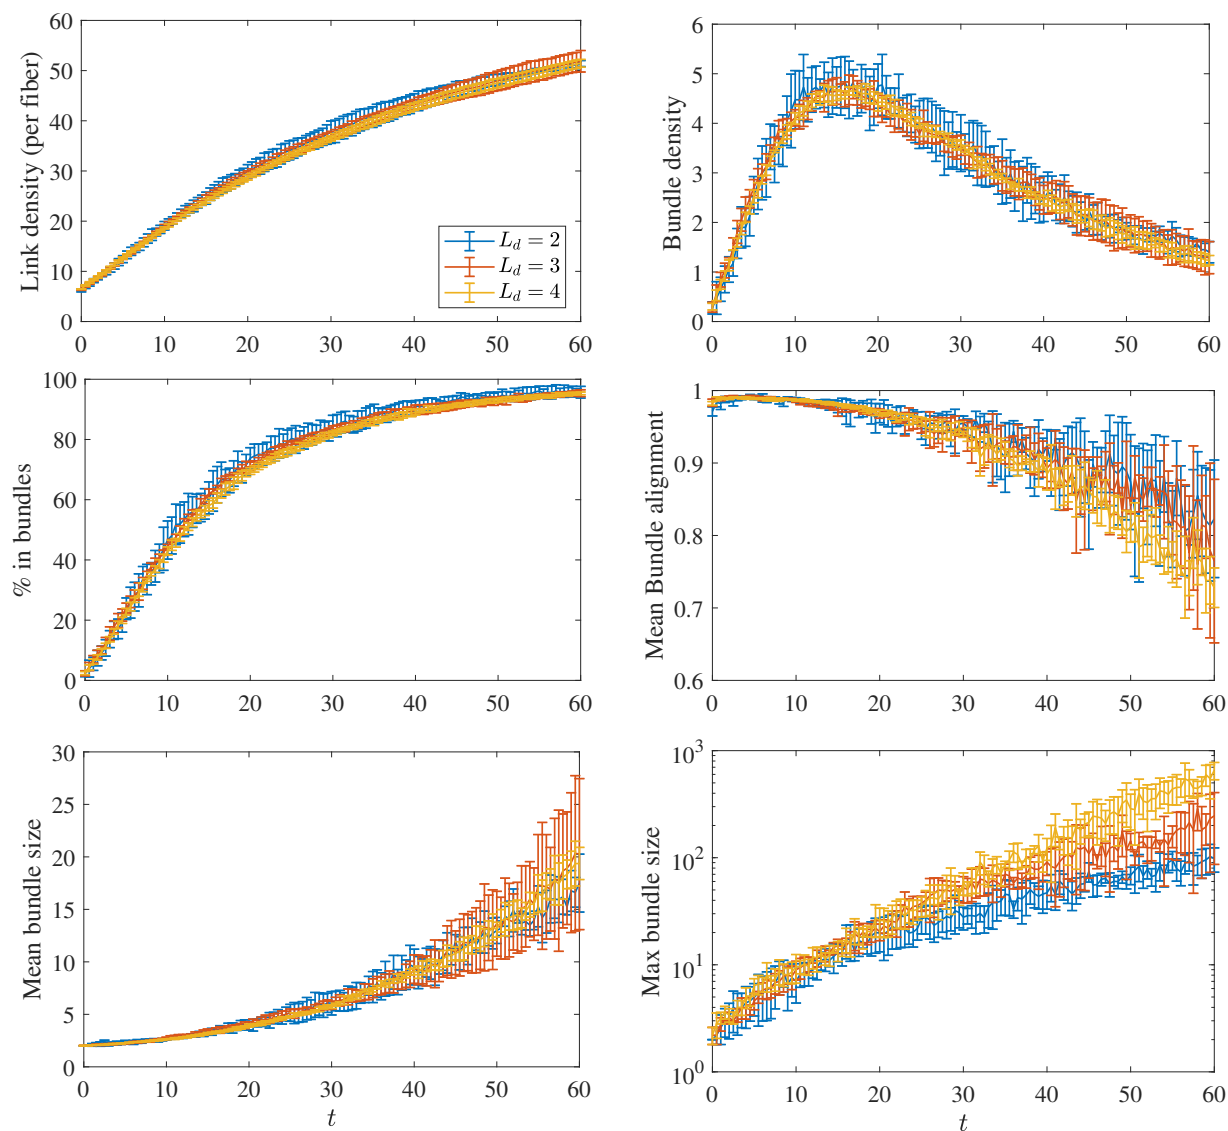


Figure 12.2: Link density, bundle density, % of fibers in bundles, mean bundle alignment, mean bundle size, and maximum bundle size, in a network of semiflexible ($\kappa = 0.07 \text{ pN}\cdot\mu\text{m}^2$) non-Brownian filaments with initial mesh size $\ell_m = 0.2 \mu\text{m}$. We show curves with different domain sizes (in μm) to establish that the statistics are repeatable in larger systems. The only statistic which is not repeatable is the maximum bundle size after many of the filaments have collapsed into one bundle ($t \gtrsim 40 \text{ s}$).

aligns the fibers, making further cross linking faster.

This process plays out on a larger scale in snapshots from the simulations, shown in Fig. 12.1 at $t = 5, 10, 20$, and 40 seconds. The initial stage of bundling (first two snapshots) is characterized by bundles of a few straight, aligned filaments, which is similar to the experimentally-observed composite bundle state [203] and the three-filament bundle shown at the top of Fig. 12.1. Later times (bottom two frames) show coalescence of these smaller bundles into larger bundles, with some curvature appearing in the fibers in the final frame. By $t = 40$ seconds, there are only a few bundles made of coalesced smaller bundles, and the network resembles the experimentally-observed clustered bundle state [203], which approaches the low energy state consisting of a single aligned bundle [218].

To quantify our observations, in Fig. 12.3 we plot the mean link density ($2C/(LF)$, see Section 11.4.1), bundle density (B/L_d^3), percentage of fibers in bundles, mean bundle alignment parameter, and mean and maximum bundle size throughout the bundling process for three values of fiber stiffness: $\kappa = 0.07$ (the value for actin), $\kappa = 0.007$ (fibers ten-fold less stiff), and $\kappa = \infty$ (rigid filaments). In all systems, we see the number of links per fiber grow in time to approach the maximum of $[\Delta s_u/c_w] \times L/\Delta s_u = 80$, while the bundle density in all systems exhibits a peak around a critical time $\tau_c \approx 16$ seconds. At this time, the other panels of Fig. 12.3 tell us that 60% of the fibers are already in bundles, which have a mean alignment parameter larger than 0.9. Figure 12.4 gives a more precise look at the composition of the bundles, which are the same for the three values of stiffness when $t \leq \tau_c$: at $t = \tau_c$, most ($> 50\%$) of the fibers are in bundles of size 11 or less, with a small percentage ($< 10\%$) in bundles of size 10 – 20, and the other 40% of the filaments not in bundles at all. Thus, a time τ_c into the bundling process, most of the fibers are in small, highly aligned bundles, as we see in the snapshots in Fig. 12.1, and the dynamics up to this point are roughly independent of the fiber stiffness. Based on Fig. 12.1, we can also think of τ_c as the time required to reach the composite bundle state. For this system of

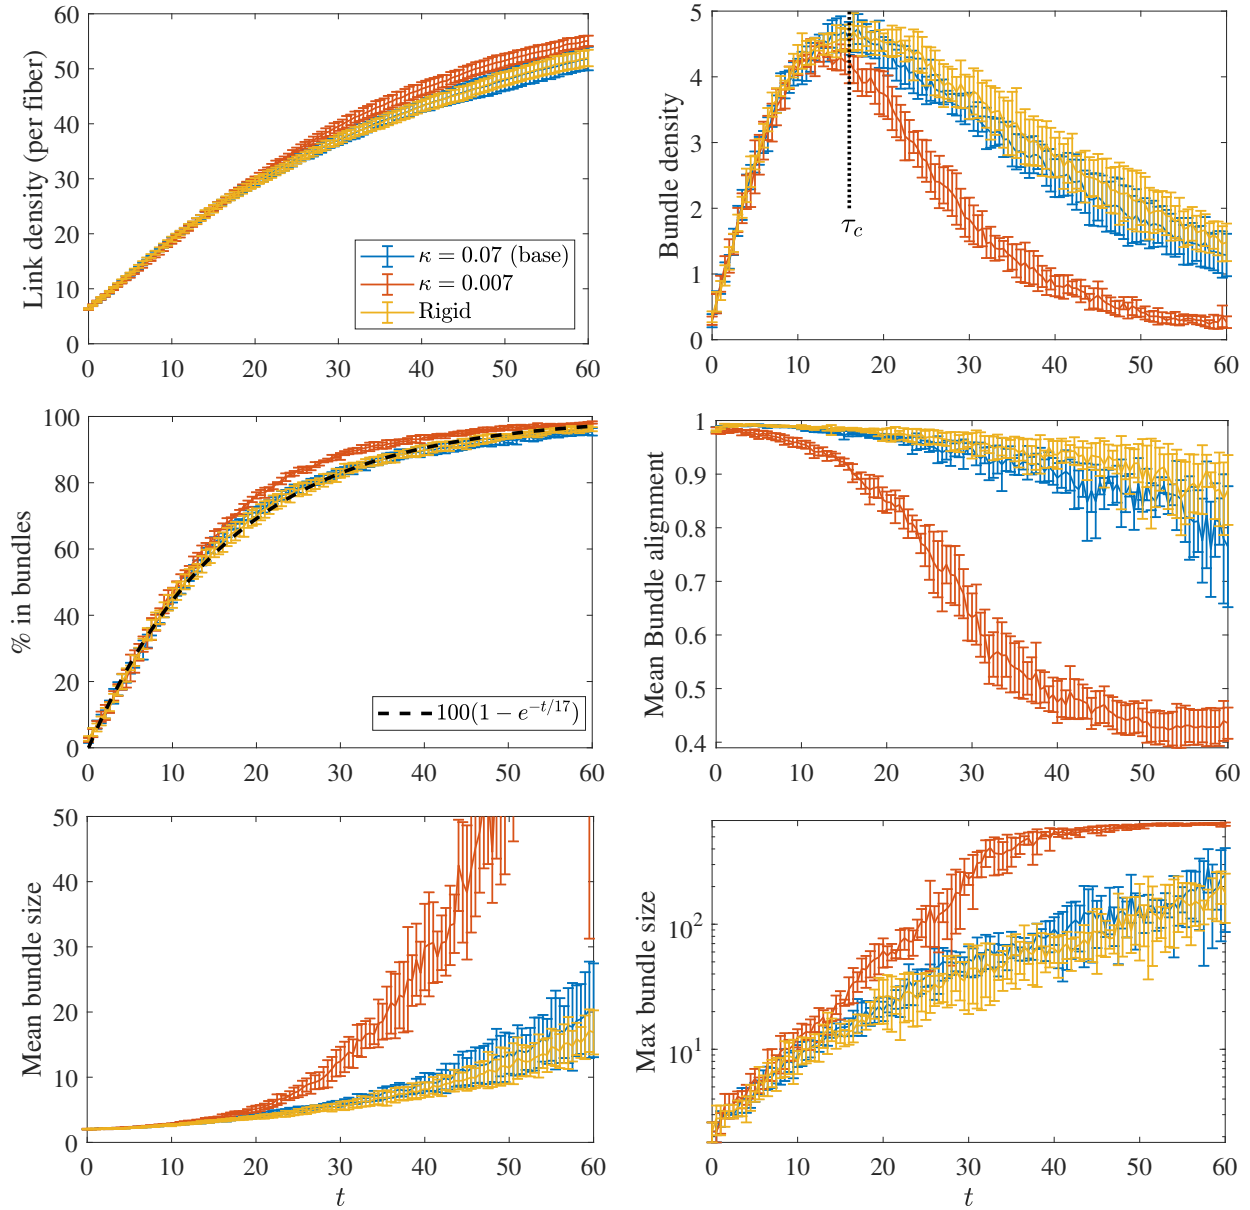


Figure 12.3: Statistics for the bundling process, where we compare the base parameters ($\kappa = 0.07$ pN $\cdot\mu\text{m}^2$, blue) with the systems with smaller bending stiffness ($\kappa = 0.007$, orange) and rigid fibers (yellow, $\kappa \rightarrow \infty$). After $\tau_c \approx 16$ s, the bundling dynamics for the less stiff fibers are significantly faster. Fibers with similar bending stiffness to actin are well-approximated by rigid fibers.

non-Brownian filaments, Fig. 12.3 shows $\tau_c \approx 16$ s corresponds to the timescale of increase of the percentage of fibers in bundles (see middle left frame), meaning it is also the timescale on which the fibers' rotational degrees of freedom are arrested or constrained.

After the coalescence time, we see a transition to coalescence of bundles, and the flexibility of the fibers comes into play. Figure 12.3 shows that the number of bundles is declining and the mean bundle alignment is dropping for $t > \tau_c$, which implies that bundles are forming with non-aligned fibers. The mean and maximum bundle sizes also start to grow, which again means that small bundles are coming together to form the larger ones we see in the bottom row of Fig. 12.1. Figure 12.4 shows that by $t = 60$ seconds, at least 75% of the fibers are in bundles of size 30 or larger. It is in this stage where the flexibility of the fibers can become important: when $\kappa = 0.007$ (fibers ten-fold less stiff than actin), coalescence of bundles occurs faster than in systems with $\kappa = 0.07$ or systems with rigid fibers, since in the former case the fibers are more compliant and can be linked together more easily by deforming. That said, when $\kappa = 0.07$ (persistence length $17 \mu\text{m}$), Figs. 12.3 and 12.4 show that the dynamics throughout the bundling process are well-approximated by rigid filaments. This analysis is of course limited by the fiber length we've chosen: in particular, we have shown that, in the absence of Brownian motion, rigid filaments are a good approximation to semiflexible actin filaments for fibers of length $\leq 1 \mu\text{m}$, which are most common in vivo. The approximation will be worse as the filament length gets larger.

To simplify our analysis of parameter changes, in the next three sections we will consider rigid fluctuating filaments, which we can study by using a modified version of the temporal integrator in Chapter 8 where the kinematic matrix \mathbf{K} is replaced by a rigid body kinematic matrix \mathbf{K}_r . That is, we consider the rigid filament dynamics

$$\partial_t \mathbf{X} = \mathbf{K}_r \mathbf{N}_r \mathbf{K}_r^T \mathbf{F}_{\text{ext}} + \sqrt{2k_B T} \mathbf{K}_r \mathbf{N}_r^{1/2} \mathcal{W}, \quad (12.1)$$

where \mathbf{K}_r is the kinematic matrix for a *rigid* fiber which acts on a 6-vector $(\boldsymbol{\Omega}, \mathbf{U}_{\text{MP}})$

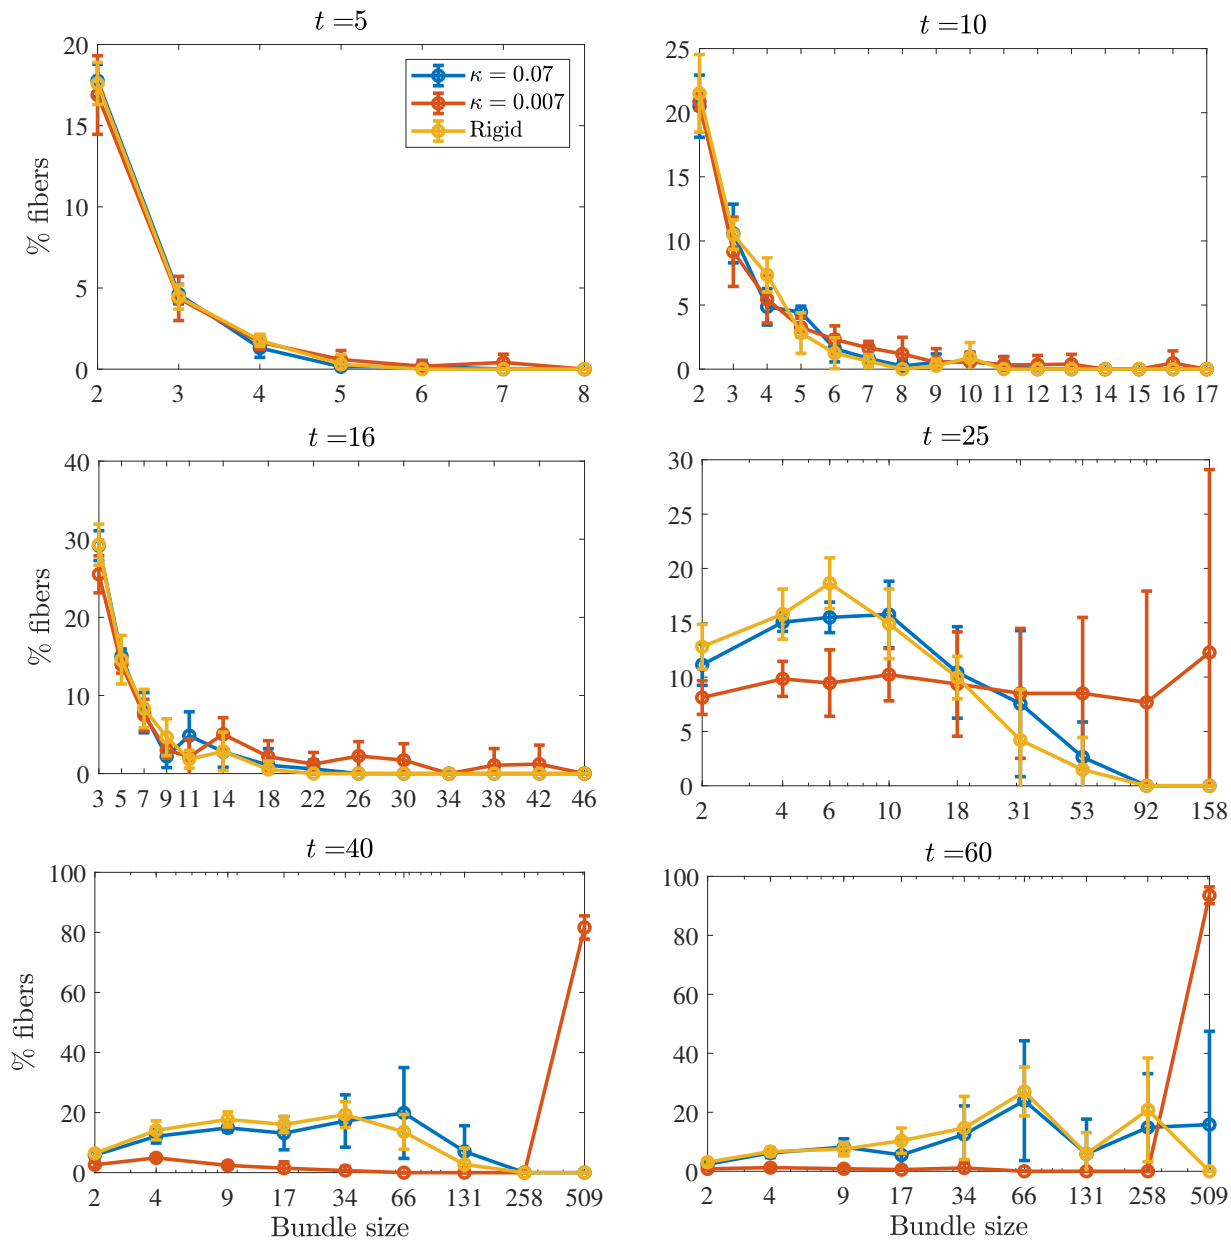


Figure 12.4: Bundle sizes over time in the system with $F = 675$ semiflexible filaments in a domain of size $L_d = 3 \mu\text{m}$ ($\tau_c \approx 16$). We show the percentage of fibers that are in bundles of various sizes over time for $\kappa = 0.07$ pN· μm^2 (blue), $\kappa = 0.007$ pN· μm^2 (orange), and rigid fibers (yellow). For times $t = 25, 40$, and 60 seconds, the x coordinate reflects the center of a histogram bin with logarithmically-scaled width. At $t = 5$ s, 25% of the filaments are in bundles of sizes 2 or 3, while most of the other fibers are not in bundles. At $t = \tau_c = 16$ s, about 50% of the fibers are in bundles of size 10 or less, with a small percentage in larger bundles, and the rest not in bundles at all (this is the composite bundle state). For semiflexible fibers with $\kappa = 0.07$ and rigid fibers, about 75% of the fibers are in bundles of size 30 or larger by $t = 60$, while for fibers with $\kappa = 0.007$ the entire suspension has coalesced together by $t = 40$ seconds.

to give velocity in an analogous way to (6.28). The 6×6 rigid-body mobility is $\mathbf{N}_r = \left(\mathbf{K}_r^T \widetilde{\mathbf{M}}^{-1} \mathbf{K}_r \right)^\dagger$, exactly as in (8.3). Because the fibers are rigid and straight, there are no stochastic drift terms [235], and so we can apply the algorithm in Section 8.2.3 to simulate rigid fibers, but without steps 1–3.

12.2 Thermal fluctuations speed up bundling

An important quantity in the case of rigid fibers is the time for a fiber to diffuse across a mesh size. In our initial set-up, the fibers are spaced approximately ℓ_m apart, and they first must find each other in order to cross link and begin the bundling process. The theoretical translational diffusion coefficient of a straight fiber, derived in [235], can be written in terms of the 3×3 translational mobility matrix \mathbf{N}_{tt} for rigid body motions (the trans-trans block of the 6×6 matrix \mathbf{N}_r) as

$$D_t = \frac{k_B T}{3} \text{tr}(\mathbf{N}_{\text{tt}}) \approx \frac{k_B T}{3\mu L} 1.67, \quad (12.2)$$

where the last equality gives the result for a fiber aspect ratio of $\epsilon = 0.004$, which we obtain from intra-fiber hydrodynamics (note that this estimate accounts for the anisotropy of the fiber, since \mathbf{N}_{tt} has an eigenvalue in the parallel direction which is twice as large as the perpendicular directions). The mean square displacement of the fiber center is then $\langle r^2(t) \rangle = 6D_t t$. Substituting the parameters in Table 11.1, we obtain $6D_t \approx 0.13 \mu\text{m}^2/\text{s}$, and thus the time to diffuse a mesh size is given by $\tau_m = \ell_m^2 / (6D_t) \approx \ell_m^2 / 0.13 \text{ s}$. Since diffusion promotes mixing of the suspension and gives more opportunities for cross linking, our expectation is that thermal fluctuations should speed up the transition from the homogeneous meshwork to the composite bundle state, where bundles are made of a few fibers which must be close enough together to cross link. This assumes that the CL concentration is large enough for links to bind as soon as fibers are close enough together; we will analyze this assumption in

Section 12.3.

To understand how thermal diffusion affects the bundling process, in Fig. 12.5 we plot the statistics both with (orange) and without (blue) thermal fluctuations. We see that the entire process is faster with diffusion, as we might expect, but the degree of acceleration changes before and after τ_c . Before τ_c , the process with diffusion is significantly faster than without; for instance, it takes about 3 seconds for 50% of the fibers to be in bundles with diffusion, while without diffusion it takes 12 seconds (Fig. 12.5 inset), which is a difference of a factor of 4. Indeed, the critical bundling time $\tau_c \approx 4$ seconds with diffusion, while we have already seen $\tau_c \approx 16$ seconds without diffusion, so that the difference is again a factor of 4. For $t > \tau_c$, when bundles start to coalesce, the difference is only a factor of 2; an exponential fit to the decaying bundle density gives a constant of 20 seconds for simulations with diffusion and 40 seconds for simulations without diffusion.

A similar relationship holds when we look not at the number of bundles (which depends on F_B , the minimum number of fibers forming a bundle), but the percentage of fibers in bundles, which is independent of F_B and shown in the middle left frame of Fig. 12.5. Unlike in the non-Brownian case, where a single timescale fits the data, the Brownian case requires two timescales for fitting, which are 3 seconds (which is close, but not equal to $\tau_c \approx 4$ s), and 13 seconds. In this case, the new fast timescale of 3 seconds reflects the ability of *Brownian* filaments to freely diffuse translationally and rotationally early in the simulation. Later in the simulation, the filaments are arrested, and the timescale on which filaments enter bundles approaches that of non-Brownian filaments, 17 seconds. This provides more evidence for our two stage model of bundling, where thermal fluctuations make more of a difference in the first stage when fibers are less constrained by cross links. Equivalently, entropic effects (Brownian motion of fibers and cross-linker stretching) are more important at early times, while at later times energetic effects trap the fibers in the clustered bundle state.

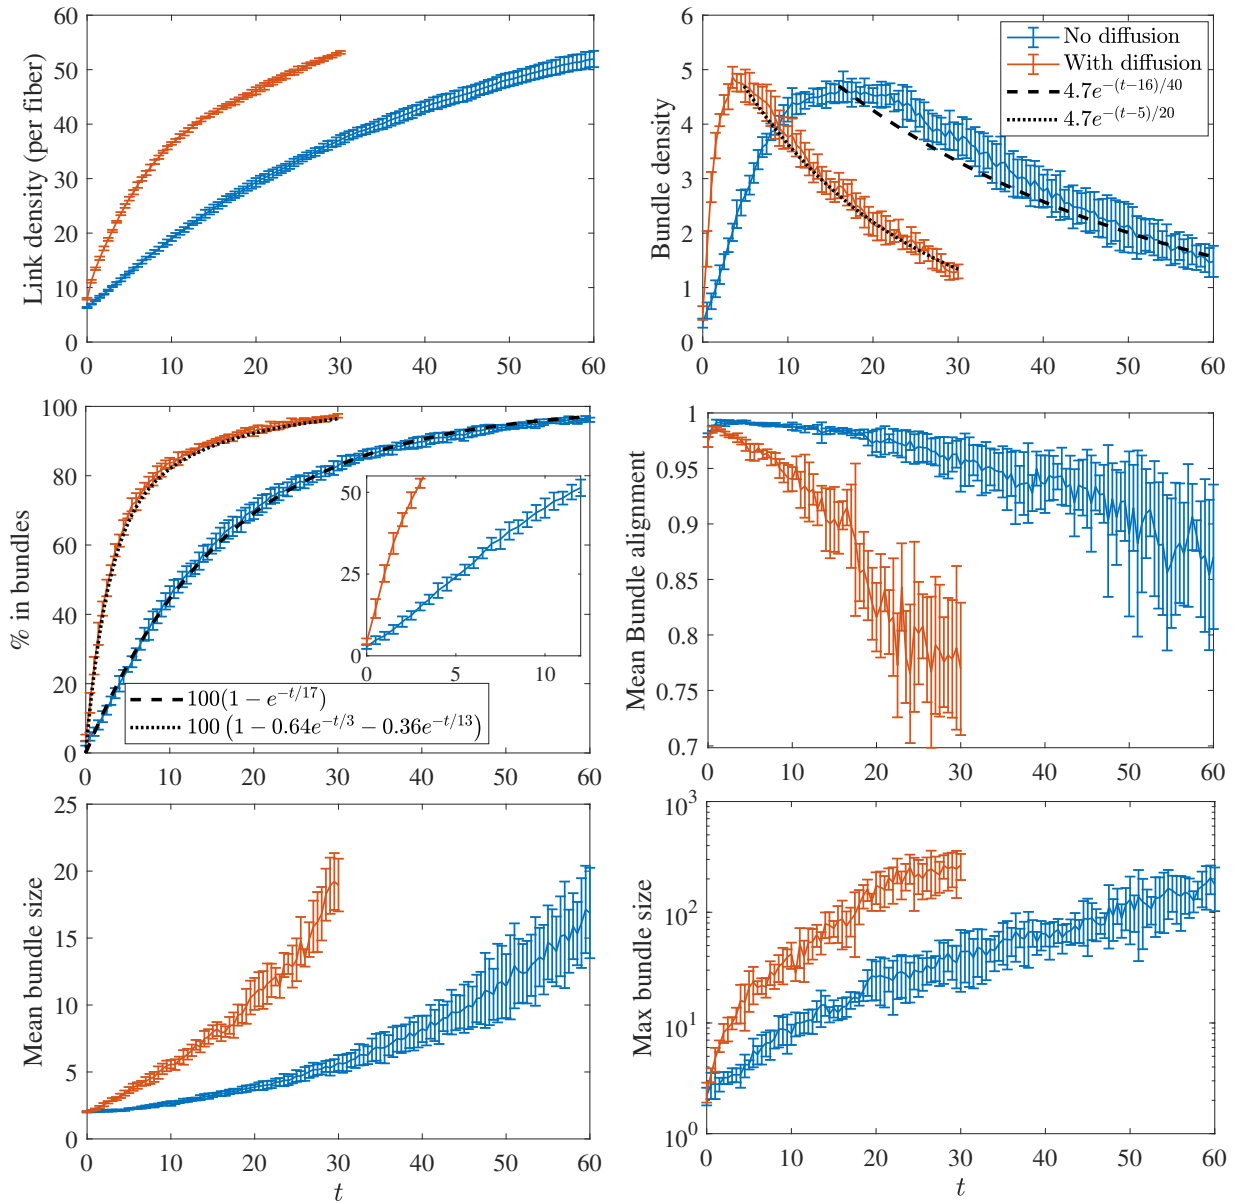


Figure 12.5: Statistics for the bundling process with and without thermal fluctuations. The blue lines show the results without thermal movement, while the orange lines show the results with translational and rotational diffusion. Here we use $\Delta t = 10^{-4}$; we have verified that the statistical noise exceeds the time stepping error for this time step size. The peak in the bundle density occurs at $\tau_c \approx 16$ for systems without diffusion, while for systems with diffusion it occurs at $\tau_c \approx 4$ seconds. These results should be taken with a grain of salt, since with Brownian fluctuations the CL binding model also changes (Section 9.1.2). Because we also widen the range at which we allow cross linking in this case (from $\sqrt{k_B T/K_c}$ to $2\sqrt{k_B T/K_c}$), it turns out that the bundling process is faster even without filament fluctuations.

