





RESEARCH ARTICLE | APRIL 20 2023

Bending fluctuations in semiflexible, inextensible, slender filaments in Stokes flow: Toward a spectral discretization

Ondrej Maxian  ; Brennan Sprinkle ; Aleksandar Donev 



J. Chem. Phys. 158, 154114 (2023)

<https://doi.org/10.1063/5.0144242>

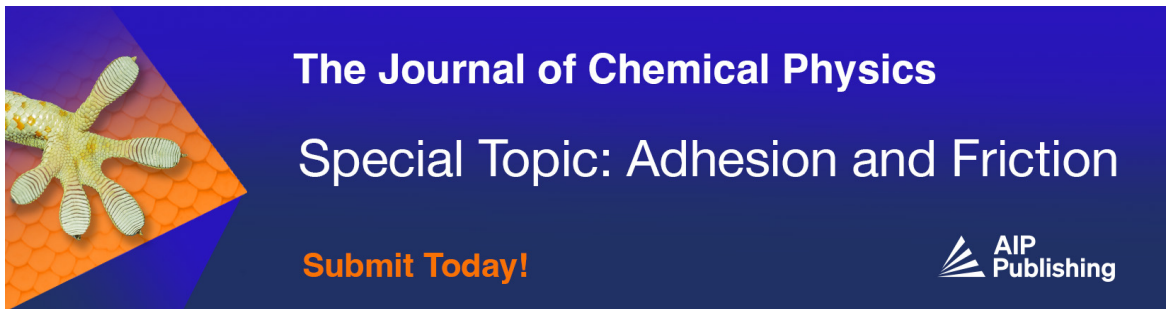


View
Online





Export
Citation

CrossMark



The Journal of Chemical Physics
Special Topic: Adhesion and Friction

Submit Today!



Bending fluctuations in semiflexible, inextensible, slender filaments in Stokes flow: Toward a spectral discretization

Cite as: J. Chem. Phys. 158, 154114 (2023); doi: 10.1063/5.0144242

Submitted: 28 January 2023 • Accepted: 30 March 2023 •

Published Online: 20 April 2023



View Online



Export Citation



CrossMark

Ondrej Maxian,^{1,a)}  Brennan Sprinkle,²  and Aleksandar Donev¹ 

AFFILIATIONS

¹Department of Mathematics, Courant Institute, New York University, New York, New York 10012, USA

²Department of Applied Mathematics and Statistics, Colorado School of Mines, Golden, Colorado 80401, USA

^{a)}Author to whom correspondence should be addressed: om759@nyu.edu

ABSTRACT

Semiflexible slender filaments are ubiquitous in nature and cell biology, including in the cytoskeleton, where reorganization of actin filaments allows the cell to move and divide. Most methods for simulating semiflexible inextensible fibers/polymers are based on discrete (bead-link or blob-link) models, which become prohibitively expensive in the slender limit when hydrodynamics is accounted for. In this paper, we develop a novel coarse-grained approach for simulating fluctuating slender filaments with hydrodynamic interactions. Our approach is tailored to relatively stiff fibers whose persistence length is comparable to or larger than their length and is based on three major contributions. First, we discretize the filament centerline using a coarse non-uniform Chebyshev grid, on which we formulate a *discrete* constrained Gibbs–Boltzmann (GB) equilibrium distribution and overdamped Langevin equation for the evolution of unit-length tangent vectors. Second, we define the hydrodynamic mobility at each point on the filament as an integral of the Rotne–Prager–Yamakawa kernel along the centerline and apply a spectrally accurate “slender-body” quadrature to accurately resolve the hydrodynamics. Third, we propose a novel midpoint temporal integrator, which can correctly capture the Ito drift terms that arise in the overdamped Langevin equation. For two separate examples, we verify that the equilibrium distribution for the Chebyshev grid is a good approximation of the blob-link one and that our temporal integrator for overdamped Langevin dynamics samples the equilibrium GB distribution for sufficiently small time step sizes. We also study the dynamics of relaxation of an initially straight filament and find that as few as 12 Chebyshev nodes provide a good approximation to the dynamics while allowing a time step size two orders of magnitude larger than a resolved blob-link simulation. We conclude by applying our approach to a suspension of cross-linked semiflexible fibers (neglecting hydrodynamic interactions between fibers), where we study how semiflexible fluctuations affect bundling dynamics. We find that semiflexible filaments bundle faster than rigid filaments even when the persistence length is large, but show that semiflexible *bending* fluctuations only further accelerate agglomeration when the persistence length and fiber length are of the same order.

Published under an exclusive license by AIP Publishing. <https://doi.org/10.1063/5.0144242>

I. INTRODUCTION

The closer we look at biological systems, the more we find slender filaments performing important work. These filaments are responsible for cell motility and division in prokaryotes^{1,2} and eukaryotes,^{3–6} which makes them indispensable for processes such as wound healing, stem cell differentiation, and organism growth. In cells, three kinds of filaments can be distinguished: microtubules, which are sufficiently stiff as to behave deterministically;⁷ intermediate filaments, which have small persistence length and are

consequently found in entangled networks *in vivo*;^{8–10} and actin filaments, which have visible bending fluctuations driven by thermal forces, but have sufficient stiffness to maintain a relatively smooth appearance of the filament centerline.¹¹

We are motivated here by actin filaments, which have been shown to take on a range of morphologies when combined with cross-linking proteins *in vivo*,¹² *in vitro*,¹³ and *in silico*.^{14,15} Our interest, in particular, is how thermal fluctuations, hydrodynamic interactions, and cross-linking compete or cooperate with each other to determine the steady state morphology and stress–strain behavior

of actin networks. While there have been a number of coarse-grained theories examining this question,^{16–18} there is still a need for detailed simulation of each of the microscopic components in the system so that assumptions made in deriving coarse-grained theories can be validated and tested more rigorously. This was the idea behind our previous work on *deterministic* actin filaments,¹⁵ which looked at how each of the system components contributes to the behavior of the cross-linked actin network on short and long timescales. Our ultimate goal is to extend the study of Ref. 15 to *fluctuating* (Brownian) actin filaments so that we can determine how thermal fluctuations affect the viscoelastic gel behavior.

With this goal in mind, we turn here to the simulation of semi-flexible filaments interacting through a viscous medium. 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 that 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.^{19–23} In this case, the chain is represented by a series of beads or blobs connected by springs that 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 (Ref. 24, 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 [see Ref. 25 for the limit on modern graphics processing units (GPUs)].

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 (Ref. 24, Sec. 4) (we will use the lesser-used “blob-link” terminology because of the coupling with “blob-based” hydrodynamics²⁶). 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.^{27,28} An alternative view, which is the 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.²⁹

The main issue with the blob-link model in our context arises when we account for the hydrodynamics of the chain. Denoting the position of bead i by $\mathbf{X}_{[i]}$, hydrodynamic interactions give us a mobility matrix $\tilde{\mathbf{M}}(\mathbf{X})$, which describes the relationship between the forces on the beads \mathbf{F} and their velocities \mathbf{U} via $\mathbf{U} = \tilde{\mathbf{M}}\mathbf{F}$. For blob-link or bead-spring models in unbounded domains, the

obvious choice for the mobility matrix is to use a pairwise hydrodynamic kernel between the blobs/beads. One such kernel is the Rotne–Prager–Yamakawa (RPY) tensor,^{30,31} which is obtained by approximately solving the mobility problem for a pair of spheres of radius \hat{a} an arbitrary distance apart in the fluid. That is, the 3×3 mobility block associated with beads i and j in a blob-link model is

$$\tilde{\mathbf{M}}_{\{i\}\{j\}} = \tilde{\mathbf{M}}_{\text{RPY}}(\mathbf{X}_{\{i\}}, \mathbf{X}_{\{j\}}; \hat{a}), \quad (1)$$

where the 3×3 matrix $\tilde{\mathbf{M}}_{\text{RPY}}(\mathbf{x}, \mathbf{y}; \hat{a})$ is defined in (B2). While this kernel and others like it neglect interactions involving more than two beads, it is commonly used in polymer physics to describe the hydrodynamics of a chain^{32–35} because the RPY kernel is symmetric positive definite for any locations of the two particles. This fact can also be seen by using an immersed boundary formulation to justify (1) (Ref. 36, Sec. 2.1). One contribution of this paper will be to develop a temporal integrator for a fluctuating blob-link chain with mobility given by (1).

As we will show, however, the blob-link model breaks down in the limit of very slender fibers, such as actin filaments, which have a radius of $a = 4 \text{ nm}$ ³⁷ and can have lengths L of a few micrometers *in vivo* and tens of micrometers *in vitro*.^{38,39} This set of parameters gives aspect ratios on the order of $\epsilon = a/L = 10^{-3}$, making the hydrodynamics multiscale in nature. By using a blob-link model with blobs of radius $\sim a$, we are implicitly marrying ourselves to the resolution of this smallest lengthscale. If we wanted to accurately simulate very flexible filaments or filaments that are nearly touching, then there would be no way around this. However, for smoother filament shapes and non-dense suspensions where we are interested in looking at the bulk behavior, a heavy price is paid: since the hydrodynamics we are really interested in occur on lengthscales $L = a/\epsilon$, the number of required blobs to resolve the hydrodynamics of a slender filaments is approximately $L/a \sim 1/\epsilon$.^{40,41} In addition to being prohibitively expensive in the spatial discretization, this model also presents a problem for temporal integration since having 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, a sensible resolution in the slender limit is to take a continuum limit of the discrete model so that 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 (Refs. 42 and 43; see Fig. 1). Likewise, the links become a continuous function $\mathbb{T}(s) = \partial_s \mathbb{X}(s)$, and the force on each link becomes a force density $\mathbf{f}(s)$ that is defined everywhere along the filament. The hydrodynamic mobility $\mathcal{M}[\mathbb{X}(\cdot)]$ that maps force density to velocity can then be defined as the integral operator,

$$\mathbf{U}(s) = \int_0^L \mathbf{M}_{\text{RPY}}(\mathbb{X}(s), \mathbb{X}(s'); \hat{a}) \mathbf{f}(s') ds' := (\mathcal{M}\mathbf{f})(s). \quad (2)$$

In Ref. 36, Appendix A, we showed that the RPY integral (2) is asymptotically equivalent to slender body theory (SBT), which is the more commonly used mobility for continuum curves in Stokes




FIG. 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 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 2 shows a more detailed look at this discretization.

flow,^{44–47} if the relationship between the true filament radius a and the regularization radius \hat{a} is given by

$$\hat{a} = \frac{e^{3/2}}{4} a \approx 1.1204a. \quad (3)$$

This statement is only true for the translation–translation component of the RPY kernel; see Ref. 48 for how this choice of radius performs in the context of rotation, as well as Ref. 49 for the corresponding analysis for regularized Stokeslets,⁵⁰ and Ref. 40 for a numerical analysis for the immersed boundary method.⁵¹

To simulate this continuum model, we still ultimately need to choose collocation points at which to represent the fiber positions and tangent vectors. Because we are interested in actin filaments, for which the fiber shapes are relatively smooth, an attractive discretization is a spectral one, where the collocation points are chosen from a Chebyshev grid in the arclength s . The motivation for using Chebyshev points comes from classical numerical analysis; since Chebyshev polynomials give spectrally accurate polynomial interpolants, the error in approximating smooth fiber shapes decreases exponentially in the number of nodes, similar to using a Fourier basis for a closed loop fiber (Ref. 52, Chap. 10). In addition, the well-conditioned property of the Chebyshev interpolation matrix allows us to obtain globally accurate interpolants to the fiber shape, which are simpler than breaking the fiber into panels. In previous work on deterministic filaments,^{36,53} we developed numerical methods for such global Chebyshev discretizations⁵³ and accompanying quadrature schemes (Ref. 36, Appendix G) for integral (2) on the spectral grid. Using these schemes, we showed that the number of collocation points required to resolve integral (2) is (roughly) independent of the fiber aspect ratio ϵ (Ref. 36, Sec. 4.4), which makes this mobility definition, and the spectral grid that accompanies it, an appealing choice for *smooth* slender filaments. In this paper, we

will conduct the same analysis for *fluctuating* (Brownian) semiflexible filaments, showing that there are indeed significant advantages to using this model of the mobility over the discrete model, at least for very slender filaments.

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 $\mathbb{T}(s)$ 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 show in Appendix C 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. 22, 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 *discrete* orientations of the tangent vectors and do not converge in the continuum limit. This issue has either not been treated in previous continuum methods for filament Brownian dynamics^{54,55} or been partially sidestepped by making the filament extensible through a penalty force.^{56,57} The latter approach, which has characterized a number of finite element methods for biopolymers,^{58–61} 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. 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 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. However, by using mobility (2) 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.

Seeing that this is, to our knowledge, the first paper 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 the 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 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 radius), many quantities should be approximated well since scales 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 the conclusion of this paper laying out much of the work that remains.

II. LANGEVIN EQUATION FOR SEMIFLEXIBLE FILAMENTS

In this section, we formulate the overdamped Langevin equation for the evolution of a discrete inextensible semiflexible filament. As discussed in the Introduction, we propose a fully discrete spectral discretization composed of tangent vectors τ , which are defined on a Chebyshev grid of size N , and node positions X , which are defined on a Chebyshev grid of size $N + 1$. Once we formulate the discrete filament model, we can write the Gibbs–Boltzmann probability distribution for the fiber energy. The overdamped Langevin equation then follows once we define the dynamics of how the fibers evolve deterministically.

A. Discrete fiber model

Let us first formulate a spectral discretization of an inextensible filament, which is well-suited for fluctuating hydrodynamics. As discussed in the Introduction and shown in Fig. 2, we choose to track a discrete collection of tangent vectors τ , which evolve as rigid rods by rotating on the unit sphere. To completely define the fiber, we need to also track the fiber midpoint X_{MP} . These two quantities give a set of node positions X , which in turn define a smooth polynomial interpolant $\mathbb{X}(s)$ for the fiber centerline. We use this interpolant to compute elastic energy and, importantly, hydrodynamic interactions via (2). Once we define a discretization, we can postulate the constrained Gibbs–Boltzmann distribution in terms of the tangent vectors τ . This distribution depends on the fiber elastic energy \mathcal{E}_{bend} , which we discretize in the final part of this section.

1. Discretization

We define a spectral discretization of an inextensible filament by transferring the blob-link discretization concept onto a non-uniform Chebyshev grid (see Fig. 2). We begin with a collection of N unit-length tangent vectors $\tau = \{\tau_{\{p\}}\}_{p=1}^N$, with $\tau_{\{p\}} \cdot \tau_{\{p\}} = 1$ for $p = 1, \dots, N$, on a type 1 Chebyshev grid [i.e., a Chebyshev grid that does not include the endpoints; here and throughout this paper, τ refers to a $3N \times 1$ column-stacked vector of each 3×1 tangent vector $\tau_{\{p\}}$; the (scalar) p th entry of this vector is denoted as τ_p]. To obtain the fiber position, we oversample τ to a (type 2, endpoints-included) Chebyshev grid of size $N + 1$ using the evaluation matrix $E_{N \rightarrow N+1}$ and then integrate the result exactly using the Chebyshev integration matrix D_{N+1}^\dagger (pseudo-inverse of the differentiation matrix D_{N+1}). This can be written as

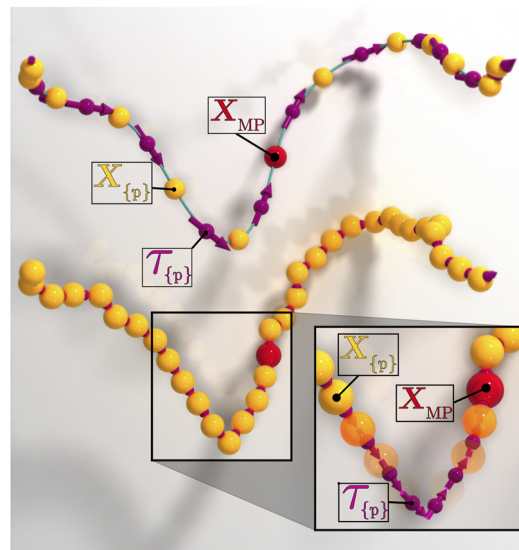


FIG. 2. Transferring the fiber discretization for the blob-link model (bottom) to a spectral grid (top). We show the tangent vector locations as magenta vectors visible through the “x-ray” inset—for the blob-link discretization, these are the points exactly between the beads, while for the spectral discretization, they are defined on a type 1 Chebyshev grid with N points ($N = 12$ here). The positions, which can be obtained from τ and 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 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.

$$X = \left(D_{N+1}^\dagger E_{N \rightarrow N+1} \quad B \right) \begin{pmatrix} \tau \\ X_{MP} \end{pmatrix} := \mathcal{X} \bar{\tau}. \quad (4)$$

The matrix B is such that the midpoint of X on the grid of size $N + 1$ is X_{MP} (when N is even, 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 X are defined on a grid of size $N + 1$, whereas τ 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 \mathcal{X} , which is done by differentiating X on the $N + 1$ point grid and then down-sampling to the grid of size N via the matrix $E_{N+1 \rightarrow N}$. The midpoint is determined from the $N + 1$ nodes via sampling the $N + 1$ -term interpolating polynomial at the midpoint using the matrix $E_{N+1 \rightarrow MP}$. Together, this gives

$$\mathcal{X}^{-1} = \begin{pmatrix} E_{N+1 \rightarrow N} D_{N+1} \\ E_{N+1 \rightarrow MP} \end{pmatrix}, \quad (5)$$

which is the actual inverse of \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 as we have done previously,^{36,53} in which case information is lost when converting the unit-length tangent vectors to the positions

X . For the deterministic examples we studied in Refs. 36 and 53, the fibers were relatively smooth, and so the lost (high-frequency) information had a negligible effect. This is no longer the case for Brownian filaments, and so the method here is required to correctly track 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 conditions on τ .