To show that the network morphology has not changed when we add thermal movement, in Fig. 12.6 we show the networks at $t \approx \tau_c$ without and with thermal fluctuations. The composite bundle network morphology at τ_c is similar between the two, which demonstrates that fluctuations speed up the pace of bundling without changing the types of bundles that evolve.

The results of this section are actually harder to interpret when we remember that including Brownian fluctuations changed two things in our model: first, the filaments do move faster because of fluctuations. But secondly, the binding model also changes, as when we include fluctuations we shift the binding kinetics so they are in detailed balance, as described in Section 9.1.2. Because we also widen the range at which we allow cross linking in this case (from $\sqrt{k_B T / K_c}$ to $2\sqrt{k_B T / K_c}$), it turns out that the bundling process is faster even without filament fluctuations. Indeed, it is not really consistent in the first place to have a model which has Brownian fluctuations for the CLs and not the filaments. Therefore, we will now compare simulations *with* fluctuations with different parameters.

In subsequent sections, we will analyze how the timescale τ_c that we use to quantify the speed of bundling depends on the microscopic parameters. While the precise value of τ_c depends on the parameter F_B (the minimum number of filaments in a “bundle”), Fig. 12.5 shows that this timescale can roughly capture the initial growth rate of the percentage of fibers in bundles, which is independent of F_B . Since τ_c is easier to measure by looking at the peak bundle density (and is in principle easier to measure experimentally through microscopy) than by fitting a double-exponential curve (which is an ill-conditioned problem for larger timescales), we will use the bundle density maximum as the definition of τ_c . Of course, making F_B larger increases τ_c , as we show in Fig. 12.7 by setting $F_B = 5$, but the ratio of τ_c between the Brownian and non-Brownian system remains the same. However, increasing F_B will cause us to miss the initial stage of bundling where two-filament bundles form and the fibers’ rotational degrees of freedom are arrested, so we will use $F_B = 2$




Figure 12.6: Representative snapshots of a network of rigid fibers without (left) and with (right) Brownian motion. Both snapshots are taken at $t = \tau_c$, which is 16 seconds for simulations without fluctuations and 4 seconds for simulations with fluctuations. The networks are qualitatively the same.

henceforth.

12.3 Dependence on actin and CL concentration

Our conclusion that thermal fluctuations significantly accelerate the initial stage of the bundling process is dependent on having a sufficient concentration of CLs. While thermal fluctuations undoubtedly increase the frequency of fibers coming close enough together for cross linking, the bundling process still must be initiated via binding of a CL. Consequently, in this section we consider a range of values of mesh size (actin concentration) and k_{on} (CL attachment rate, which is proportional to CL concentration), to get a more complete picture of how the critical bundling time τ_c depends on these parameters. In particular, we will consider mesh sizes $\ell_m = 0.2$ ($F = 675$, $L_d = 3 \mu\text{m}$), 0.4 ($F = 400$, $L_d = 4 \mu\text{m}$), and $0.8 \mu\text{m}$ ($F = 338$, $L_d = 6 \mu\text{m}$), and single-end binding rates $k_{\text{on}} = 1.25$, 5 (the base value), and




Figure 12.7: Comparing Brownian and non-Brownian filaments with a minimum of $F_B = 5$ filaments per bundle. When we increase to $F_B = 5$ filaments per bundle, we observe the same characteristic growth and decay as with $F_B = 2$ (Fig. 12.5), with the peak occurring three-fold faster in simulations with Brownian dynamics and the later dynamics being accelerated by a factor of about two (from a timescale of 60 s to 25 s). These confirm the qualitative (and quantitative) conclusions from Fig. 12.5.

$20 (\mu\text{m}\cdot\text{s})^{-1}$. By changing the rate at which a single CL end binds to a fiber, we effectively vary the CL concentration.

Figure 12.8 shows the resulting evolution of the bundle density for the nine different systems, as well as the resulting critical bundling time τ_c . For systems with large k_{on} , where binding is essentially instantaneous once filaments come close enough together, $\tau_c \approx 3$ s for the small enough mesh sizes of $\ell_m = 0.2$ and $0.4 \mu\text{m}$. Once the mesh size increases to $0.8 \mu\text{m}$, the bundling time increases, but only to about 4.5 s (see inset of Fig. 12.8). Thus τ_c is not a strong function of mesh size for larger k_{on} , which implies that the process for large k_{on} , where there is always sufficient cross-linker available for binding, is primarily limited by cross-linking dynamics (alignment of filaments), with diffusion (across the mesh size) playing only a secondary role.

Let us now consider the case of slower k_{on} . In this case, filaments could come close enough to link together, but diffuse away before a CL can actually bind them. As a result of this,

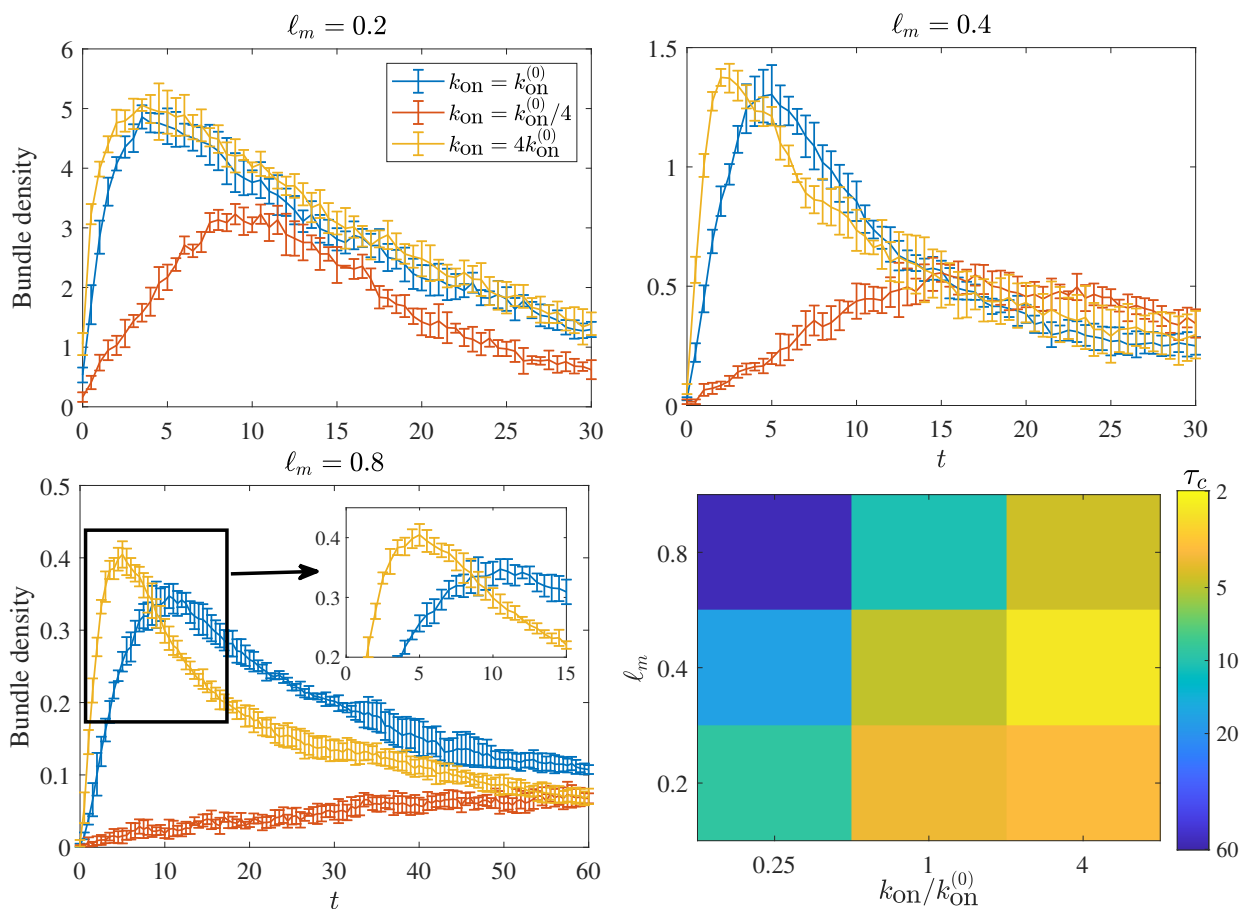


Figure 12.8: Bundling time scales for a range of initial mesh sizes ℓ_m and binding rate k_{on} . The first three frames show the trajectory of the bundle density for the different mesh sizes, where blue lines denote our base value of $k_{\text{on}} = k_{\text{on}}^{(0)} = 5 \text{ } (\mu\text{m}\cdot\text{s})^{-1}$, red lines denote $k_{\text{on}}/k_{\text{on}}^{(0)} = 1/4$, and yellow lines $k_{\text{on}}/k_{\text{on}}^{(0)} = 4$. The bottom right frame shows the dependence of the critical bundling time τ_c on ℓ_m and k_{on} .




Figure 12.9: Snapshots of the network at $t = \tau_c$ with initial mesh size $\ell_m = 0.4 \mu\text{m}$ ($F = 400$ filaments of length $L = 1$ in a domain of size $L_d = 4$) with $k_{\text{on}}/k_{\text{on}}^{(0)} = 1/4$ (left, $\tau_c \approx 17$), $k_{\text{on}}/k_{\text{on}}^{(0)} = 1$ (middle, $\tau_c \approx 5$), and $k_{\text{on}}/k_{\text{on}}^{(0)} = 4$ (right, $\tau_c \approx 2.5$). A smaller k_{on} (smaller CL concentration) gives fewer but larger bundles at $t = \tau_c$, as well as a smaller percentage of fibers in bundles.

the bundling process is slowed, and in fact the peak bundle density drops. Indeed, as shown in Fig. 12.9, networks with smaller k_{on} (lower CL concentration) contain larger bundles at $t = \tau_c$ than those with larger k_{on} (higher concentration). Upon reducing k_{on} , two filaments finding each other becomes the limiting step in the bundling process. This causes a slow growth of the bundle curve, and a bias towards larger bundles, which build up at a faster rate (relative to τ_c), and the process is rate-limited by two-filament bundle formation. The scaling of τ_c at small k_{on} (left column of the bottom right panel in Fig. 12.8) is reminiscent of a diffusion-limited process, as it increases from 9 s to 17 s to 56 s as the mesh size doubles, scaling approximately as ℓ_m^2 as the mesh size increases. In some sense, diffusion is actually a hindrance to bundling, since fibers that are close to each other diffuse away before a CL can bind them together.

We note that a roughly constant bundling time can be achieved by decreasing k_{on} as the mesh size decreases (moving from the top right to the bottom left of the bottom right panel in Fig. 12.8). This implies that the relative concentration of CL required for a particular bundled state decreases with the mesh size, as has been found experimentally [206]. When the mesh size is smaller, the filaments are in contact for longer, and so it is less important

that a CL be available immediately to bind them together. By contrast, filaments in larger-mesh-size systems are only in contact for a brief time, so relatively more CLs are necessary to ensure that these filaments are linked when they come into contact with each other.

12.4 Faster bundling for shorter filaments

We will now explore the dependence of the critical bundling time τ_c on the fiber length. Experimentally, it has been shown that shorter filaments bundle faster [202, 209], but it is still unclear whether this is due to thermal movements, cross-linking kinetics, or some combination of both. In this section, we show that the experimental results can only be reproduced if we consider thermal movements, so that cross-linking kinetics are not responsible for the speedup in bundling. We use a fixed mesh size of $\ell_m = 0.2 \mu\text{m}$, which translates to $F = 675$ filaments of length $L = 1 \mu\text{m}$ in a domain of size $L_d = 3 \mu\text{m}$, and $F = 400$ filaments of length $L = 0.5 \mu\text{m}$ in a domain of size $L_d = 2 \mu\text{m}$.

In Fig. 12.10, we show how the bundle density, percentage of fibers in bundles, and mean bundle size evolve for the two different filament lengths both (a) without and (b) with actin diffusion. In Fig. 12.10(a), we see that in the absence of Brownian motion the behavior in the two systems is similar, with the peak bundle density occurring in both cases around $\tau_c \approx 15$ seconds. Furthermore, there is only a mild difference in the percentage of fibers in bundles over time. The mean bundle size is at most twice larger for the system with shorter filaments, but we would expect this since the filaments are twice as short and there are twice as many of them if ℓ_m is fixed.

In Section 12.2, we showed that Brownian motion speeds up the bundling process by promoting mixing and more near-contacts of filaments. In particular, we saw that the time for a filament with length $L = 1 \mu\text{m}$ to diffuse a mesh size of $\ell_m = 0.2 \mu\text{m}$ is $\tau_m \approx 0.30$ s, so that filaments can find each other rapidly and begin the bundling process. In the case of filaments with $L = 0.5 \mu\text{m}$, our thermal diffusion coefficient (12.2) scales log-linearly with

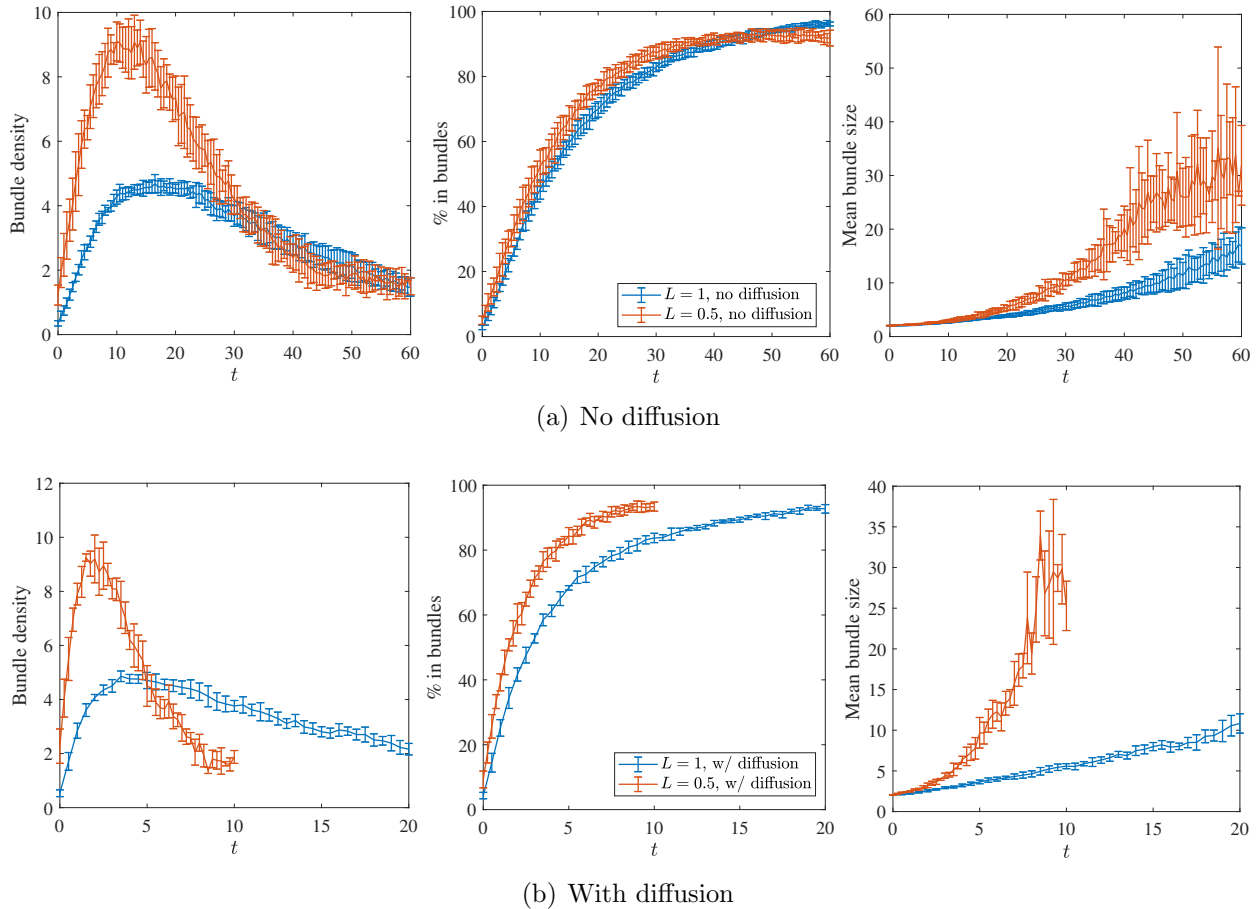


Figure 12.10: Effect of changing filament length for rigid fibers with and without Brownian motion, with constant initial mesh size $\ell_m = 0.2 \mu\text{m}$. (a) Without fiber diffusion, we show the statistics for filaments of length $L = 0.5 \mu\text{m}$ (red) and $L = 1 \mu\text{m}$ (blue), where we observe dynamics occurring on a similar timescale, especially in the initial stage ($t \leq \tau_c \approx 20$) of bundling. (b) When we add fiber diffusion, the bundling process for $L = 0.5 \mu\text{m}$ (red) is significantly faster than $L = 1 \mu\text{m}$ (blue), because filaments can diffuse faster.

the fiber length, so that it takes $\tau_m = 0.17$ s to diffuse a mesh size of $\ell_m = 0.2 \mu\text{m}$. We might expect, therefore, that at least the initial stages of the bundling process will be sped up by a factor of 2.

Figure 12.10(b) shows that this is indeed the case. For $\ell_m = 0.2 \mu\text{m}$, the bundle density peak occurs around $\tau_c \approx 2$ seconds when $L = 0.5 \mu\text{m}$, while with $L = 1 \mu\text{m}$ it occurs around $\tau_c \approx 4$ seconds, so it appears that bundling time with thermal motion scales linearly with filament length, which is in (approximate) accordance with the scaling of the translational diffusion coefficient. The faster bundling behavior also manifests itself in the link density and percentage of fibers in bundles, where we see that systems with shorter filaments reach a number of links or percentage of fibers about twice as fast. For instance, 80% of the fibers are in bundles by $t \approx 4$ seconds in the $L = 0.5 \mu\text{m}$ case, while with $L = 1 \mu\text{m}$ the 80% mark is not reached until about $t \approx 8$ seconds.

12.5 Semiflexible fibers

So far, we included fluctuations by modeling the filaments as rigid, and made statements about the role of thermal fluctuations in bundling dynamics for networks of rigid filaments¹. We now complete our investigation by studying the role of semiflexible fluctuations in the bundling process.

We use the cross linking forces from (9.10), having verified that the smooth vs. nonsmooth forcing assumption makes little difference in the overall bundling dynamics. These forces at each time step become additional forces in the Langevin equation (8.12), so that at each time step we solve (for each fiber independently)

$$\partial_t \mathbf{X} = \mathbf{K} \mathbf{N} \mathbf{K}^T (-\mathbf{L} \mathbf{X} + \mathbf{F}_{\text{ext}}) + k_B T \partial_{\mathbf{X}} \cdot (\mathbf{K} \mathbf{N} \mathbf{K}^T) + \sqrt{2k_B T} \mathbf{K} \mathbf{N}^{1/2} \boldsymbol{\mathcal{W}}. \quad (12.3)$$

¹The reason for this is that we actually published the study [C] on bundling in rigid fibers prior to developing an integrator for semiflexible fibers in [F] (see p. 2 for references).

by replacing $-\mathbf{L}\mathbf{X}^n$ with $-\mathbf{L}\mathbf{X}^n + \mathbf{F}_{\text{ext}}(\mathbf{X}^n)$ on the right hand side of the saddle point solve (8.33) (that is, we treat the spring forces explicitly in time).

The question we will examine is how the behavior changes as we increase ℓ_p/L , so we will leave all parameters in Table 11.1 constant except the bending stiffness κ . This includes $k_B T = 4.1 \times 10^{-3}$ pN $\cdot\mu\text{m}$, which corresponds to the thermal energy at room temperature. For our spatial and temporal discretization, we use $N = 12$ and $\Delta t = 10^{-4}$ s over all simulations, having verified that doubling the number of points and halving the time step size does not change the results within statistical error. Our explicit treatment of the forces from the CLs, whose base stiffness of $K_c = 10$ pN/ μm is an order of magnitude estimate for the effective stiffness of α -actinin [165, 209], limits the time step size. In particular, resolving the spring dynamics automatically resolves the equilibrium statistical mechanics.²

12.5.1 Visualizing the bundling process

The same fundamental bundling process plays out in networks of semiflexible filaments, as shown in Fig. 12.11, where we show snapshots of the bundling process at different time points for three different orders of magnitude of ℓ_p/L (these plots show 200 filaments with $L_d = 2$). The top plots ($t = 1.5$ s) correspond to initial stage of bundling, where there are many bundles of a few filaments, while subsequent plots begin to show coalescence of the bundles. There are a few takeaways here: first, we see that the bundle morphology looks qualitatively different as we decrease ℓ_p/L , with smaller persistence length having more curved fibers and therefore more curved bundles. Furthermore, the smaller persistence length bundles agglomerate faster, and at a given time they appear more clumped (especially $t = 6$ s). It is not clear, however, to what extent these differences in bending deformations are driven by

²If we convert the time step sizes from Fig. 8.5 to the units we use here, we need a time step size of 3×10^{-3} s, 1.4×10^{-3} s, and 1.2×10^{-3} s to resolve the equilibrium fluctuations for $\ell_p/L = 1, 10$, and 100, respectively. All of these are at least an order of magnitude larger than that required to resolve the explicit treatment of CLs.

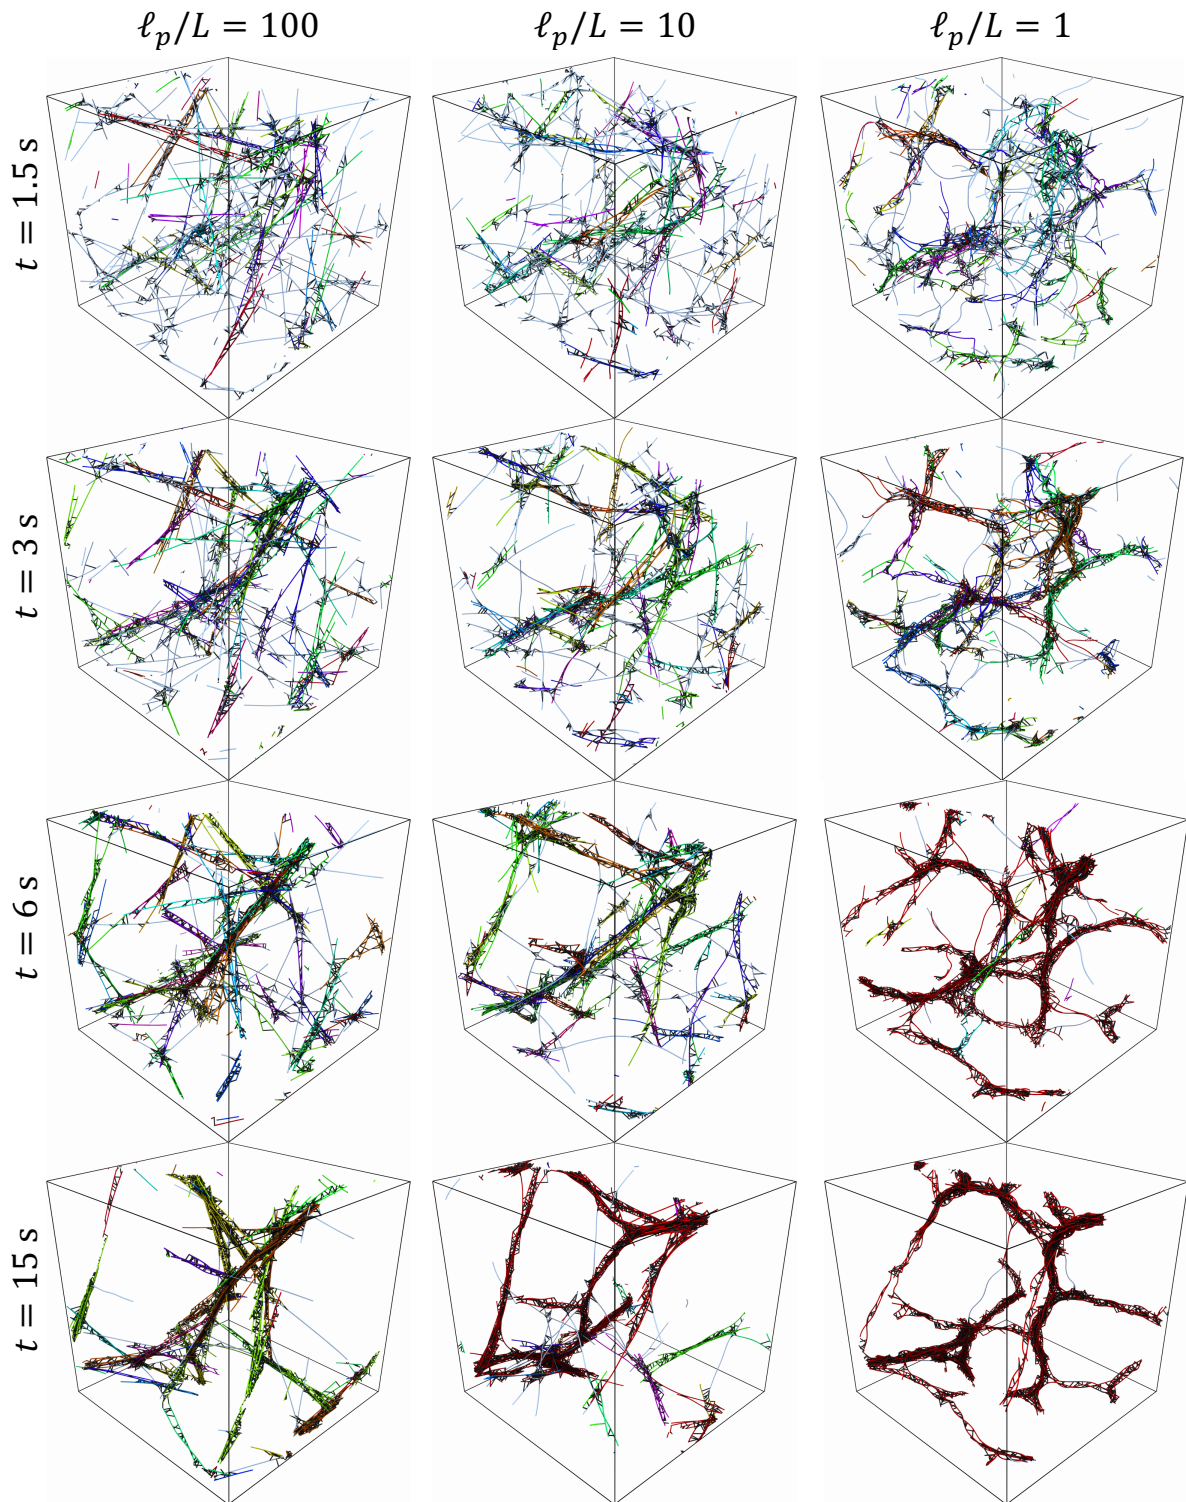


Figure 12.11: Snapshots of the bundling process in networks of varying fiber stiffness. Colored fibers are actin filaments (colored by bundle), while black fibers are the cross linkers. The bundling process is fastest for the most flexible fibers, where the final morphology shows curved bundles.

CL forces vs. the thermal bending fluctuations.³ Indeed, we did show in Section 12.2 that agglomeration (the second stage of the bundling process) happens faster for non-fluctuating filaments that are less stiff, since they are able to be bent easier by the cross linkers. Thus the main question here is whether semiflexible bending fluctuations themselves speed up (or slow down) the bundling process, or whether CL forces dominate.

12.5.2 Quantifying the role of semiflexible bending fluctuations

To get at this question, we need to dissect the evolution of the actin filaments in a cross-linked network into the three possible ways they can move: action by CL forces, thermal rotation and translation (keeping the fiber shape fixed), and bending fluctuations. So far, we studied the first two of these, showing in particular that deterministic filaments (those that can only move by CL forces) behave the same as rigid ones when $\ell_p/L \sim 10$. In this section, we can finally consider the full system, and how the third possible motion, transverse bending fluctuations, impacts bundling. It will be important, however, to separate these fluctuations from the rotational and translational diffusion that we have considered previously. Since our temporal integrator in Section 8.2.3 does not distinguish between the two kinds of fluctuations, for comparison we consider an alternative model where the only fluctuations are translational and rotational diffusion, as if the fibers were rigid in their current configuration. We consider a dynamics which takes the deterministic part from the semiflexible Langevin equation (12.3), and the random part from the rigid fiber Langevin equation (12.1)

$$\partial_t \mathbf{X} = \mathbf{K} \mathbf{N} \mathbf{K}^T (-\mathbf{L} \mathbf{X} + \mathbf{F}_{\text{ext}}) + \sqrt{2k_B T} \mathbf{K}_r \mathbf{N}_r^{1/2} \mathcal{W}. \quad (12.4)$$

³A similar question was studied in [236] for microtubule networks, where the authors showed that large *nonthermal* forces combine with polymerization dynamics to generate bent microtubule shapes in cells.

Our rationale is that if such semiflexible rigid body motion (“SF-RBD”) simulations give the same results as those with semiflexible bending fluctuations (“SF-Bend”), then bending fluctuations are not important to the bundling process. Note however that the dynamics (12.4) are problematic for stiffer fibers, since in that case the pseudo-inversion in the calculation of $\mathbf{N}_r = \left(\mathbf{K}_r^T \widetilde{\mathbf{M}}^{-1} \mathbf{K}_r\right)^\dagger$ involves singular values which are close to zero. We find that the dynamics for SF-RBD filaments with $\ell_p/L = 100$ are quite sensitive to the tolerance we use, and so we do not report them here. We find it sufficient to only consider up to $\ell_p/L = 10$, for which we obtain dynamics that are not sensitive to the tolerance (see Fig. 12.12). For $\ell_p/L = 100$, fibers with semiflexible bending fluctuations behave almost identically to rigid fibers (see Fig. 12.12), which we have already shown we can simulate without difficulty.

We begin by looking at the differences in bundling dynamics between *semiflexible* filaments and the rigid filaments we considered in Fig. 12.3. This is actually the same test we performed in Fig. 12.3, but this time we consider filament fluctuations in addition to CL forces. The bundle density (number of bundles divided by periodic cell volume) and percent of fibers in bundles over time are shown using darker colors in Fig. 12.12, where we consider $\ell_p/L = 1, 10$, and 100, and compare the results to the case when the fibers are actually rigid, the dynamics of which are governed by the overdamped Ito Langevin equation (12.1).

We observe almost complete overlap between the trajectory for $\ell_p/L = 100$ and that for rigid fibers, (except for the bundle density curve at late times, which is when cross linkers exert extremely large forces on the filaments). This is not a surprise given that we see bundles of *straight* filaments when $\ell_p/L = 100$ in Fig. 12.11, but it does showcase the ability of our temporal integrator to remain accurate in the stiff limit, in contrast to that for SF-RBD filaments. Dropping to $\ell_p/L = 10$, the pictures in Fig. 12.11 show curved bundles, and the plots in Fig. 12.12 for $\ell_p/L = 10$ show significant deviations from rigid fibers, especially in the later stages of bundling. The curves for $\ell_p/L = 1$ do not even match rigid fibers at early times, which indicates that the semiflexible fluctuations impact the first stage of




Figure 12.12: Quantifying the dynamics of bundling in semiflexible fiber networks, and determining the role of bending fluctuations. We use the bundle density (left, number of bundles per unit volume) and percent of fibers in bundles (right) as metrics for bundling dynamics. We compare the semiflexible filaments with $\ell_p/L = 1$ (blue), $\ell_p/L = 10$ (red), and $\ell_p/L = 100$ (green) to the rigid filament networks we have been studying (black), finding that $\ell_p/L = 100$ is “stiff enough” to be rigid. We then separate the effect of bending fluctuations from CL forces and rigid body diffusion by comparing our semiflexible filament simulations (SF-Bend) to filaments which only fluctuate by rigidly rotating and translating (SF-RBD, lighter colors, see text for details). The CL forces and rigid body fluctuations are sufficient to account for the speed-up in bundling when $\ell_p/L = 10$, but not when $\ell_p/L = 1$.

bundle formation, in contrast to larger persistence lengths where the fluctuations appear to only accelerate later stages of bundle agglomeration. Since filaments are weakly cross-linked at early times, these results suggest that semiflexible bending fluctuations are accelerating the bundling process when $\ell_p/L \sim 1$. For $\ell_p/L \gtrsim 10$, the deviations from rigid fibers come only when the fibers are strongly cross linked, suggesting that CL forces combine with fiber flexibility to accelerate bundling.

To make this statement more precise, we compare simulations with (SF-Bend) and without (SF-RBD) thermal bending forces using lighter colors in Fig. 12.12. When $\ell_p/L = 10$, we see identical dynamics between SF-Bend and SF-RBD filaments, which means that bending fluctuations contribute minimally to the bundling process for $\ell_p/L \geq 10$, and demonstrates that the curvature of the bundles we see in Fig. 12.11 when $\ell_p/L = 10$ is indeed driven pri-

marily by cross linking forces. When $\ell_p/L = 1$, by contrast, we see faster bundling dynamics with SF-Bend filaments than with SF-RBD filaments, and we also see bundles in Fig. 12.11 that appear to have wavy spatial shapes. This implies that thermal bending fluctuations can accelerate bundling both in the initial and later stages, but only when the persistence length is comparable to the contour length of the fiber, in which case the transverse fluctuations effectively increase the probability that a CL (which can only stretch a finite amount) can bind two filaments. However, since actin filaments have persistence length on the order 10–20 μm [18], we can conclude that the bundling dynamics of filaments with length 1–2 μm are *not* significantly impacted by semiflexible bending fluctuations.

12.6 Influence of steric interactions

We now include steric interactions using the formulation in Section 9.2. We will simulate bundling of 200 fibers of length $L = 1 \mu\text{m}$ and persistence length $\ell_p = 2 \mu\text{m}$ in a periodic box of size $L_d = 2 \mu\text{m}$, and run bundling dynamics with various values of the time step. Because the fibers are not too slender, $\epsilon = 0.004$, our choice is to use a global upsampling scheme with $1/\epsilon$ points per fiber.

Figure 12.13 shows three sets of snapshots of the bundled networks at times (from left to right) $t = 2, 4$, and 8 seconds. The top row is a simulation without steric interactions, where we see lots of overlaps between the fibers (shown in black, and defined by oversampling the Chebyshev polynomials to $1/\epsilon$ points and looking for pairs of points $2a$ or less apart). In the middle row, we include sterics without changing the time step from the cross-linking limited $\Delta t = 10^{-4}$. This results in a trajectory which looks okay at early times, but has overlaps at later times, and eventually looks completely unreasonable (see the soup of fibers at $t = 8$). Decreasing the time step to $\Delta t = 2 \times 10^{-5}$ gives nicer results: we see fibers which have a limited number of overlaps, and which form bundles. The main difference we see between the no sterics and converged sterics simulation is the spacing in the bundles: there is clearly




Figure 12.13: Snapshots of the bundling process at $t = 2$ (left), 4 (middle), and 8 seconds (right), with different models of steric interactions. Top row: no sterics. Middle row: sterics with $\Delta t = 10^{-4}$. Bottom row: sterics with $\Delta t = 2 \times 10^{-5}$. Fiber overlaps are identified by oversampling the Chebyshev polynomials to $1/\epsilon$ points and looking for pairs of points $2a$ or less apart, and shown using black circles.

more space between the fibers in bundles when we include steric interactions. The amount of spacing is a function of our imposed sterics model, which is itself a soft approximation of a rigid interaction potential, so it is not necessarily the “truth.” The simulation with sterics also appears to be slightly behind (less bundles at a given time) the simulation without sterics, which is not a surprise since it has been previously speculated that sterics would slow down the initial parts bundling process [218, Note 23], with fibers being held apart at later times by cross linkers.

We now take our qualitative analysis a step further and look at the statistics of the bundling process with and without steric interactions in Fig. 12.14. Our first observation is that, in the first 2 or so seconds, all simulations with steric interactions give essentially the same dynamics. Furthermore, those dynamics are *slower* than the dynamics without sterics, which is expected from the pictures in Fig. 12.13. However, the fundamental trajectory of the curves does not change. The most noteworthy change with steric interactions is the displacement of the fibers at later times, which drops by about 10% when steric interactions are included (and the time step is small enough for trajectories to be converged; for $\Delta t = 5 \times 10^{-5}$ we see larger displacements at later times which comes from the explicit treatment of sterics with a time step size that is too large). The 10% figure is actually quite small, and suggests that the “kinetic arrest” of bundled filaments reported in [210] is driven primarily by cross linking and not sterics.

In Fig. 12.15, we conclude by looking at the number of contacts over time between the $1/\epsilon = 250$ uniform points on the filaments (this is the mean across 5 simulations). Here we normalize by the number of contacts *without* sterics, so that we can look at the reduction when sterics is included. As would be obviously expected, the number of contacts drops as the time step decreases. For the two time step sizes for which we see convergence of the statistics in Fig. 12.14 (yellow and purple), the number of contacts is reduced to 1 – 2% of the number of contacts without sterics. This ought to be acceptable for our purposes, since




Figure 12.14: Bundling statistics without steric interactions (blue, $\Delta t = 10^{-4}$) and with steric interactions (red, $\Delta t = 5 \times 10^{-5}$, yellow, $\Delta t = 2 \times 10^{-5}$, and purple $\Delta t = 10^{-5}$). We plot our typical statistics (link density, bundle density, percent of fibers in bundles, and mean bundle alignment), and then plot the mean and maximum displacement over a time of 0.02. These plots show that the code with steric interactions converges (to within statistical error bars for $\Delta t \leq 2 \times 10^{-5}$) to a bundling trajectory that is slightly slower than that without.




Figure 12.15: The number of contacts between the 250 uniform points on each fiber in simulations with various time step sizes (red, $\Delta t = 5 \times 10^{-5}$, yellow, $\Delta t = 2 \times 10^{-5}$, and purple $\Delta t = 10^{-5}$). Here a contact is defined as two uniform points coming within a fiber diameter of each other. When the time step size is small enough, there are very few contacts (purple), while when it is too large the number of contacts grows in time.

with Brownian motion we still expect to see some contacts, as our steric potential is finite and not a hard-core potential.

12.7 Bundling, turnover times, and the steady state

Because we define the bundle density in terms of bundles of an arbitrary number of filaments ($F_B = 2$), the precise value of the timescale τ_c that we obtain is also somewhat arbitrary. However, if we increase the number of filaments required for a bundle to $F_B = 5$, Fig. 12.7 shows that the peak in the bundle density occurs about a factor of 2 later in both Brownian and non-Brownian filament simulations. We therefore postulate that the *ratio* $\tau_c^{(A)}/\tau_c^{(B)}$ between systems A and B is a meaningful quantity, approximately independent of the definition of τ_c , and can be used to predict the steady state morphology in systems with fiber turnover.

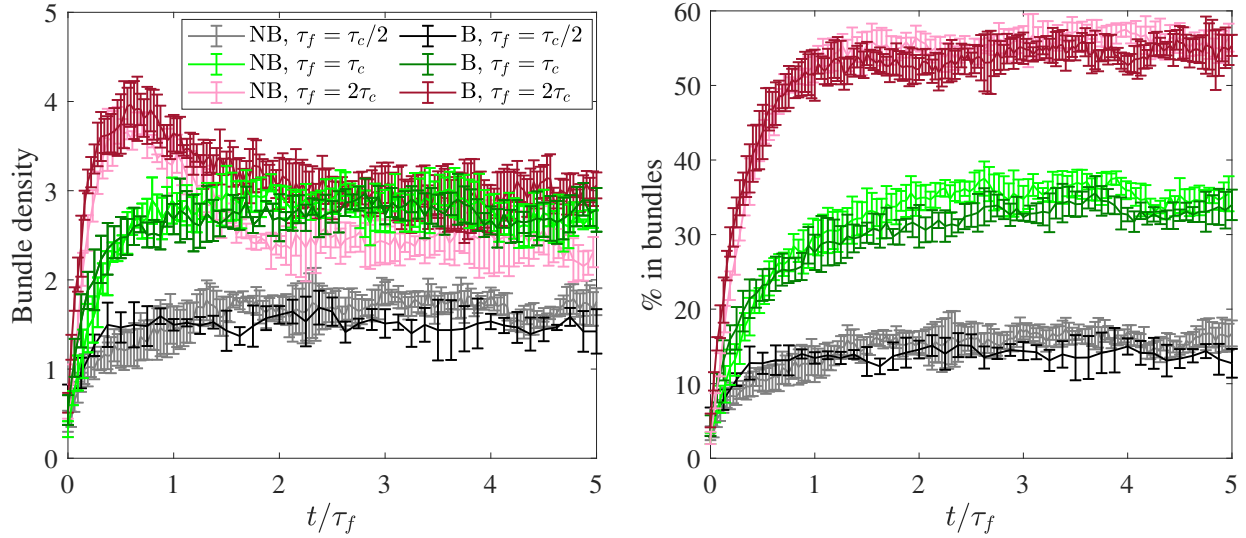


Figure 12.16: Steady state morphologies for systems with turnover. We introduce filament turnover with mean filament lifetime τ_f and observe the steady state bundle density (left) and percentage of fibers in bundles (right) for $\tau_f/\tau_c = 1/2$ (black), 1 (green), and 2 (red). Note that using a constant τ_f/τ_c in the two systems ensures $\tau_f^{(B)}/\tau_f^{(NB)} = \tau_c^{(B)}/\tau_c^{(NB)}$. Using both non-Brownian (lighter colors, $\tau_c \approx 16$ s) and Brownian (darker colors, $\tau_c \approx 4$ s) filaments, we show that the steady state bundling statistics are roughly the same when τ_f/τ_c is matched.

To test this, we introduce filament turnover with mean filament lifetime τ_f and fix τ_f as a function of τ_c , so that the ratio of the turnover times equals the ratio of the bundling times between the Brownian (B) and non-Brownian (NB) cases, $\tau_f^{(B)}/\tau_f^{(NB)} = \tau_c^{(B)}/\tau_c^{(NB)}$, or, equivalently, $\tau_f^{(NB)}/\tau_c^{(NB)} = \tau_f^{(B)}/\tau_c^{(B)}$. In Fig. 12.16 we vary the ratio τ_f/τ_c between 0.5 and 2 and plot the bundle density and percentage of fibers in bundles as they evolve to a steady state in each case. Despite the system of Brownian filaments having much faster bundling dynamics than the system of non-Brownian filaments, the morphology of the steady state is the same in the Brownian and non-Brownian cases, as is shown in the snapshots of Fig. 12.17. We will make use of this in the next chapter, where we choose a ratio τ_f/τ_c , generate a steady state, and perform rheological measurements on the steady states to determine how the mechanical properties change with morphology.

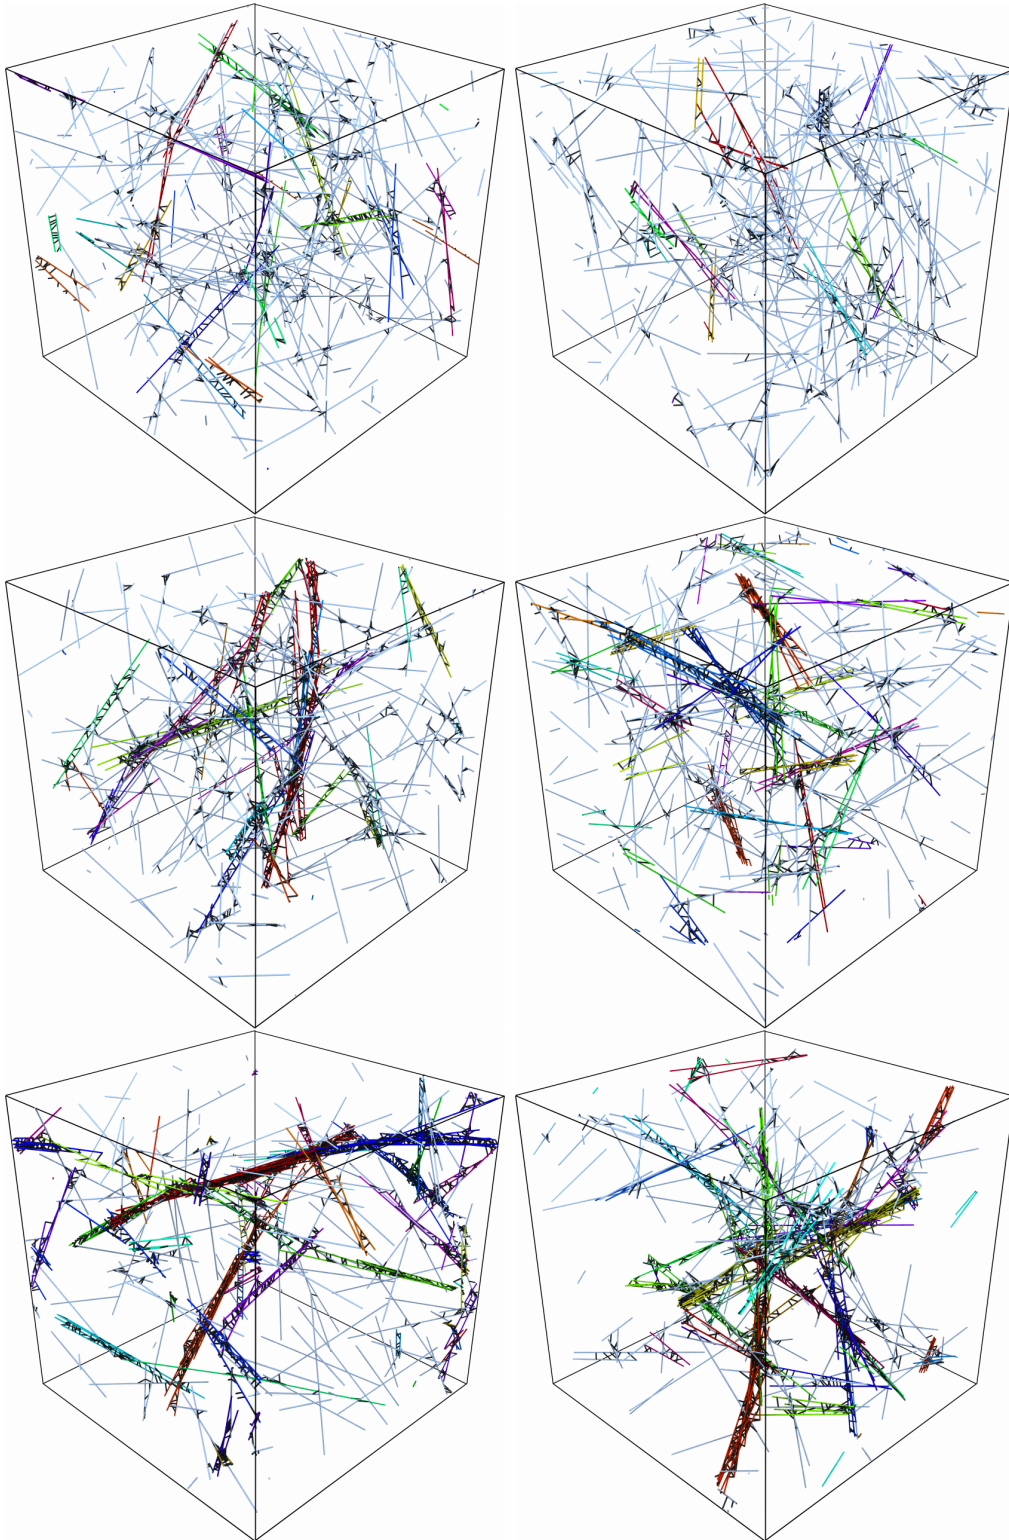


Figure 12.17: Steady states with fiber turnover for (top to bottom) $\tau_f = \tau_c/2$, $\tau_f = \tau_c$, and $\tau_f = 2\tau_c$ for a system with $L_d = 2$ and $\ell_m = 0.2$. The left column is for non-Brownian fibers ($\tau_c = 16$ seconds) and the right column is for Brownian ones ($\tau_c = 4$ seconds).

Chapter 13

Rheology of transiently cross-linked actin networks

Despite, and maybe even because of, the large volume of experimental and theoretical studies on cross-linked actin gel mechanics, a number of questions about the rheology of cross-linked networks still remain open. We focus on three of them: What role does the morphology of the network play in the mechanical response? Could there be a significant contribution from nonlocal hydrodynamic interactions? And can we inform a simple and practical spring-and-dashpot model by detailed microscopic simulations?

This chapter is built around answering these questions in *deterministic* suspensions (no Brownian motion). Since previous experimental and modeling studies show that the mechanical behavior of actin networks is most affected by the presence of short, stiff CLs which induce morphological changes in the network, we will focus our analysis on CLs with spring stiffness $10 \text{ pN}/\mu\text{m}$ and a 50 nm rest length, thereby mimicking α -actinin. We show that changing the turnover time τ_f with all other parameters fixed gives two model systems: an isotropically cross-linked network, or homogeneous meshwork, and a composite network of bundles embedded in the meshwork (here we follow the classification scheme of [203]). Unlike previous simulation studies [166, 218], the networks that we consider here are in a dynamic

steady state, which allows us to precisely quantify their steady-state mechanical properties. In particular, we use small amplitude oscillatory shear (SAOS) rheology over a range of biologically-relevant frequencies (0.01 Hz to 100 Hz)¹ to show that the networks have three characteristic timescales on which the links go from rigid to flexible, permanent to dynamic, and viscoelastic to purely viscous. It is on this last timescale that remodeling of the network occurs.

This chapter is, to our knowledge, the first to incorporate many-body hydrodynamic interactions (HIs) in the study of transiently cross-linked actin networks. Here we find results that depend on the timescale under consideration. At high frequency, we show that the viscous modulus scales linearly with ω , and that accounting for nonlocal hydrodynamics simply changes this coefficient by a fixed percentage. This seemingly uninteresting regime still requires more analysis, however, since thermal fluctuations (Brownian motion), which are not included here, are expected to contribute significantly for high frequencies. In contrast, at low frequency, we show that nonlocal hydrodynamic interactions lead to a significant *decrease* in the moduli for structures with a significant degree of bundling, and that the decrease can be explained by the disturbance flows inside a bundle. We conclude this chapter by using our knowledge of the timescales and role of nonlocal hydrodynamic interactions to suggest coarse-grained continuum models for the transiently cross-linked actin gel.

13.1 Permanent cross linkers

Before understanding the dynamics of transiently linked networks, it will be useful to look at the viscous and elastic moduli for a network of $F = 700$ fibers and 8400 CLs. In this section only, we use a slightly different set of parameters. We set the fiber length $L = 2 \mu\text{m}$

¹“Biologically relevant” timescales are from ~ 0.1 sec (the fastest relevant microscopic changes in the network structure, like unbinding of a CL [186]) to tens of seconds (turnover time, [237, 238]). The motility and division cycles for fast cells are on the scale of minutes (hundreds of seconds) [11]. Because of this, in vivo rheological experiments often focus on the range from 0.1 to 100 s [186] – we follow their example.

with aspect ratio $\epsilon = 10^{-3}$ and bending stiffness $\kappa = 0.01 \text{ pN} \cdot \mu\text{m}^2$, and use a box of size $L_d = 4 \mu\text{m}$. For the fluid viscosity, we use $\mu = 1 \text{ Pa}\cdot\text{s}$. The CLs have rest length $\ell = 0.5 \mu\text{m}$ and spring constant $K_c = 1 \text{ pN}/\mu\text{m}$. In order to avoid transient behavior, we first find a steady state configuration of the network by initializing straight fibers with CLs and running the system forward in time using local drag and $\Delta t = 0.005 \text{ s}$ without any background flow. After $t = 2500$ seconds, the maximum L^2 norm of the fiber velocity is approximately $4 \times 10^{-6} \mu\text{m}/\text{s}$, which indicates a near steady state.

In order to measure the steady state viscous and elastic moduli, we must wait for some intrinsic time on which the network reaches a new steady state in the shear flow. This relaxation time scale, which we denote by τ_r , combines the cross-linker and fiber relaxation time scales. For cross-linkers, a characteristic time scale of link relaxation is $\tau_{\text{CL}} = \mu\ell/K_c = 0.5$ seconds for our parameters. For fibers, we assume that the network is sufficiently constrained that the length scale on which the fibers can relax is the mesh size $\ell_m \sim \sqrt{L_d^3/(FL)} \approx 0.21$, or characteristic distance between filaments.² Thus a characteristic time scale of fiber relaxation is $\tau_F = \mu\ell_m^4/\kappa \approx 0.2$ seconds, and our expectation is for τ_r to be on the order $0.1 - 1$ second.

To measure the time scale τ_r more precisely, we start with the steady state configurations, turn on a shear flow (7.50) with $\omega = 0.2\pi \text{ rad/s}$ and $\dot{\gamma}_0 = 0.02\pi \text{ 1/s}$, and run for one cycle (until $T = 10$ seconds) with $\Delta t = 0.005 \text{ s}$. We then turn off the shear flow and measure the velocity of the fibers for another 5 seconds. We track the mean L^2 fiber velocity, given by

$$\bar{v}(t) = \frac{1}{F} \sum_{i=1}^F v^{(i)}(t), \quad \text{where} \quad (13.1)$$

$$v^{(i)}(t) = \left(\sum_{p=1}^N \left\| \mathbf{X}_p^{(i)}(t) - \mathbf{X}_p^{(i)}(t + 0.05) \right\|^2 w_p \right)^{1/2},$$

²Another possibility is to assume that the fiber relaxation length scale is governed by the distance between two CLs on each fiber. Since there are an average of 24 CL connections on each fiber (12F CLs, each of which binds to 2 fibers), this gives a distance $\ell_c = 2/24 \approx 0.08 \mu\text{m}$, and a time scale $\mu\ell_c^4/\kappa$ of about 0.005 seconds, which is much faster than that measured in our numerical experiments.




Figure 13.1: Decay of the mean fiber velocity to a steady state after being sheared for one cycle with $\omega = 0.2\pi$ (permanent cross linkers). We plot the mean fiber velocity, given by (13.1) and normalized by $\bar{v}(t = 0)$, over five seconds and compare the result to the double-exponential fit $0.64e^{-t/0.36} + 0.36e^{-t/2.39}$ to estimate the network relaxation time scale τ_r .

and normalize by $\bar{v}(0)$ to obtain the exponential-like decay shown in Fig. 13.1. The time scale of relaxation to steady state is $\tau_r \approx 0.5 - 2$ s, with the best fit being a sum of two exponentials $\tau_1 = 0.4$ s and $\tau_2 = 2.4$ s. These relaxation time scales are in line with our physical estimates of $\tau_F = 0.2$ s and $\tau_{CL} = 0.5$ s.

Thus in order to measure the steady state moduli, we wait one second or one cycle (whichever is longer) prior to measuring the stress (and moduli) over three cycles of shear flow. We use a maximum strain $\gamma_0 = \dot{\gamma}_0/\omega = 0.1$ to stay in the linear regime (data not shown), and we give frequencies in Hz. To initialize for a given ω , we use the final network configuration from the previous (next smallest) frequency. A sample configuration of the fibers in the network, taken with $\omega = 1$ Hz at the point of maximum strain, is shown in Fig. 13.2. Because the CLs are permanent here, they do not have the chance to thermally ratchet the fibers into bundles, and so the network architecture is relatively homogeneous.

Figure 13.3(a) shows the steady-state elastic and viscous moduli when the dynamics of the network are computed with fully nonlocal hydrodynamics. Figure 13.3(a) also shows




Figure 13.2: Steady state fiber configurations for the network of 700 fibers and 8400 CLs with $\omega = 1$ Hz and $g = 0.1$. We show (left) a three-dimensional snapshot of all the fibers in the unit cell, (middle) a view along the z axis, and (right) a snapshot of all the fibers (white/blue) and links (black) bound to a single fiber (orange) located near the center of the simulation cell.

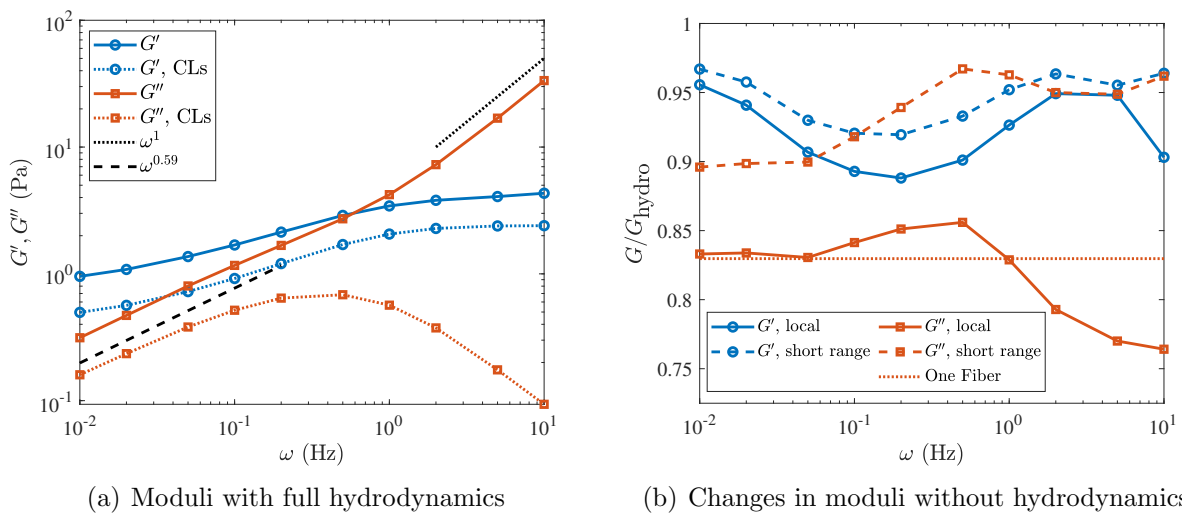


Figure 13.3: Elastic and viscous moduli for a permanently cross-linked network of 700 fibers and 8400 CLs. (a) Values of the elastic modulus G' (blue circles) and viscous modulus G'' (orange squares, this excludes the piece due to the background fluid stress $\sigma_{21}^{(\mu)}$ defined in (11.6)) for a gel of 700 fibers and 8400 CLs. We use fully nonlocal hydrodynamics to compute the dynamics and moduli. We show the contribution of the CL stress (11.9) to the moduli as dotted lines. (b) Fraction of G' (blue circles) and G'' (orange squares) that can be recovered when the dynamics are computed using local drag (solid lines) and short-ranged hydrodynamics only (i.e., only intra-fiber but no inter-fiber hydrodynamics). We also show the scaling of the stress for a single fiber in a shear flow as a dotted orange line.

separately the contribution to the moduli of the CL stress (11.9), which is 5 – 10 times more elastic than viscous. There is a clear transition in both of the moduli for $\omega \approx 0.5 - 1$ Hz which can be understood using the characteristic time scales in the problem. Since the characteristic relaxation time scale $\tau_r \approx 0.5 - 2$ s, the behavior of the moduli can be divided into three regimes: low frequency (background flow time scale $\tau_\omega \gg \tau_r$), medium frequency ($\tau_\omega \approx \tau_r$), and high frequency ($\tau_\omega \ll \tau_r$), where $\tau_\omega = \omega^{-1}$.