While we motivated (4) using a spectral discretization, these equations also hold for the standard blob-link discretization;³⁵ 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 maps \mathcal{X} and \mathcal{X}^{-1} have been specified.

2. Gibbs-Boltzmann distribution

Now that we have introduced our discrete fiber model, we can write the Gibbs-Boltzmann equilibrium distribution as a function of the chain degrees of freedom $\tilde{\tau}$. Letting $\mathcal{E}_{\text{bend}}(\tilde{\tau})$ denote the discrete bending energy of the fiber, which we discretize in Sec. II A 3, we take the Gibbs-Boltzmann distribution to be

$$dP_{\text{eq}}(\tilde{\tau}) = Z^{-1} \exp(-\mathcal{E}_{\text{bend}}(\tilde{\tau})/k_B T) d\mu_0(\tilde{\tau}),$$

$$P_{\text{eq}}(\tilde{\tau}) = Z^{-1} \exp(-\mathcal{E}_{\text{bend}}(\tilde{\tau})/k_B T) \prod_{p=1}^N \delta(\tau_{\{p\}}^T \tau_{\{p\}} - 1). \quad (6)$$

What we mean by the product of δ functions is really that the base measure $d\mu_0(\tilde{\tau})$ in the Gibbs-Boltzmann distribution (6) corresponds to the tangent vectors $\tau_{\{p\}}$ being independently uniformly distributed on the unit sphere for $p = 1, \dots, N$. For blob-link chains, the tangent vectors $\tau_{\{p\}}$ are uniformly spaced, as shown in Fig. 2, and this distribution follows naturally from their inextensibility and statistical independence in a freely jointed chain, in which case (6) defines the standard worm-like chain model.

In our spectral discretization, the tangent vectors $\tau_{\{p\}}$ are defined on a Chebyshev grid, and this model lacks the physical motivation that characterizes the blob-link model (see Fig. 2). 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 $\mathbb{T}(s)$ to perform infinitely fast Brownian motion in s on the unit sphere; thus, only (6) 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 $\mathbb{T} \cdot \mathbb{T} - 1$ 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).⁶² In the first step toward a more mathematically justified approach, we postulate (6) as a reasonable guess, but do not claim any precise sense of convergence of (6) for a spectral grid to the equilibrium distribution of a continuum worm-like chain. We do show (in Sec. IV A and Appendix D) 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. Forthcoming work will justify (6) for spectral chains through the theory of coarse graining.

3. Bending energy and forces

We now turn to the evaluation of the bending energy $\mathcal{E}_{\text{bend}}$ and the resulting force and force density that it generates on the fiber centerline for free fibers. In continuum, the fiber bending or curvature energy is given by the squared L^2 norm of the fiber curvature vector,

$$\mathcal{E}_{\text{bend}}[\mathbb{X}(\cdot)] = \frac{\kappa}{2} \int_0^L \partial_s^2 \mathbb{X}(s) \cdot \partial_s^2 \mathbb{X}(s) ds. \quad (7)$$

Now, there are two competing ideas for how to generate a force (or force density) from this energy. In the continuum perspective we have used in the past,^{36,53} we take a functional derivative of $\mathcal{E}_{\text{bend}}$ in continuum to yield a force density (inside the integral) $\mathbf{f}(s) = -\kappa \partial_s^4 \mathbb{X}$, subject to the boundary conditions $\partial_s^2 \mathbb{X}(0, L) = \partial_s^2 \mathbb{X}(0, L) = \mathbf{0}$. In a spectral method, we discretize this force density using rectangular spectral collocation (RSC),^{63,64} which does *not* guarantee that the total force and torque are zero on a nonsmooth fiber, since the force obtained does not come from differentiating a discrete energy functional.⁵¹ As a result, we cannot write down the equilibrium distribution we would expect the Brownian dynamics to obey, even for extensible fibers.

Because of this, it is useful to discretize the energy functional (7) directly and then use the matrix that results to compute forces, implementing the boundary conditions naturally. On a spectral grid, we can put the energy functional (7) in the form of an energy function $\mathcal{E}_{\text{bend}}(\mathbf{X}) = \frac{1}{2} \mathbf{X}^T \mathbf{L} \mathbf{X}$, where

$$\mathbf{L} = \kappa (\mathbf{E}_{N_x \rightarrow 2N_x} \mathbf{D}^2)^T \mathbf{W}_{2N_x} (\mathbf{E}_{N_x \rightarrow 2N_x} \mathbf{D}^2) := \kappa (\mathbf{D}^2)^T \tilde{\mathbf{W}} \mathbf{D}^2 \quad (8)$$

is a matrix that takes two derivatives of \mathbf{X} on a grid of size $N_x = N + 1$ and then does an inner product of those derivatives on a $2N_x$ grid using Clenshaw-Curtis weights matrix \mathbf{W}_{2N_x} . The inner product on the $2N_x$ grid is computed with an upsampled representation $\mathbf{E}_{N_x \rightarrow 2N_x} \mathbf{X}$, where the upsampling matrix $\mathbf{E}_{N_x \rightarrow 2N_x}$ is applied by using the points \mathbf{X} to form the Chebyshev interpolant $\mathbb{X}(s)$ and then evaluating the interpolant on the upsampled grid. The purpose of this 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

$$\tilde{\mathbf{W}} = \mathbf{E}_{N_x \rightarrow 2N_x}^T \mathbf{W}_{2N_x} \mathbf{E}_{N_x \rightarrow 2N_x}. \quad (9)$$

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

$$\mathcal{E}_{\text{bend}}(\tilde{\tau}) = \frac{1}{2} \tilde{\tau}^T \mathcal{X}^T \mathbf{L} \mathcal{X} \tilde{\tau} := \frac{1}{2} \tilde{\tau}^T \mathbf{L}_\tau \tilde{\tau}. \quad (10)$$

The force on the fiber nodes is obtained by differentiating the energy function,

$$\mathbf{F} = -\frac{\partial \mathcal{E}_{\text{bend}}}{\partial \mathbf{X}} = -\mathbf{L} \mathbf{X} = -\mathbf{L} \mathcal{X} \tilde{\tau}. \quad (11)$$

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 (11) by

$$\mathbf{f} = \tilde{\mathbf{W}}^{-1} \mathbf{F} \quad (12)$$

and use the Chebyshev interpolant $f(s)$ as a continuum force density. Note that the matrix $\tilde{\mathbf{W}}$, 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 (9) 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 (Ref. 65, Sec. 6), i.e., the function that satisfies $\langle \mathbf{X}, \mathbf{f} \rangle = \mathbf{X}^T \mathbf{F}$.⁶⁶

B. Dynamics

If we substitute energy (10) into the Gibbs–Boltzmann distribution (6), we see that our goal is to write an overdamped Langevin equation that is in detailed balance with respect to the distribution

$$P_{\text{eq}}(\tilde{\boldsymbol{\tau}}) = Z^{-1} \exp\left(-\tilde{\boldsymbol{\tau}}^T \mathbf{L} \tilde{\boldsymbol{\tau}} / (2k_B T)\right) \prod_{p=1}^N \delta(\boldsymbol{\tau}_{\{p\}}^T \boldsymbol{\tau}_{\{p\}} - 1) \quad (13)$$

and includes hydrodynamic interactions between points on the filament. To do this, we first need to discuss how inextensible filaments evolve *deterministically*. This begins with a description of the kinematics of discrete inextensible filaments, which are analogous to that of a series of elastically interacting rigid rods. We then describe the evaluation of hydrodynamic interactions and formulate the equations of motion in the deterministic setting. The overdamped Langevin equation can be obtained from these deterministic considerations and (13) by following a standard formulation.

1. Kinematics

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

$$\partial_t \boldsymbol{\tau} = \boldsymbol{\Omega} \times \boldsymbol{\tau} := -\mathbf{C}[\boldsymbol{\tau}] \boldsymbol{\Omega}, \quad (14)$$

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:

$$\mathbf{C} \boldsymbol{\tau} = -\mathbf{C}^T \boldsymbol{\tau} = \mathbf{0}, \quad \partial_{\boldsymbol{\tau}} \cdot \mathbf{C}^T = \mathbf{0}. \quad (15)$$

The first equation of Eq. (15) 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 the j th entry

$$\left(\partial_{\boldsymbol{\tau}} \cdot \mathbf{C}^T\right)_j := \frac{\partial}{\partial \boldsymbol{\tau}_k} C_{jk}^T = \frac{\partial}{\partial \boldsymbol{\tau}_k} C_{kj} = 0 \quad (16)$$

because the k th row of \mathbf{C} has no entries that depend on $\boldsymbol{\tau}_k$; here and throughout this paper, 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 (14) is analogous to the time evolution of the unit quaternion describing a rigid body's orientation [Ref. 29, Eqs. (5)–(8)], which can be represented as a unit vector on the unit 4 sphere.

The evolution of the tangent vectors according to (14) automatically implies that the evolution of the positions \mathbf{X} is given by

$$\partial_t \boldsymbol{\mathcal{X}} = \mathbf{X} \begin{pmatrix} -\mathbf{C} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \boldsymbol{\Omega} \\ U_{\text{MP}} \end{pmatrix} := \boldsymbol{\mathcal{X}} \tilde{\mathbf{C}} \boldsymbol{\alpha} := \mathbf{K} \boldsymbol{\alpha}. \quad (17)$$

Note the analogy with (4), 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

$$\mathbf{K}^{-1} = \begin{pmatrix} \mathbf{C} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \boldsymbol{\mathcal{X}}^{-1} = \tilde{\mathbf{C}}^T \boldsymbol{\mathcal{X}}^{-1}. \quad (18)$$

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 space of \mathbf{C} .⁵⁷ Thus, \mathbf{K}^{-1} is not a true inverse of \mathbf{K} , but rather

$$\mathbf{K}^{-1} \mathbf{K} = \begin{pmatrix} -\mathbf{C}^2 & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \rightarrow \tilde{\mathbf{C}} \mathbf{K}^{-1} \mathbf{K} = \tilde{\mathbf{C}}, \quad (19)$$

i.e., $\mathbf{K}^{-1} \mathbf{K}$ acts like the identity when applied to $\tilde{\mathbf{C}}$ from the right. The last equality holds because 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 (19) also means, in practice, is that

$$\mathbf{K}^{-1} \mathbf{K} \boldsymbol{\alpha} = \mathbf{K}^{-1} \mathbf{K} \begin{pmatrix} \boldsymbol{\Omega} \\ U_{\text{MP}} \end{pmatrix} = \begin{pmatrix} \boldsymbol{\Omega}^\perp \\ U_{\text{MP}} \end{pmatrix}, \quad (20)$$

where $\boldsymbol{\Omega}_{\{p\}}^\perp = (\mathbf{I} - \boldsymbol{\tau}_{\{p\}} \boldsymbol{\tau}_{\{p\}}) \boldsymbol{\Omega}_{\{p\}}$ (once again, here, we see $\mathbf{K}^{-1} \mathbf{K}$ projecting off the parallel parts of $\boldsymbol{\Omega}$ since the p th block of $\mathbf{K}^{-1} \mathbf{K}$ is $\mathbf{I}_3 - \boldsymbol{\tau}_{\{p\}} \boldsymbol{\tau}_{\{p\}}$). Thus, because the fiber evolves according to (14), it makes no difference whether we use $\boldsymbol{\alpha}$ or $\mathbf{K}^{-1} \mathbf{K} \boldsymbol{\alpha}$ for the fiber velocities. It is for this reason that we use this discretization for \mathbf{K} , as opposed to those from our prior work,^{36,53} which viewed $\boldsymbol{\Omega}$ as a continuous function of s , and were consequently plagued by aliasing errors in trying to recover $\boldsymbol{\Omega}$ from $\mathbf{K} \boldsymbol{\alpha}$ [Ref. 53, Eq. (110)].

While evolution (14) describes how the tangent vectors evolve in continuous time, in discrete time, the update $\boldsymbol{\tau}^{(n+1)} = \boldsymbol{\tau}^{(n)} - \Delta t \mathbf{C} \boldsymbol{\Omega}$ does *not* preserve the unit-length constraint. Thus, in order to keep the dynamics on the constraint in discrete time, we will solve for an angular velocity $\boldsymbol{\Omega}$ and then evolve each tangent vector $\boldsymbol{\tau}_{\{p\}}^{(n)}$ by rotating by $\boldsymbol{\Omega}_{\{p\}}$,

$$\boldsymbol{\tau}_{\{p\}}^{(n+1)} = \text{rotate}\left(\boldsymbol{\tau}_{\{p\}}^{(n)}, \boldsymbol{\Omega}_{\{p\}} \Delta t\right), \quad (21)$$

the explicit formula for which is given in Ref. 53, Eq. (111). We then update the midpoint \mathbf{X}_{MP} via $\mathbf{X}_{\text{MP}}^{(n+1)} = \mathbf{X}_{\text{MP}}^{(n)} + \Delta t \mathbf{U}_{\text{MP}}^{(n)}$. In deterministic methods,^{46,53} rotating the tangent vectors in this way keeps the dynamics on the constraint without needing to introduce penalty parameters⁴⁶ that lead to additional stiffness in temporal integration.

2. Hydrodynamic mobility

We now discuss the discretization of the mobility matrix $\tilde{\mathbf{M}}$ that relates the force on the nodal points to their velocities. In this paper, we will not consider hydrodynamic interactions between multiple filaments. Thus, the matrix $\tilde{\mathbf{M}}$ can be formed separately as a dense matrix on each fiber, and any operations with $\tilde{\mathbf{M}}$, such as inversions and square roots, done directly by storing the eigenvalue (or Cholesky if $\tilde{\mathbf{M}} \geq \mathbf{0}$) decomposition. Extensions to multiple filaments interacting hydrodynamically, in which case $\tilde{\mathbf{M}}$ cannot be formed densely, will be covered in future work.

In Ref. 36, we developed an efficient quadrature scheme for (2), but this mobility acts on force *densities* by applying a matrix \mathbf{M} to obtain the relationship $\mathbf{U} = \mathbf{M}\mathbf{f}$ (we use the notation $\tilde{\mathbf{M}}$ for the matrix mapping force to velocity and \mathbf{M} for the matrix mapping force density to velocity). Thus, to use our quadrature scheme, we first need to first convert force to force density using (12), which gives $\mathbf{U} = \mathbf{M}\tilde{\mathbf{W}}^{-1}\mathbf{F}$. If we were to compute the quadratures *exactly*, for instance, by upsampling to a fine grid, $\mathbf{M}\tilde{\mathbf{W}}^{-1}$ should be a symmetric positive definite matrix since the work dissipated in the fluid

$$\int_0^L \mathbf{f}(s) \cdot \mathbf{U}(s) ds = \mathbf{f}^T \tilde{\mathbf{W}}\mathbf{U} = \mathbf{f}^T \tilde{\mathbf{W}}\mathbf{M}\mathbf{f} \quad (22)$$

is always positive, which implies that $\tilde{\mathbf{W}}\mathbf{M}$ (and $\tilde{\mathbf{W}}^{-1}\tilde{\mathbf{W}}\mathbf{M}\tilde{\mathbf{W}}^{-1} = \mathbf{M}\tilde{\mathbf{W}}^{-1}$) is an SPD matrix. Indeed, upsampling to a fine grid gives us a reference SPD mobility,

$$\tilde{\mathbf{M}}_{\text{ref}} = \tilde{\mathbf{W}}^{-1} \mathbf{E}_u^T \mathbf{W}_u \tilde{\mathbf{M}}_{\text{RPY},u} \mathbf{W}_u \mathbf{E}_u \tilde{\mathbf{W}}^{-1}, \quad (23)$$

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 $\tilde{\mathbf{W}}^{-1}$. We then extend the force density to the upsampled grid and integrate it against the RPY kernel there (the matrix $\tilde{\mathbf{M}}_{\text{RPY},u}$ 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 $\tilde{\mathbf{M}}_{\text{ref}}$, 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 . Note that while we write (23) in a collocation perspective, it can be shown that the same fundamental matrix appears in a Galerkin method, where we would expect an automatically SPD matrix.