1. In the low frequency regime ($\omega < 0.1$ Hz), τ_ω is the longest time scale and the system is in a constant quasi-static state. If the frequency is small enough, the network has the opportunity to relax at every instant, and it therefore behaves more like an elastic solid where the links constrain the network. As in an elastic solid, the elastic modulus is unchanged with frequency and changes very little (less than 10%, as shown in Fig. 13.3(b)) when nonlocal hydrodynamics is dropped. In this regime, the viscous modulus scales like $G'' = \omega^{0.59}$; the reason for this particular scaling is not obvious to us.
2. In the mid frequency regime ($0.1 \leq \omega \leq 1$), fibers and CLs can deform and relax on the time scale of the background flow and the dynamics involve both an elastic and viscous response. In this regime, $G'' \approx G'$, as shown in Fig. 13.3(a), and the change in the elastic modulus G' due to both changes in frequency (Fig. 13.3(a)) and the inclusion of hydrodynamics (Fig. 13.3(b)) is maximal.
3. In the high frequency regime ($\omega > 1$), the background flow dominates the dynamics, the network is essentially fixed on the time scale of the shear flow, and it oscillates back and forth as a viscous fluid would. Figure 13.3(a) shows that for $\omega \gg 1/\tau_r$, the viscous modulus scales linearly with ω , as would happen for a pure viscous fluid. In this regime, the viscous modulus decreases by as much as 25% when nonlocal hydrodynamics is not included, with higher frequencies giving larger decreases. The larger the frequency, the farther the network is from its quasi steady-state and the more important dynamics

are for determining the viscous modulus.

Generally speaking, the changes in the viscous modulus when hydrodynamics is switched off are attributable to a reduction in stress on each fiber individually. The changes in the viscous modulus at most frequencies can in fact be explained by considering an isolated straight fiber that makes an angle θ with the x axis, $\boldsymbol{\tau}(s) = (\cos \theta, \sin \theta, 0)$. We put a background shear flow on the fiber $\mathbf{u}_0(x, y, z) = (y, 0, 0)$ and compute the resulting constraint force density $\boldsymbol{\lambda}$ on the fiber using both local drag and nonlocal hydrodynamics. Averaging over θ , we obtain a mean difference in the corresponding stress of $\sim 17\%$, which is plotted as a dotted line in Fig. 13.3(b), and matches the change in the network's viscous modulus when nonlocal hydrodynamics is dropped, except at the largest frequencies. This implies that the change in the network's viscous modulus comes primarily from an increase in the stress on each fiber individually when the intra-fiber, but not necessarily inter-fiber, nonlocal hydrodynamics is included in the mobility calculation. We confirm this in Fig. 13.4, which shows that including nonlocal hydrodynamics causes the magnitude of the constraint forces $\boldsymbol{\lambda}$ on the fibers in the network to increase without changing their positions substantially.

While Fig. 13.3(b) shows that there is a change in the viscous modulus of as much as 25% when the dynamics are computed by local drag, it also shows that including short-ranged hydrodynamics only, i.e., including intra-fiber but not inter-fiber hydrodynamics, gives a viscous modulus that is at least 90% of the one computed with inter-fiber hydrodynamics. The changes in the elastic modulus when nonlocal hydrodynamics is dropped are smaller (at most 10%), but not explainable simply by adding only intra-fiber hydrodynamics. Indeed, for $\omega = 0.5$, when the change in the elastic modulus due to hydrodynamics is near maximal, we find that the short-ranged intra-fiber hydrodynamics explains only one third of the change, with the rest coming from inter-fiber hydrodynamics. Physically, the elastic modulus is related to the interactions of the fibers with CLs, and the network is sufficiently connected that the CLs transmit stress across multiple links in the network, and so long-




Figure 13.4: Subset of fibers inside the unit cell, colored by the L^2 norm of the constraint forces $\lambda^{(i)}$ for $\omega = 1$ Hz at $t = 3.5$ s ($g = 0$). Normalization is the maximum L^2 norm of $\lambda^{(i)}$ in the system, which is on the fiber indicated by the arrow. We show the case of local drag at left and nonlocal hydrodynamics at right, noting that the inclusion of nonlocal hydrodynamics increases the norm of the constraint forces (which is the dominant cause of the increase in stress), without changing the fiber positions significantly.

ranged hydrodynamics plays a role.

Our main findings for permanent networks can be summarized as follows: there exists *one* critical time scale τ_r , which is on the order of a second for our parameters. On time scales longer than τ_r , the CLs are constantly in a steady state, as they would be in an elastic solid, and the network is more elastic than viscous ($G' > G''$). On time scales shorter than τ_r , the network does not have time to respond to (penalize) strain deformations, and the network is more viscous than elastic ($G'' > G'$). On time scales comparable to τ_r , the network is equally viscous and elastic. The more viscous the network behavior and the farther the network is from a quasi-steady state, the more nonlocal hydrodynamics impacts the moduli. Indeed, if a modulus G' or G'' changes substantially with changes in frequency, then we expect dynamics, including whether they are computed with nonlocal or local hydrodynamics, to matter. Interestingly, we find that intra-fiber nonlocal hydrodynamics dominates the contributions

System	τ_f (s)	Varied params	\bar{C}	\bar{B} (μm^{-3})	% in B
Meshwork	5	—	5.8	1.0	8.5
Meshwork, $\kappa/10$	5	$\kappa = 0.007$	5.6	0.93	8.1
B-In-M	10	—	12.5	2.2	23
B-In-M, $2k$	5.7	$k_{\text{on}} = 10, k_{\text{on},s} = 100,$ $k_{\text{off}} = k_{\text{off},s} = 2$	12.5	2.1	21
B-In-M, 10μ	16	$\mu = 1$	12.2	1.8	17

Table 13.1: Description of the systems we consider with transient CLs and their respective computed mean link and bundle densities. These data are from rheology simulations using $\omega = 1$. Here \bar{C} is the average number of links attached to each fiber, \bar{B} is the bundle density in bundles per μm^3 , and “% in B” is the percentage of the fibers that are in a bundle in the dynamic steady state.

of nonlocal hydrodynamics to the viscous, but not the elastic, modulus.

In the next section, we will see how the fundamental system timescales change when the cross linkers are transient (new timescale = rate of breaking) and the fibers are turning over (new timescale = polymerization rate). Our expectation is that there will be at least three timescales in this new system, which we will see is roughly the case.

13.2 Steady states with transient CLs

We first show that varying the turnover time τ_f of the F fibers creates a set of dynamic steady states with varying degrees of bundling. We also verify that the statistics are independent of the system size by considering three possible cubic periodic unit cells of different edge length L_d , all with the same mesh size (computed assuming that the fibers are distributed isotropically) $\ell_m = \sqrt{L_d^3/(FL)} \approx 0.2 \mu\text{m}$. The fibers have length $L = 1 \mu\text{m}$ throughout, so we will consider domains of size $L_d = 2 \mu\text{m}$ ($F = 200$ fibers), $L_d = 3 \mu\text{m}$ ($F = 675$ fibers), and $L_d = 4 \mu\text{m}$ ($F = 1600$ fibers). We will show that simulations of a homogeneous meshwork ($\tau_f = 5$ seconds) are suitably carried out in the smallest of the three systems, while runs where bundling is more considerable require the next-largest system ($L_d = 3 \mu\text{m}$).

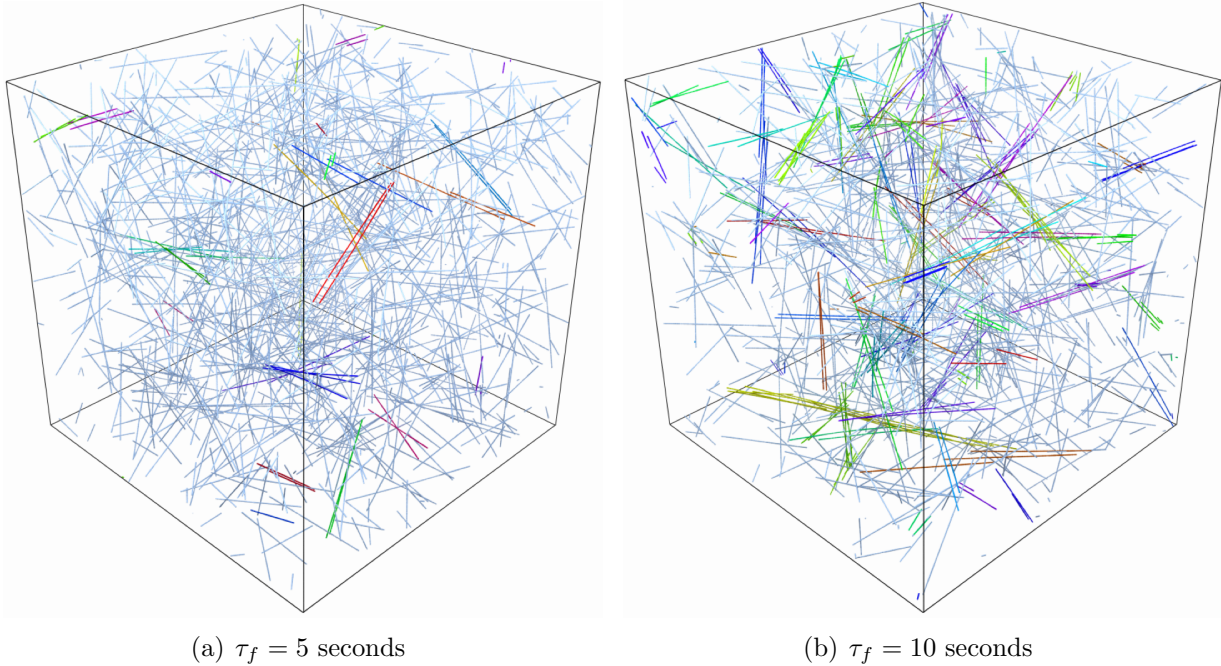


Figure 13.5: The actin gel for (a) $\tau_f = 5$ seconds and (b) $\tau_f = 10$ seconds, both shown on a domain with edge length $L_d = 3 \mu\text{m}$. Fibers in the same bundle are colored with the same color. In (a), we observe a more homogeneous mesh. In this case, we use $L_d = 2 \mu\text{m}$ in most simulations. In (b), we observe multiple bundles embedded into a mesh, and we use $L_d = 3 \mu\text{m}$ in most simulations.

13.2.1 Homogeneous meshwork morphology

Beginning with turnover time $\tau_f = 5$ s, we perform a steady state run to $10\tau_f = 50$ s using the three domain sizes. As shown in Fig. 13.5(a), the network is made primarily of fibers oriented isotropically (as can be verified by computing an order parameter of the type [211, Eq. (21)]), with a few bundles of at most two to three aligned fibers. We quantify this more precisely in Table 13.1, where we see that there is on average one bundle per μm^3 , and that less than 10% of the fibers are in bundles. Because there are only a small number of bundles and no permanent structures over long timescales in this system, we classify this system as a *homogeneous meshwork*. We will report results for it using a domain size of $L_d = 2 \mu\text{m}$ (we have verified that using $L_d = 3 \mu\text{m}$ does not change the results significantly; see Fig. 13.6).

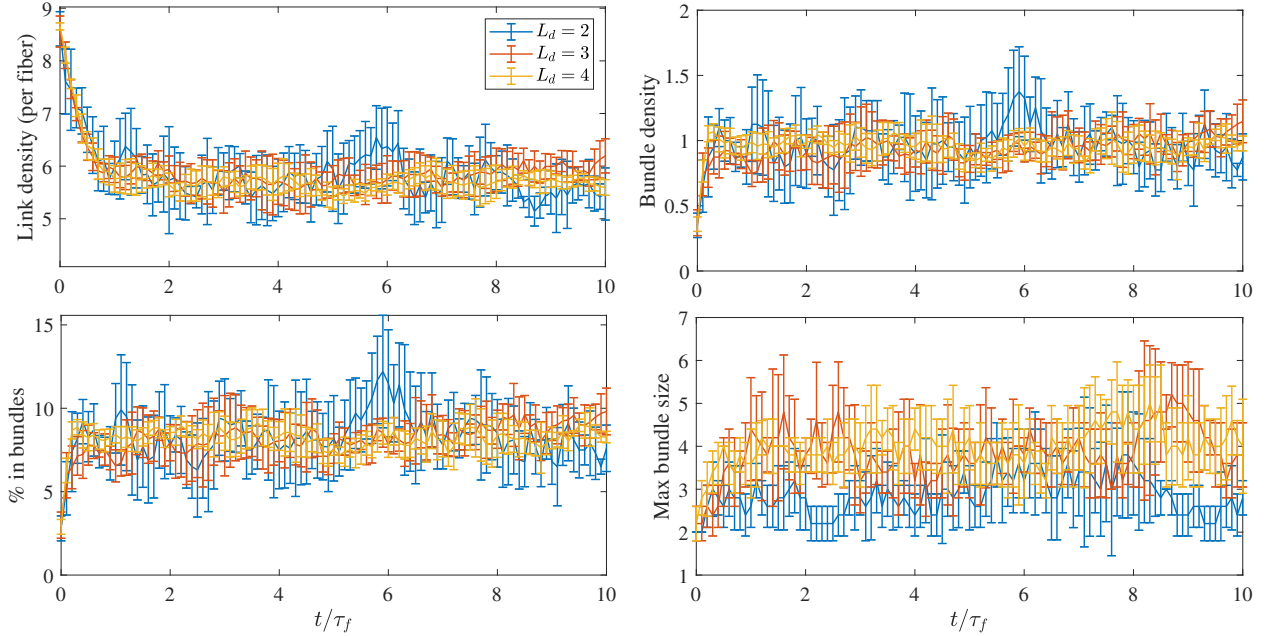


Figure 13.6: Number of links per fiber (top left), bundle density (top right, number of bundles / L_d^3), % of fibers in bundles (bottom left), and maximum bundle size (bottom right) for simulations from an initialized isotropic state with to the dynamic steady state for $\tau_f = 5$ seconds using three different system sizes: $L_d = 2$ (200 fibers, blue), $L_d = 3$ (675 fibers, red), and $L_d = 4$ (1600 fibers, yellow). Error bars are two standard errors over five trials.

13.2.2 Bundles embedded in meshwork morphology (B-In-M)

In Fig. 13.5(b), we show snapshots from the dynamic steady state with a (doubled) turnover time of $\tau_f = 10$ s. We observe significantly more bundling and fiber alignment, as well as bundles with several (four to six) fibers in them. In Table 13.1, we see that the steady-state link density has more than doubled from the 5 s turnover case, and that the percentage of fibers in bundles has gone from about 10% for $\tau_f = 5$ s to 25% for $\tau_f = 10$ s. The fluctuations of link density in the smaller system ($L_d = 2$) are quite large (about 20%, see Fig. 13.7), which makes averaging too inaccurate. For this reason, for $\tau_f = 10$ seconds we run a larger system with $L_d = 3 \mu\text{m}$, which we show in Fig. 13.5(b). We use the $L_d = 4$ system with 1600 fibers to verify that the finite system size has little effect; see, for example, Table 13.2.

Because this system has a significant number, but not more than half, of the fibers in bundles, and because the maximum bundle size is still at most ten fibers, we refer to this

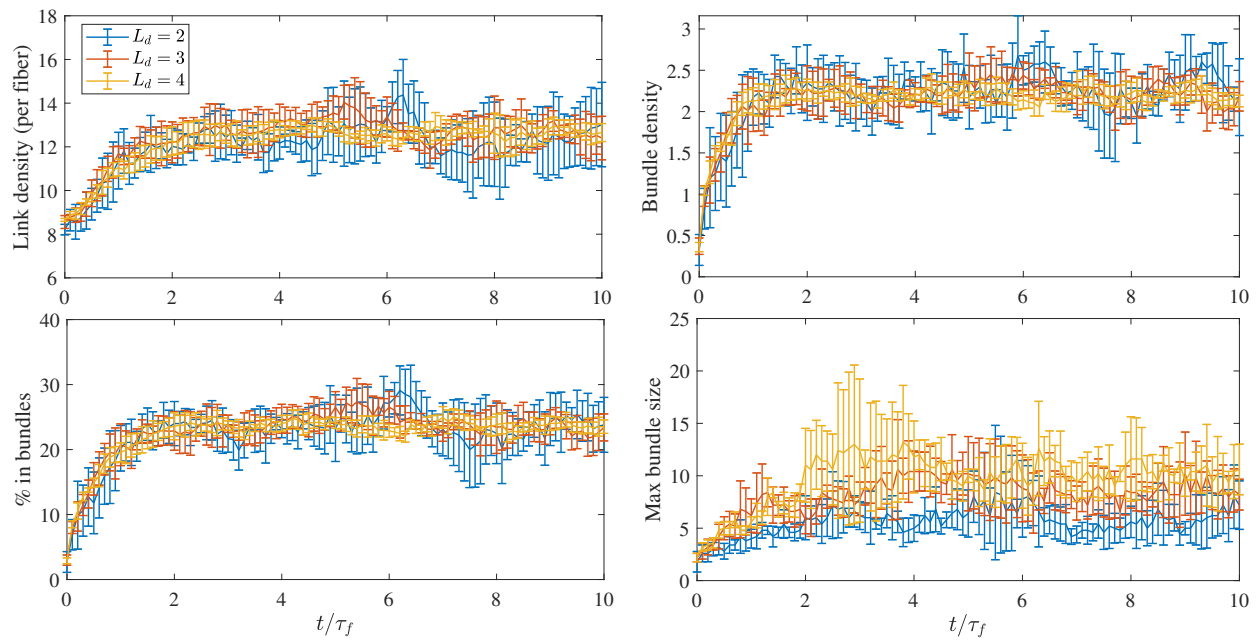


Figure 13.7: Number of links per fiber (top left), bundle density (top right, number of bundles / L_d^3), % of fibers in bundles (bottom left), and maximum bundle size (bottom right) for simulations from an initialized isotropic state with to the dynamic steady state for $\tau_f = 10$ seconds using three different system sizes: $L_d = 2$ (200 fibers, blue), $L_d = 3$ (675 fibers, red), and $L_d = 4$ (1600 fibers, yellow). Error bars are two standard errors over five trials.

system as a *bundles embedded in meshwork (B-In-M)* morphology, where there are small bundles of a few fibers embedded in an otherwise homogeneous mesh.

13.2.3 Other systems (parameter variations)

In addition to varying the morphology through the turnover time, we also look at systems with the same morphologies, but different microscopic parameters. Table 13.1 gives our description of these systems. For a homogeneous meshwork, we consider $\tau_f = 5$ s again, but this time with ten times smaller bending stiffness, so that $\kappa = 0.007$ pN $\cdot\mu\text{m}^2$. Table 13.1 shows that the morphology in this case is the same as the homogeneous meshwork with $\kappa = 0.07$ pN $\cdot\mu\text{m}^2$, so changing the bending stiffness has a minimal impact on the morphology. Indeed, the fibers remain relatively straight even with the smaller bending stiffness (see Fig. 13.8), just as they are straight in both the B-In-M and homogeneous meshwork systems in Fig. 13.5. This implies that the degree of cross-linking/network connectedness is insufficient to induce filament bending, even in B-In-M geometries. The reason for this is the fibers that are heavily cross linked are in parallel bundles, and the other fibers have so few cross linkers attached that bending is not energetically favorable.

We also consider B-In-M systems with larger viscosity, $\mu = 1$ Pa \cdot s instead of 0.1, and another system where we double the rates of link binding and unbinding. Table 13.1 shows how we vary the turnover time in these systems to obtain a B-In-M morphology. When the viscosity is increased by a factor of 10, we only need to increase the turnover time by a factor of 1.5 to obtain a similar morphology (match \bar{C} as closely as possible). Meanwhile, doubling the rates forces us to cut the turnover time almost in half to obtain the same morphology. This provides our first indication that the morphology, or the timescale on which the network organizes itself into bundles, is strongly dependent on the link binding and unbinding rates and weakly dependent on the underlying fluid viscosity. This combination of behavior suggests we are operating in a regime where the CLs are attached for long enough

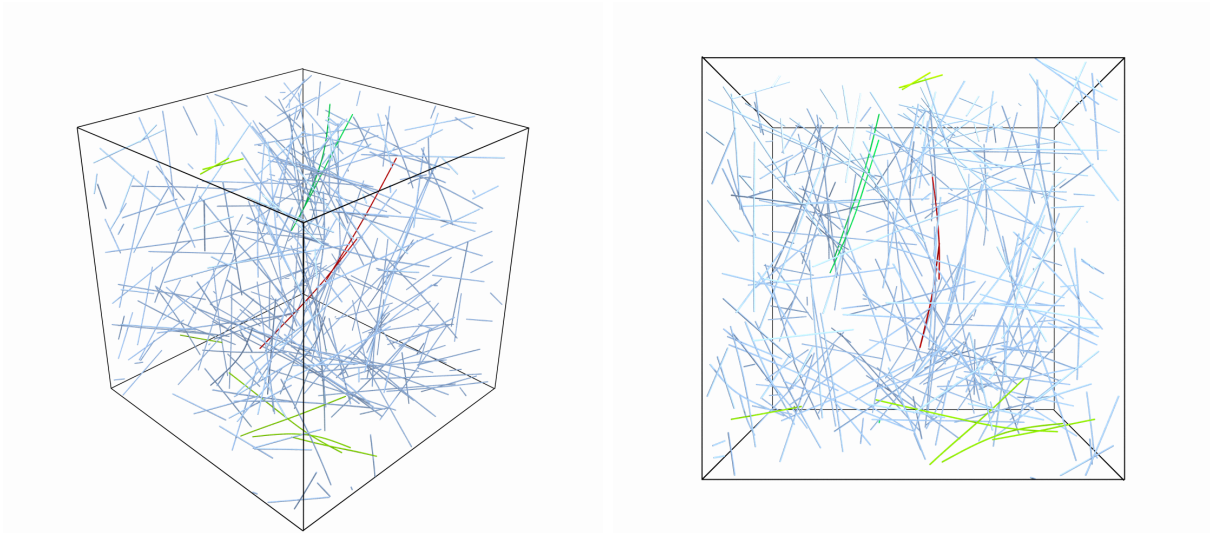


Figure 13.8: Side (left) and top (right) view of the homogeneous meshwork with $\kappa = 0.007 \text{ pN}\cdot\mu\text{m}^2$, which is 1/10 of the stiffness of actin. Despite the smaller bending stiffness, the fibers still appear relatively straight, with the exception of some of the fibers in bundles.

to move the fibers into a quasi-steady state (since using twice the link turnover rate speeds up the process by a factor of two and making the dynamics slower by changing viscosity has little effect).

13.2.4 The linear response regime (LRR)

Our goal for the rheology experiments is to measure the shear moduli of the network in its dynamic steady state, and in the linear response regime (LRR). In this section we briefly demonstrate how to find the LRR and what happens when we go beyond it, using the B-In-M geometry as an example.

Table 13.2 shows the steady state statistics and mean viscoelastic moduli for the B-In-M system with varying strains γ . In Table 13.2 and throughout this paper, we report the viscous modulus G'' *without* the contribution from the viscosity of the Newtonian solvent $G''_{\text{bkgnd}} = 2\pi\omega\mu$. We observe from the table that the LRR is $\gamma \lesssim 0.05$, since using $\gamma > 0.05$ leads to significantly larger moduli. In particular, Table 13.2 shows that larger strains of $\gamma \geq 0.10$ induce higher bundle densities and the formation of more links. Furthermore, the

L_d (μm)	γ	\bar{C}	B (μm^{-3})	% in B	G' (Pa)	G'' (Pa)
3	0	12.6	2.3	24.1		
3	0.01	12.5	2.2	23.2	0.37	0.64
3	0.025	12.5	2.2	23.5	0.36	0.63
3	0.05	12.8	2.2	23.9	0.39	0.67
3	0.10	13.2	2.3	25.1	0.52	0.77
3	0.15	13.9	2.4	26.6	0.79	0.96
3	\pm	0.4	0.1	0.8	0.04	0.03
4	0.025	12.7	2.2	23.8	0.37	0.66
4	0.05	12.8	2.2	24.0	0.42	0.70
4	\pm	0.1	0.1	0.5	0.02	0.02

Table 13.2: Network structure statistics for the runs over 120 seconds for the B-In-M network with $\omega = 1$ Hz. We give the mean for each statistic, with uncertainties on the measurements (\pm row) equal to two standard deviations across five trials. The uncertainties in the viscoelastic moduli for $\gamma = 0.15$ are larger, $G' = 0.79 \pm 0.09$ Pa and $G'' = 0.96 \pm 0.12$ Pa. The comparable values of the statistics and moduli for $L_d = 3 \mu\text{m}$ and $L_d = 4 \mu\text{m}$ demonstrate that the finite-size errors are smaller than the statistical errors.

links that do form in the larger-strain systems are themselves more strained, as the mean square link strain at $\gamma = 0.1$ and $\gamma = 0.15$ is 10% and 25% higher, respectively, than in the linear regime. The combination of these two factors gives large increases in the moduli when we reach the nonlinear regime. The viscoelastic moduli for larger strains are also subject to significantly larger uncertainties (e.g., the uncertainty for $\gamma = 0.15$ is at least double the uncertainty for $\gamma = 0.1$). The larger uncertainties come from fluctuations in the underlying microstructure, as the fluctuations in the link density increase by 50% from $\gamma = 0.01$ to $\gamma = 0.15$.

Another way we can confirm that we are in the LRR is by looking at the Fourier spectrum of the stress. Specifically, we will write the stress as

$$\sigma(t) = \sum_{k=-\infty}^{\infty} \hat{\sigma}(k) e^{2\pi i k \omega t} \quad (13.2)$$

and use the discrete Fourier transform to look at the amplitude of the coefficients $\hat{\sigma}(k)$ for various γ . In Fig. 13.9, we show the spectrum $|\hat{\sigma}(k)|$ for integer values of k with $\tau_f = 10$




Figure 13.9: Spectrum of the stress for the B-In-M network. The coefficients $\hat{\sigma}(k)$ are defined in (18) in the main text. We consider various maximum strain values $\gamma = 0.01$ (green), 0.025 (purple), 0.05 (yellow), 0.1 (orange), and 0.15 (blue). We see the emergence of a $k = 3$ nonlinear harmonic for larger strains, suggesting a nonlinear response.

s, where we observe a peak at the obvious location of $k = 1$. It is the emergence of higher harmonics for larger strains, however, that tells when we leave the LRR. From Fig. 13.9, we see that $\gamma = 0.1$ is definitely *not* in the linear regime for the B-In-M geometry, since it has a $k = 3$ Fourier coefficient that is only about ten times smaller than the one for $k = 1$. Increasing γ to 0.15 , we see an even stronger response from the $k = 3$ harmonic, and the emergence of the $k = 2$ harmonic as well.

13.2.5 Timescales of stress relaxation

We first seek to gain some understanding of the timescales in the system with a stress relaxation test, snapshots from which are shown in Fig. 13.10. We simulate for 0.25 seconds using $\omega = 1$ Hz, so that the system ends at its maximum strain γ (the results are largely independent of both the strain and shear rate used to obtain it). In Fig. 13.11, we show the decay of the stress, with $t = 0$ denoting the point at which the system reaches the maximum strain. In addition to our dynamic network geometries, for which we average over 30 trials,

we consider *permanently* linked, interconnected networks of the type considered in Section 13.1, where $\bar{C} = 25$ ensures the network is well above the rigidity percolation threshold (over 95% of the fibers are in one connected component, where two fibers are connected if they are linked by at least one CL). These networks, for which we *do not* include fiber turnover, give smoother stress profiles, and so we only average over 5 trials.