When we use our efficient quadrature scheme to approximate (23), the matrix that results is $\tilde{\mathbf{M}} = \mathbf{M}\tilde{\mathbf{W}}^{-1}$, which is not guaranteed to be symmetric positive definite. Indeed, in Ref. 36, Sec. 4.4.1, 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 define the mobility as the symmetric matrix,

$$\tilde{\mathbf{M}} = \frac{1}{2} (\mathbf{M}\tilde{\mathbf{W}}^{-1} + \tilde{\mathbf{W}}^{-1}\mathbf{M}^T). \quad (24)$$

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 (23). In our tests, our quadrature-based mobility, which is slightly modified for the case of random filaments as described in Appendix B 2, gives the same dynamics for a relaxing filament as the reference mobility (23); see Appendix B 4 a. Thus, the error made in quadrature, symmetrization, and eigenvalue truncation is small, in practice, even though the quadrature is only spectrally accurate for smooth fibers.

In Appendix B 4, we also compare the slender body quadrature mobility to two other commonly used mobilities for slender body hydrodynamics: direct RPY-based quadrature on the Chebyshev grid and a local slender body theory.^{54,55} Our results show that there are advantages to using the special quadrature both in terms of resolving spatial scales as $\hat{\epsilon} \rightarrow 0$ and in terms of the time step required for temporal accuracy. With that in mind, Fig. 3 shows how we use the fiber geometry and special quadrature to apply the mobility $\tilde{\mathbf{M}}\mathbf{F}$.

3. Deterministic evolution

We are now ready to put the pieces of our discretization together to obtain the deterministic evolution of the fiber centerline. 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 [Ref. 53, Eq. (54)],

$$\begin{aligned} \mathbf{K}\boldsymbol{\alpha} &= \tilde{\mathbf{M}}(-\mathbf{L}\mathbf{X} + \boldsymbol{\Lambda}), \\ \mathbf{K}^T\boldsymbol{\Lambda} &= \mathbf{0}, \end{aligned} \quad (25)$$

where \mathbf{K} , $\tilde{\mathbf{M}}$, \mathbf{L} , and \mathbf{K}^T are all square matrices of size $3N_x = 3(N + 1)$ and $\boldsymbol{\Lambda}$ is a constraint force. This is a reformulation of Ref. 53, Eq. (54)

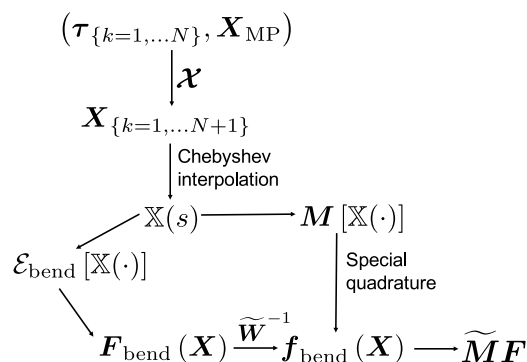


FIG. 3. Summary of how we use the fiber configuration to compute the forces \mathbf{F} and mobility $\tilde{\mathbf{M}}$. Given $(\boldsymbol{\tau}, \mathbf{X}_{\text{MP}})$, we apply the map \mathcal{X} from (4) to obtain the Chebyshev polynomial $\mathbb{X}(s)$, which we then use to compute the bending energy $\mathcal{E}_{\text{bend}}$ in (7), force \mathbf{F} in (11), and force density \mathbf{f} in (12). The velocity $\tilde{\mathbf{M}}\mathbf{F}$ is computed by applying special quadrature to the force density \mathbf{f} ; see Sec. II B 2 (because the quadrature matrix might have negative eigenvalues, we form $\tilde{\mathbf{M}}$ explicitly to compute and truncate its eigenvalues).

in terms of *force* rather than force density as the mobility matrix $\tilde{\mathbf{M}}(\mathbf{X})$ relates the velocity of the filament centerline to the *forces* at the nodal points. As shown in Appendix A, the deterministic dynamics (25) can also be obtained by minimizing a constrained Lagrangian function, which implies that they are dissipative and have the structure of a gradient descent flow.

Using a Schur complement approach, the Lagrange multipliers Λ can be eliminated from (25) to yield

$$\partial_t \mathbf{X} = \mathbf{K}\alpha = -(\mathbf{K}\mathbf{N}\mathbf{K}^T)\mathbf{L}\mathbf{X}, \quad (26)$$

where

$$\mathbf{N} = (\mathbf{K}^T \tilde{\mathbf{M}}^{-1} \mathbf{K})^\dagger \quad (27)$$

is the mobility matrix projected onto the space of inextensible motions. We can transform this equation to obtain the evolution of $\tilde{\boldsymbol{\tau}}$ by substituting $\mathbf{X} = \mathcal{X}\tilde{\boldsymbol{\tau}}$ into (26) to obtain

$$\begin{aligned} \partial_t \tilde{\boldsymbol{\tau}} &= -\mathcal{X}^{-1}(\mathbf{K}\mathbf{N}\mathbf{K}^T)\mathbf{L}\mathcal{X}\tilde{\boldsymbol{\tau}} = -(\tilde{\mathbf{C}}\mathbf{N}\tilde{\mathbf{C}}^T)\mathcal{X}^T\mathbf{L}\mathcal{X}\tilde{\boldsymbol{\tau}} \\ &= -(\tilde{\mathbf{C}}\mathbf{N}\tilde{\mathbf{C}}^T)\mathbf{L}_\tau \tilde{\boldsymbol{\tau}}. \end{aligned} \quad (28)$$

We note that these equations apply for continuous time. In discrete time, we will solve (25) for $\alpha = (\Omega, \mathbf{U}_{\text{MP}})$ and then update the tangent vectors on the unit sphere by $\Omega\Delta t$ using the nonlinear rotation (21).

4. Overdamped Langevin equation in $\tilde{\boldsymbol{\tau}}$

The deterministic equation (28) 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 [Ref. 29, Eq. (12)] in the rigid body case, we have the overdamped Ito Langevin equation,

$$\partial_t \tilde{\boldsymbol{\tau}} = -(\tilde{\mathbf{C}}\mathbf{N}\tilde{\mathbf{C}}^T)\mathbf{L}_\tau \tilde{\boldsymbol{\tau}} + k_B T \partial_{\tilde{\boldsymbol{\tau}}} \cdot (\tilde{\mathbf{C}}\mathbf{N}\tilde{\mathbf{C}}^T) + \sqrt{2k_B T \tilde{\mathbf{C}}\mathbf{N}^{1/2}} \mathcal{W}, \quad (29)$$

describing the evolution of $\tilde{\boldsymbol{\tau}}$. Here, $\mathcal{W}(t)$ is a collection of white noise processes (the formal derivative of Brownian motions), the divergence with respect to $\tilde{\boldsymbol{\tau}}$ is defined in (16), and $\mathbf{N}^{1/2}$ satisfies the fluctuation–dissipation relation,

$$\mathbf{N}^{1/2}(\mathbf{N}^{1/2})^T = \mathbf{N}. \quad (30)$$

The arguments in Ref. 29, Sec. II B 2 can be used to show that (29) is time-reversible with respect to the Gibbs–Boltzmann equilibrium distribution (13). As described there, because we formulate (29) with respect to $\tilde{\boldsymbol{\tau}}$ 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 (30), can be applied by adding Brownian noise with covariance $\tilde{\mathbf{M}}$ to the right-hand side of saddle point system (25). In an Euler–Maruyama discretization, this corresponds to solving the saddle point system,

$$\begin{aligned} \mathbf{K}\alpha &= \tilde{\mathbf{M}}(-\mathbf{L}\mathbf{X} + \Lambda) + \sqrt{\frac{2k_B T}{\Delta t}} \tilde{\mathbf{M}}^{1/2} \boldsymbol{\eta}, \\ \mathbf{K}^T \Lambda &= \mathbf{0}, \end{aligned} \quad (31)$$

where $\boldsymbol{\eta}$ is an i.i.d. vector of standard normal random variables. As discussed at length in Ref. 68, Sec. II(B), solving this saddle point system gives

$$\begin{aligned} \partial_t \mathbf{X} &= \mathbf{K}\alpha = -(\mathbf{K}\mathbf{N}\mathbf{K}^T)\mathbf{L}\mathbf{X} + \sqrt{\frac{2k_B T}{\Delta t}} (\mathbf{K}\mathbf{N}\mathbf{K}^T) \tilde{\mathbf{M}}^{-1/2} \boldsymbol{\eta} \\ &= -(\mathbf{K}\mathbf{N}\mathbf{K}^T)\mathbf{L}\mathbf{X} + \sqrt{\frac{2k_B T}{\Delta t}} \mathbf{K}\mathbf{N}^{1/2} \boldsymbol{\eta}, \end{aligned} \quad (32)$$

giving $\mathbf{N}^{1/2} = \mathbf{N}\mathbf{K}^T \tilde{\mathbf{M}}^{-1/2}$. 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 $\tilde{\mathbf{M}}^{1/2} \boldsymbol{\eta}$. For hydrodynamics that is localized to each fiber, we do this using the eigenvalue decomposition of $\tilde{\mathbf{M}}$, which already must be computed for the purposes of eigenvalue truncation (see Sec. II B 2). 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 $\tilde{\mathbf{M}}^{1/2}$ can be computed via the Lanczos algorithm⁶⁹ or the positively split Ewald (PSE) method.⁷⁰ While the case of nonlocal hydrodynamics will not be treated in this paper, the saddle point method is a useful foundation for future work.

5. Overdamped Langevin equation in \mathbf{X}

We now use (29) to derive an overdamped Langevin equation in terms of \mathbf{X} . If we multiply (29) on the left by \mathbf{X} and expand $\mathbf{L}_\tau = \mathcal{X}^T \mathbf{L} \mathcal{X}$, we obtain

$$\begin{aligned} \partial_t \mathbf{X} &= -(\mathbf{K}\mathbf{N}\mathbf{K}^T)\mathbf{L}\mathbf{X} + k_B T \partial_{\tilde{\boldsymbol{\tau}}} \cdot (\mathbf{K}\mathbf{N}\tilde{\mathbf{C}}^T) \\ &\quad + \sqrt{2k_B T \mathbf{K}\mathbf{N}^{1/2}} \mathcal{W}. \end{aligned} \quad (33)$$

We use the chain rule to write differentiation with respect to \mathbf{X} as

$$\frac{\partial f}{\partial \mathbf{X}_k} = \frac{\partial f}{\partial \tilde{\boldsymbol{\tau}}_p} \frac{\partial \tilde{\boldsymbol{\tau}}_p}{\partial \mathbf{X}_k} = \frac{\partial f}{\partial \tilde{\boldsymbol{\tau}}_p} \mathcal{X}_{pk}^{-1}.$$

We now rewrite the divergence in (33) as

$$\begin{aligned} \frac{\partial}{\partial \tilde{\boldsymbol{\tau}}_j} (\mathbf{K}\mathbf{N}\tilde{\mathbf{C}}^T)_{ij} &= \frac{\partial}{\partial \mathbf{X}_k} (\mathbf{K}\mathbf{N}\tilde{\mathbf{C}}^T)_{ij} \frac{\partial \mathbf{X}_k}{\partial \tilde{\boldsymbol{\tau}}_j} \\ &= \frac{\partial}{\partial \mathbf{X}_k} (\mathbf{K}\mathbf{N}\tilde{\mathbf{C}}^T)_{ij} \mathcal{X}_{kj} = \frac{\partial}{\partial \mathbf{X}_k} (\mathbf{K}\mathbf{N}\mathbf{K}^T)_{ik} \end{aligned} \quad (34)$$

so that the Ito equation (29) could equivalently be formulated in terms of \mathbf{X} as might be expected from the deterministic equation (26),

$$\begin{aligned} \partial_t \mathbf{X} &= -(\mathbf{K}\mathbf{N}\mathbf{K}^T)\mathbf{L}\mathbf{X} + k_B T \partial_{\mathbf{X}} \cdot (\mathbf{K}\mathbf{N}\mathbf{K}^T) \\ &\quad + \sqrt{2k_B T \mathbf{K}\mathbf{N}^{1/2}} \mathcal{W} \end{aligned} \quad (35)$$

$$\doteq -\tilde{\mathbf{N}}\mathbf{L}\mathbf{X} + \sqrt{2k_B T \tilde{\mathbf{N}}} \circ \tilde{\mathbf{N}}^{-1/2} \mathcal{W}, \quad (36)$$

where $\tilde{\mathbf{N}} = \mathbf{K}\mathbf{N}\mathbf{K}^T$ and the second equality denotes that paths of (35) and (36) have the same probability distribution. Equation (36) is the Langevin equation written in a split Stratonovich–Ito²⁹ or kinetic⁷¹ 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 [cf. Ref. 29, 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 Eq. (29) for $\tilde{\boldsymbol{\tau}}$, knowing that the one in \mathbf{X} can be obtained by this simple transformation. Note that (35) is much simpler than the overdamped equations derived previously for bead-link models.^{27,28}

III. TEMPORAL INTEGRATION

In this section, we discuss our temporal integrator for the overdamped Langevin equation (29). The scheme, which is in the spirit of the Fixman method⁷² and similar to that of Westwood *et al.* for rigid bodies,⁷³ 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 and 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 (29). We first separate it into three terms,

$$\frac{\partial}{\partial \tilde{\boldsymbol{\tau}}_j} (\tilde{\mathbf{C}}\mathbf{N}\tilde{\mathbf{C}}^T)_{ij} = (\partial_j \tilde{\mathbf{C}}_{ik} \mathbf{N}_{kp} \tilde{\mathbf{C}}_{pj}^T + \tilde{\mathbf{C}}_{ik} (\partial_j \mathbf{N}_{kp}) \tilde{\mathbf{C}}_{pj}^T + \tilde{\mathbf{C}}_{ik} \mathbf{N}_{kp} (\partial_j \tilde{\mathbf{C}}_{pj}^T), \quad (37)$$

where ∂_j is shorthand for $\partial/\partial \tilde{\boldsymbol{\tau}}_j$. As shown in Ref. 29, Sec. III(A), rotating the tangent vectors at every time step n using the Euler–Maruyama method,

$$\boldsymbol{\tau}^{(n+1)} = \text{rotate} \left(\boldsymbol{\tau}^{(n)}, \sqrt{2k_B T \Delta t} (\mathbf{N}^{(n)})^{1/2} \boldsymbol{\eta}^{(n)} + \mathcal{O}(\Delta t) \right), \quad (38)$$

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 (37) is zero by (16). Therefore, our schemes simply need to generate the additional drift term,

$$(k_B T) \tilde{\mathbf{C}}_{ik} (\partial_j \mathbf{N}_{kp}) \tilde{\mathbf{C}}_{pj}^T \Delta t, \quad \text{i.e.,} \quad (k_B T) \tilde{\mathbf{C}} (\nabla_{\tilde{\boldsymbol{\tau}}} \mathbf{N} : \tilde{\mathbf{C}}^T) \Delta t. \quad (39)$$

Because $\text{rotate}(\boldsymbol{\tau}^{(n)}, \Delta t \boldsymbol{\Omega}) = \boldsymbol{\tau}^{(n)} + \Delta t \mathbf{C} \boldsymbol{\Omega} - (\boldsymbol{\Omega} \cdot \boldsymbol{\Omega}) \Delta t^2 \boldsymbol{\tau}^{(n)} / 8 + \mathcal{O}(\Delta t^3)$, the term (39) corresponds to $\mathcal{O}(\Delta t^2)$ to a rotate procedure by an angle $\Delta t (k_B T) \nabla_{\tilde{\boldsymbol{\tau}}} \mathbf{N} \tilde{\mathbf{C}}^T$. Thus, our task will be to design numerical methods to produce the stochastic drift term in $\boldsymbol{\Omega}$,

$$\text{Drift} = (k_B T) \nabla_{\tilde{\boldsymbol{\tau}}} \mathbf{N} : \tilde{\mathbf{C}}^T, \quad \text{i.e.,} \quad \text{Drift}_k = (k_B T) (\partial_j \mathbf{N}_{kp}) \tilde{\mathbf{C}}_{jp}, \quad (40)$$

in expectation, which will give (39) after rotation over time step size Δt .