Figure 13.11 gives us two pieces of information that are important for our analysis going forward. First, we observe that statically-linked networks can store elastic energy, since $g_0 := \sigma_{21}(t \rightarrow \infty) > 0$, while all of the energy dissipates in transiently-linked networks. Second, both permanent and dynamic networks have multiple relaxation timescales, as most of the stress ($\approx 80\%$ for dynamic networks) relaxes over the first 0.2 s or so, with the remaining part taking seconds to relax. In Fig. 13.11, we show a two-timescale decay curve that approximately, but not exactly, matches the decay of the stress in both statically- and transiently-linked networks. Given that one of the time constants is on the order 0.1 s and the other is on the order 1 s, Fig. 13.11 demonstrates that both slow and fast timescales exist in cross-linked networks. Our rheological experiments will give more precise estimates for the individual timescales, but it is worth noting that the dissipation of *all* of the stress over a $\mathcal{O}(1s)$ timescale implies that the longest timescale in the system cannot be much longer than five seconds. This is not surprising since the filament turnover time is of this order, but is in sharp contrast to modeling studies without turnover, where stress takes hundreds of seconds to dissipate [218, Fig. 4(a)].

One observation we can make from Fig. 13.11 is that the long-time behavior is similar for all dynamic networks, so the long timescale behavior is roughly independent of viscosity. We made this latter observation once before when we saw that a similar turnover time can be used to generate a B-In-M morphology for larger viscosity systems. We will look at the short-timescale behavior more closely in our rheological tests, which we discuss next.

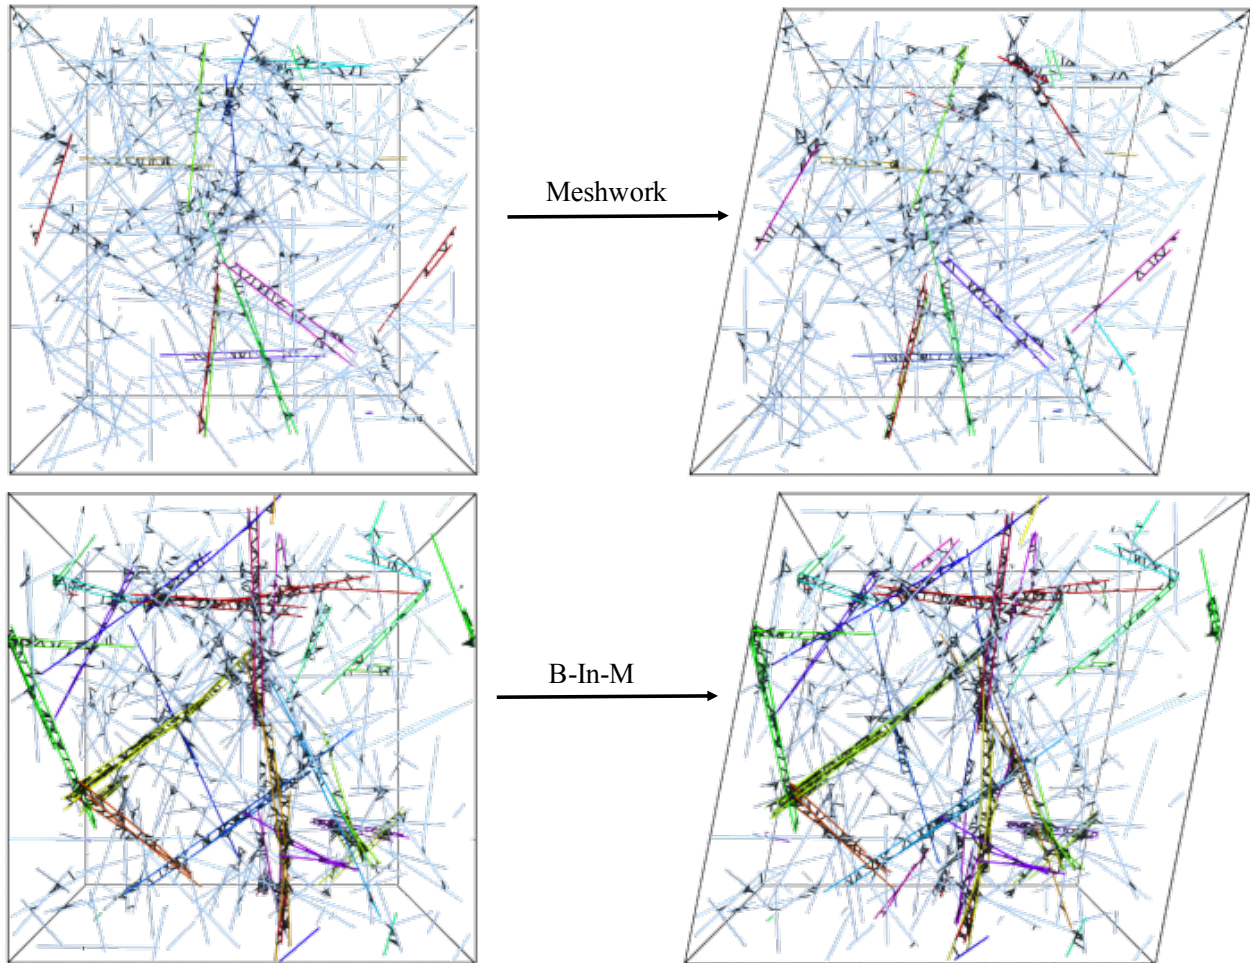


Figure 13.10: Snapshots from the stress relaxation test in a homogeneous meshwork (top) and B-In-M geometry (bottom). We begin with an un-sheared unit cell at left, then shear the network until it reaches a maximum strain (20% in this case, shown at right), after which we turn off the shear and measure the relaxation of the stress. As in Fig. 13.5, the colored fibers are in bundles, and the CLs are shown in black. For the B-In-M geometry, these snapshots are from a smaller domain ($L_d = 2 \mu\text{m}$) than we typically use so that we can also see the CLs.




Figure 13.11: Normalized stress profiles over time in the stress relaxation test. We consider four different systems: a permanent network (blue), for which the stress relaxes to a nonzero value ($g_0 \approx 0.07$ Pa after accounting for normalization), and three dynamic networks, for which the stress relaxes to a value of zero (orange is the homogeneous meshwork, yellow the B-In-M morphology, and purple the B-In-M morphology with ten times larger viscosity). To illustrate the point that there are multiple intrinsic relaxation timescales in the system, we show a double-exponential curve which approximately matches the decay of stress for both permanent and transient CLs.

13.3 Viscoelastic moduli

In Fig. 13.12, we plot the elastic and viscous moduli as a function of frequency for the five different systems defined in Table 13.1. We also include the data for a permanent, interconnected network of the type considered in Section 13.1, which we show with and without the constant elastic element $g_0 = G'(\omega \rightarrow 0)$ determined in Fig. 13.11. For the system with ten times larger viscosity, we show the data with time/frequency rescaled by a factor of ten, so that the value for $\omega = 1$ Hz is mapped to $\omega = 10$ Hz. These data are for the local drag model ((2.24) without the finite part integral), since it is faster to simulate; we will consider the effect of nonlocal hydrodynamics in Section 13.4.

Examining the data, we make the following observations:

- On very short timescales (shorter than 0.05 s, or frequencies larger than 20 Hz), the viscous moduli all approach the viscous-fluid scaling of $\approx (0.13 \text{ Pa}\cdot\text{s}) \omega$. The timescale on which this occurs is directly proportional to the viscosity, as the rescaled G'' curve for the larger viscosity system makes the transition at the same time as the other B-In-M curves. In addition, the elastic moduli for B-In-M networks are about twice as large as those of homogeneous meshwork curves at high frequency, so the elastic modulus at short times is proportional to the link density.
- On very long timescales (several seconds, low frequencies), the viscous modulus again scales linearly with frequency, but with a larger slope ($G'' \approx (0.3 \text{ Pa}\cdot\text{s}) \omega$), which indicates that the links have become viscous on those timescales. The elastic modulus is much smaller than the viscous modulus on these timescales, and both of the moduli scale nonlinearly with the link density, as B-In-M morphologies have moduli about 4 times as large as homogeneous meshworks.
- There is an intermediate timescale of about 1 second on which the moduli for the B-In-M meshwork (yellow) diverge from those for the system with twice the bind-

ing/unbinding rate (purple), and on which both of these curves diverge from the light blue permanent network curve. In particular, for $\omega \leq 2$ Hz the elastic and viscous modulus are both smaller for the system with faster link turnover. This timescale is related to (in fact, it is exactly equal to) $1/k_{\text{off}}$, the timescale on which individual links bind and unbind. The rescaled curve with larger viscosity also diverges from the other B-In-M curves on this intermediate timescale, which indicates that a timescale has been introduced (by the dynamic linking) that does *not* scale with viscosity.

We will explore each of these fast, slow, and intermediate timescales in subsequent sections. The picture we sketch is of a rigid network for timescales shorter than the fastest intrinsic timescale $\tau_1 \sim 0.1$ s; on these timescales the links are predominantly elastic and the viscous modulus is the same as if the fibers were in a fluid without CLs. On timescales longer than τ_1 but shorter than the intermediate timescale $\tau_2 \sim 1$ s, the links still appear permanent, but they provide additional viscosity by deforming the network to its steady state in response to flow deformations. On timescales longer than τ_2 , but shorter than the longest timescale $\tau_3 \sim 5$ s, the links come off before they are able to fully respond to the deformations induced by the background flow, and the slope of the elastic and viscous modulus changes relative to the one observed in permanent networks. Finally, on timescales longer than τ_3 , the network is almost completely viscous, and the viscosity is controlled by the morphology.

13.3.1 Short timescales: elastic links and viscous fibers

Our hypothesis is that the shortest timescale in the system, τ_1 , dictates when the links make a negligible contribution to the viscous modulus. For timescales $\tau < \tau_1$, the network is essentially frozen: the fibers are rigid, the links are static, and the fibers and links contribute exclusively to the viscous and elastic modulus, respectively. Specifically, as $\omega \rightarrow \infty$, we expect that the links will make a constant contribution to G' and that G'' will scale like $\eta_0\omega$, i.e., as a viscous fluid, with the viscous coefficient η_0 independent of the number of links




Figure 13.12: Elastic (left) and viscous (right) modulus for the five systems described in Tables 11.1 and 13.1 (see Table 13.1 for what parameters are varied in each system). For the B-In-M morphology with ten times the viscosity ($\mu = 1$ Pa·s), we show the data with time rescaled by a factor of ten. We also show the results for permanently-cross-linked networks, and for the elastic modulus the dashed light blue curve is the remaining elastic modulus when $g_0 \approx 0.07$ Pa (measured in Fig. 13.11) is subtracted off.

in the system. In Fig. 13.12, we show that $\eta_0 \approx 0.13$ Pa·s if ω is given in Hz. We have confirmed that this scaling continues for frequencies all the way up to 1000 Hz.

To verify that the viscous modulus at short timescales is dominated by the fibers only (independent of the number and nature of the links), we compare to the theoretical shear viscosity enhancement η_0^* for an isotropic suspension of rigid cylindrical fibers, where the mobility is computed by the local drag approximation. For dilute suspensions, the enhanced viscosity is given to order $(\log \epsilon)^{-2}$ in terms of the fiber density f and aspect ratio ϵ as [82, 239, 240]

$$\frac{\eta_0^*}{\mu} = \frac{\pi f}{45} \left(-\frac{1}{2 \ln \epsilon} \left(\frac{-\ln \epsilon + 0.640}{-\ln \epsilon - 0.5} \right) \right). \quad (13.3)$$

For our parameters ($\epsilon = 0.004$, $f = FL/L_d^3 = 25$), we obtain $\eta_0^*/\mu \approx 0.194$ for cylindrical fibers.³ Recalling that $G'' = 2\pi\omega\eta_0$ for a viscous fluid, our scaling in Fig. 13.12 shows that

³We obtained (13.3) by dropping the finite part integral in [82, Eq. (7.1)], which gives 1/2 instead of 3/2 in the denominator of [82, Eq. (8.13)]. Including intra-fiber hydrodynamics in the mobility changes the 0.5 in (13.3) back to 1.5. Substituting our parameters and recomputing, we get $\eta_0^*/\mu = 0.242$ with intra-fiber hydrodynamics, which means the local drag approximation gives only 80% of the correct viscosity (see Section 13.4).

we obtain a similar value of $\eta_0/\mu = 0.13/(2\pi \times 0.1) \approx 0.21$.

It is important to separate the viscous contribution from the *links*, which approaches zero at short timescales, from that due to the *fibers*, which is infinite as $\omega \rightarrow \infty$. For this reason, we define


$$G''_F = 2\pi\eta_0\omega \quad (13.4)$$

as the viscous modulus coming from the fibers. In Fig. 13.13(a), we examine the behavior of the moduli due to the links, G' and $G'' - G''_F$, rescaled by the average number of links in the system \bar{C} . After rescaling time in the system with larger viscosity, we see that the data for $\omega > 10$ Hz can be collapsed onto a single curve for both the elastic and viscous modulus. Thus on timescales shorter than $\tau_1 \sim 0.1$ s, each link behaves as an elastic element with strength independent of morphology. Note that the elastic modulus is more than 2 – 3 times larger than the viscous modulus (coming from the links) at high frequencies.

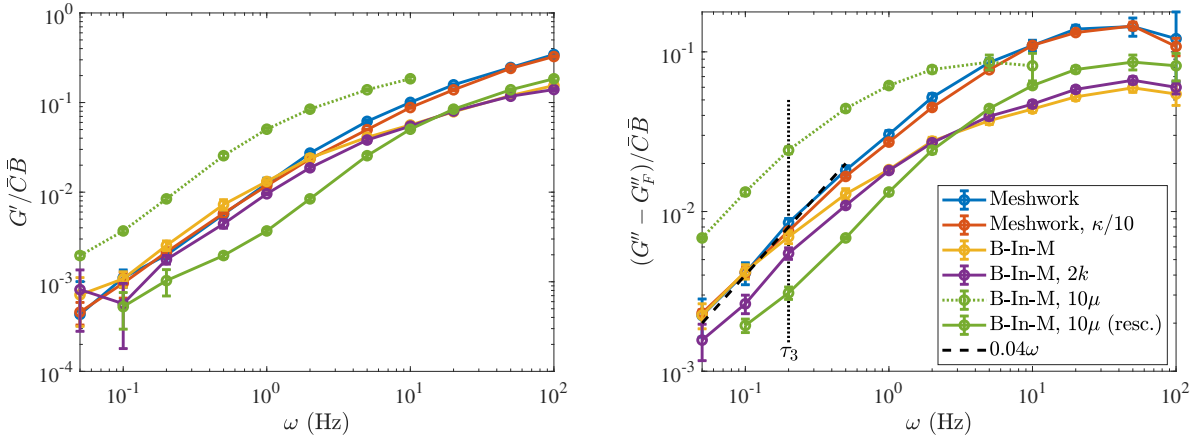
13.3.2 Viscosity on long timescales

Let us now transition to the opposite limit of long timescales. On long timescales, Fig. 13.13(a) shows that each system is about five times more viscous than elastic (e.g., compare the values at $\omega = 0.2$ Hz), and that normalizing by \bar{C} does not place the curves on top of each other. We say that these “long timescales” are longer than τ_3 , the longest intrinsic timescale in the system, or the timescale on which the network totally “remodels” itself. Our goal in this section is two-fold: first, we would like to find a new rescaling that better matches the curves at low frequencies, and second, we would like to estimate τ_3 , which we do by finding the regime in which G'' is a linear function of ω for each system.

Figure 13.13(b) shows our attempt to accomplish both of these goals. For both the viscous and elastic modulus, we rescale by the mean link density \bar{C} multiplied by the mean number of bundles \bar{B} , the idea being that bundles make a stronger contribution to the moduli at low frequency. We see a better match at low frequencies than in Fig. 13.13(a), which implies that



(a) Scaling by number of links (short timescales)



(b) Scaling by number of links and bundles (long timescales)

Figure 13.13: Elastic (left) and viscous (right) modulus *due to the CLs* for the five systems described in Tables 11.1 and 13.1 (see Table 13.1 for what parameters are varied in each system). To obtain a viscous modulus due to the links alone, we subtract the component due to the fibers G''_F defined in (13.4). (a) We normalize by the link density \bar{C} and see that the curves all match at short timescales on the order $\tau_1 \approx 0.02$ s (after we also rescale time by viscosity). (b) We normalize by the link density multiplied by the bundle density. For the system with ten times larger viscosity, we also include the raw data (not rescaled) as a dotted green line. In the viscous modulus plot on the right, we show a linear slope as a dashed black line and define $\tau_3 \approx 5$ s as the end of the low-frequency linear regime in G'' . This timescale is roughly independent of viscosity.

bundles have a stronger effect on the mechanical behavior at long timescales than at short ones. In particular, the homogeneous meshworks (blue and orange) and B-In-M (yellow) curves all scale onto the same curve for frequencies $\omega \leq \omega_3 \approx 0.2 \text{ Hz} = 1/\tau_3$, but the curves with larger viscosity (green) and faster link (un)binding (purple) do not scale as well. In both of these cases, there is a change in a timescale shorter than τ_3 , as we discuss in the next section.

To estimate τ_3 , we determine the regime where the viscous modulus G'' scales *linearly* with ω and equate τ_3 with the end of the linear region. From Fig. 13.13(b), we see that $\tau_3 \approx 5$ seconds, and that the linear regime endures longer for homogeneous meshworks, which implies that they start becoming purely viscous at shorter timescales than systems with bundles (or, equivalently, these networks “remodel” themselves faster).

To understand exactly how much viscosity is being provided by the links and bundles in the different systems, we use the slope in Fig. 13.13(b) to obtain

$$G''_{\text{CL}} := G'' - G''_F \approx 0.04\bar{C}\bar{B}\omega \quad \text{for } \tau > \tau_3. \quad (13.5)$$

Since \bar{C} has units links per fiber and \bar{B} has units bundles per μm^3 , 0.04 in this equation has units $\text{Pa}\cdot\text{s}/\mu\text{m}^3$. To extract a viscosity, we write $G''_{\text{CL}} = 2\pi\mu_{\text{CL}}\omega$, so that in the B-In-M system with $\bar{C} = 12.5$ and $\bar{B} = 2 \mu\text{m}^{-3}$ we obtain $\mu_{\text{CL}}/\mu = (0.04 \times 12.5 \times 2)/(2\pi \times 0.1) = 1.6$. This is eight times more viscosity than the fibers alone (0.2) and 1.6 times the background fluid viscosity. For the homogeneous meshwork with $\bar{C} = 5.75$ and $\bar{B} = 1 \mu\text{m}^{-3}$, the additional viscosity is $\mu_{\text{CL}}/\mu = (0.04 \times 5.75 \times 1)/(2\pi \times 0.1) = 0.4$. This is twice the viscosity of the fibers but 2.5 times smaller than that of the background fluid. Comparing the two morphologies confirms some of the prior intuition that bundles embedded in meshworks provide stronger resistance to flow than homogeneous meshworks [205].

13.3.3 Intermediate timescales

So far, we have addressed the two extreme limits of the suspension behavior. Over short timescales, $\tau < \tau_1$, the entire network is rigid, the CLs are purely elastic, and the additional viscosity is the same as that due to the fibers alone. Over long timescales, $\tau > \tau_3$, all of the energy in the network is dissipated. Bundles break up, there is no elastic modulus, and the morphology of the network, and the amount by which links can deform it, determines the viscous behavior.

There is also an intermediate timescale τ_2 at which *some of* the links, but not the entire network, start to behave viscously. While we expect $\tau_2 \sim 1/k_{\text{off}}$, this could change based on the network parameters, as the same value of k_{off} acts differently depending on the speed of network deformation. To better understand the timescale τ_2 , we compare the elastic modulus with dynamic links to the one for statically-linked networks (by starting from the same equilibrium structure and keeping the links fixed). We do this for five seconds or five cycles, and measure the modulus over the last three cycles.

Figure 13.14(a) shows the results for our B-In-M structures. The solid lines show the elastic modulus with dynamic linking, and the dotted lines show the elastic modulus with permanent linking. The timescale τ_2 is when the dynamic linking elastic modulus diverges significantly from the permanent link modulus, so that the links come off before they can provide their maximum elastic response. From Fig. 13.14(a), we see that τ_2 is obviously smaller when we increase the rate of link turnover (the purple curve diverges from the dotted yellow one faster than the solid yellow one does); this is expected from $\tau_2 \sim 1/k_{\text{off}}$, which is 1 s for the yellow curve and 0.5 s for the purple one.

Morphology has no impact on the timescale τ_2 , as we obtain the same characteristic time $\tau_2 \approx 1$ s for the homogeneous meshwork (not shown, because the results are indistinguishable) as the B-In-M morphology. Thus the morphology only appears to have a strong influence on the longest timescale τ_3 .

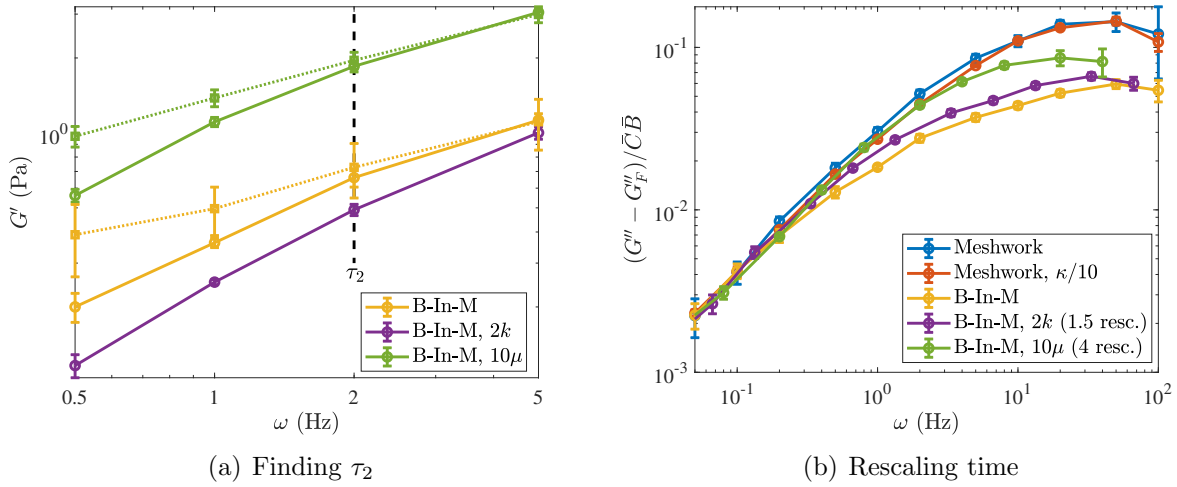


Figure 13.14: (a) Elastic modulus at medium frequencies. For each network type (indicated in the legend), we compare the solid line, which has the elastic modulus with dynamic links, to the dotted line, which is the elastic modulus with permanent links. The intermediate timescale $\tau_2 \approx 0.5$ s is the inverse of the frequency where the data start to diverge. We consider only bundle-in-mesh morphologies using our standard parameters (yellow), twice the link turnover rate (purple, but the static link reference is still the dotted yellow), and ten times the viscosity (green). (b) Rescaling time to get all of the viscous modulus curves on the same plot. This is the same plot as Fig. 13.13(b) (right panel), but now we rescale the time for the larger viscosity green curve by a factor of 4 instead of 10 ($\omega \rightarrow 4\omega$), and we rescale the time for the faster link turnover purple curve by 1.5 ($\omega \rightarrow \omega/1.5$). This demonstrates that the data do not scale simply with the parameters at medium and low frequencies, when multiple timescales are involved.

It is harder to determine the effect of viscosity on τ_2 . On the one hand, Fig. 13.14(a) shows the true data for G' vs. ω (without scaling time by μ), so it appears that the dynamic curve diverges from the permanent one at about the same timescale regardless of viscosity. On the other hand, for a fixed k_{off} the difference between the solid and dashed curves at $\omega = 1$ Hz is larger for larger viscosity (green). Thus the timescale τ_2 depends on the interaction of dynamic linking and deformation in a nontrivial way.

To illustrate this complicated dependency, we show that different rescalings are required when the viscosity (dynamics) or link turnover rates change from the base parameters. We recall from Fig. 13.13(b) that the larger viscosity (green) and faster link turnover (purple) B-In-M systems do not follow the same normalization as the other systems. Figure 13.14(b) shows that scaling time by a factor of 4 (and not the expected 10) puts the green larger viscosity curve onto the rest of the curves at low frequencies, while scaling time by a factor of 1.5 (and not the expected 2) puts the curve with twice the link turnover rate on top of the others as well. The reason for these particular scalings is not obvious to us, other than that they fall somewhere between the expected scaling and unity.

13.4 Role of nonlocal hydrodynamic interactions

We have seen that the network behaves very differently on short timescales, where each link behaves elastically independent of the network morphology, than on long timescales, where the system is purely viscous and the morphology exerts a significant influence on the resistance to flow. In this section, we show that nonlocal hydrodynamic interactions significantly decrease the elastic and viscous moduli in B-In-M morphologies. We show that the decrease is most significant at low frequencies, i.e., the frequencies where we already know morphology has a strong influence on the behavior. We explain the decrease in two ways: first, nonlocal hydrodynamic interactions slow down the bundling process, causing less bundles to form for a fixed turnover time, and second, they create entrainment flows that

reduce the stress inside of the bundles that do form.

We consider only the two characteristic morphologies in this section without varying any other parameters from those given in Table 11.1. For each of the morphologies, we compute the viscoelastic moduli with nonlocal hydrodynamics, then compute them again using local drag and intra-fiber hydrodynamics. In Fig. 13.15, we plot the error/fraction of the moduli recovered with the various mobility operators. For a homogeneous meshwork system (blue lines), we see that the elastic modulus is the same within 10 – 20% whether we use local drag, intra-fiber, or full hydrodynamics. At low frequencies, the viscous modulus is also the same within 10% regardless of the hydrodynamic model used, but there is an obvious error in the viscous modulus when the local drag model is used at high frequencies. This discrepancy of about 20% is also present in permanent networks and in B-In-M morphologies, and in all cases it can be mostly recovered by adding only *intra-fiber* hydrodynamics to the mobility. The 20% difference between local drag and intra-fiber hydrodynamics is similar to the theoretical estimate discussed in footnote 3, and to the results we obtained in Section 13.1. There appears to be a small increase (at most 5%) in the viscous modulus when we switch from intra-fiber hydrodynamics to fully nonlocal hydrodynamics, which would be more significant at smaller mesh sizes [239].

The more interesting changes with nonlocal hydrodynamics come at low frequencies in the B-In-M geometries. Unlike with the homogeneous meshwork, the local drag model makes about a 50% overestimation in the elastic modulus and a 30% overshoot in the viscous modulus at low frequencies for B-In-M morphologies, with the largest change coming at $\omega = 0.5$ Hz, or a timescale of 2 seconds. We have previously seen that morphology has a strong influence on the moduli at these timescales, so we explore the impact of nonlocal hydrodynamics on morphology next.




Figure 13.15: Proportion of the (left) elastic and (right) viscous modulus that is recovered using various mobility approximations. We compute the modulus using full hydrodynamics, then plot the fraction of it recovered using local drag or intra-fiber hydrodynamics. We show the results for the homogeneous meshwork in blue and the B-In-M morphology in orange. For each line color, a solid line shows the results for local drag and a dotted line shows the results for intra-fiber hydrodynamics. Intra-fiber hydrodynamics cannot explain the deviations in the elastic and viscous modulus at low frequency for the B-In-M geometry.

13.4.1 Rescaling of time cannot explain long-timescale moduli

The first possible explanation for the decrease in G' and G'' with nonlocal hydrodynamics at low frequencies is a change in the network structure. In this section, we show that, while there are less bundles and links in the dynamic steady state with full (intra- and inter-fiber) hydrodynamics for a given turnover time, this by itself cannot entirely explain the decrease in the moduli.

To do this, we focus on the frequency where full hydrodynamics matters the most as a percentage. For $\omega = 0.5$ Hz, we measure the moduli with intra-fiber hydrodynamics as $G'_{\text{IF}} = 0.18 \pm 0.02$ Pa, $G''_{\text{IF}} = 0.42 \pm 0.01$ Pa, while the moduli with full hydrodynamics are $G'_H = 0.127 \pm 0.004$ Pa, $G''_H = 0.35 \pm 0.01$ Pa. As shown in Table 13.3, the steady state link and bundle density are 12.3 and $2.2 \mu\text{m}^{-3}$, respectively, with intra-fiber hydrodynamics, while with full hydrodynamics they are 11.7 and $2.0 \mu\text{m}^{-3}$. Thus at its dynamic steady state, the network has on average fewer links and fewer bundles when we simulate with full hydrodynamics than when we simulate with intra-fiber hydrodynamics. This is because

τ_f (s)	Mobility	\bar{C}	\bar{B} (μm^{-3})	% in B	G' (Pa)	$G'' - \eta_0\omega$ (Pa)
10	H	11.7 ± 0.3	2.03 ± 0.04	20.9 ± 0.6	0.127 ± 0.004	0.27 ± 0.01
9.3	IF	11.5 ± 0.4	2.07 ± 0.07	21.2 ± 1.1	0.17 ± 0.01	0.33 ± 0.03
10	IF, μ^*	11.5 ± 0.3	2.01 ± 0.07	20.1 ± 0.8	0.19 ± 0.02	0.38 ± 0.03
10	IF	12.3 ± 0.2	2.15 ± 0.05	22.6 ± 0.7	0.18 ± 0.02	0.34 ± 0.01

Table 13.3: Statistics for system with $\omega = 0.5$ Hz and varying turnover times and mobility models (IF stands for intra-fiber hydrodynamics where from Table 13.4 we obtain $\eta_0/\mu = 1.6$, and H stands for full, nonlocal hydrodynamics, where $\eta_0/\mu = 1.7$). In the second and third row, we attempt to tune the turnover time or viscosity ($\mu = \mu^* = 0.14$ Pa·s) so that the steady state morphology with intra-fiber hydrodynamics matches the steady state morphology with hydrodynamics and $\tau_f = 10$ s. Even after doing this, the elastic and viscous modulus are still significantly higher with the intra-fiber mobility.

disturbance flows in bundling fibers (fibers moving towards each other) *oppose* the direction of motion, which slows down the bundling process. Thus when we fix a turnover time, there are on average fewer bundles (and fewer links) when the bundling process is slower.