A. Implicit methods

We first motivate our temporal discretization of (29) by considering the discretization of the unconstrained *linearized* stochastic differential equation (SDE),

$$\partial_t \mathbf{X} = -\tilde{\mathbf{M}}[\mathbf{X}_0] \mathbf{L} \Delta \mathbf{X} + \mathbf{U}_B, \quad (41)$$

where \mathbf{X}_0 is an equilibrium position, $\Delta \mathbf{X} = \mathbf{X} - \mathbf{X}_0$, the matrix \mathbf{L} discretizes the energy $\mathcal{E}_{\text{bend}} = \Delta \mathbf{X}^T \mathbf{L} \Delta \mathbf{X}$, and \mathbf{U}_B is the Brownian velocity given by fluctuation–dissipation balance as $\sqrt{2k_B T} (\tilde{\mathbf{M}}[\mathbf{X}_0])^{1/2} \boldsymbol{\mathcal{W}}(t)$. Because this SDE is unconstrained, the equilibrium covariance of $\Delta \mathbf{X}$ is known from statistical mechanics,

$$\mathbb{E}[\Delta \mathbf{X} \Delta \mathbf{X}^T] = k_B T \mathbf{L}^{-1}. \quad (42)$$

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 (42) for arbitrary Δt when applied to (41). To do this, we follow the analysis of Ref. 74, Sec. III(B) to derive the steady state covariance for a given temporal integrator.

We consider an implicit–explicit method for (41) of the form

$$\mathbf{X}^{(n+1)} = \mathbf{X}^{(n)} + \Delta t \left(-\tilde{\mathbf{M}} \mathbf{L} (\mathbf{X}^{(n+1)} + (1-c)\mathbf{X}^{(n)}) + \mathbf{U}_B^{(n)} \right),$$

which can be rearranged to yield

$$(\mathbf{I} + c \Delta t \tilde{\mathbf{M}} \mathbf{L}) \mathbf{X}^{(n+1)} = \left[(\mathbf{I} - (1-c) \Delta t \tilde{\mathbf{M}} \mathbf{L}) \mathbf{X}^{(n)} + \Delta t \mathbf{U}_B^{(n)} \right].$$

We now take an outer product of the two sides of the equation and then substitute the desired covariance from (42), which at steady state is independent of the time step n , to obtain the matrix equation,

$$2k_B T \Delta t \tilde{\mathbf{M}} + k_B T (2c-1) \Delta t^2 \tilde{\mathbf{M}} \mathbf{L} \tilde{\mathbf{M}} = \Delta t^2 \mathbb{E}[\mathbf{U}_B \mathbf{U}_B^T].$$

To obtain the exact covariance for $c=1$ (backward Euler), we therefore set

$$\mathbf{U}_B = \sqrt{\frac{2k_B T}{\Delta t}} \left(\tilde{\mathbf{M}}^{1/2} \boldsymbol{\eta} + \sqrt{\frac{\Delta t}{2}} \tilde{\mathbf{M}} \mathbf{L}^{1/2} \tilde{\boldsymbol{\eta}} \right), \quad (43)$$

where $\tilde{\boldsymbol{\eta}}$ 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 $\sqrt{2k_B T / \Delta t} \tilde{\mathbf{M}}^{1/2} \boldsymbol{\eta}$. While this choice has been preferred for other applications,⁷⁴ 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 (43) 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.

B. Midpoint scheme

We can now present our “midpoint” method, which can produce the drift term (40) 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.*⁷³ for rigid body suspensions, but differs in the ways we detail in Appendix C 2 c. The method as presented here is optimized for dense linear algebra in the sense that it requires only two mobility evaluations per time step if those mobilities can be stored as dense matrices. See Appendix C 2 b

for modifications when the mobility cannot be stored as a dense matrix.

At each time step n , we perform the following steps:

1. Compute a rotation rate for the tangent vectors based on the Brownian velocity $\tilde{\mathbf{M}}^{1/2}\boldsymbol{\eta}^{(n)}$,

$$\boldsymbol{\alpha}^{(n,*)} = \sqrt{\frac{2k_B T}{\Delta t}} \left(\mathbf{K}^{(n)} \right)^{-1} \left(\tilde{\mathbf{M}}^{(n)} \right)^{1/2} \boldsymbol{\eta}^{(n)} = \begin{pmatrix} \boldsymbol{\Omega}^{(n,*)} \\ \mathbf{U}_{\text{MP}}^{(n,*)} \end{pmatrix}, \quad (44)$$

where \mathbf{K}^{-1} is defined in (18).

2. Rotate the tangent vectors by $\boldsymbol{\Omega}^{(n,*)}\Delta t/2$ to generate a new configuration,

$$\begin{aligned} \tilde{\boldsymbol{\tau}}^{(n+1/2,*)} &= \begin{pmatrix} \text{rotate}(\boldsymbol{\tau}^{(n)}, (\Delta t/2)\boldsymbol{\Omega}^{(n,*)}) \\ \mathbf{X}_{\text{MP}}^{(n)} + (\Delta t/2)\mathbf{U}_{\text{MP}}^{(n,*)} \end{pmatrix} \quad (45) \\ &= \sqrt{\frac{k_B T \Delta t}{2}} \tilde{\mathbf{C}}^{(n)} \left(\mathbf{K}^{(n)} \right)^{-1} \left(\tilde{\mathbf{M}}^{(n)} \right)^{1/2} \\ &\quad \times \boldsymbol{\eta}^{(n)} + \mathcal{O}(\Delta t^{3/2}). \end{aligned} \quad (46)$$

3. Evaluate the mobility $\tilde{\mathbf{M}}^{(n+1/2,*)}$ and use it to compute the additional drift velocity using the random finite difference (RFD)⁷⁵ with $\delta \sim \sqrt{\Delta t}$,

$$\begin{aligned} \mathbf{U}_{\text{MD}}^{(n)} &= \sqrt{\frac{2k_B T}{\Delta t}} \left(\tilde{\mathbf{M}}^{(n+1/2,*)} - \tilde{\mathbf{M}}^{(n)} \right) \\ &\quad \times \left(\left(\tilde{\mathbf{M}}^{(n)} \right)^{-1/2} \right)^T \boldsymbol{\eta}^{(n)}. \end{aligned} \quad (47)$$

This term might be impractical for large systems because it is based on solving a resistance problem to obtain $(\tilde{\mathbf{M}}^{-1/2})^T$. Appendix C 2 b has an alternative approach, which generates the same drift term via an RFD in which $\tilde{\mathbf{M}}$ only has to be applied rather than inverted.

4. To obtain the tangent vector rotation rates, solve the saddle point system

$$\begin{aligned} \begin{pmatrix} -\tilde{\mathbf{M}} & (\mathbf{I} + c\Delta t\tilde{\mathbf{M}}\mathbf{L})\mathbf{K} \\ \mathbf{K}^T & \mathbf{0} \end{pmatrix}^{(n+1/2,*)} \begin{pmatrix} \boldsymbol{\Lambda} \\ \boldsymbol{\alpha} \end{pmatrix}^{(n+1/2)} \\ = \begin{pmatrix} -\tilde{\mathbf{M}}^{(n+1/2,*)}\mathbf{L}\mathbf{X}^{(n)} + \mathbf{U}_B^{(n)} + \mathbf{U}_{\text{MD}}^{(n)} \\ \mathbf{0} \end{pmatrix} \end{aligned} \quad (48)$$

for $\boldsymbol{\Lambda}^{(n+1/2)}$ and $\boldsymbol{\alpha}^{(n+1/2)} = (\boldsymbol{\Omega}^{(n+1/2)}, \mathbf{U}_{\text{MP}}^{(n+1/2)})$. The Brownian velocity \mathbf{U}_B is defined in (43), and the first part of it is used in (44) to generate the midpoint configuration, i.e., the same $\boldsymbol{\eta}$ is used in steps 1, 3, and 4.

5. Update the fiber via (21),

$$\tilde{\boldsymbol{\tau}}^{(n+1)} = \begin{pmatrix} \text{rotate}(\boldsymbol{\tau}^{(n)}, \Delta t\boldsymbol{\Omega}^{(n+1/2)}) \\ \mathbf{X}_{\text{MP}}^{(n)} + \Delta t\mathbf{U}_{\text{MP}}^{(n+1/2)} \end{pmatrix}. \quad (49)$$

Solving (48) yields (to leading order in Δt)

$$\begin{aligned} \boldsymbol{\alpha}^{(n+1/2)} &= \left(\mathbf{N}\mathbf{K}^T \tilde{\mathbf{M}}^{-1} \right)^{(n+1/2,*)} \\ &\quad \times \left(-\tilde{\mathbf{M}}^{(n+1/2,*)}\mathbf{L}\mathbf{X} + \mathbf{U}_B^{(n)} + \mathbf{U}_{\text{MD}}^{(n)} \right). \end{aligned} \quad (50)$$

In Appendix C 2, we show that using this value of $\boldsymbol{\alpha}$ in the nonlinear update (50) generates the drift term (40) in expectation. Thus, after the rotation, we obtain dynamics consistent with (29).

In this paper, we discretize filaments with at most 30 Chebyshev nodes and consider hydrodynamics on one filament at a time. We therefore use direct solvers for the saddle point system (48) in our implementation, which is available (along with python files for all dynamic examples in this paper) at <https://github.com/stochasticHydroTools/SlenderBody>. To implement the midpoint method efficiently for a blob-link discretization, where there are many blobs even on a single filament, we need to use an iterative solver for (48). In fact, if such solvers are accessible, then the midpoint temporal integrator can be applied as is using the positively split Ewald method⁷⁰ to generate random displacements with covariance $\tilde{\mathbf{M}}$, as has been done for suspensions of rigid bodies.⁶⁸ Further details will be provided elsewhere; here, we will only use such an implementation of the blob-link discretization to compare to our spectral results.

IV. EQUILIBRIUM STATISTICAL MECHANICS

This section is devoted to the equilibrium statistical mechanics of fluctuating inextensible filaments. We focus first on comparing the Gibbs-Boltzmann distribution (6) for blob-link and spectral chains through Markov Chain Monte Carlo (MCMC) calculations and 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. We focus on two examples: small fluctuations of a filament held near a curved base state, and free fibers. The former example is advantageous because it allows us to linearize both the inextensibility constraint (for static calculations) and the SDE (35) (for dynamics) and then break the dynamics into a set of modes. We can then compare the variance of each mode between the blob-link and spectral discretizations. Free fibers, however, are far more common and relevant in practice, and so they are our focus in this section. We relegate the details on the curved filament to Appendix D and here give only a few summary statements to point to the parallels between the two examples.

For free fibers, we will use the end-to-end distance,

$$r(t) = \frac{1}{L} \|\mathbb{X}(s=0, t) - \mathbb{X}(s=L, t)\|, \quad (51)$$

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 approximately⁷⁶

$$\begin{aligned} G(r) &= \frac{1}{Z} \sum_{\ell=1}^{\infty} \frac{1}{(\ell_p^* (1-r))^{3/2}} \\ &\quad \times \exp\left(-\frac{(\ell-1/2)^2}{\ell_p^* (1-r)}\right) H_2\left(\frac{\ell-1/2}{\sqrt{\ell_p^* (1-r)}}\right) r^2, \end{aligned} \quad (52)$$

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 (52) 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).

A. Quantifying the Gibbs-Boltzmann distribution with MCMC sampling

We first examine if the equilibrium distribution (6) for a freely fluctuating *spectral* filament approximates well that of a blob-link chain. To do this, we use MCMC to sample the Gibbs-Boltzmann measure,

$$d\pi(\mathbf{X}) \propto \exp\left(-\frac{\mathcal{E}(\mathbf{X})}{k_B T}\right) d\mu_0(\mathbf{X}) := \tilde{\pi}(\mathbf{X}) d\mu_0(\mathbf{X}), \quad (53)$$

where $\mathcal{E}(\mathbf{X}) = \frac{1}{2}\mathbf{X}^T \mathbf{L} \mathbf{X}$ and the base measure $d\mu_0(\mathbf{X})$ is defined in (6). We generate a proposed configuration $\tilde{\mathbf{X}}$ by randomly rotating the tangent vectors (see Appendix D 2 c for details), which preserves the base measure. In this case, the probability of acceptance is simply the Metropolis factor,

$$p_{\text{acc}} = \max\left(\frac{\tilde{\pi}(\Delta\tilde{\mathbf{X}})}{\tilde{\pi}(\Delta\mathbf{X})}, 1\right). \quad (54)$$

We use this MCMC procedure 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. 4. 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 toward 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 $\langle \boldsymbol{\tau}(s + \Delta s) \cdot \boldsymbol{\tau}(s) \rangle$ 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 $\boldsymbol{\tau}$ is defined (see Fig. 2). 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 5 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.

In Appendix D 2, we perform a similar analysis for a filament undergoing small thermal fluctuations around a curved base state. By breaking the dynamics into a set of modes of the linearized covariance matrix, we show that the spectral method with $N = 12$ nodes can successfully give the correct variance of the first ten modes (with a larger error for $\ell_p/L = 1$ than for $\ell_p/L \gtrsim 10$), while $N = 24$

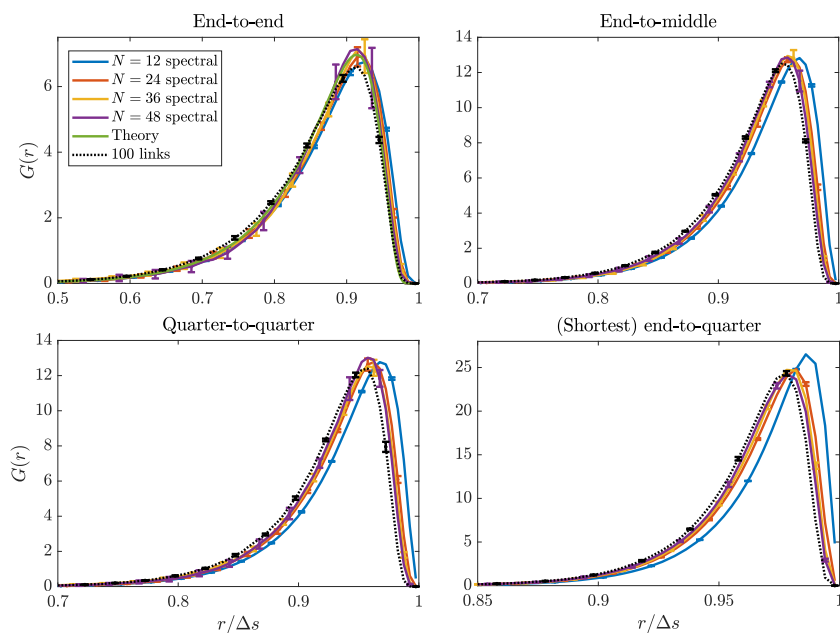


FIG. 4. MCMC results for a freely fluctuating spectral fiber with $\ell_p/L = 1$. We compute distances along the fiber, including end-to-end ($\Delta s = L$, top left), end-to-middle ($\Delta s = L/2$, top right), quarter-to-quarter ($\Delta s = L/2$, bottom left), and the nearest end-to-quarter ($\Delta s = L/4$, 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 (52). 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 .

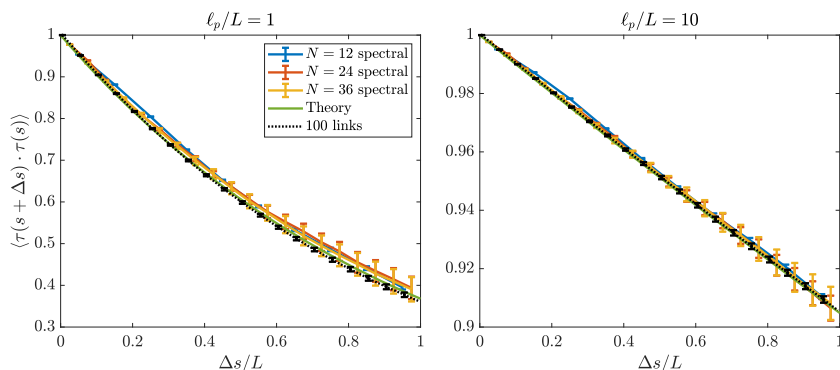


FIG. 5. Tangent vector correlation function $\langle \tau(s + \Delta s) \cdot \tau(s) \rangle$, computed from MCMC sampling. We show the data for a blob-link discretization with 100 links in black and $N = 12, 24,$ and 36 spectral nodes in blue, red, and yellow. The theory $e^{-\Delta s/\ell_p}$ is shown in green.

and $N = 36$ are sufficient to give the correct variance of the first 25 modes (see Fig. 18). Thus, whether we consider small or large fluctuations, the Gibbs–Boltzmann distribution (6) for *spectral* filaments is a good approximation of the more physical one for blob-link chains.