To compensate for the changes in structure, we drop the turnover time or increase the viscosity for intra-fiber hydrodynamics simulations so that the steady state morphology matches the morphology with hydrodynamics as closely as possible. We then measure the moduli for these new steady states. In Table 13.3, we show our attempt to match the statistics for intra-fiber hydrodynamics with smaller turnover times with those for hydrodynamics with $\tau_f = 10$ s. Comparing $\tau_f = 10$ s with full hydrodynamics with $\tau_f = 9.3$ s with intra-fiber (IF) hydrodynamics, we see the IF simulations have larger moduli, even when we closely match the steady state link density \bar{C} and bundle density \bar{B} . The same is true when we attempt to rescale time by increasing the background fluid viscosity (second to last line in Table 13.3), as in that case the larger values of G' and G'' persist despite a decrease in link and bundle density. So nonlocal hydrodynamics must both change the bundle morphology and lower the stress for a fixed morphology.

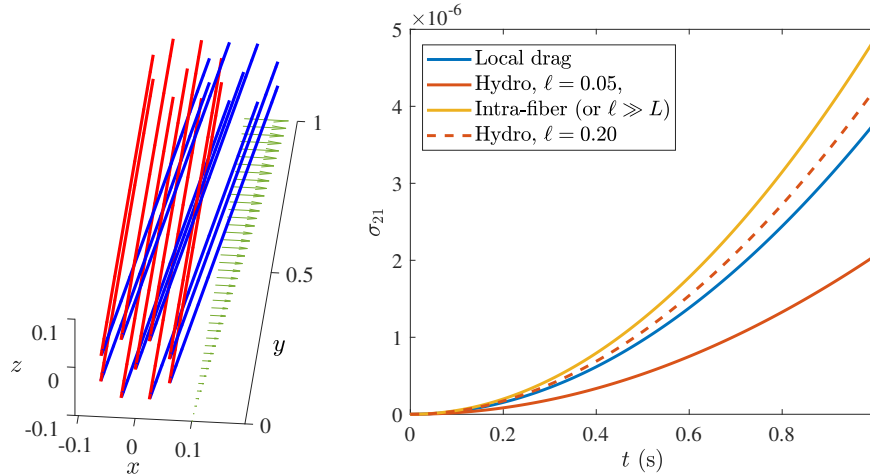


Figure 13.16: Reduction of stress in bundles explains smaller moduli with hydrodynamics. (Left) We manufacture a bundle geometry without CLs by placing nine fibers of length $L = 1 \mu\text{m}$ (red) in and around an octagon with side length ℓ and straining with constant rate $\dot{\gamma} = 0.1 \text{ s}^{-1}$ until $t = 1 \text{ s}$ (blue). (Right) The resulting stress evolution for different hydrodynamic models. The local drag (blue) and intra-fiber (yellow) results are independent of ℓ , while the stress for full hydrodynamics (red) depends strongly on ℓ . For $\ell = 0.05 \mu\text{m}$ (solid red, the simulation parameters), there is a significant decrease in stress which comes from the entrainment of the fibers in each other’s flow fields. For $\ell = 0.20 \mu\text{m}$ (dashed orange), the decrease is minimal; note that for full hydrodynamics with $\ell \gg L$ we would recover the “intra-fiber” curve.

13.4.2 Bundled fibers: when flow reduces stress

The reason for the decrease in moduli with inter-fiber hydrodynamics for a fixed morphology has to do with the flows inside a bundle. If the fibers are packed tightly within a bundle, then we expect their disturbance flows to have a strong influence on each other. When a fiber in a bundle moves with the bulk fluid, the disturbance flow it creates is in the same direction as the bulk motion, so the other fibers in the bundle will naturally move as well. This effect, which is not present when we use local drag or intra-fiber mobility, explains the reduction in stress for a fixed rate of strain.

Figure 13.16 establishes this more rigorously. We consider a set of nine fibers with $L = 1 \mu\text{m}$ *without* cross-links. The fibers are arranged in a regular octagon with side length ℓ with another fiber centered at the origin. We apply a constant straining flow with strength $\dot{\gamma} = 0.1 \text{ s}^{-1}$ and measure the stress over the first second in Fig. 13.16. The solid red line shows the




Figure 13.17: Our continuum model, informed by the timescales discussed in previous sections. We use three Maxwell elements with timescales τ_1 , τ_2 , and τ_3 , all in parallel with a viscous dashpot to describe the network. The viscous dashpot η_0 represents the high frequency viscosity of the permanently cross-linked fiber suspension. The first Maxwell element has timescale $\tau_1 \approx 0.02$ seconds associated with it, and represents the relaxation of the fibers to a transient elastic equilibrium (the networks before and after relaxation are shown to the left and right of this Maxwell element; the relaxing fibers are shown in blue); on this timescale, the links are effectively static. The second Maxwell element, with timescale $\tau_2 \approx 0.5$ s, represents the unbinding of some links (shown more transparent than the others) and the appearance of new links (orange) – compare the networks to the left and right of this Maxwell element. The third Maxwell element with timescale $\tau_3 \approx 5$ s represents network remodeling (compare the networks to the left and right of this element); for timescales larger than τ_3 , some of the fibers (shown in green) and links (orange) turn over and the network completely remodels from the initial state.

stress with full hydrodynamics and the simulation parameter of $\ell = 0.05 \mu\text{m}$, so that the fibers are very close. In this case, we see that the stress with full hydrodynamics is slightly more than half of that with local drag (which is independent of ℓ). Similar to our simulation result for cross-linked gels, the stress decrease is *not* explained by intra-fiber hydrodynamics, which increases the stress from local drag by the theoretically predicted 20 – 25%. Thus the decrease in stress with nonlocal hydrodynamics comes from the nonlocal flows induced by the fibers on each other. We demonstrate this in Fig. 13.16 by spacing the fibers farther apart ($\ell = 0.2 \mu\text{m}$) and showing that the stress with full hydrodynamics increases towards the intra-fiber curve.

13.5 A generalized Maxwell continuum framework

In order to enable whole-cell modeling, it is important to introduce a continuum model for the passively cross-linked actin gel. We are motivated by the behavior of the network: at short timescales, it scales as a viscous fluid with a constant elastic modulus (a dashpot in parallel with a spring), while at long timescales it is purely viscous. Combining these two, we introduce a generalized Maxwell model of the type shown in Fig. 13.17. We have the viscous dashpot of strength η_0 in parallel with three Maxwell elements, each of which has a strength g_i and an associated relaxation timescale τ_i , on which the element goes from elastic (for timescales shorter than τ_i) to viscous (longer than τ_i). The total elastic and viscous (in excess of the solvent viscosity) modulus are given by sums of the viscous dashpot and each Maxwell element [233]

$$G'_{\text{GM}}(\omega; \mathbf{p}) = \sum_{i=1}^3 g_i \frac{(\omega\tau_i)^2}{1 + \omega^2\tau_i^2}, \quad G''_{\text{GM}}(\omega; \mathbf{p}) = \eta_0\omega + \sum_{i=1}^3 g_i \frac{\omega\tau_i}{1 + \omega^2\tau_i^2}, \quad (13.6)$$

where we have denoted the 7 parameters by a vector \mathbf{p} . As discussed in [241, 242], we can obtain the fitting parameters (g_i, τ_i, η_0) by maximizing the log-likelihood function

$$L = -\frac{1}{2} \sum_{k=1}^K \left[\left(\frac{1}{\delta'_k} (G'_k - G'_{\text{GM}}(\omega_k; \mathbf{p})) \right)^2 + \left(\frac{1}{\delta''_k} (G''_k - G''_{\text{GM}}(\omega_k; \mathbf{p})) \right)^2 \right], \quad (13.7)$$

where δ'_k and δ''_k are the uncertainties in $G'(\omega_k)$ and $G''(\omega_k)$, respectively, and K is the number of frequencies studied. Weighting the observations by uncertainty will cause the fit to give more weight to more certain measurements. The uncertainty in the fit $\hat{\mathbf{p}}$ is estimated by the square root of the diagonal entries of the inverse of the Fisher information matrix,

$$\delta p_i = 2\sqrt{\mathbf{F}_{ii}^{-1}(\hat{\mathbf{p}})}, \quad \text{where} \quad \mathbf{F}_{ij}(\mathbf{p}) = -\frac{\partial^2 L}{\partial p_i \partial p_j}(\hat{\mathbf{p}}). \quad (13.8)$$

so that $p_i \pm \delta p_i$ is a 95% confidence interval for p_i .

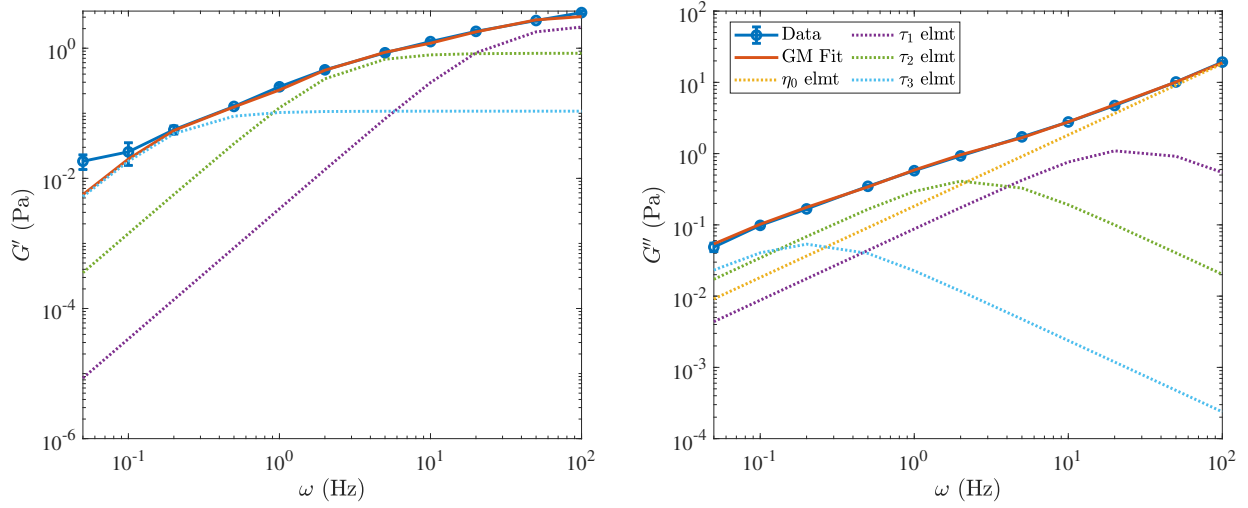


Figure 13.18: Fitting the model of Fig. 13.17 (the moduli (13.6)) to the data from the B-In-M system. The data are shown in blue, with the total fit shown in red. The dotted lines show the contribution of each element in Fig. 13.17 to the total fit. Here $\tau_1 = 0.04$ s is the fastest timescale shown in purple, $\tau_2 = 0.41$ s is the intermediate timescale shown in green, $\tau_3 = 4.5$ s is the longest timescale shown in light blue. The yellow dotted line shows the contribution of the pure viscous element.


In Fig. 13.19(a), we verify that using three (as opposed to two or four) Maxwell elements is indeed the best choice for the bundle-embedded meshwork with $\tau_f = 10$ s and full hydrodynamics (the other systems are similar). As detailed in [242, 243], we determine this by increasing the number of Maxwell elements until the fit stops improving substantially and we start to overfit the data, leading to ill-conditioning. In addition to comparing the fit with three timescales to that with two or four, in Fig. 13.19(a) we also compare to a fit using a continuum of timescales, which is suggested by physical theories [193, 217] and used in prior studies [186]. In this approach, the discrete values of g_i in (13.6) are replaced by a continuous function $g(\tau)$, which is integrated (instead of summed) on $0 \leq \tau \leq \tau_{\max}$. If we take the common choice of $g(\tau) = g_0\tau^{-\alpha}$ [244], we obtain

$$G'_{\text{CM}}(\omega; \mathbf{p}) = \int_0^{\tau_{\max}} g_0 \tau^{-\alpha} \frac{(\omega\tau)^2}{1 + \omega^2\tau^2} d\tau, \quad G''_{\text{CM}}(\omega; \mathbf{p}) = \eta_0\omega + \int_0^{\tau_{\max}} g_0 \tau^{-\alpha} \frac{\omega\tau}{1 + \omega^2\tau^2} d\tau, \quad (13.9)$$


where $(g_0, \alpha, \tau_{\max}, \eta_0)$ are the fitting parameters. In Fig. 13.19(a), we show that the fit to the data with this continuum spectrum of timescales is not substantially better than the three timescale generalized Maxwell fit. Our preference is the three timescale fit, since in this case we can assign physical to meaning to each of the timescales, as we do in Sections 13.3.1–13.3.3. Given that adding more discrete timescales leads to ill-conditioning, the continuum approach, where there are infinitely many timescales (and in fact infinitely many for choices $g(\tau)$), is even more poorly conditioned, and in that case it is difficult to assign a physical meaning to the fitting parameters.

The three-timescale fit that we obtain for the B-In-M system with full hydrodynamics, along with the contribution of each of the separate Maxwell elements, is shown in Fig. 13.18. Admittedly, the elastic modulus data $G'(\omega)$ at small frequencies *do not* fit the generalized Maxwell model since they do not decay as ω^2 . This suggests that the material response is likely more complicated than just three Maxwell elements. At the same time, we have already shown that $G''(\omega)$ scales linearly with ω at small frequencies, and thus that the viscous modulus, which dominates the elastic modulus at low frequencies, does fit the generalized Maxwell assumption. Our choice is therefore to fit only G'' , and not G' , for frequencies less than 0.1 Hz (10 s timescales).

Table 13.4 gives the fitting parameters for all of the systems we have considered with multiple options for the mobility. For almost all systems, we have $\tau_1 \approx 0.03$ s, $\tau_2 \approx 0.3$ s, and $\tau_3 \approx 3$ s, which is in line with our estimates of these timescales in Figs. 13.13 and 13.14(a). The only exception is the system with higher viscosity; in this case we have already shown that the timescale τ_1 scales with viscosity to become about 0.3 s, while the timescale τ_2 remains about 0.5 s. The result is that τ_1 and τ_2 blend into a single timescale, and a two-timescale model is the best choice for the data, as we show in Fig. 13.19(b). Comparing B-In-M and meshwork geometries, we observe that the long timescale τ_3 is either longer (comparing B-In-M and meshwork with full hydrodynamics) or its contribution in (13.6) is



(a) B-In-M



(b) B-In-M, 10μ

Figure 13.19: Fits to the elastic (left) and viscous (right) modulus data for (a) The B-In-M geometry with full hydrodynamics and (b) The B-In-M geometry with ten times larger viscosity and local drag mobility. We consider 1 (red), 2 (yellow), 3 (purple), and 4 (green) Maxwell modes in parallel with a viscous dashpot. The elastic modulus is only fit for $\omega \geq 0.1$ Hz. The B-In-M system with $\mu = 0.1$ Pa·s is best fit with three elements, while the higher viscosity system ($\mu = 1$ Pa·s) is best fit with only two elements, since in (b) the three and four element Maxwell models have a timescale larger than 100 seconds, indicating overfitting. For (a), the dashed line shows the fit to a continuous spectrum of timescales (13.9), which yields $g_0 = 0.22$, $\alpha = 1.4$, $\tau_{\max} = 3.5$, and $\eta_0 = 0.16$, and is not a better fit than the 3 timescale generalized Maxwell model.

System	Mobility	(τ_1, τ_2, τ_3) (s)	(g_1, g_2, g_3) (Pa)	η_0 (Pa·s)
Meshwork	LD	(0.032, 0.31, 4.2)	(1.56, 0.44, 0.03)	0.13
Meshwork, $\kappa/10$	LD	(0.024, 0.15, 1.5)	(1.34, 0.45, 0.07)	0.13
Meshwork	Intra-fiber	(0.022, 0.17, 1.4)	(1.49, 0.59, 0.10)	0.16
Meshwork	Full hydro	(0.025, 0.21, 2.1)	(1.42, 0.54, 0.05)	0.17
B-In-M	LD	(0.026, 0.28, 2.7)	(2.91, 1.26, 0.30)	0.13
B-In-M	Intra-fiber	(0.030, 0.28, 2.4)	(2.77, 1.17, 0.29)	0.16
B-In-M	Full hydro	(0.039, 0.41, 4.5)	(2.26, 0.84, 0.11)	0.18
B-In-M, $2k$	LD	(0.034, 0.31, 2.7)	(2.54, 1.10, 0.17)	0.13
B-In-M, 10μ	LD	(0.49, —, 4.3)	(3.13, —, 0.38)	1.33

Table 13.4: Parameters in the three-timescale generalized Maxwell model shown in Fig. 13.17. Refer to Table 13.1 for a detailed description of each system. The uncertainty in each measurement is about 25%, with the exception of η_0 , which has an uncertainty of about 5%.

stronger (comparing B-In-M and meshwork with local drag), which is consistent with our observation in Section 13.3.2 that B-In-M geometries have longer remodeling times and a higher resistance to deformation. Finally, we notice how the viscosity η_0 scales directly with μ and increases by $\approx 30\%$ when we account for intra-fiber or nonlocal hydrodynamics, in accordance with our observations in Sections 13.3.1 and 13.4.

Chapter 14

Discussion and conclusions

This part of the dissertation was substantially more applied than the previous parts. In it, we used the methodology we developed in Parts I and II to simulate the dynamics of cross-linked actin networks. We focused on two main problems: the dynamic process by which cross linkers ratchet filaments into bundles (Chapter 12) and how morphology and hydrodynamic interactions impact the macroscopic viscoelastic behavior (Chapter 13).

We found that, even without thermal movements, actin filaments can still bundle, as filaments that are initially close enough are linked together at small patches with CLs. These CLs pull fibers together and align them, thereby allowing more CLs to bind to other sections of the fibers. What results initially, for times smaller than the critical bundling coalescence time τ_c , is a collection of bundles with a few highly aligned filaments, also called a composite bundle state [203]. For times larger than τ_c , these bundles coalesce into larger bundles using a similar mechanism as that for individual fibers, and a clustered bundle state forms. Our critical bundling timescale τ_c thus describes the initial time at which networks transition from the composite bundle state to the clustered bundle state. In networks with fiber turnover, a clustered bundle steady state is only possible if the turnover time is much larger than τ_c .

We quantified the role of diffusion throughout the bundling process, finding that it has a

larger impact in the initial stages of bundling, when the filaments are not severely constrained by CLs and can move freely to find each other. We associated this stage with $t < \tau_c$, and showed that adding thermal fluctuations decreases τ_c from 16 seconds to 4 seconds, although much of this could be driven by our change in the cross linking model to achieve detailed balance. We showed that the stage when bundles coalesce ($t > \tau_c$) is less affected by thermal diffusion (sped up by a factor of 2), since at that stage the filaments are constrained by CLs, which are more involved in bundle coalescence. This complements the observation in [209, 212] that bundling occurs faster in a fluid-like environment, where filaments can move freely prior to kinetic arrest.

At first glance, the order of magnitude of τ_c that we obtained seems shorter than the characteristic bundling time obtained experimentally, which is generally reported to be on the order of minutes [230]. The comparison is difficult, however, since experimental times generally include polymerization, and the bundling timescale in experiments is defined by the onset of the clustered bundle steady state, which is much later than the composite bundle state where we define τ_c . Nevertheless, the most instructive comparison is between our work and [209, Fig. 4], which shows experimentally that the addition of 10% nucleates (which speeds up the polymerization process) gives a saturated bundled state after 100 seconds of polymerization and bundling, where the bundles are made of at least 15 – 30 filaments and are spaced some 10 – 20 μm apart. Given this observation, and the fact that bundling slows down over time, it is not difficult to imagine that the transition from the homogeneous state to the composite bundle state could take place on the order of 5 – 10 seconds after cross linkers are added to a system of (polymerized and capped) actin filaments.

While diffusion of fibers speeds up the bundling process, we showed that it must be combined with a sufficient concentration of CLs for rapid bundling to occur. In particular, we showed that a high concentration of CLs (high CL binding rate) can induce bundling for filaments of any mesh size, with a critical bundling time τ_c that depends only weakly on the

mesh size. By contrast, when the concentration of CLs is small, bundling is more difficult for any fixed mesh size, and gets near impossible as the mesh size increases, as near-fiber contacts become less frequent. This is in accordance with a number of experimental papers [202, 206] which find that bundling requires a critical CL concentration. In addition, because the fibers are in contact for a short time at larger mesh sizes, the system must be saturated with CLs for bundling to proceed at a reasonable rate. This saturation is less important at smaller mesh sizes, where fiber pairs come into contact more frequently. Translating our results to experimental parameters, we find that the ratio of the crosslinker concentration to the F-actin concentration that is needed for a particular bundling time scale decreases as the actin concentration increases, which is in accordance with existing experimental observations [206, Fig. 3].

As already mentioned, one of the drawbacks of some experimental studies is the sensitivity of the bundling time to the rate of actin polymerization. For example, it is shown in [209, Fig. 4(d)] that polymerization kinetics make an order of magnitude difference in the bundling kinetics. While simultaneous polymerization and bundling also occurs *in vivo*, our study here allowed us to divorce bundling and polymerization by focusing on a fixed filament length. By doing this, we showed that shorter filaments bundle faster exclusively because they can diffuse faster, because without thermal fluctuations we saw no difference in the bundling kinetics between short and long filaments. This clarifies why shorter actin filaments are able to associate more rapidly into bundles without the presence of a background actin mesh [202, 212].

We also examined how semiflexible bending fluctuations affect the bundling process, which is a novel contribution to the biophysics literature. Here we showed that the bundling of fluctuating actin filaments, which have persistence length on the order $10 \mu\text{m}$ [18], significantly differs from that of rigid fibers, which is in contrast to the conclusion we drew when filaments did not fluctuate (and only moved by CL forces). However, when $\ell_p/L \gtrsim 10$, we

showed that these differences can be explained by a reduced model where actin filaments rotate and translate randomly, and bend via CL forces. In this case the combination of translational and rotational rigid-body diffusion with fiber flexibility and CL forces drives the observed differences from rigid fiber dynamics. Semiflexible bending fluctuations only accelerate bundling dynamics when $\ell_p/L \approx 1$, in which case the differences from rigid fibers can be observed even in the regime when the fibers are weakly cross linked.

By applying our steric repulsion algorithm to bundling filaments, we showed that the observed “kinetic arrest” of bundling fibers [209, 210] is driven primarily by cross linking, with an only 10% change in the fiber displacement when we account for steric interactions. In fact, contrary to intuition, our simulations show that actin fibers can reorient into a percolated network even when the initial mesh size is smaller than the filament length. An interesting future direction would be to see how, if at all, this conclusion changes when we also account for the electrostatic interactions between actin filaments, since actin is a polyelectrolyte made of charged molecules [245].

There are, of course, other timescales that we could have examined in the bundling process. Our choice to focus on the timescale τ_c was motivated by our observation that the steady state morphology of cross-linked actin networks is driven by a competition between bundling (which occurs on timescale τ_c) and filament turnover (which occurs on timescale τ_f). While it is intuitively obvious that increasing the turnover timescale τ_f will produce a steady state with more bundles, it is fair to ask whether the ratio τ_f/τ_c alone controls the steady state morphology, or if some other microscopic parameters come into play. In Fig. 12.16, we showed that for turnover times $\tau_f = \tau_c/2$, τ_c , and $2\tau_c$, the gel evolves to a steady state where the bundle density and percentage of fibers in bundles depend primarily on the ratio τ_f/τ_c , for either Brownian or non-Brownian fibers (recall that τ_c differs by a factor of 4 for these two cases). Snapshots in Fig. 12.17 show little qualitative difference between the network morphology of the Brownian and non-Brownian steady states for a fixed τ_f/τ_c .

We took advantage of this finding in Chapter 13, where we used the fixed values $\tau_f/\tau_c \approx 0.3$ and $\tau_f/\tau_c \approx 0.6$ to generate two steady state structures for rheological analysis. We saw that the former structure is a homogeneous meshwork, while the latter contains bundles embedded in a mesh. Despite the small differences in these two structures (one has 10% of fibers in bundles, the other 20%), the differences in their behavior, especially at long timescales, was remarkable.

For gels with permanent cross linkers, we found one timescale which dominates the behavior, while gels with transient cross linkers had three. On the fastest time scale, ~ 0.03 seconds (which is the one we found for permanent networks), actin fibers relax viscously, locally, and rapidly to the transient elastic equilibrium generated by the network configuration. On intermediate timescales, ~ 0.5 seconds, the CLs turn over, generating new transient elastic equilibria, and finally on the slow time scale, $\sim 3 - 5$ seconds, the network undergoes global remodeling through a combination of filament and CL turnover. One of the most useful practical results of the simulations is that very complex cross-linked gel mechanics can be approximated with the generalized Maxwell mechanical circuit shown in Fig. 13.17, in which three Maxwell elements in parallel correspond to elastoviscous gel deformations on the three characteristic time scales, in addition to the effective viscosity of the actin fiber suspension in the background fluid.

Our simulations predict that at small frequencies (time scales longer than a few seconds), the effective elastic and viscous behaviors are controlled by the CL mechanics (as opposed to fiber mechanics). In this regime, an effectively viscous mechanical response to deformation dominates relatively weak elastic behavior. At low frequencies, the dominant viscous response originates from the CLs stretching and deforming the fibers (in addition to the significant viscosity of the background fluid); the fibers by themselves contribute little to the net viscosity. The mechanical moduli scale nonlinearly with the CL density, because the CLs within actin bundles respond to shear differently than the CLs in the actin mesh.

Thus, the gel’s morphology exerts a significant influence on the actin resistance to flow at long timescales.

At high frequencies, or timescales shorter than one tenth of a second, we found that the effective elastic modulus of the network is higher than (but on the same order of magnitude as) the effective viscous modulus (once we removed the viscous scaling generated by the fibers). The elasticity simply originates from multiple elastic springs of CLs that can be considered static on the short time scale. The greatest contribution to the net viscosity is that of the background fluid, with the fiber suspension contributing about 25% additional viscosity. At moderate frequencies around 1 Hz (intermediate timescale, between 0.1 and 3 seconds), the mechanical behavior of the crosslinked network is complex, combining comparable contributions of all mechanical factors described for low and high frequencies.

Our model allowed us, for the first time, to estimate quantitatively the role of nonlocal inter-fiber hydrodynamic interactions on the actin gel rheology. While these interactions have little effect in a homogeneous cross-linked actin meshwork, they significantly decrease the elastic and viscous moduli of heterogeneous B-In-M gels. We showed that this decrease is most significant at low frequencies, i.e., on timescales of a few seconds, when the network morphology has a strong influence on its mechanical behavior. At these frequencies, in the presence of bundles, the hydrodynamic interactions cause 30–50% downward corrections to the viscoelastic moduli. Two factors are responsible for this effect: first, hydrodynamic interactions slow down the bundling process, so fewer bundles form, and second, nonlocal hydrodynamics creates entrainment flows that reduce the stress inside of the bundles that do form.

While almost all of our simulations fell in the linear viscoelastic regime, we did observe a strain hardening effect, with the network resistance increasing beyond the linear response for higher-amplitude deformations. This effect has been observed before experimentally [188], which has prompted a variety of explanations and computational studies. Both Mulla et

al. [193] and Kim et al. [246], for instance, propose that strain hardening is due to the mechanics of individual filaments, which stiffen nonlinearly in response to applied stress. Because our goal was to measure the viscoelastic moduli of the networks in their dynamic steady state, we did not systematically explore the viscoelastic behavior at large strains or perform measurements at zero frequency. That said, our simulations do offer two interesting insights: first, the mechanical behavior becomes nonlinear at just 10% or greater strains, which is a typical result for in vitro networks [247], and second, there are more bundles (which are more resistive to strain) at larger strain rates, and so at least one reason for the nonlinearity is the mechanosensitive character of the actin gel morphology. Our conclusion is identical to the modeling study [248], which found that strain combined with dynamic cross-linking induces significant morphological changes in the network, leading to strain hardening [248, Fig. 3]. In fact, it was recently shown that actin filament mobility is a necessary condition for strain hardening and mechanical hysteresis [249].

In the linear response regime, our results are in line with most of the prior experimental and modeling studies. We found viscoelastic moduli on the order 0.1 – 10 Pa [189, 202], with the links becoming viscous on long timescales and elastic on short ones [186]. Our measurements showed that decreasing the fiber bending stiffness decreased the viscoelastic moduli, but by a slower than linear rate [167] (in our study, decreasing κ by a factor of ten reduced the moduli by at most 30%). We also found that the elastic stress in the gel (from the CLs exerting force on the fibers) is much larger than the viscous drag stress over long times, which agrees with a previous computational study [246].

14.1 Future directions

What all modeling studies, including ours, suffer from is that only few features of the actin gel are simulated, with many factors, forces, and processes either ignored or approximated crudely. There is, of course, a good reason for this: even a very limited caricature of the

gel exhibits complex dynamics; the fact that many experimental studies paint very different pictures of the gel mechanics likely stems from the same biological complexity. We propose that to understand the full complexity, one has to keep adding dynamic processes to previous simpler benchmark models and examine the changes that new dynamics bring. Thus, our next step will be to use the methodology of Chapter 8 to add thermal forces responsible for random bending of the fibers. For a one micron long fiber with persistence length close to 20 microns, the thermal bending causes lateral fluctuations of no more than 30 to 40 nm, which is significantly less than the mesh size of the networks we investigated, so we do not expect that our results will change drastically. Indeed, a previous computational study in permanent networks showed that thermal fluctuations only impact the mechanical moduli when the average distance between cross links is greater than the persistence length of the fibers [167, Fig. 7(c)], which is a regime we are far outside of here.

Despite this speculation, the fact remains that our study was unable to reproduce a number of experimental observations in transiently cross-linked actin networks that have been traced to the underlying thermal fluctuations of the filaments. For instance, we report a viscous modulus that decays as for a viscous fluid ($G'' \sim \omega$ at low frequencies), whereas a number of recent experimental, modeling, and theoretical studies [192, 195, 217, 218] report an $\omega^{1/2}$ dependence of both moduli at low frequencies. It is important to note, however, that fiber turnover, which was not accounted for in these previous studies, changes the behavior on timescales larger than the turnover time, as does our use of shorter filaments relative to [218]. In the high-frequency range, our results show viscous-fluid scaling of G'' at large ω , but experiments [188] and models [166] have shown an $\omega^{3/4}$ dependence at high frequencies. Now that we have performed a benchmark study which shows that a *non-Brownian* cross-linked actin network exhibits $G'' \sim \omega$ for large ω , it will be interesting to see if adding thermal bending modes [217, 218] reproduces the scaling relation $G'' \sim \omega^{3/4}$, and how these relations are affected by hydrodynamic interactions. In concert with this, we will consider more

bundled morphologies, where the flexibility of the filaments, and therefore their transverse bending fluctuations, may become more important. Similarly, we will test how explicit simulation of translational and rotational diffusion and (un)binding of individual CLs [218], as well as force dependence of CL kinetics, affect these results. A key component to any of these steps is to find an expression for the additional suspension stress due to Brownian motion, which, although available in some of the engineering literature [250, 251], is so far lacking in theoretical understanding from an applied stochastic analysis perspective. It is for this reason that our rheological simulations have yet to consider Brownian motion.

We failed to observe a local maximum or minimum in the viscous modulus at intermediate frequencies, as was reported in [166, 194]. One explanation for this is a difference in parameters. In the computational study [166], the unloaded CL unbinding rate is 0.1 s^{-1} , and the CL stiffness is 10 pN/nm ($10,000 \text{ pN}/\mu\text{m}$). In this parameter regime, the CLs *always* make the most of their mechanical contribution to the network mechanics while bound, since they cross-link the fibers for seconds and are extremely stiff. Indeed, [166, Fig. 6] shows that the peak in G'' , when the timescale of the driving frequency matches the link unbinding rate, becomes less sharp and almost disappears when k_{off} increases to the order of one second, as it is in our study here. This suggests that the local maximum is only observed when the relaxation of the network is many orders of magnitude faster than the link binding rate. This issue clearly requires more investigation, and in the future we will explore how using a force-dependent unbinding rate, as is done in [166], impacts the results, and whether this assumption is responsible for the peak.

We also hope to place our model of cross-linker dynamics on more rigorous footing by comparing it to a model that actually tracks the diffusion, binding, and unbinding of individual CLs. Other modeling studies addressed bundling in more complex systems, for example formation of unipolar bundles from a branched actin network [252] and bundling in the presence of a mix of CLs and myosin molecular motors [163, 174]. Interestingly, the appearance

of the bundles in these more complex systems [174], which form when CL concentration is above a threshold value [253], resemble those predicted by our model without motors. Another level of complexity is limits on bundle sizes due to chirality effects [254] and long-range electrostatic repulsion between the filaments (reviewed in [255]). Elastic twisting of actin filaments has been posited to contribute to elastic energy storage [211] and to the emergence of chirality in cross-linked gels [256], so it would be interesting to use the methodology of Chapter 6 to see how twist elasticity affects the suspension behavior. In addition, the gels *in vivo* are very heterogeneous in the sense that there is a distribution of fiber lengths and turnover times, and a diversity of CL types co-existing with myosin motors [171, 212, 257]. How network heterogeneity and active molecular force generation affect the gel mechanics is another important question. Finally, in this work it was too difficult for us to simulate the experimental steady state clustered bundled morphologies, since we simulated actin filament lengths of 1 μm and the observed steady states have bundles separated by hundreds of microns [206, 208]. More efficient, GPU-based, simulation techniques might enable the efficient simulation of even larger systems.

Bibliography

- [1] AT Chwang and Th Y Wu. A note on the helical movement of micro-organisms. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 178(1052):327–346, 1971.
- [2] Howard C Berg and Robert A Anderson. Bacteria swim by rotating their flagellar filaments. *Nature*, 245(5425):380, 1973.
- [3] Christopher Brennen and Howard Winet. Fluid mechanics of propulsion by cilia and flagella. *Annual Review of Fluid Mechanics*, 9(1):339–398, 1977.
- [4] Eric Lauga and Thomas R Powers. The hydrodynamics of swimming microorganisms. *Reports on Progress in Physics*, 72(9):096601, 2009.
- [5] Sookkyung Lim and Charles S Peskin. Fluid-mechanical interaction of flexible bacterial flagella by the immersed boundary method. *Physical Review E*, 85(3):036307, 2012.
- [6] Alexander M Maier, Cornelius Weig, Peter Oswald, Erwin Frey, Peer Fischer, and Tim Liedl. Magnetic propulsion of microswimmers with dna-based flagellar bundles. *Nano letters*, 16(2):906–910, 2016.
- [7] Tătulea-Codrean Maria and Eric Lauga. Geometrical constraints on the tangling of bacterial flagellar filaments. *Scientific Reports (Nature Publisher Group)*, 10(1), 2020.

- [8] Miguel Angel Bibbó. *Rheology of semiconcentrated fiber suspensions*. PhD thesis, Massachusetts Institute of Technology, 1987.
- [9] Michael B Mackaplow and Eric SG Shaqfeh. A numerical study of the rheological properties of suspensions of rigid, non-brownian fibres. *J. Fluid Mech.*, 329:155–186, 1996.
- [10] RR Sundararakumar and Donald L Koch. Structure and properties of sheared fiber suspensions with mechanical contacts. *Journal of Non-Newtonian Fluid Mechanics*, 73(3):205–239, 1997.
- [11] Bruce Alberts, Alexander Johnson, Julian Lewis, Martin Raff, Keith Roberts, and Peter Walter. *Molecular biology of the cell*. Garland Science, 2002.
- [12] Bernd Wagner, Rainer Tharmann, Ilka Haase, Markus Fischer, and Andreas R Bausch. Cytoskeletal polymer networks: the molecular structure of cross-linkers determines macroscopic properties. *Proceedings of the National Academy of Sciences*, 103(38):13974–13978, 2006.
- [13] David A Head, Alex J Levine, and FC MacKintosh. Deformation of cross-linked semi-flexible polymer networks. *Physical review letters*, 91(10):108102, 2003.
- [14] Wylie W Ahmed and Timo Betz. Dynamic cross-links tune the solid–fluid behavior of living cells. *Proceedings of the National Academy of Sciences*, 112(21):6527–6528, 2015.
- [15] Ehssan Nazockdast, Abtin Rahimian, Denis Zorin, and Michael Shelley. A fast platform for simulating semi-flexible fiber suspensions applied to cell mechanics. *J. Comput. Phys.*, 329:173–209, 2017.
- [16] Michael J Shelley. The dynamics of microtubule/motor-protein assemblies in biology and physics. *Annual Review of Fluid Mechanics*, 48:487–506, 2016.

- [17] Sayantan Dutta, Reza Farhadifar, Wen Lu, Robert Blackwell, David B Stein, Margot Lakonishok, Vladimir I Gelfand, Stanislav Shvartsman, Michael Shelley, et al. Self-organized intracellular twistors. *arXiv preprint arXiv:2304.02112*, 2023.
- [18] Frederick Gittes, Brian Mickey, Jilda Nettleton, and Jonathon Howard. Flexural rigidity of microtubules and actin filaments measured from thermal fluctuations in shape. *The Journal of cell biology*, 120(4):923–934, 1993.
- [19] Joseph B Keller and Sol I Rubinow. Slender-body theory for slow viscous flow. *Journal of Fluid Mechanics*, 75(4):705–714, 1976.
- [20] Robert E Johnson. An improved slender-body theory for stokes flow. *Journal of Fluid Mechanics*, 99(2):411–431, 1980.
- [21] Robert Edward Johnson. *Slender-body theory for Stokes flow and flagellar hydrodynamics*. PhD thesis, California Institute of Technology, 1977.
- [22] Anna-Karin Tornberg and Michael J Shelley. Simulating the dynamics and interactions of flexible fibers in stokes flows. *Journal of Computational Physics*, 196(1):8–40, 2004.
- [23] GK Batchelor. The stress system in a suspension of force-free particles. *Journal of fluid mechanics*, 41(3):545–570, 1970.
- [24] Yoichiro Mori, Laurel Ohm, and Daniel Spirn. Theoretical justification and error analysis for slender body theory with free ends. *Archive for Rational Mechanics and Analysis*, 235(3):1905–1978, 2020.
- [25] Yoichiro Mori and Laurel Ohm. Accuracy of slender body theory in approximating force exerted by thin fiber on viscous fluid. *Studies in Applied Mathematics*, 2021.
- [26] Yoichiro Mori, Laurel Ohm, and Daniel Spirn. Theoretical justification and error

- analysis for slender body theory. *Communications on Pure and Applied Mathematics*, 73(6):1245–1314, 2020.
- [27] Lyndon Koens and Eric Lauga. The boundary integral formulation of stokes flows includes slender-body theory. *Journal of Fluid Mechanics*, 850, 2018.
- [28] Mohit Garg and Ajeet Kumar. A slender body theory for the motion of special cosserat filaments in stokes flow. *Mathematics and Mechanics of Solids*, page 10812865221083323, 2022.
- [29] Lyndon Koens and Eric Lauga. Slender-ribbon theory. *Physics of Fluids*, 28(1):013101, 2016.
- [30] Neeraj S Borker and Donald L Koch. Slender body theory for particles with non-circular cross-sections with application to particle dynamics in shear flows. *Journal of Fluid Mechanics*, 877:1098–1133, 2019.
- [31] DK Srivastava. Slender body theory for stokes flow past axisymmetric bodies: a review article. *International Journal of Applied Mathematics and Mechanics*, 8(15):14–39, 2012.
- [32] Thomas Götz. *Interactions of fibers and flow: asymptotics, theory and numerics*. dissertation. de, 2001.
- [33] Thomas R Powers. Dynamics of filaments and membranes in a viscous fluid. *Reviews of Modern Physics*, 82(2):1607, 2010.
- [34] Constantine Pozrikidis et al. *Boundary integral and singularity methods for linearized viscous flow*. Cambridge university press, 1992.
- [35] Panayiota Katsamba, Sébastien Michelin, and Thomas D Montenegro-Johnson. Slen-

- der phoretic theory of chemically active filaments. *Journal of Fluid Mechanics*, 898, 2020.
- [36] Debasish Das and Eric Lauga. Computing the motor torque of escherichia coli. *Soft matter*, 14(29):5955–5967, 2018.
- [37] Xiumei Liu and Gerald H Pollack. Mechanics of f-actin characterized with microfabricated cantilevers. *Biophysical journal*, 83(5):2705–2715, 2002.
- [38] Basile Audoly and Yves Pomeau. *Elasticity and Geometry: From hair curls to the non-linear response of shells*. Oxford University Press, 2010.
- [39] Harishankar Manikantan and David Saintillan. Subdiffusive transport of fluctuating elastic filaments in cellular flows. *Physics of Fluids*, 25(7):073603, 2013.
- [40] Harishankar Manikantan and David Saintillan. Effect of flexibility on the growth of concentration fluctuations in a suspension of sedimenting fibers: Particle simulations. *Physics of Fluids*, 28(1):013303, 2016.
- [41] Mehdi Jabbarzadeh and Henry C Fu. A numerical method for inextensible elastic filaments in viscous fluids. *Journal of Computational Physics*, page 109643, 2020.
- [42] Cornelius Lanczos. *The variational principles of mechanics*. Courier Corporation, 2012.
- [43] Ellis Harold Dill. Kirchhoff’s theory of rods. *Archive for History of Exact Sciences*, pages 1–23, 1992.
- [44] Reinhard Vogel. The bacterial flagellum: Modeling the dynamics of the elastic filament and its transition between polymorphic helical forms. 2012.
- [45] Mordecai B Rubin. *Cosserat theories: shells, rods and points*, volume 79. Springer Science & Business Media, 2013.

- [46] Carlos Floyd, Haoran Ni, Ravinda Gunaratne, Radek Erban, and Garegin A Papoian. On stretching, bending, shearing and twisting of actin filaments i: Variational models. *arXiv preprint arXiv:2112.01480*, 2021.
- [47] Simon F Schoeller, Adam K Townsend, Timothy A Westwood, and Eric E Keaveny. Methods for suspensions of passive and active filaments. *Journal of Computational Physics*, 424:109846, 2021.
- [48] Timothy A Westwood and Eric E Keaveny. Coordinated motion of active filaments on spherical surfaces. *Physical Review Fluids*, 6(12):L121101, 2021.
- [49] Sookkyung Lim, Anca Ferent, X Sheldon Wang, and Charles S Peskin. Dynamics of a closed rod with twist and bend in fluid. *SIAM Journal on Scientific Computing*, 31(1):273–302, 2008.
- [50] Sarah D Olson, Sookkyung Lim, and Ricardo Cortez. Modeling the dynamics of an elastic rod with intrinsic curvature and twist using a regularized stokes formulation. *Journal of Computational Physics*, 238:169–187, 2013.
- [51] Kenta Ishimoto and Eamonn A Gaffney. An elastohydrodynamical simulation study of filament and spermatozoan swimming driven by internal couples. *IMA Journal of Applied Mathematics*, 83(4):655–679, 2018.
- [52] Frank TM Nguyen and Michael D Graham. Impacts of multiflagellarity on stability and speed of bacterial locomotion. *Physical Review E*, 98(4):042419, 2018.
- [53] Michael Reichert. *Hydrodynamic interactions in colloidal and biological systems*. PhD thesis, 2006.
- [54] Hirofumi Wada and Roland R Netz. Non-equilibrium hydrodynamics of a rotating filament. *EPL (Europhysics Letters)*, 75(4):645, 2006.

- [55] Hirofumi Wada and Roland R Netz. Model for self-propulsive helical filaments: kink-pair propagation. *Physical review letters*, 99(10):108102, 2007.
- [56] Weicheng Huang and M Khalid Jawed. Numerical simulation of bundling of helical elastic rods in a viscous fluid. *Computers & Fluids*, page 105038, 2021.
- [57] Yi Man, William Page, Robert J Poole, and Eric Lauga. Bundling of elastic filaments induced by hydrodynamic interactions. *Physical Review Fluids*, 2(12):123101, 2017.
- [58] Lyndon Koens. Tubular-body theory for viscous flows. *Physical Review Fluids*, 7(3):034101, 2022.
- [59] Howard C Berg. *E. coli in Motion*. Springer, 2004.
- [60] Charles S Peskin. Flow patterns around heart valves: a numerical method. *J. Comput. Phys.*, 10(2):252–271, 1972.
- [61] Sune Lomholt and Martin R Maxey. Force-coupling method for particulate two-phase flow: Stokes flow. *J. Comput. Phys.*, 184(2):381–405, 2003.
- [62] Ricardo Cortez. The method of regularized stokeslets. *SIAM Journal on Scientific Computing*, 23(4):1204–1225, 2001.
- [63] Ricardo Cortez, Lisa Fauci, and Alexei Medovikov. The method of regularized stokeslets in three dimensions: analysis, validation, and application to helical swimming. *Physics of Fluids*, 17(3):031504, 2005.
- [64] Jens Rotne and Stephen Prager. Variational treatment of hydrodynamic interaction in polymers. *J. Chem. Phys.*, 50(11):4831–4837, 1969.
- [65] Eligiusz Wajnryb, Krzysztof A Mizerski, Pawel J Zuk, and Piotr Szymczak. Generalization of the rotne–prager–yamakawa mobility and shear disturbance tensors. *Journal of Fluid Mechanics*, 731, 2013.

- [66] Jason E Butler and Eric SG Shaqfeh. Brownian dynamics simulations of a flexible polymer chain which includes continuous resistance and multibody hydrodynamic interactions. *The Journal of chemical physics*, 122(1):014901, 2005.
- [67] Victor A Beck and Eric SG Shaqfeh. Ergodicity breaking and conformational hysteresis in the dynamics of a polymer tethered at a surface stagnation point. *The Journal of chemical physics*, 124(9):094902, 2006.
- [68] Jianghui Wang, Emilio J Tozzi, Michael D Graham, and Daniel J Klingenberg. Flipping, scooping, and spinning: Drift of rigid curved nonchiral fibers in simple shear flow. *Physics of Fluids*, 24(12):123304, 2012.
- [69] Thomas T Bringley and Charles S Peskin. Validation of a simple method for representing spheres and slender bodies in an immersed boundary method for stokes flow on an unbounded domain. *Journal of Computational Physics*, 227(11):5397–5425, 2008.
- [70] Ricardo Cortez and Michael Nicholas. Slender body theory for stokes flows with regularized forces. *Communications in Applied Mathematics and Computational Science*, 7(1):33–62, 2012.
- [71] Laurel Ohm. Remarks on regularized stokeslets in slender body theory. *Fluids*, 6(8):283, 2021.
- [72] Boan Zhao and Lyndon Koens. Regularized stokeslets lines suitable for slender bodies in viscous flow. *Fluids*, 6(9):335, 2021.
- [73] MR Maxey and BK Patel. Localized force representations for particles sedimenting in stokes flow. *International journal of multiphase flow*, 27(9):1603–1626, 2001.
- [74] Charles S Peskin. The immersed boundary method. *Acta Numer.*, 11:479–517, 2002.

- [75] Andrew M Fiore and James W Swan. Rapid sampling of stochastic displacements in brownian dynamics simulations with stresslet constraints. *The Journal of chemical physics*, 148(4):044114, 2018.
- [76] Bakytzhan Kallemov, Amneet Bhalla, Boyce Griffith, and Aleksandar Donev. An immersed boundary method for rigid bodies. *Communications in Applied Mathematics and Computational Science*, 11(1):79–141, 2016.
- [77] James Lighthill. Flagellar hydrodynamics. *SIAM review*, 18(2):161–230, 1976.
- [78] PJ Zuk, E Wajnryb, KA Mizerski, and P Szymczak. Rotne–prager–yamakawa approximation for different-sized particles in application to macromolecular bead models. *Journal of Fluid Mechanics*, 741, 2014.
- [79] Dhairya Malhotra and Alex Barnett. Efficient convergent boundary integral methods for slender bodies. *In preparation*. Software obtained from <https://github.com/dmalhotra/test-slenderbody>.
- [80] Anna-Karin Tornberg. Accurate evaluation of integrals in slender-body formulations for fibers in viscous flow. *arXiv preprint arXiv:2012.12585*, 2020.
- [81] Eric E Keaveny and Michael J Shelley. Applying a second-kind boundary integral equation for surface tractions in stokes flow. *Journal of Computational Physics*, 230(5):2141–2159, 2011.
- [82] GK Batchelor. Slender-body theory for particles of arbitrary cross-section in stokes flow. *Journal of Fluid Mechanics*, 44(3):419–440, 1970.
- [83] Vasily Kantsler and Raymond E Goldstein. Fluctuations, dynamics, and the stretch-coil transition of single actin filaments in extensional flows. *Physical review letters*, 108(3):038103, 2012.

- [84] Miklós Bergou, Max Wardetzky, Stephen Robinson, Basile Audoly, and Eitan Grispun. Discrete elastic rods. In *ACM SIGGRAPH 2008 papers*, pages 1–12. 2008.
- [85] Wanho Lee, Yongsam Kim, Charles S Peskin, and Sookkyung Lim. A novel computational approach to simulate microswimmers propelled by bacterial flagella. *Physics of Fluids*, 33(11):111903, 2021.
- [86] Benjamin J Walker, Kenta Ishimoto, and Eamonn A Gaffney. Hydrodynamic slender-body theory for local rotation at zero reynolds number. *Physical Review Fluids*, 8(3):034101, 2023.
- [87] Wen Yan, Saad Ansari, Adam Lamson, Matthew A Glaser, Robert Blackwell, Meredith D Betterton, and Michael Shelley. Toward the cellular-scale simulation of motor-driven cytoskeletal assemblies. *Elife*, 11:e74160, 2022.
- [88] Sune Lomholt and Martin R Maxey. Force-coupling method for particulate two-phase flow: Stokes flow. *Journal of Computational Physics*, 184(2):381–405, 2003.
- [89] Florencio Balboa Usabiaga, Bakytzhan Kallemov, Blaise Delmotte, Amneet Bhalla, Boyce Griffith, and Aleksandar 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.
- [90] Boyce E Griffith, Richard D Hornung, David M McQueen, and Charles S Peskin. An adaptive, formally second order accurate version of the immersed boundary method. *J. Comput. Phys.*, 223(1):10–49, 2007.
- [91] Ondrej Maxian and Charles S Peskin. An immersed boundary method with subgrid resolution and improved numerical stability applied to slender bodies in stokes flow. *SIAM J. Sci. Comput.*, 42(4):847–868, 2020.

- [92] Daniel R Ladiges, A Nonaka, K Klymko, GC Moore, JB Bell, SP Carney, AL Garcia, SR Natesh, and A Donev. Discrete ion stochastic continuum overdamped solvent algorithm for modeling electrolytes. *Physical Review Fluids*, 6(4):044309, 2021.
- [93] Daniel R Ladiges, Jailun G Wang, Ishan Srivastava, A Nonaka, JB Bell, SP Carney, AL Garcia, and A Donev. Modeling electrokinetic flows with the discrete ion stochastic continuum overdamped solvent algorithm. *Physical Review E*, 106(3):035104, 2022.
- [94] Minghao W Rostami and Sarah D Olson. Kernel-independent fast multipole method within the framework of regularized stokeslets. *Journal of Fluids and Structures*, 67:60–84, 2016.
- [95] Helge I Andersson, Elena Celledoni, Laurel Ohm, Brynjulf Owren, and Benjamin K Tapley. An integral model based on slender body theory, with applications to curved rigid fibers. *Physics of Fluids*, 33(4):041904, 2021.
- [96] Harishankar Manikantan and David Saintillan. Subdiffusive transport of fluctuating elastic filaments in cellular flows. *Physics of Fluids*, 25(7):073603, 2013.
- [97] Y-N Young. Hydrodynamic interactions between two semiflexible inextensible filaments in stokes flow. *Physical Review E*, 79(4):046317, 2009.
- [98] Anna-Karin Tornberg and Katarina Gustavsson. A numerical method for simulations of rigid fiber suspensions. *Journal of Computational Physics*, 215(1):172–196, 2006.
- [99] Charles W Wolgemuth, Thomas R Powers, and Raymond E Goldstein. Twirling and whirling: Viscous dynamics of rotating elastic filaments. *Physical Review Letters*, 84(7):1623, 2000.
- [100] Sookkyung Lim and Charles S Peskin. Simulations of the whirling instability by the immersed boundary method. *SIAM Journal on Scientific Computing*, 25(6):2066–2083, 2004.

- [101] Wanho Lee, Yongsam Kim, Sarah D Olson, and Sookkyung Lim. Nonlinear dynamics of a rotating elastic rod in a viscous fluid. *Physical Review E*, 90(3):033012, 2014.
- [102] Isaac R Bruss, Heena K Mutha, Katherine Stoll, Brent Collins, Vinh Nguyen, David JD Carter, Michael P Brenner, and Kasey J Russell. Twirling, whirling, and tensioning: Plectoneme formation and suppression in flexible filaments. *Physical Review Research*, 1(3):032020, 2019.
- [103] Andrew M Fiore, Florencio Balboa Usabiaga, Aleksandar Donev, and James W Swan. Rapid sampling of stochastic displacements in brownian dynamics simulations. *The Journal of chemical physics*, 146(12):124116, 2017.
- [104] Anna-Karin Tornberg. Fft based spectral ewald methods as an alternative to fast multipole methods.
- [105] Raul P Perez. Universally adaptable multiscale molecular dynamics (uammd). <https://github.com/RaulPPelaez/UAMMD>, 2021.
- [106] C Barbier, R Dendievel, and D Rodney. Numerical study of 3d-compressions of entangled materials. *Computational Materials Science*, 45(3):593–596, 2009.
- [107] Taeyoon Kim, W Hwang, and RD Kamm. Computational analysis of a cross-linked actin-like network. *Experimental Mechanics*, 49(1):91–104, 2009.
- [108] Robert D Groot. Mesoscale simulation of semiflexible chains. i. endpoint distribution and chain dynamics. *The Journal of Chemical Physics*, 138(22):224903, 2013.
- [109] Fernando Vargas-Lara and Jack F Douglas. Fiber network formation in semi-flexible polymer solutions: an exploratory computational study. *Gels*, 4(2):27, 2018.
- [110] I Pincus, A Rodger, and J Ravi Prakash. Dilute polymer solutions under shear flow:

- comprehensive qualitative analysis using a bead-spring chain model with a fene-fraenkel spring. *arXiv preprint arXiv:2206.01870*, 2022.
- [111] Camilo Cruz, Francisco Chinesta, and Gilles Regnier. Review on the brownian dynamics simulation of bead-rod-spring models encountered in computational rheology. *Archives of Computational Methods in Engineering*, 19(2):227–259, 2012.
- [112] Venkata Siva Krishna, Praphul Kumar, Bharatkumar Sharma, and Indranil Saha Dalal. Petascale brownian dynamics simulations of highly resolved polymer chains with hydrodynamic interactions using modern gpus. *arXiv preprint arXiv:2208.06559*, 2022.
- [113] F. Balboa Usabiaga, B. Kallemov, B. Delmotte, A. P. S. Bhalla, B. E. 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, 2016. Software available at <https://github.com/stochasticHydroTools/RigidMultiblobsWall>.
- [114] David C Morse. Theory of constrained brownian motion. *Advances in Chemical Physics*, 128:65–189, 2003.
- [115] Hans C Öttinger. *Stochastic processes in polymeric fluids: tools and examples for developing simulation algorithms*. Springer Science & Business Media, 2012.
- [116] S. DeLong, F. Balboa Usabiaga, and A. Donev. Brownian dynamics of confined rigid bodies. *J. Chem. Phys.*, 143(14):144107, 2015. Software available at <https://github.com/stochasticHydroTools/RigidMultiblobsWall>.
- [117] Enrico Grazi. What is the diameter of the actin filament? *FEBS letters*, 405(3):249–252, 1997.
- [118] Paul A Janmey, Soren Hvidt, J Käs, Dietmar Lerche, Anthony Maggs, Erich Sackmann, Manfred Schliwa, and Thomas P Stossel. The mechanical properties of

- actin gels. elastic modulus and filament motions. *Journal of Biological Chemistry*, 269(51):32503–32513, 1994.
- [119] James L McGrath, Eric A Osborn, Yanik S Tardy, C Forbes Dewey, and John H Hartwig. Regulation of the actin cycle in vivo by actin filament severing. *Proceedings of the National Academy of Sciences*, 97(12):6532–6537, 2000.
- [120] Michael J Shelley and Tetsuji Ueda. The nonlocal dynamics of stretching, buckling filaments. *Advances in multi-fluid flows*, pages 415–425, 1996.
- [121] Michael J Shelley and Tetsuji Ueda. The stokesian hydrodynamics of flexing, stretching filaments. *Physica D: Nonlinear Phenomena*, 146(1-4):221–245, 2000.
- [122] Kai Liu, John Lowengrub, and Jun Allard. Efficient simulation of thermally fluctuating biopolymers immersed in fluids on 1-micron, 1-second scales. *Journal of computational physics*, 386:248–263, 2019.
- [123] Paul M Ryan and Charles W Wolgemuth. A finite volume algorithm for the dynamics of filaments, rods, and beams. *Journal of Computational Physics*, page 111375, 2022.
- [124] Christian J Cyron and Wolfgang A Wall. Finite-element approach to brownian dynamics of polymers. *Physical Review E*, 80(6):066704, 2009.
- [125] Yuan Lin, X Wei, J Qian, KY Sze, and VB Shenoy. A combined finite element-langevin dynamics (fem-ld) approach for analyzing the mechanical response of biopolymer networks. *Journal of the Mechanics and Physics of Solids*, 62:2–18, 2014.
- [126] CJ Cyron and WA Wall. Numerical method for the simulation of the brownian dynamics of rod-like microstructures with three-dimensional nonlinear beam elements. *International journal for numerical methods in engineering*, 90(8):955–987, 2012.