B. Sampling with the midpoint temporal integrator

We now discuss how the midpoint integrator of Sec. III B can also give samples from the Gibbs–Boltzmann distribution (6). 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 (35) around a certain state and compute a set of eigenmodes and associated timescales for the dynamics. This analysis, which we carry out in Appendix D 3 a, is a discrete version of that carried out by Kantsler and Goldstein⁷⁷ in continuum, the difference being that the mobility in the latter case was approximated by local drag so that the calculations could be done semi-analytically. For free filaments, the largest timescale in the problem is associated with the first “fundamental”

bending mode,⁷⁷ which we show in the inset of Fig. 19. The timescale associated with this mode is roughly

$$\tau_{\text{fund}} = 0.003 \frac{4\pi\mu L^4}{\kappa \ln(\hat{\epsilon}^{-1})}, \quad (55)$$

and so we will report time in units of τ_{fund} . There is a slight complication, however, as Fig. 19 shows that the linearized timescales for two different $\hat{\epsilon}$ do *not* collapse onto the same curve when rescaled by the estimate of (55). In fact, the expected log scaling, which comes from slender body theory,^{44,45} 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).

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

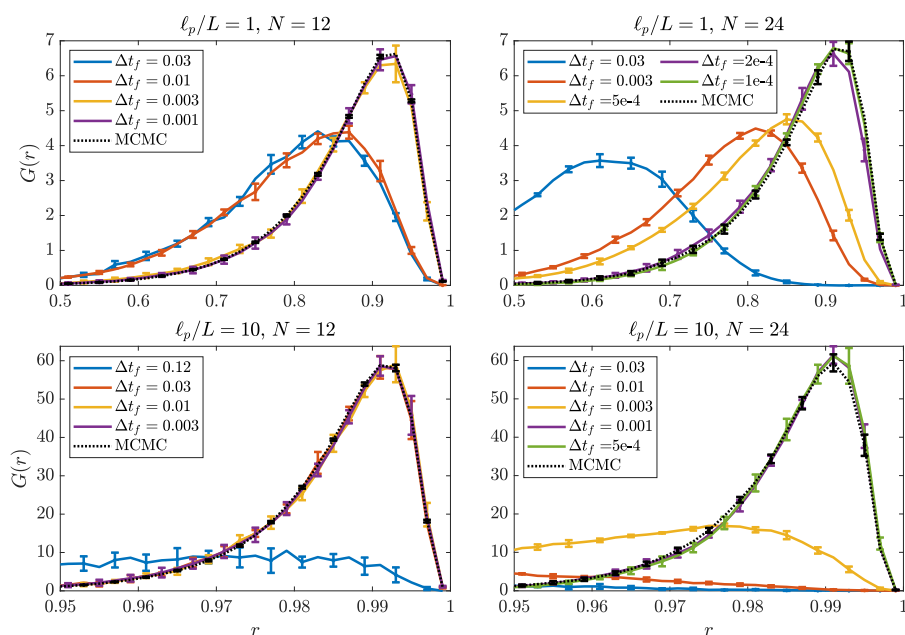


FIG. 6. Distribution of end-to-end distance for the midpoint temporal integrator discussed in Sec. III B with the RPY mobility and $\hat{\epsilon} = 10^{-3}$. Time step sizes are reported as $\Delta t_f = \Delta t/\tau_{\text{fund}}$, where τ_{fund} is the slowest system relaxation time, defined in (55). 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.

(2) 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 6 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. In Appendix D 4, we re-interpret this in terms of modes, finding that we need to resolve roughly twice as many modes when we double N . The scaling of the timescale of each mode $\tau_k \sim k^{-4}$ then implies a decrease in the time step size of 16, which points to the limitations of our temporal integrator for larger N .

Switching our focus to κ , Fig. 6 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. In terms of modal analysis, the number of modes we need to resolve decreases by about 3 for every factor of 10 increase in κ (see Appendix D 4), hence the increase in relative time step size. However, since the timescale τ_{fund} (and the timescale of each of the modes) scales like $1/\kappa$, the net effect of this behavior is no change in the absolute time step size required for accurate equilibrium statistics.

Appendix D 4 also shows how the required time step size changes with $\hat{\epsilon}$, although the analysis is less straightforward since there is no simple rescaling of time in this case. Our results show that the number of modes we need to resolve increases weakly as $\hat{\epsilon}$ decreases so that our time step size drops by a factor of roughly 5 when we drop from $\hat{\epsilon} = 10^{-2}$ to $\hat{\epsilon} = 10^{-3}$.

V. 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 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 $\hat{\epsilon}$. For a direct blob-link discretization, it has already been established that $1/\hat{\epsilon}$ beads are required to resolve hydrodynamics^{40,41} although one could use an asymptotic theory, such as slender body theory (SBT), to model the hydrodynamics approximately.^{54,55} However, in our spectral discretization, we can resolve deterministic hydrodynamics with $\mathcal{O}(1)$ points (i.e., N independent of $\hat{\epsilon}$) (Ref. 36, Sec. 4.4); in this section (and Appendix B 4), 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.^{77–80} 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. 6. As discussed in Ref. 79, 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.⁸⁰

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,

$$t^* = \frac{t}{\bar{t}}, \quad \bar{t} = 0.0008 \frac{4\pi\mu L^4}{\kappa \ln(\hat{\epsilon}^{-1})} \approx 0.27\tau_{\text{fund}}, \quad r^* = \frac{r - \bar{r}}{1 - \bar{r}}, \quad (56)$$

where \bar{r} is the mean end-to-end distance computed in Sec. IV A 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. 19), meaning that all 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. 7 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 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 (MSD) 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 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 $(2k_B T \bar{t}) \text{trace}(N_{\text{tt}})$, where N_{tt} is the 3×3 matrix




FIG. 7. 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 (56)]. In both cases, realizations of the fibers at different times are artificially staggered along the direction of the arrows for visual clarity. See the supplementary material for an animated version of this figure.

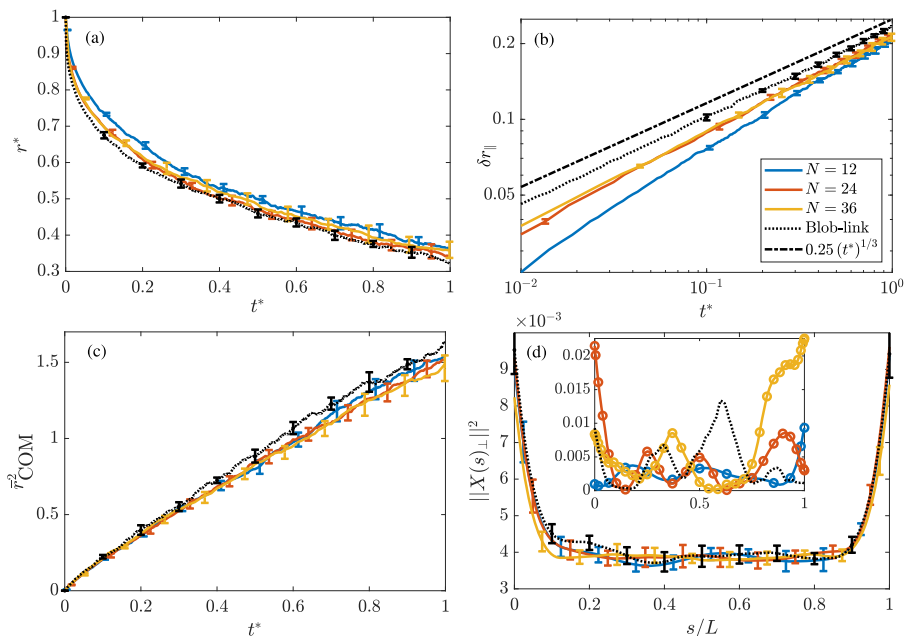


FIG. 8. Comparing the spectral and blob-link trajectories for relaxation of a straight filament to its equilibrium fluctuations. We set $\ell_p/L = 1$ and $\hat{\epsilon} = 10^{-2}$ and simulate until $t^* = 1$, as defined in (56), comparing spectral discretizations with $N = 12$ (blue), $N = 24$ (red), 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 Sec. IV A (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 the 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.

relating forces on a rigid fiber to its translational velocity. The normalized displacement is denoted as \bar{r}_{COM}^2 .

To separate the error in the *dynamics* from that of equilibrium statistical mechanics, in Fig. 8(a), we normalize r^* by the average \bar{r} obtained from MCMC in Sec. IV A. 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(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 $N = 12$ (respectively, $N = 24$) uses a time step that is two (respectively, one) orders of magnitude larger than that of the blob-link chains, as well as a number of collocation points that is an order of magnitude fewer.

To compare our results to theory, in Fig. 8(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 Ref. 79, 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 Ref. 79, 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(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 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 9 shows the results of this test in the end-to-end distance [compare to Figs. 8(a) and 8(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.⁷⁹ Without blob-link data, it is difficult to say 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




FIG. 9. 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 Figs. 8(a) and 8(b), and the difference between the two discretizations is similar to the $\hat{\epsilon} = 10^{-2}$ case, as desired.

slender fibers with the same number of points, independent of the fiber aspect ratio.

VI. BUNDLING OF TRANSIENTLY CROSS-LINKED SEMIFLEXIBLE FILAMENTS

In previous work,⁸¹ some of us examined the dynamics of bundling in transiently cross-linked actin networks. These networks form an important part of the cytoskeleton, and their arrangement into bundled structures is a topic of interest for cell division and motility.^{4,82} As such, they have also been reconstituted *in vitro*,^{13,83–85} and the focus of our previous study⁸¹ was to compare our computational results to those observed experimentally in Ref. 13. In Ref. 81, 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 filaments, leaving semiflexible filaments for future work. Having developed a temporal integrator for semiflexible filaments, we are now ready to complete our investigation by studying the role of semiflexible fluctuations in bundling of cross-linked fiber networks. We emphasize that there are *no* hydrodynamic interactions between different fibers in these simulations.

To simulate cross-linking, we couple the filament model developed here with a Markov chain describing the transiently bound cross-linkers (CLs). Our model of cross-linking is laid out in detail in Refs. 15 and 81, so here, we present only a summary. Each filament is divided into N_u uniformly spaced binding sites separated by distance $\Delta s_u = L/(N_u - 1)$. Assuming that the CLs diffuse rapidly relative to the filaments, the binding of one end of a CL to one of these sites can be approximated by a single rate k_{on} with units $1/(\text{length} \times \text{time})$. Once the first end is bound, the second end can bind to a nearby filament with rate

$$k_{on,s}(\ell'_k) = k_{on,s}^0 \exp\left(-\frac{K_c}{2} \frac{(\ell'_k - \ell_c)^2}{k_B T}\right), \quad (57)$$

where ℓ'_k is the deformed length of the CL (distance between the pair of binding sites), ℓ_c is the rest length of the CL, and K_c is its stiffness. The relationship (57) ensures that the CL dynamics are in detailed balance, that is, the links are passive and do not consume energy. To efficiently search for nearby pairs of filaments, we limit ℓ'_k to two standard deviations of the Gaussian (57), that is, we only

search for pairs of binding sites $2\sqrt{k_B T/K_c}$ apart. Each of the binding reactions has an associated unbinding (reverse) reaction with a rate on the order $1/s^{86}$ so that there are a total of four possible reactions, which are simulated using a version of the standard Stochastic simulation/Gillespie algorithm.^{15,87,88}

We use a time splitting algorithm to update the filaments and cross-linkers in sequence. At each time step, we take a step Δt of the stochastic simulation algorithm with the filament positions fixed. This gives pairs of binding sites that are bound together and, consequently, a force $\mathbf{F}^{(CL)}$ exerted on the corresponding filament pairs. In previous work (Ref. 53, Sec. 6.1), we spread this force as a smoothed delta function around the cross-linker binding location, ensuring smoothness of the cross-linking force density and the subsequent fiber shapes. Since the smoothness assumption does not apply to fluctuating filaments, it is more physical to use instead the spring cross-linking energy

$$\mathcal{E}_{CL} = \frac{K_c}{2} \left(\left\| \mathbb{X}^{(i)}(s_i^*) - \mathbb{X}^{(j)}(s_j^*) \right\| - \ell_c \right)^2 \quad (58)$$

between points $\mathbb{X}^{(i)}(s_i^*)$ (on fiber i) and $\mathbb{X}^{(j)}(s_j^*)$ (on fiber j). If we introduce the matrix \mathbf{R}_u , which resamples the Chebyshev interpolant \mathbb{X} at uniformly spaced binding sites, this energy can be rewritten in terms of the Chebyshev collocation points $\mathbf{X}^{(i)}$ and $\mathbf{X}^{(j)}$ and the resulting force computed by differentiating the energy with respect to \mathbf{X} . The final expression for the force $\mathbf{F}^{(CL)}$ at point p for a CL attached to binding site k then becomes the standard force for a spring (equal and opposite at the two fibers i and j) multiplied by the (k, p) entry of \mathbf{R}_u . These forces at each time step become additional forces in the Langevin equation (35) so that at each time step, we solve (for each fiber independently)

$$\partial_t \mathbf{X} = \mathbf{K} \mathbf{N} \mathbf{K}^T \left(-\mathbf{L} \mathbf{X} + \mathbf{F}^{(CL)} \right) + k_B T \partial_X \cdot \left(\mathbf{K} \mathbf{N} \mathbf{K}^T \right) + \sqrt{2k_B T \mathbf{K} \mathbf{N}^{1/2}} \mathcal{W} \quad (59)$$

by replacing $-\mathbf{L} \mathbf{X}^{(n)}$ with $-\mathbf{L} \mathbf{X}^{(n)} + \mathbf{F}^{(CL)}(\mathbf{X}^{(n)})$ on the right-hand side of the saddle point solve (48) (that is, we treat the spring forces explicitly in time). A quantitative comparison of simulations with the smoothed forcing from Ref. 53, Sec. 6 and the new energy-based forcing from (58) shows little difference between the two models.

Throughout this section, we will use the parameters given in Ref. 81, Table 1. Just as in that study, we consider filaments with

initial mesh size $0.2 \mu\text{m}$, which corresponds to 200 filaments of length $L = 1 \mu\text{m}$ in a periodic domain of edge length $L_d = 2 \mu\text{m}$ and 675 filaments of length $L = 1 \mu\text{m}$ in a periodic domain of edge length $L_d = 3 \mu\text{m}$ (as discussed in Ref. 81, the results are repeatable as we increase the domain size until the structure begins to collapse into 1 or 2 bundles). The question we will examine is how the behavior changes as we increase ℓ_p/L , so we will leave all parameters constant except the bending stiffness κ . This includes $k_B T = 4.1 \times 10^{-3} \text{ pN } \mu\text{m}$, which corresponds to the thermal energy at room temperature. For our spatial and temporal discretization, we use $N = 12$ and $\Delta t = 10^{-4} \text{ s}$ 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 $K_c = 10 \text{ pN}/\mu\text{m}$ is an order of magnitude estimate for the effective stiffness of α -actinin,^{13,89} limits the time step size. In particular, resolving the spring dynamics automatically resolves the equilibrium statistical mechanics.⁹⁰

A. Visualizing the bundling process

In previous work,⁸¹ we showed that bundling of filaments occurs via a thermal zipper mechanism, where cross-linkers stretch to bind nearby pairs of filaments and then contract to their rest length. This contraction pulls the filaments closer together, which allows for binding of additional cross-linkers. The resulting equilibrium configuration contains filaments, which are aligned in parallel and spaced roughly a distance equal to the cross-linker rest length (Ref. 81, Fig. 1). In Ref. 81, we showed that this small-scale ratcheting mechanism also leads to large scale bundling for networks of rigid diffusing filaments and that the bundling process roughly occurs in two stages: first, individual filaments come together into small bundles of a few filaments. Then, the bundles begin to coalesce, forming bundles of bundles and eventually one very large bundle. To quantify this process, we map the filament connections to a connected graph, where a connection in the graph exists when at least two CLs spaced $L/4$ apart connect the two filaments. The “bundles” are then connected regions in this graph. We track over time the bundle density, defined as the number of bundles per unit volume, and see a peak before the bundles start to coalesce. As discussed in Ref. 81, our definition of bundle density, while arbitrary, is a good way to compare the dynamics across multiple systems.