- [127] Christian J Cyron, Kei W Müller, Andreas R Bausch, and Wolfgang A Wall. Micromechanical simulations of biopolymer networks with finite elements. *Journal of Computational Physics*, 244:236–251, 2013.
- [128] Lloyd N Trefethen. *Spectral methods in MATLAB*, volume 10. Siam, 2000.
- [129] Martin Kröger, Jorge Ramirez, and Hans Christian Öttinger. Projection from an atomistic chain contour to its primitive path. *Polymer*, 43(2):477–487, 2002.
- [130] Jared L Aurentz and Lloyd N Trefethen. Block operators and spectral discretizations. *SIAM Review*, 59(2):423–446, 2017.
- [131] Tobin A Driscoll and Nicholas Hale. Rectangular spectral collocation. *IMA Journal of Numerical Analysis*, 36(1):108–132, 2015.
- [132] Dunzhu Li, Michael Gurnis, and Georg Stadler. Towards adjoint-based inversion of time-dependent mantle convection with nonlinear viscosity. *Geophysical Journal International*, 209(1):86–105, 2017.
- [133] Zewen Shen and Kirill Serkh. Is polynomial interpolation in the monomial basis unstable? *arXiv preprint arXiv:2212.10519*, 2023.
- [134] Benjamin J Walker and Eamonn A Gaffney. Regularised non-uniform segments and efficient no-slip elastohydrodynamics. *Journal of Fluid Mechanics*, 915, 2021.
- [135] Lorenzo Pareschi and Giovanni Russo. Implicit–explicit runge–kutta schemes and applications to hyperbolic systems with relaxation. *Journal of Scientific computing*, 25(1):129–155, 2005.
- [136] Steven Delong, Boyce E Griffith, Eric Vanden-Eijnden, and Aleksandar Donev. Temporal integrators for fluctuating hydrodynamics. *Physical Review E*, 87(3):033302, 2013.

- [137] Olinde Rodrigues. *Des lois géométriques qui régissent les déplacements d'un système solide dans l'espace: et de la variation des coordonnées provenant de ces déplacements considérés indépendamment des causes qui peuvent les produire.* 1840.
- [138] James W Swan and John F Brady. Simulation of hydrodynamically interacting particles near a no-slip boundary. *Physics of Fluids*, 19(11):113306, 2007.
- [139] NJ De Mestre and WB Russel. Low-reynolds-number translation of a slender cylinder near a plane wall. *Journal of Engineering Mathematics*, 9(2):81–91, 1975.
- [140] Ludvig af Klinteberg and Alex H Barnett. Accurate quadrature of nearly singular line integrals in two and three dimensions by singularity swapping. *BIT Numerical Mathematics*, pages 1–36, 2020.
- [141] Ludvig af Klinteberg, Davoud Saffar Shamshirgar, and Anna-Karin Tornberg. Fast ewald summation for free-space stokes potentials. *Research in the Mathematical Sciences*, 4(1):1–32, 2017.
- [142] Joar Bagge and Anna-Karin Tornberg. Fast ewald summation for stokes flow with arbitrary periodicity. *arXiv preprint arXiv:2210.01255*, 2022.
- [143] Hideki Kobayashi and Ryoichi Yamamoto. Implementation of lees–edwards periodic boundary conditions for direct numerical simulations of particle dispersions under shear flow. *The Journal of chemical physics*, 134(6):064110, 2011.
- [144] Dean R Wheeler, Norman G Rowley, and Richard L Fuller. Non-equilibrium molecular dynamics simulation of the shear viscosity of liquid methanol: adaptation of the ewald sum to lees-edwards boundary conditions. *Molecular Physics*, 92(1):55–62, 1997.
- [145] Hidenori Hasimoto. On the periodic fundamental solutions of the stokes equations and their application to viscous flow past a cubic array of spheres. *Journal of Fluid Mechanics*, 5(2):317–328, 1959.

- [146] Brennan Sprinkle, Florencio Balboa Usabiaga, Neelesh A. Patankar, and Aleksandar Donev. Large scale Brownian dynamics of confined suspensions of rigid particles. *J. Chem. Phys.*, 147(24):244103, 2017. Software available at <https://github.com/stochasticHydroTools/RigidMultiblobsWall>.
- [147] Edmond Chow and Yousef Saad. Preconditioned krylov subspace methods for sampling multivariate gaussian distributions. *SIAM Journal on Scientific Computing*, 36(2):A588–A608, 2014.
- [148] M. Hütter and H.C. Öttinger. Fluctuation-dissipation theorem, kinetic stochastic integral and efficient simulations. *J. Chem. Soc., Faraday Trans.*, 94(10):1403–1405, 1998.
- [149] Marshall Fixman. Simulation of polymer dynamics. i. general theory. *The Journal of Chemical Physics*, 69(4):1527–1537, 1978.
- [150] Timothy A Westwood, Blaise Delmotte, and Eric E Keaveny. A generalised drift-correcting time integration scheme for brownian suspensions of rigid particles with arbitrary shape. *Journal of Computational Physics*, 467:111437, 2022.
- [151] S. Delong, B. E. Griffith, E. Vanden-Eijnden, and A. Donev. Temporal Integrators for Fluctuating Hydrodynamics. *Phys. Rev. E*, 87(3):033302, 2013.
- [152] S. Delong, F. Balboa Usabiaga, R. Delgado-Buscalioni, B. E. Griffith, and A. Donev. Brownian dynamics without Green’s functions. *J. Chem. Phys.*, 140(13):134110, 2014. Software available at <https://github.com/stochasticHydroTools/FIB>.
- [153] Sander L Poelert and Amir A Zadpoor. Analytical and numerical methods for capturing the thermal fluctuations of semiflexible polymers. *Macromolecular theory and simulations*, 21(6):357–371, 2012.

- [154] Jan Wilhelm and Erwin Frey. Radial distribution function of semiflexible polymers. *Physical review letters*, 77(12):2581, 1996.
- [155] Benedikt Obermayer, Wolfram Möbius, Oskar Hallatschek, Erwin Frey, and Klaus Kroy. Freely relaxing polymers remember how they were straightened. *Physical Review E*, 79(2):021804, 2009.
- [156] Panagiotis Dimitrakopoulos. Longitudinal relaxation of initially straight flexible and stiff polymers. *Physical review letters*, 93(21):217801, 2004.
- [157] KM Zero and R Pecora. Rotational and translational diffusion in semidilute solutions of rigid-rod macromolecules. *Macromolecules*, 15(1):87–93, 1982.
- [158] Sophie Marbach and Christopher E Miles. Coarse-grained dynamics of transiently-bound fast linkers. *arXiv preprint arXiv:2212.08777*, 2022.
- [159] Daniel T Gillespie. Stochastic simulation of chemical kinetics. *Annu. Rev. Phys. Chem.*, 58:35–55, 2007.
- [160] David F Anderson. A modified next reaction method for simulating chemical systems with time dependent propensities and delays. *The Journal of chemical physics*, 127(21):214107, 2007.
- [161] Aleksandar Donev, Chiao-Yu Yang, and Changho Kim. Efficient reactive brownian dynamics. *The Journal of chemical physics*, 148(3):034103, 2018.
- [162] Michael P Allen and Dominic J Tildesley. *Computer simulation of liquids*. Oxford university press, 2017.
- [163] Tamara Carla Bidone, Wonyeong Jung, Daniel Maruri, Carlos Borau, Roger D Kamm, and Taeyoon Kim. Morphological transformation and force generation of active cytoskeletal networks. *PLoS computational biology*, 13(1):e1005277, 2017.

- [164] Denis James Evans, Debra Joy Searles, and Stephen Rodney Williams. *Fundamentals of classical statistical thermodynamics: dissipation, relaxation, and fluctuation theorems*. John Wiley & Sons, 2016.
- [165] Shimin Le, Xian Hu, Mingxi Yao, Hu Chen, Miao Yu, Xiaochun Xu, Naotaka Nakazawa, Felix M Margadant, Michael P Sheetz, and Jie Yan. Mechanotransmission and mechanosensing of human alpha-actinin 1. *Cell reports*, 21(10):2714–2723, 2017.
- [166] X Wei, C Fang, B Gong, J Yao, J Qian, and Y Lin. Viscoelasticity of 3d actin networks dictated by the mechanochemical characteristics of cross-linkers. *Soft Matter*, 2021.
- [167] Taeyoon Kim, Wonmuk Hwang, Hyungsuk Lee, and Roger D Kamm. Computational analysis of viscoelastic properties of crosslinked actin networks. *PLoS Comput Biol*, 5(7):e1000439, 2009.
- [168] David S Courson and Ronald S Rock. Actin cross-link assembly and disassembly mechanics for α -actinin and fascin. *Journal of Biological Chemistry*, 285(34):26350–26357, 2010.
- [169] Adam B Johnston, Agnieszka Collins, and Bruce L Goode. High-speed depolymerization at actin filament ends jointly catalysed by twinfilin and srv2/cap. *Nature cell biology*, 17(11):1504–1511, 2015.
- [170] Hao Yuan Kueh, Guillaume T Charras, Timothy J Mitchison, and William M Brieher. Actin disassembly by cofilin, coronin, and aip1 occurs in bursts and is inhibited by barbed-end cappers. *The Journal of cell biology*, 182(2):341–353, 2008.
- [171] Angelika Manhart, Aleksandra Icheva, Christophe Guerin, Tobias Klar, Rajaa Boujemaa-Paterski, Manuel Thery, Laurent Blanchoin, and Alex Mogilner. Recon-

- stitution of the equilibrium state of dynamic actin networks. *bioRxiv*, page 437806, 2018.
- [172] Helge Holden, Kenneth H Karlsen, and Knut-Andreas Lie. *Splitting methods for partial differential equations with rough solutions: Analysis and MATLAB programs*, volume 11. European Mathematical Society, 2010.
- [173] Brennan Sprinkle, Ernest B Van Der Wee, Yixiang Luo, Michelle M Driscoll, and Aleksandar Donev. Driven dynamics in dense suspensions of microrollers. *Soft Matter*, 16(34):7982–8001, 2020.
- [174] Konstantin Popov, James Komianos, and Garegin A Papoian. Medyan: Mechanochemical simulations of contraction and polarity alignment in actomyosin networks. *PLoS computational biology*, 12(4):e1004877, 2016.
- [175] Zydrunas Gimbutas and Leslie Greengard. Computational software: Simple fmm libraries for electrostatics, slow viscous flow, and frequency-domain wave propagation. *Communications in Computational Physics*, 18(2):516–528, 2015.
- [176] Dag Lindbo and Anna-Karin Tornberg. Fast and spectrally accurate summation of 2-periodic stokes potentials. *arXiv preprint arXiv:1111.1815*, 2011.
- [177] Wen Yan and Michael Shelley. Universal image systems for non-periodic and periodic stokes flows above a no-slip wall. *Journal of Computational Physics*, 375:263–270, 2018.
- [178] Andrew M Fiore and James W Swan. Rapid sampling of stochastic displacements in brownian dynamics simulations with stresslet constraints. *The Journal of chemical physics*, 148(4):044114, 2018.
- [179] Anna Broms and Anna-Karin Tornberg. A barrier method for contact avoiding particles in stokes flow. *arXiv preprint arXiv:2301.01666*, 2023.

- [180] Wen Yan, Huan Zhang, and Michael J Shelley. Computing collision stress in assemblies of active spherocylinders: Applications of a fast and generic geometric method. *The Journal of chemical physics*, 150(6):064109, 2019.
- [181] Zachary Ferguson, Minchen Li, Teseo Schneider, Francisca Gil-Ureta, Timothy Langlois, Chenfanfu Jiang, Denis Zorin, Danny M Kaufman, and Daniele Panozzo. Intersection-free rigid body dynamics. *ACM Transactions on Graphics*, 40(4), 2021.
- [182] Mohammad RK Mofrad. Rheology of the cytoskeleton. *Annual Review of Fluid Mechanics*, 41:433–453, 2009.
- [183] Bernhard Hinner, Markus Tempel, Erich Sackmann, Klaus Kroy, and Erwin Frey. Entanglement, elasticity, and viscous relaxation of actin solutions. *Physical Review Letters*, 81(12):2614, 1998.
- [184] Daniel H Wachsstock, WH Schwartz, and Thomas D Pollard. Affinity of alpha-actinin for actin determines the structure and mechanical properties of actin filament gels. *Biophysical journal*, 65(1):205, 1993.
- [185] O Lieleg, Mireille Maria Anna Elisabeth Claessens, Y Luan, and AR Bausch. Transient binding and dissipation in cross-linked actin networks. *Physical review letters*, 101(10):108101, 2008.
- [186] Loïc Chaubet, Abdullah R Chaudhary, Hossein K Heris, Allen J Ehrlicher, and Adam G Hendricks. Dynamic actin cross-linking governs the cytoplasm’s transition to fluid-like behavior. *Molecular biology of the cell*, 31(16):1744–1752, 2020.
- [187] Daniel H Wachsstock, WH Schwarz, and TD Pollard. Cross-linker dynamics determine the mechanical properties of actin gels. *Biophysical journal*, 66(3):801–809, 1994.
- [188] ML Gardel, Jennifer Hyunjong Shin, FC MacKintosh, L Mahadevan, P Matsudaira,

- and DA Weitz. Elastic behavior of cross-linked and bundled actin networks. *Science*, 304(5675):1301–1305, 2004.
- [189] ML Gardel, Jennifer Hyunjong Shin, FC MacKintosh, L Mahadevan, PA Matsudaira, and DA Weitz. Scaling of f-actin network rheology to probe single filament elasticity and dynamics. *Physical review letters*, 93(18):188102, 2004.
- [190] O Müller, HE Gaub, M Bärmann, and E Sackmann. Viscoelastic moduli of sterically and chemically cross-linked actin networks in the dilute to semidilute regime: measurements by oscillating disk rheometer. *Macromolecules*, 24(11):3111–3120, 1991.
- [191] R Ruddies, WH Goldmann, G Isenberg, and E Sackmann. The viscoelasticity of entangled actin networks: the influence of defects and modulation by talin and vinculin. *European biophysics journal*, 22(5):309–321, 1993.
- [192] Norman Y Yao, Daniel J Becker, Chase P Broedersz, Martin Depken, Frederick C MacKintosh, Martin R Pollak, and David A Weitz. Nonlinear viscoelasticity of actin transiently cross-linked with mutant α -actinin-4. *Journal of molecular biology*, 411(5):1062–1071, 2011.
- [193] Yuval Mulla, FC MacKintosh, and Gijsje H Koenderink. Origin of slow stress relaxation in the cytoskeleton. *Physical review letters*, 122(21):218102, 2019.
- [194] O Lieleg, KM Schmoller, Mireille Maria Anna Elisabeth Claessens, and AR Bausch. Cytoskeletal polymer networks: viscoelastic properties are determined by the microscopic interaction potential of cross-links. *Biophysical journal*, 96(11):4725–4732, 2009.
- [195] Norman Y Yao, Chase P Broedersz, Martin Depken, Daniel J Becker, Martin R Pollak, Frederick C MacKintosh, and David A Weitz. Stress-enhanced gelation: A dynamic nonlinearity of elasticity. *Physical review letters*, 110(1):018103, 2013.

- [196] David C Morse. Viscoelasticity of concentrated isotropic solutions of semiflexible polymers. 1. model and stress tensor. *Macromolecules*, 31(20):7030–7043, 1998.
- [197] Ben Fabry, Geoffrey N Maksym, James P Butler, Michael Glogauer, Daniel Navajas, and Jeffrey J Fredberg. Scaling the microrheology of living cells. *Physical review letters*, 87(14):148102, 2001.
- [198] Ben Fabry, Geoffrey N Maksym, James P Butler, Michael Glogauer, Daniel Navajas, Nathan A Taback, Emil J Millet, and Jeffrey J Fredberg. Time scale and other invariants of integrative mechanical behavior in living cells. *Physical Review E*, 68(4):041914, 2003.
- [199] F Ziemann, J Rädler, and E Sackmann. Local measurements of viscoelastic moduli of entangled actin networks using an oscillating magnetic bead micro-rheometer. *Biophysical journal*, 66(6):2210–2216, 1994.
- [200] Prince E Rouse Jr. A theory of the linear viscoelastic properties of dilute solutions of coiling polymers. *The Journal of Chemical Physics*, 21(7):1272–1280, 1953.
- [201] Simon L Freedman, Cristian Suarez, Jonathan D Winkelman, David R Kovar, Gregory A Voth, Aaron R Dinner, and Glen M Hocky. Mechanical and kinetic factors drive sorting of f-actin cross-linkers on bundles. *Proceedings of the National Academy of Sciences*, 116(33):16192–16197, 2019.
- [202] KE Kasza, CP Broedersz, GH Koenderink, YC Lin, W Messner, EA Millman, F Nakamura, TP Stossel, FC MacKintosh, and DA Weitz. Actin filament length tunes elasticity of flexibly cross-linked actin networks. *Biophysical journal*, 99(4):1091–1100, 2010.
- [203] Oliver Lieleg, Mireille MAE Claessens, and Andreas R Bausch. Structure and dynamics of cross-linked actin networks. *Soft Matter*, 6(2):218–225, 2010.

- [204] Yiider Tseng, Kwang M An, and Denis Wirtz. Microheterogeneity controls the rate of gelation of actin filament networks. *Journal of Biological Chemistry*, 277(20):18143–18150, 2002.
- [205] Yiider Tseng, Benjamin W Schafer, Steven C Almo, and Denis Wirtz. Functional synergy of actin filament cross-linking proteins. *Journal of Biological Chemistry*, 277(28):25609–25616, 2002.
- [206] KM Schmoller, O Lieleg, and AR Bausch. Structural and viscoelastic properties of actin/filamin networks: cross-linked versus bundled networks. *Biophysical journal*, 97(1):83–89, 2009.
- [207] Li Hou, Katherine Luby-Phelps, and Frederick Lanni. Brownian motion of inert tracer macromolecules in polymerized and spontaneously bundled mixtures of actin and filamin. *The Journal of cell biology*, 110(5):1645–1654, 1990.
- [208] Oliver Lieleg, Kurt M Schmoller, Christian J Cyron, Yuxia Luan, Wolfgang A Wall, and Andreas R Bausch. Structural polymorphism in heterogeneous cytoskeletal networks. *Soft Matter*, 5(9):1796–1803, 2009.
- [209] Tobias T Falzone, Martin Lenz, David R Kovar, and Margaret L Gardel. Assembly kinetics determine the architecture of α -actinin crosslinked f-actin networks. *Nature communications*, 3(1):1–9, 2012.
- [210] Giulia Foffano, Nicolas Levernier, and Martin Lenz. The dynamics of filament assembly define cytoskeletal network morphology. *Nature communications*, 7(1):1–8, 2016.
- [211] Rui Ma and Julien Berro. Structural organization and energy storage in crosslinked actin assemblies. *PLoS computational biology*, 14(5):e1006150, 2018.
- [212] Kimberly L Weirich, Shiladitya Banerjee, Kinjal Dasbiswas, Thomas A Witten, Suriyanarayanan Vaikuntanathan, and Margaret L Gardel. Liquid behavior of cross-linked

- actin bundles. *Proceedings of the National Academy of Sciences*, 114(9):2131–2136, 2017.
- [213] Marcus Tempel, Gerhard Isenberg, and Erich Sackmann. Temperature-induced sol-gel transition and microgel formation in α -actinin cross-linked actin networks: a rheological study. *Physical Review E*, 54(2):1802, 1996.
- [214] FC MacKintosh, Josef Käs, and PA Janmey. Elasticity of semiflexible biopolymer networks. *Physical review letters*, 75(24):4425, 1995.
- [215] H Isambert and AC Maggs. Dynamics and rheology of actin solutions. *Macromolecules*, 29(3):1036–1040, 1996.
- [216] Adrian F Pegoraro, Paul Janmey, and David A Weitz. Mechanical properties of the cytoskeleton and cells. *Cold Spring Harbor perspectives in biology*, 9(11):a022038, 2017.
- [217] Chase P Broedersz, Martin Depken, Norman Y Yao, Martin R Pollak, David A Weitz, and Frederick C MacKintosh. Cross-link-governed dynamics of biopolymer networks. *Physical review letters*, 105(23):238101, 2010.
- [218] Kei W Müller, Robijn F Bruinsma, Oliver Lieleg, Andreas R Bausch, Wolfgang A Wall, and Alex J Levine. Rheology of semiflexible bundle networks with transient linkers. *Physical review letters*, 112(23):238102, 2014.
- [219] Itamar Borukhov, Robijn F Bruinsma, William M Gelbart, and Andrea J Liu. Structural polymorphism of the cytoskeleton: a model of linker-assisted filament aggregation. *Proceedings of the National Academy of Sciences*, 102(10):3673–3678, 2005.
- [220] AG Zilman and SA Safran. Role of cross-links in bundle formation, phase separation and gelation of long filaments. *EPL (Europhysics Letters)*, 63(1):139, 2003.

- [221] Jan Kierfeld, Torsten Kühne, and Reinhard Lipowsky. Discontinuous unbinding transitions of filament bundles. *Physical review letters*, 95(3):038102, 2005.
- [222] Xueping Yu and AE Carlsson. Kinetics of filament bundling with attractive interactions. *Biophysical journal*, 87(6):3679–3689, 2004.
- [223] Hélene Karcher, Jan Lammerding, Hayden Huang, Richard T Lee, Roger D Kamm, and Mohammad R Kaazempur-Mofrad. A three-dimensional viscoelastic model for cell deformation with experimental verification. *Biophysical journal*, 85(5):3336–3349, 2003.
- [224] Srboľjub M Mijailovich, Milos Kojic, Miroslav Zivkovic, Ben Fabry, and Jeffrey J Fredberg. A finite element model of cell deformation during magnetic bead twisting. *Journal of Applied Physiology*, 93(4):1429–1436, 2002.
- [225] Wanda Strychalski, Calina A Copos, Owen L Lewis, and Robert D Guy. A poroelastic immersed boundary method with applications to cell biology. *Journal of Computational Physics*, 282:77–97, 2015.
- [226] Calina A Copos and Robert D Guy. A porous viscoelastic model for the cell cytoskeleton. *The ANZIAM Journal*, 59(4):472–498, 2018.
- [227] Robert D Guy, Toshiyuki Nakagaki, and Grady B Wright. Flow-induced channel formation in the cytoplasm of motile cells. *Physical Review E*, 84(1):016310, 2011.
- [228] Nicolas Desprat, Alain Richert, Jacqueline Simeon, and Atef Asnacios. Creep function of a single living cell. *Biophysical journal*, 88(3):2224–2233, 2005.
- [229] Katherine Luby-Phelps. Cytoarchitecture and physical properties of cytoplasm: volume, viscosity, diffusion, intracellular surface area. In *International review of cytology*, volume 192, pages 189–221. Elsevier, 1999.

- [230] Rudolf K Meyer and Ueli Aebi. Bundling of actin filaments by alpha-actinin depends on its molecular length. *The Journal of cell biology*, 110(6):2013–2024, 1990.
- [231] Philip A Kuhlman, Jacqueline Ellis, David R Critchley, and Clive R Bagshaw. The kinetics of the interaction between the actin-binding domain of α -actinin and f-actin. *FEBS letters*, 339(3):297–301, 1994.
- [232] Jingyuan Xu, Denis Wirtz, and Thomas D Pollard. Dynamic cross-linking by α -actinin determines the mechanical properties of actin filament networks. *Journal of Biological Chemistry*, 273(16):9570–9576, 1998.
- [233] Faith A Morrison et al. *Understanding rheology*. Oxford University Press, USA, 2001.
- [234] Frédéric Eghiaian, Annafrancesca Rigato, and Simon Scheuring. Structural, mechanical, and dynamical variability of the actin cortex in living cells. *Biophysical journal*, 108(6):1330–1340, 2015.
- [235] Masato Makino and Masao Doi. Brownian motion of a particle of general shape in newtonian fluid. *Journal of the Physical Society of Japan*, 73(10):2739–2745, 2004.
- [236] Clifford P Brangwynne, FC MacKintosh, and David A Weitz. Force fluctuations and polymerization dynamics of intracellular microtubules. *Proceedings of the National Academy of Sciences*, 104(41):16128–16133, 2007.
- [237] Minakshi Guha, Mian Zhou, and Yu-li Wang. Cortical actin turnover during cytokinesis requires myosin ii. *Current biology*, 15(8):732–736, 2005.
- [238] Julie A Theriot and Timothy J Mitchison. Actin microfilament dynamics in locomoting cells. *Nature*, 352(6331):126–131, 1991.
- [239] Eric SG Shaqfeh and Glenn H Fredrickson. The hydrodynamic stress in a suspension of rods. *Physics of Fluids A: Fluid Dynamics*, 2(1):7–24, 1990.

- [240] Michael B Mackaplow and Eric SG Shaqfeh. A numerical study of the rheological properties of suspensions of rigid, non-brownian fibres. *Journal of Fluid Mechanics*, 329:155–186, 1996.
- [241] H Henning Winter. Analysis of dynamic mechanical data: inversion into a relaxation time spectrum and consistency check. *Journal of Non-Newtonian Fluid Mechanics*, 68(2-3):225–239, 1997.
- [242] Ian McDougall, Nese Orbey, and John M Dealy. Inferring meaningful relaxation spectra from experimental data. *Journal of Rheology*, 58(3):779–797, 2014.
- [243] M Baumgaertel and HH Winter. Determination of discrete relaxation and retardation time spectra from dynamic mechanical data. *Rheologica Acta*, 28(6):511–519, 1989.
- [244] A Ya Malkin. The use of a continuous relaxation spectrum for describing the viscoelastic properties of polymers. *Polymer Science Series A*, 48(1):39–45, 2006.
- [245] Ernesto Alva, Annitta George, Lorenzo Brancaleon, and Marcelo Marucho. Hydrodynamic and polyelectrolyte properties of actin filaments: Theory and experiments. *Polymers*, 14(12):2438, 2022.
- [246] Taeyoon Kim, Margaret L Gardel, and ED Munro. Determinants of fluidlike behavior and effective viscosity in cross-linked actin networks. *Biophysical journal*, 106(3):526–534, 2014.
- [247] Jonathan Stricker, Tobias Falzone, and Margaret L Gardel. Mechanics of the f-actin cytoskeleton. *Journal of biomechanics*, 43(1):9–14, 2010.
- [248] Jan A Åström, PB Sunil Kumar, Ilpo Vattulainen, and Mikko Karttunen. Strain hardening, avalanches, and strain softening in dense cross-linked actin networks. *Physical Review E*, 77(5):051913, 2008.

- [249] Danielle R Scheff, Steven A Redford, Chatipat Lorpaiboon, Sayantan Majumdar, Aaron R Dinner, and Margaret L Gardel. Actin filament alignment causes mechanical hysteresis in cross-linked networks. *Soft Matter*, 17(22):5499–5507, 2021.
- [250] G Bossis and JF Brady. The rheology of brownian suspensions. *The Journal of chemical physics*, 91(3):1866–1874, 1989.
- [251] John F Brady. Brownian motion, hydrodynamics, and the osmotic pressure. *The Journal of chemical physics*, 98(4):3335–3341, 1993.
- [252] Le Yang, David Sept, and AE Carlsson. Energetics and dynamics of constrained actin filament bundling. *Biophysical journal*, 90(12):4295–4304, 2006.
- [253] Aravind Chandrasekaran, Arpita Upadhyaya, and Garegin A Papoian. Remarkable structural transformations of actin bundles are driven by their initial polarity, motor activity, crosslinking, and filament treadmilling. *PLoS computational biology*, 15(7):e1007156, 2019.
- [254] Gregory M Grason. Colloquium: Geometry and optimal packing of twisted columns and filaments. *Reviews of Modern Physics*, 87(2):401, 2015.
- [255] Jörg Schnauß, Tina Händler, and Josef A Käs. Semiflexible biopolymers in bundled arrangements. *Polymers*, 8(8):274, 2016.
- [256] Yee Han Tee, Tom Shemesh, Visalatchi Thiagarajan, Rizal Fajar Hariadi, Karen L Anderson, Christopher Page, Niels Volkmann, Dorit Hanein, Sivaraj Sivaramakrishnan, Michael M Kozlov, et al. Cellular chirality arising from the self-organization of the actin cytoskeleton. *Nature cell biology*, 17(4):445–457, 2015.
- [257] Kimberly L Weirich, Samantha Stam, Edwin Munro, and Margaret L Gardel. Actin bundle architecture and mechanics regulate myosin ii force generation. *Biophysical Journal*, 120(10):1957–1970, 2021.