The same fundamental process plays out in networks of semiflexible filaments, as shown in Fig. 10, 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 ($t = 1.5 \text{ s}$) correspond to the 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 $t = 6 \text{ s}$). It is not clear, however, to what extent these differences in bending deformations are driven by CL forces vs the thermal bending fluctuations (a similar question was studied in Ref. 7 for microtubule networks, where the authors showed that large *nonthermal* forces combine with polymerization dynamics to generate bent microtubule shapes in cells). Indeed, we did

show in previous work⁸¹ 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.

B. 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. In previous work,⁸¹ 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 the rigid ones when $\ell_p/L \sim 10$. We used this assumption to justify neglect of thermal bending fluctuations for *actin* filaments of length $1 \mu\text{m}$, for which $\ell_p/L \approx 17$.¹¹ In fact, we neglected any bending and treated the filaments as rigid so that they evolved under both CL forces and thermal rotation and translation. We showed that translational and rotational diffusion accelerates the bundling process significantly since there is more mixing and filaments are able to find each other faster (assuming that there are sufficiently many CLs to link two filaments that are close together).

In this work, 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 Sec. III B 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. That is, instead of (59), we consider the dynamics

$$\partial_t \mathbf{X} = \mathbf{K} \mathbf{N} \mathbf{K}^T (-\mathbf{L} \mathbf{X} + \mathbf{F}^{(\text{CL})}) + \sqrt{2k_B T} \mathbf{K} \mathbf{N}_r^{1/2} \mathcal{W}, \quad (60)$$

where \mathbf{K}_r is the kinematic matrix for a *rigid* fiber, which acts on a 6-vector $(\boldsymbol{\Omega}, \mathbf{U}_{\text{MP}})$ to give velocity in an analogous way to (17). The 6×6 rigid-body mobility is $\mathbf{N}_r = (\mathbf{K}_r^T \tilde{\mathbf{M}}^{-1} \mathbf{K}_r)^{-1}$,⁹¹ exactly as in (27). Our rationale for (60) is that if such semiflexible rigid body motion (“SF-RBD”) simulations give the same results as those with semiflexible bending (“SF-Bend”) fluctuations, then bending fluctuations are not important to the bundling process.

To advance the dynamics of SF-RBD filaments, we first take a step of the stochastic simulation algorithm for the CLs (as for SF-Bend fibers) and then solve (60) using a splitting scheme. The splitting scheme is to first perform a rigid body rotation and translation to add the random term in (60) [see Ref. 81, Eq. (15)] and then compute the cross-linking forces and perform a deterministic saddle point solve to capture the deterministic term in (60).

We begin by looking at the differences in bundling dynamics between *semiflexible* filaments and the rigid filaments we considered in Ref. 81. This is actually the same test we performed in Ref. 81, 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. 11, where we consider $\ell_p/L = 1, 10$, and 100 and

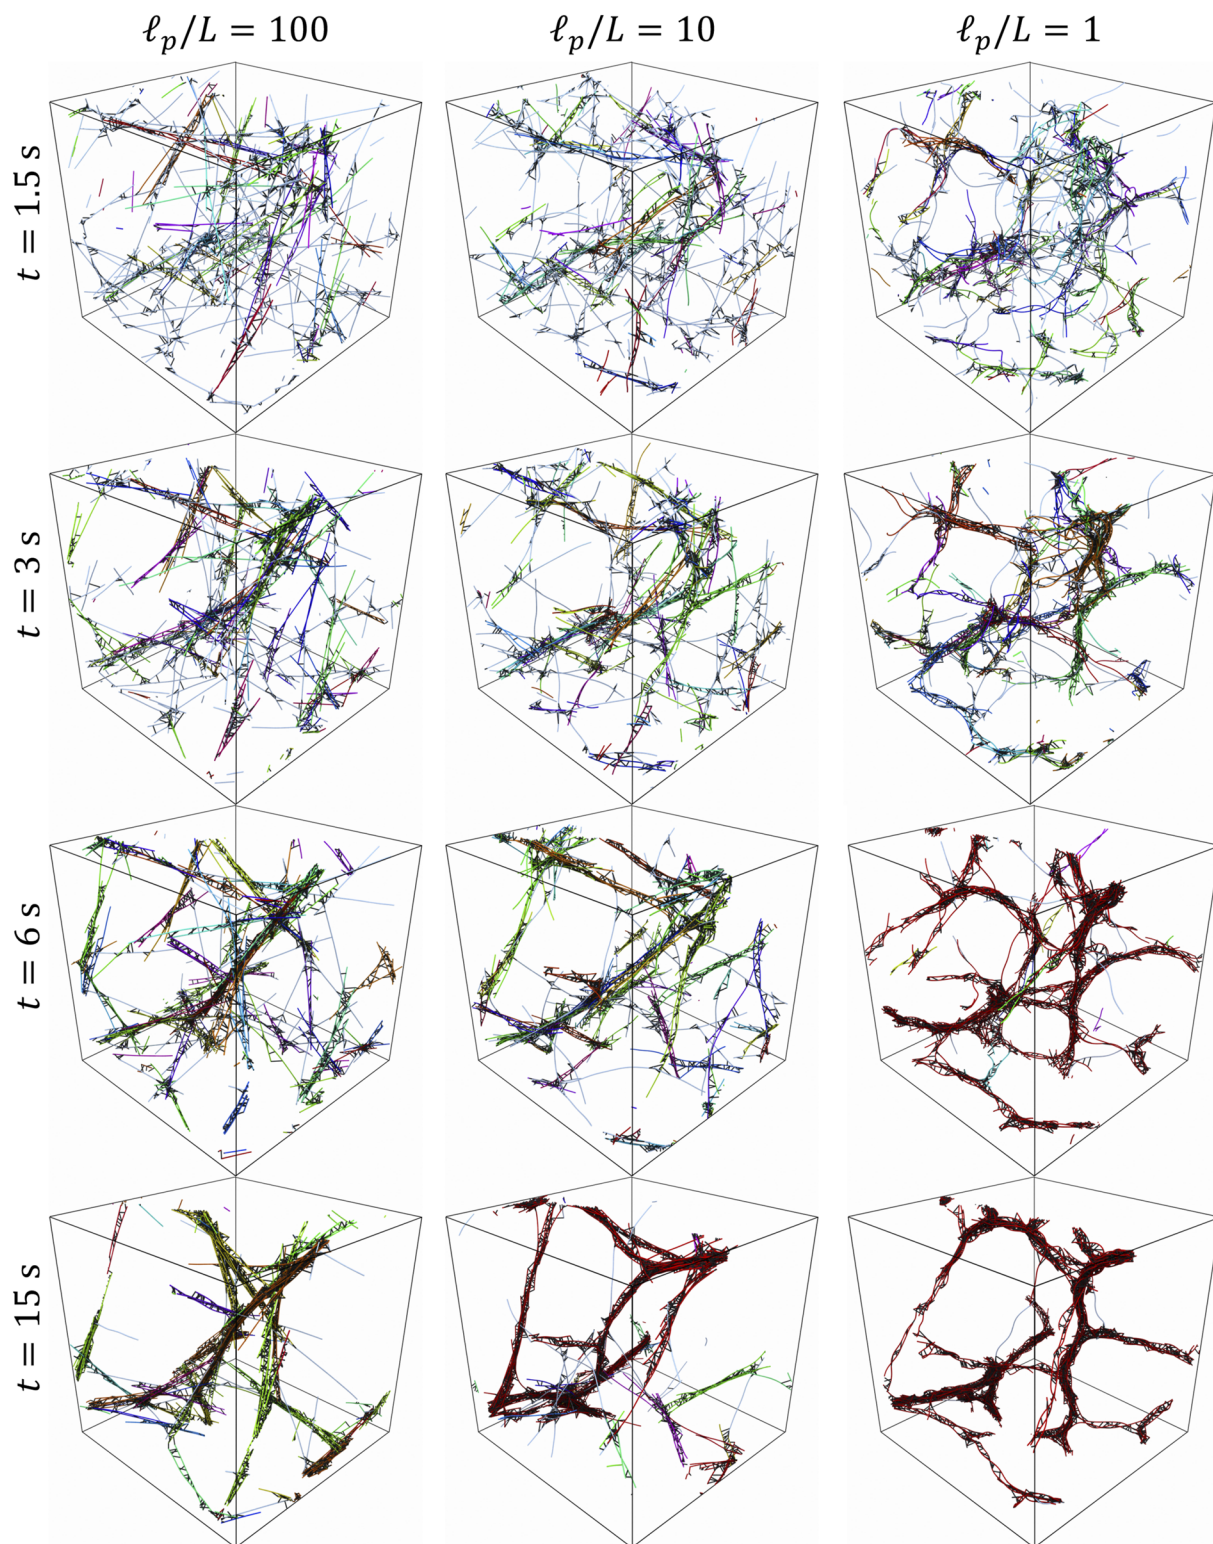


FIG. 10. Snapshots of the bundling process in networks of varying fiber stiffness. Colored fibers are actin filaments (colored by bundle), while black fibers are the cross-linkers. The bundling process is fastest for the most flexible fibers, where the final morphology shows curved bundles.




FIG. 11. 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 (see Ref. 81 for discussion of these). 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 simulated in Ref. 81 (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 the 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 $\ell_p/L = 1$. See the [supplementary material](#) for animations comparing bundling dynamics in SF-Bend and SF-RBD (or rigid) filaments.

compare the results to the case when the fibers are actually rigid, the dynamics of which are governed by the overdamped Ito Langevin equation,

$$\partial_t \mathbf{X} = \mathbf{K}_r \mathbf{N}_r \mathbf{K}_r^T \left(-\mathbf{LX} + \mathbf{F}^{(\text{CL})} \right) + \sqrt{2k_B T \mathbf{K}_r \mathbf{N}_r^{1/2}} \mathbf{W}. \quad (61)$$

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. 10, but it does show the ability of our temporal integrator to remain accurate in the stiff limit (this is not the case for SF-RBD dynamics.⁹¹ Dropping to $\ell_p/L = 10$, the pictures in Fig. 10 show curved bundles, and the plots in Fig. 11 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 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. 11. 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 $\ell_p/L \gtrsim 10$ and demonstrates that the curvature of the bundles we see in Fig. 10 when $\ell_p/L = 10$ is indeed driven primarily 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. 10 that appear to have wavy spatial shapes. This implies that thermal bending fluctuations can accelerate bundling in both 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 ,¹¹ we can conclude that the bundling dynamics of filaments with length 1–2 μm are *not* significantly impacted by semiflexible bending fluctuations.

VII. CONCLUSIONS

This paper represents a first step in applying spectral methods to fluctuating inextensible filaments in Stokes flow. While the advantages of spectral methods for simulating *smooth* fibers are well known,^{47,53,92} this paper is to our knowledge the first attempt to use Chebyshev polynomials to simulate fibers that are inherently nonsmooth due to Brownian bending fluctuations. Our main motivation for doing this comes from hydrodynamics: our slender body quadrature scheme developed in Ref. 36 allows us to compute the hydrodynamic mobility using $\mathcal{O}(1)$ points per filament (with respect to the aspect ratio ϵ), as opposed to the $1/\epsilon$ points required in traditional blob-link (bead-link) methods.^{35,40,41} The quadrature scheme works on a spectral grid since it requires a global interpolating function for the fiber centerline. Thus, to include fluctuations, we postulated the Gibbs–Boltzmann distribution (6) on the spectral grid of N nodes and then constructed spatially discrete overdamped Langevin dynamics that is in detailed balance with respect to that distribution.

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, such as 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 (6) 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, which are better suited for simulations where the small $\mathcal{O}(\hat{a})$ lengthscales need to be resolved.

Our approach to the spatial discretization as presented in Sec. II A seems straightforward, but was actually the product of several months of trial and error. The reason is that the nonsmoothness of thermal fluctuations exposes all the weaknesses of a particular spectral discretization, especially when that discretization fails to properly track high-frequency modes. It is therefore vital that our discretization be free from aliasing errors. In particular, it is important that we be able to recover the original tangent vector rotation rates from the fiber velocity field by “inverting” the kinematic matrix and thereby develop temporal integrators that capture the right drift terms in expectation. 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 such as 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).

We showed that our spectral method with slender body RPY quadrature can reduce the number of nodes required to resolve hydrodynamics and also substantially increase the time step size required for accuracy relative to both blob-link methods and spectral methods with other hydrodynamic models (such as local slender body theory). This is an important feature because 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. By breaking the dynamics into a set of modes, we showed that to generate the correct equilibrium statistics, the temporal integrator must resolve a number of modes that scales linearly with N . Seeing as the timescales of the modes scale like k^{-4} for large node index k (see Fig. 19), this means that our 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, such as that used by Schoeller *et al.* in Ref. 35. Stochastic exponential Lie integrators have not yet been developed (to our knowledge), and so they represent an interesting avenue of exploration for future work.

We ended this paper by applying our algorithm to the bundling dynamics of transiently cross-linked fluctuating actin filaments. When the fibers do not turnover, we showed previously⁸¹ that there is nothing stopping them from forming large bundles through a thermal ratcheting mechanism where cross-linkers (CLs) stretch to bind filaments and then contract to pull them closer together and allow additional CL bindings. Here, we showed that the bundling of fluctuating actin filaments, which have persistence length on the order $10 \mu\text{m}$,¹¹ significantly differs from that of rigid fibers, which is in contrast to the conclusion we drew previously when filaments did not fluctuate (and only moved by CL forces).⁸¹ 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.

Our study of bundling dynamics, like our previous work,⁸¹ did not consider hydrodynamic interactions between filaments. In previous work,¹⁵ we showed that such interactions slow down the bundling process for non-fluctuating fibers. An interesting extension of the work here would be to see if such a conclusion also holds for fluctuating filaments. Simulating hydrodynamic interactions with fluctuations is in fact possible in (log) linear time with the midpoint temporal integrator developed here once suitable iterative solvers are developed. For the blob-link discretization, this has already been done by some of us for triply periodic systems and will be shared in separate work. For spectral methods, our previous work,⁵³ which was restricted to the deterministic setting, relied on splitting the mobility into intra- and inter-fiber hydrodynamics and then computing the second piece fast using oversampled quadrature on a GPU by the Positively Split Ewald (PSE) method.^{15,70} This approach might break down for fluctuating fibers unless we can guarantee that each piece is separately symmetric positive definite (SPD).⁷⁰ A simple solution would be to use oversampled quadrature for the mobility on *all fibers* (which Appendix B 4 a shows is what our special quadrature approximates) so that \mathbf{M} and its square root can be applied fast using the PSE method. Other possible approaches are to use a Galerkin approach, which we have shown is consistent

with the discretization presented here [see (23)], to define a consistent SPD mobility matrix or use the positive splitting of the RPY kernel in the PSE method to split the action of the mobility into a near field special quadrature and a far-field fast summation. After we formulate the mobility, it will be interesting to see how many tangent vectors N are required to resolve the hydrodynamics since in the case of bundled suspensions, the lengthscale of filament interaction decreases over time. Presumably, more than $N = 12$ tangent vectors would be required to resolve the individual hydrodynamic interactions, but whether the error made in using a small number of points affects the overall suspension behavior (e.g., the bundling time) significantly is hard to predict.

From an application standpoint, the rheology of transiently cross-linked actin networks with fiber turnover is an exciting open problem in biological physics. Recent experiments have shown that the stress relaxation in the actin cytoskeleton is slow, with a viscous modulus that scales as $\omega^{1/2}$ on long timescales.⁹³ Since then, there have been a number of physical theories seeking to explain this behavior,^{18,94–96} all of which are based to some degree on a coarse-grained relationship between the cross-linker distribution and the stress relaxation timescales. Thus, the ability to simulate many-fiber suspensions for long times is vital to justify more rigorously the assumptions made by continuum theories and study the role of Brownian motion and bending fluctuations in stress relaxation in the actin cytoskeleton. For the bundled networks in Sec. VI, we have already shown¹⁵ that both intra- and inter-fiber hydrodynamic interactions change the viscoelastic moduli in the deterministic setting, and it is therefore interesting to see how or if this conclusion changes with fluctuations. Doing these calculations will require care, however, since in rheology, we are looking for *quantitative* values of stress, which includes a contribution from Brownian motion [Ref. 97, Eq. (3.169)] that needs to be formulated carefully for overdamped inextensible fibers, and computed efficiently numerically. We have already seen how fluctuating fiber numerics can be quite sensitive to spatial and temporal discretizations. When we combine this with other (non-equilibrium) dynamics, such as cross-linker attachment/detachment and oscillatory shear flow, we will need extra care to ensure that the results are not dominated by discretization artifacts. The stress also has a contribution from steric interactions, which are neglected here but can be added in future work by approximating each fiber with $\mathcal{O}(1)$ spherocylinders,⁹⁸ which resolve the smallest lengthscale on which the fiber bends (roughly ℓ_p). Robust quantitative results will give us the ability to connect microscale Brownian filament hydrodynamics to macroscopic cytoskeletal behavior.

SUPPLEMENTARY MATERIAL

See the [supplementary material](#) for an animation of fiber relaxation for blob-link and spectral chains and also the side-by-side comparisons of bundling dynamics for $\ell_p/L = 100$ vs rigid fibers and SF-Bend fibers vs SF-RBD fibers when $\ell_p/L = 10, 1$.

ACKNOWLEDGMENTS

We thank Eric Vanden-Eijnden, Pep Español, Ravi Jagadeeshan, Isaac Pincus, and Hans C. Ottinger for useful discussions about the continuum vs discrete inextensible chain. We also thank

Raul P. Pelaez for help with the implementation of the blob-link algorithm and Alex Mogilner for reviewing the section on actin bundling dynamics. Ondrej Maxian was supported by the NYU Dean's Dissertation Fellowship and the NSF via Grant No. GRFP/DGE-1342536. This work was also supported by the National Science Foundation through Division of Mathematical Sciences Award No. DMS-2052515 and through a Research and Training Group in Modeling and Simulation under Award No. RTG/DMS-1646339.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Ondrej Maxian: Conceptualization (equal); Data curation (lead); Formal analysis (lead); Funding acquisition (equal); Investigation (lead); Methodology (equal); Software (lead); Writing – original draft (lead); Writing – review & editing (equal). **Brennan Sprinkle:** Funding acquisition (equal); Methodology (equal); Software (supporting); Validation (equal); Visualization (lead); Writing – review & editing (supporting). **Aleksandar Donev:** Conceptualization (equal); Formal analysis (supporting); Funding acquisition (equal); Methodology (equal); Project administration (lead); Supervision (lead); Writing – review & editing (lead).

DATA AVAILABILITY

The data that support the findings of this study are openly available at <https://github.com/stochasticHydroTools/SlenderBody>.

APPENDIX A: SADDLE POINT SYSTEM AS GRADIENT DESCENT DYNAMICS

In this appendix, we show that the deterministic dynamics (25) take the form of a constrained gradient descent. Let us consider a backward Euler integrator for (25) and the Lagrangian

$$\mathcal{L}[\mathbf{X}, \boldsymbol{\alpha}, \boldsymbol{\lambda}] = \frac{1}{2} \mathbf{X}^T \mathbf{L} \mathbf{X} + \frac{1}{2\Delta t} (\mathbf{X} - \mathbf{X}^{(n)})^T (\tilde{\mathbf{M}}^{(n)})^{-1} \times (\mathbf{X} - \mathbf{X}^{(n)}) + \boldsymbol{\lambda}^T \tilde{\mathbf{W}} (\mathbf{K}^{(n)} \boldsymbol{\alpha} - \frac{1}{\Delta t} (\mathbf{X} - \mathbf{X}^{(n)})). \quad (\text{A1})$$

Here, the matrices \mathbf{K} and $\tilde{\mathbf{M}}$ are evaluated at time n and, consequently, are constant with respect to \mathbf{X} . In (A1), 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 $\tilde{\mathbf{W}}$ 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,

$$\frac{\delta \mathcal{L}}{\delta \mathbf{X}} = \mathbf{0} \rightarrow -\mathbf{LX} + \tilde{\mathbf{W}}\boldsymbol{\lambda} = \frac{1}{\Delta t} \left(\tilde{\mathbf{M}}^{(n)} \right)^{-1} \left(\mathbf{X} - \mathbf{X}^{(n)} \right), \quad (\text{A2})$$

$$\frac{\delta \mathcal{L}}{\delta \boldsymbol{\lambda}} = \mathbf{0} \rightarrow \mathbf{K}^{(n)} \boldsymbol{\alpha} = \frac{1}{\Delta t} \left(\mathbf{X} - \mathbf{X}^{(n)} \right), \quad (\text{A3})$$

$$\frac{\delta \mathcal{L}}{\delta \boldsymbol{\alpha}} = \mathbf{0} \rightarrow \left(\mathbf{K}^{(n)} \right)^T \tilde{\mathbf{W}}\boldsymbol{\lambda} = \mathbf{0}. \quad (\text{A4})$$

Combining (A2) and (A3) and taking $\Delta t \rightarrow 0$, we arrive at the deterministic saddle point system,

$$\mathbf{K}\boldsymbol{\alpha} = \tilde{\mathbf{M}}(-\mathbf{LX} + \boldsymbol{\Lambda}), \quad (\text{A5})$$

$$\mathbf{K}^T \boldsymbol{\Lambda} = \mathbf{0}. \quad (\text{A6})$$

Here, we have set the force $\boldsymbol{\Lambda} = \tilde{\mathbf{W}}\boldsymbol{\lambda}$, which represents another instance of the conversion between force and force density defined in (12). In this case, it is easy to see how this conversion arises from the L^2 inner product. Equation (A6) is the principle of virtual work invoked in our previous work (Ref. 53, Sec. 3.4).

1. Adding fluctuations

With fluctuations, the same variational technique can be used to derive the saddle point system (31) by adding the work done (entropy dissipated) by the random force $\mathbf{F}_{\text{stoch}}$,

$$\left(\frac{\mathbf{X} - \mathbf{X}^{(n)}}{\Delta t} \right)^T \mathbf{F}_{\text{stoch}} \Delta t = \left(\frac{\mathbf{X} - \mathbf{X}^{(n)}}{\Delta t} \right)^T \sqrt{\frac{2k_B T}{\Delta t}} \left(\tilde{\mathbf{M}}^{(n)} \right)^{-1/2} \boldsymbol{\eta}^{(n)} \Delta t,$$

to the Lagrangian (A2). However, solving system (31) produces only the first and last term of the Langevin equation (35), and not the 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.

APPENDIX B: RPY MOBILITY

In this appendix, we give a more detailed discussion of the RPY mobility, beginning with the definition of the mobility for two particles. We then give a few details on modifications of our “slender-body” quadrature scheme³⁶ for nonsmooth filaments. Specifically, since our quadrature scheme was developed with smooth filaments in mind, we assumed that the tangent vectors obtained by differentiating \mathbb{X} on the N_x point grid have unit length. This is not necessarily the case for random filaments as the tangent vectors are unit length only when confined to the N collocation points tracking $\boldsymbol{\tau}$. We therefore have to make a few small modifications, as detailed in Appendix B 2. In Appendix B 3, we consider the error incurred in symmetrizing mobility (24) and in truncating its negative eigenvalues. We do this by looking at the errors with respect to the symmetric positive definite upsampled mobility (23). Finally, in Appendix B 4, we compare our special quadrature mobility to other mobility options for the relaxing filament in Sec. V, showing that there are significant advantages to using special quadrature from both a spatial and temporal perspective.

1. RPY kernel

To define the RPY kernel between two points \mathbf{x} and \mathbf{y} , we first let $\mathbf{r} = \mathbf{x} - \mathbf{y}$ and $r = \|\mathbf{r}\|$ with $\hat{\mathbf{r}} = \mathbf{r}/r$. The RPY kernel is based on linear combinations of the Stokeslet and doublet, which are defined as

$$\mathbf{S}(\mathbf{x}, \mathbf{y}) = \frac{1}{8\pi\mu} \left(\frac{\mathbf{I} + \hat{\mathbf{r}}\hat{\mathbf{r}}}{r} \right)$$

and

$$\mathbf{D}(\mathbf{x}, \mathbf{y}) = \frac{1}{8\pi\mu} \left(\frac{\mathbf{I} - 3\hat{\mathbf{r}}\hat{\mathbf{r}}}{r^3} \right). \quad (\text{B1})$$

The RPY kernel for an unbounded fluid is then given by^{30,31,70}

$$\tilde{\mathbf{M}}_{\text{RPY}}(\mathbf{x}, \mathbf{y}; \hat{\mathbf{a}}) = \begin{cases} \mathbf{S}(\mathbf{x}, \mathbf{y}) + \frac{2\hat{\mathbf{a}}^2}{3} \mathbf{D}(\mathbf{x}, \mathbf{y}), & r > 2\hat{\mathbf{a}}, \\ \frac{1}{8\pi\mu} \left[\left(\frac{4}{3\hat{\mathbf{a}}} - \frac{3r}{8\hat{\mathbf{a}}^2} \right) \mathbf{I} + \frac{1}{8\hat{\mathbf{a}}^2 r} \mathbf{r}\mathbf{r} \right], & r \leq 2\hat{\mathbf{a}}. \end{cases} \quad (\text{B2})$$

2. Modifications to slender body quadrature for nonsmooth filaments

A subtle point in computing the mobility on the grid of size N_x is that the quadrature schemes we use to compute (2) are based on having unit-length tangent vectors, i.e., $(\mathbf{DX})_{\{p\}} \cdot (\mathbf{DX})_{\{p\}} = 1$ for all p on the grid of size N_x . For random filaments, the length of the tangent vectors is 1 on the N point grid, but might differ substantially (by as much as 50% or so) when these polynomials are resampled to the $N_x = N + 1$ point grid. As such, we need to modify the quadrature schemes in Ref. 36, Appendix G to obtain a quadrature scheme that converges to the continuum integral (2) as the number of points becomes large.

This is quite straightforward to do. In Ref. 36, Appendix G, we break the RPY integral into three pieces: the Stokeslet on $|s - s'| > 2\hat{\mathbf{a}}$, the doublet on $|s - s'| > 2\hat{\mathbf{a}}$, and the piece for $|s - s'| \leq 2\hat{\mathbf{a}}$. The last piece is done by direct Gauss–Legendre quadrature, so there are no changes to that in this case. The first two pieces are done using special quadrature schemes, which involve subtracting the leading order singularity and integrating what remains. The leading order singularities in the case of a tangent vector that does not have the norm one are

$$\mathbf{S}_1(s, s') = \frac{1}{8\pi\mu} \left(\frac{\mathbf{I} + \widehat{\partial_s \mathbb{X}}(s) \widehat{\partial_{s'} \mathbb{X}}(s')}{\|\widehat{\partial_s \mathbb{X}}(s)\| \|s - s'\|} \right),$$

$$\mathbf{D}_1(s, s') = \frac{1}{8\pi\mu} \left(\frac{\mathbf{I} - 3\widehat{\partial_s \mathbb{X}}(s) \widehat{\partial_{s'} \mathbb{X}}(s')}{\|\widehat{\partial_s \mathbb{X}}(s)\|^3 \|s - s'\|^3} \right),$$

where the hat denotes a normalized vector. Thus, the remaining integrals are

$$\mathbf{U}^{(\text{int}, \text{S})}(s) = \int_{D(s)} \left(\mathbf{S}(\mathbb{X}(s), \mathbb{X}(s')) f(s') - \mathbf{S}_1(s, s') f(s) \right) ds',$$

$$\mathbf{U}^{(\text{int}, \text{D})}(s) = \int_{D(s)} \left(\mathbf{D}(\mathbb{X}(s), \mathbb{X}(s')) f(s') - \mathbf{D}_1(s, s') f(s) \right) ds'.$$

Here, $D(s)$ is the domain on which $|s - s'| \geq 2\hat{\mathbf{a}}$ for each s , defined in Ref. 36, Eq. (G.2).

In Ref. 36, Appendix G, we develop a special quadrature scheme, which is based on putting the integrals in the form

$$\begin{aligned} \mathbf{U}^{(\text{int},S)}(s) &= \int_{D(s)} \mathbf{g}_{\text{tt}}(s, s') \frac{(s' - s)}{|s' - s|} ds', \\ \mathbf{U}^{(\text{int},D)}(s) &= \int_{D(s)} \mathbf{g}_D(s', s) \frac{(s' - s)}{|s' - s|^3} ds'. \end{aligned}$$

Our special quadrature scheme can be implemented as described there, but with the following modified definitions for tangent vectors that do not have the norm one:

$$\begin{aligned} \mathbf{g}_{\text{tt}}(s, s') &= (\mathbf{S}(\mathbb{X}(s), \mathbb{X}(s'))f(s') - \mathbf{S}_1(s, s')f(s)) \frac{|s' - s|}{s' - s}, \\ \mathbf{g}_D(s', s) &= (\mathbf{D}(\mathbb{X}(s), \mathbb{X}(s'))f(s') - \mathbf{D}_1(s, s')f(s)) \frac{|s' - s|^3}{(s' - s)}. \end{aligned}$$

These functions are nonsingular at $s = s'$, with the finite limits

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

3. Eigenvalues of slender body quadrature mobility matrix

In this part of the appendix, we look in more detail at how to determine a systematic threshold for eigenvalue truncation of our slender body quadrature. We do this by comparing the eigenvalues of the symmetrized mobility (24) with the reference mobility

(23). We will use $N_u = 1000$ points for the reference mobility, having verified that using even more reference points gives the same results.

To motivate the truncation of the negative eigenvalues of (24), we first compare the eigenvalues of (24) (without any truncation) to those of the reference mobility (23). To do this, we fix $N = 24$ and $\ell_p/L = 1$ and consider a set of ten samples from the equilibrium distribution studied in Sec. IV A. The mean eigenvalues of $\widetilde{\mathbf{M}}_{\text{ref}}$ and our approximate $\widetilde{\mathbf{M}}$ are shown in Fig. 12, where we have repeated the test three times to generate error bars. Fortunately, as we have observed previously (Ref. 36, Sec. 4.4.1), the eigenvalues of the mobility are not particularly sensitive to the fiber shape; hence, the error bars in Fig. 12 are smaller than the symbol size. For $\hat{\epsilon} = 10^{-3}$, we observe large errors in the largest eigenvalues, which correspond to modes that are flat in the fiber interior and peak at the fiber endpoint. These modes are rapidly damped and are associated with fast timescales.

When we drop to $\hat{\epsilon} = 10^{-2}$, we begin to see negative eigenvalues in the quadrature-based mobility. These negative eigenvalues, which separate themselves from the positive eigenvalues of the reference mobility, are associated with highly oscillatory modes, which have a small Rayleigh quotient $\mathbf{v}^T \widetilde{\mathbf{M}}_{\text{ref}} \mathbf{v}$ using the reference mobility. Because modes with negative eigenvalues in the quadrature mobility are high frequency and, therefore, give small Rayleigh quotients in the reference mobility, a sensible way to truncate the eigenvalues of $\widetilde{\mathbf{M}}$ in (24) is to set any eigenvalues less than the smallest eigenvalue of $\widetilde{\mathbf{M}}_{\text{ref}}$ to that smallest eigenvalue. That is, we post-process the eigenvalues of (24) to set any initial $\lambda_k < \sigma$ to $\lambda_k = \sigma$, where σ is the smallest eigenvalue of the reference mobility matrix $\widetilde{\mathbf{M}}_{\text{ref}}$ for that N , L , and $\hat{\epsilon}$. In Fig. 12, we do this by averaging over 30 samples, but, as discussed previously, this makes little difference, and the eigenvalues for a straight filament could also be used. Thus, in practice, our Python code sets the threshold σ by computing the reference mobility (23) for a straight filament with $1/\epsilon$ upsampling points and finding the smallest eigenvalue of the resulting matrix.

4. Comparing special quadrature to other mobilities

This part of the appendix is concerned with the different options for the matrix $\widetilde{\mathbf{M}}$, which maps force to velocity in fluctuating fibers. We consider four possibilities:




FIG. 12. Eigenvalues of $\widetilde{\mathbf{M}}$ for ten random filaments, sampled from the equilibrium distribution for $\ell_p/L = 1$ and $N = 24$ (here, $\mu = 1$). We show the eigenvalues of the reference mobility (23) in blue and those of the symmetrized quadrature mobility (24) in red. The largest eigenvalues are incorrect for small $\hat{\epsilon}$ because they correspond to highly nonsmooth modes, which are flat in the fiber interior and large at the endpoints. For $\hat{\epsilon} = 10^{-2}$, we start to see negative eigenvalues in the symmetrized mobility (24).

1. RPY special quadrature as discussed in Sec. II B 2, in which we first compute the matrix \mathbf{M} , then symmetrize $\mathbf{M}\tilde{\mathbf{W}}^{-1}$, and truncate its negative eigenvalues.
2. RPY oversampled quadrature on a grid of size N_u , which is the SPD matrix (23). The main part of this computation is the action of the matrix $\tilde{\mathbf{M}}_{\text{RPY},u}$, which maps forces to velocities on the upsampled grid.
3. Direct evaluation of the RPY kernel on the spectral grid, i.e., $\tilde{\mathbf{M}} = \tilde{\mathbf{M}}_{\text{RPY}}$, 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 (12).
4. The local drag approximation to the matrix \mathbf{M} , which is a block-diagonal matrix with 3×3 blocks,

$$8\pi\mu\mathbf{M}_{\{i,i\}} = \left(\frac{a_S(s_i) + a_{\text{CLI}}(s_i)}{\|\partial_s \mathbb{X}(s_i)\|} + \frac{2\hat{a}^2}{3} \frac{a_D(s_i)}{\|\partial_s \mathbb{X}(s_i)\|^3} \right) \mathbf{I} + \left(\frac{a_S(s_i) + a_{\text{CLT}}(s_i)}{\|\partial_s \mathbb{X}(s_i)\|} - 2\hat{a}^2 \frac{a_D(s_i)}{\|\partial_s \mathbb{X}(s_i)\|^3} \right) \boldsymbol{\tau}_{\{i\}} \boldsymbol{\tau}_{\{i\}}, \quad (\text{B3})$$

where $\boldsymbol{\tau}_{\{i\}} = \partial_s \mathbb{X}(s_i) / \|\partial_s \mathbb{X}(s_i)\|$ and

$$a_S(s) = \begin{cases} \ln\left(\frac{(L-s)s}{4\hat{a}^2}\right), & 2\hat{a} < s < L - 2\hat{a}, \\ \ln\left(\frac{(L-s)}{2\hat{a}}\right), & s \leq 2\hat{a}, \\ \ln\left(\frac{s}{2\hat{a}}\right), & s \geq L - 2\hat{a}, \end{cases} \quad (\text{B4})$$

$$a_D(s) = \begin{cases} \frac{1}{4\hat{a}^2} - \frac{1}{2s^2} - \frac{1}{2(L-s)^2}, & 2\hat{a} < s < L - 2\hat{a}, \\ \frac{1}{8\hat{a}^2} - \frac{1}{2(L-s)^2}, & s \leq 2\hat{a}, \\ \frac{1}{8\hat{a}^2} - \frac{1}{2s^2}, & s \geq L - 2\hat{a}, \end{cases}$$

$$a_{\text{CLI}}(s) = \begin{cases} \frac{23}{6}, & 2\hat{a} < s < L - 2\hat{a}, \\ \frac{23}{12} + \frac{4s}{3\hat{a}} - \frac{3s^2}{16\hat{a}^2}, & s \leq 2\hat{a}, \\ \frac{23}{12} + \frac{4(L-s)}{3\hat{a}} - \frac{3(L-s)^2}{16\hat{a}^2}, & s \geq L - 2\hat{a}, \end{cases}$$

$$a_{\text{CLT}}(s) = \begin{cases} \frac{1}{2}, & 2\hat{a} < s < L - 2\hat{a}, \\ \frac{1}{4} + \frac{s^2}{16\hat{a}^2}, & s \leq 2\hat{a}, \\ \frac{1}{4} + \frac{(L-s)^2}{16\hat{a}^2}, & s \geq L - 2\hat{a}. \end{cases}$$

The derivation of this formula is given in Ref. 48, but here, we have made some modifications since the tangent vectors $\partial_s \mathbb{X}(s_i)$ might not be exactly the norm one when the filaments are Brownian. If $\|\partial_s \mathbb{X}(s_i)\| = 1$, then this formula reduces to Ref. 48, Eq. (B7). We note that this local drag theory is qualitatively the same as the more commonly used $8\pi\mu\mathbf{M}_{\{i,i\}} = \ln(\epsilon^{-2})(\mathbf{I} + \boldsymbol{\tau}_{\{i\}}\boldsymbol{\tau}_{\{i\}}) + (\mathbf{I} - 3\boldsymbol{\tau}_{\{i\}}\boldsymbol{\tau}_{\{i\}})$, which describes filaments with ellipsoidally tapered radius functions. Since our filaments are not ellipsoidally tapered, the RPY-based theory from Ref. 48 is quantitatively more accurate (the trends we describe about spatial and temporal convergence are unchanged when we use ellipsoidal SBT). As in step 1, we symmetrize the matrix $\tilde{\mathbf{M}} = \mathbf{M}\tilde{\mathbf{W}}^{-1}$ and truncate its negative eigenvalues (the matrix is not automatically symmetric) to obtain the force-velocity mobility matrix.

Throughout this appendix, we will consider the test of an initially straight single filament relaxing to its equilibrium fluctuations from Sec. V, utilizing (56) for temporal and spatial rescaling (we will rescale each trajectory by its equilibrium mean from MCMC). We will first consider only trajectories that are converged in time so that our goal is to isolate the difference between spatial discretizations. We will then zero in on potential temporal convergence problems for the mobilities that appear advantageous at first.

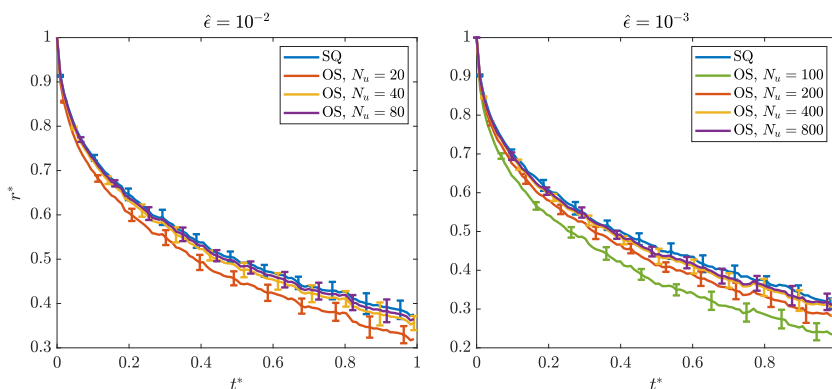


FIG. 13. End-to-end distance trajectory comparing special quadrature (SQ) with oversampled quadrature (OS) with various numbers of oversampling points N_u and a fixed number of $N = 12$ collocation points. The left plot shows $\hat{\epsilon} = 10^{-2}$, for which we can use a small number of oversampled points. The more slender $\hat{\epsilon} = 10^{-3}$ at right requires many more oversampled points, but eventually converges to the special quadrature results.

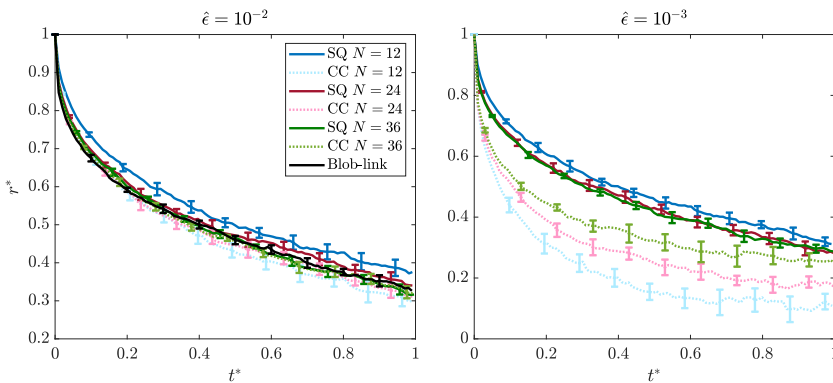


FIG. 14. End-to-end distance trajectory comparing special quadrature (SQ) with direct quadrature (CC) with various numbers of collocation points. The left plot shows $\hat{\epsilon} = 10^{-2}$, for which we can use a small number of collocation points. The more slender $\hat{\epsilon} = 10^{-3}$ at right shows the weaknesses of direct quadrature (similar to an under-resolved bead link model).

a. 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. 13, where we fix $N = 12$ and show that the number of oversampled points N_u required for the special quadrature 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.

b. Special vs direct quadrature

Let us 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 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) although this may be accidental (for this specific example).

As when we compared it to oversampled quadrature, the special quadrature scheme shines only when we drop to $\hat{\epsilon} = 10^{-3}$. This time (right panel of Fig. 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.

c. 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 $\hat{\epsilon}$.




FIG. 15. End-to-end distance trajectory comparing special quadrature (SQ) with the local drag theory (LD) 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 $\hat{\epsilon} = 10^{-3}$.

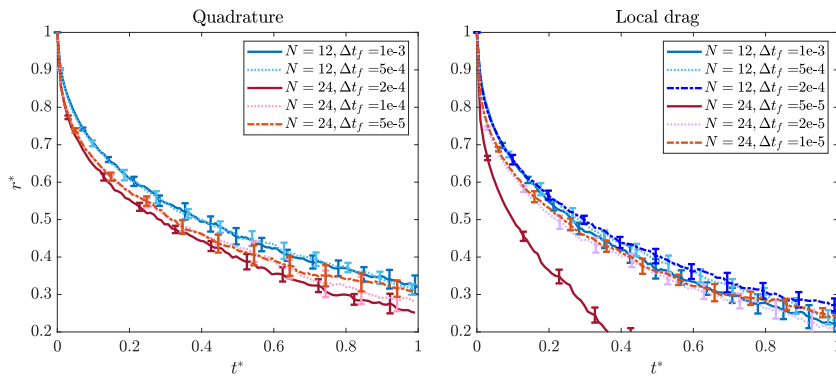


FIG. 16. Studying temporal convergence of special quadrature (left) and local drag (right). Blue colors show $N = 12$, while red colors show $N = 24$. This is for $\hat{\epsilon} = 10^{-3}$, where local drag is most effective. Time step sizes are reported in terms of $\Delta t_f = \Delta t / \tau_{\text{fund}}$, where τ_{fund} is defined in (55).

Figure 15 shows how our local drag theory (B3) 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 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. 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 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.

APPENDIX C: STOCHASTIC DRIFT TERMS

This appendix is devoted to discussing the stochastic drift term (40) in the overdamped Ito Langevin equation for the tangent vector evolution. Our first goal is to show that this drift term is vital to correctly sample from the equilibrium distribution (6). Because our midpoint temporal integrator (Sec. III B) captures the drift term in expectation by moving the fibers to the midpoint, it is not so simple to turn off the drift term and see what happens. Because of this, in Appendix C 1, we present a simple numerical scheme, which captures the drift term via an expensive random finite difference (RFD).²⁹ In that scheme, we can easily turn off the drift term and show the disastrous effects that result when it is not included (see Appendix D 5). After this, in Appendix C 2, we turn to the more

technical details of how the midpoint scheme captures the drift term (40) in expectation. We note that the midpoint scheme as presented in the main text is optimized for dense linear algebra as term (47) minimizes the number of evaluations of \tilde{M} , assuming that a dense inversion is relatively cost-efficient. Here, we present an alternative method to generate the same drift term, which applies when it is impossible to form \tilde{M} as a dense matrix. We then compare our temporal integration scheme to the related scheme of Westwood *et al.* for rigid bodies⁷³ in Appendix C 2 c.

1. 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 (29) *without* the correct drift terms does not give the correct equilibrium statistical mechanics. This scheme begins by solving the system of equations,

$$\tilde{M}^{(n)} \left(L \left(cX^{(n+1,*)} + (1-c)X^{(n)} \right) + \Lambda \right) + U_B^{(n)} = K^{(n)} \alpha, \quad (C1)$$

$$\left(K^{(n)} \right)^T \Lambda = 0,$$

where

$$X^{(n+1,*)} = X^{(n)} + \Delta t K^{(n)} \alpha \quad (C2)$$

is the linearized position at the next time step that we use in the implicit method determined by the coefficient c . Substituting (C2) into (C1), we begin by solving the saddle point system,

$$\begin{pmatrix} -\tilde{M} & (I + c\Delta t \tilde{M}L)K \\ K^T & 0 \end{pmatrix}^{(n)} \begin{pmatrix} \Lambda \\ \alpha \end{pmatrix}^{(n)} = \begin{pmatrix} -\tilde{M}LX^{(n)} + U_B^{(n)} \\ 0 \end{pmatrix}, \quad (C3)$$

where all time-dependent quantities are evaluated at time n , the Brownian velocity U_B is defined in (43), and $\alpha = (\Omega, U_{MP})$. Using the Schur complement to eliminate Λ , we obtain the solution

$$\begin{aligned} \alpha^{(n)} &= \left(-NK^T LX + \sqrt{\frac{2k_B T}{\Delta t}} N^{1/2} \eta \right)^{(n)} + \mathcal{O}(\Delta t) \\ &= \left(-N\bar{C}^T L_\tau \bar{\tau} + \sqrt{\frac{2k_B T}{\Delta t}} N^{1/2} \eta \right)^{(n)} + \mathcal{O}(\Delta t), \end{aligned} \quad (C4)$$

where the terms of order Δt come from using an implicit method. Comparing (C4) with the SDE (29), we see that we still need to capture the drift term (40). To do this, we add a random finite difference term (Ref. 29, Sec. III C) by setting

$$\tilde{\boldsymbol{\tau}}^{(n,\text{RFD})} = \begin{pmatrix} \text{rotate}(\boldsymbol{\tau}^{(n)}, \delta \boldsymbol{\eta}_{\{1:3N\}}^{(n,\text{RFD})}) \\ \mathbf{X}_{\text{MP}} + \delta L \boldsymbol{\eta}_{\{3N+1:3N+3\}}^{(n,\text{RFD})} \end{pmatrix}, \quad (\text{C5})$$

$$\boldsymbol{\alpha}^{(n,\text{RFD})} = \frac{k_B T}{\delta} \left(\mathbf{N}(\tilde{\boldsymbol{\tau}}^{(n,\text{RFD})}) - \mathbf{N}(\tilde{\boldsymbol{\tau}}^{(n)}) \right) \times \begin{pmatrix} \boldsymbol{\eta}_{\{1:3N\}}^{(n,\text{RFD})} \\ L^{-1} \boldsymbol{\eta}_{\{3N+1:3N+3\}}^{(n,\text{RFD})} \end{pmatrix}, \quad (\text{C6})$$

where $\delta \ll 1$ is a dimensionless small parameter and $\boldsymbol{\eta}^{(n,\text{RFD})}$ is generated independent of $\boldsymbol{\eta}^{(n)}$. The procedure (C5) corresponds to first rotating the tangent vectors around a random axis by a random small angle, which we follow in (C6) 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)

$$\begin{aligned} \mathbb{E}[\boldsymbol{\alpha}^{(\text{RFD})}] &= (k_B T) \partial_j (\mathbf{N}_{ik}) \mathbb{E}[\boldsymbol{\eta}_k^{(\text{RFD})} (\tilde{\boldsymbol{\tau}}^{(\text{RFD})} - \tilde{\boldsymbol{\tau}})_j] + \mathcal{O}(\delta) \\ &= (k_B T) \partial_j (\mathbf{N}_{ik}) \mathbb{E}[\boldsymbol{\eta}_k^{(\text{RFD})} \boldsymbol{\eta}_m^{(\text{RFD})}] \tilde{\mathbf{C}}_{jm} + \mathcal{O}(\delta) \\ &= (k_B T) \partial_j (\mathbf{N}_{ik}) \tilde{\mathbf{C}}_{jk}, \end{aligned}$$

which is exactly the drift term in (40). In the second equality, we use the fact that the update (C5) is $\tilde{\mathbf{C}}\boldsymbol{\eta}$ to first order in δ . Thus, the RFD term $\boldsymbol{\alpha}^{(\text{RFD})}$ captures (40) in expectation, and the fiber update

$$\begin{aligned} \boldsymbol{\tau}^{(n+1)} &= \text{rotate}(\boldsymbol{\tau}^{(n)}, \Delta t (\boldsymbol{\Omega}^{(n)} + \boldsymbol{\Omega}^{(n,\text{RFD})})), \\ \mathbf{X}_{\text{MP}}^{(n+1)} &= \mathbf{X}_{\text{MP}}^{(n)} + \Delta t (\mathbf{U}_{\text{MP}}^{(n)} + \mathbf{U}_{\text{MP}}^{(n,\text{RFD})}), \end{aligned} \quad (\text{C7})$$

is a weakly first-order accurate temporal integrator for (29).

2. Drift terms in midpoint scheme

Proceeding to the midpoint scheme in Sec. III B, we now show that the tangent vector update (50) generates the drift term (40) 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. Substituting for \mathbf{U}_B and \mathbf{U}_{MD} , we expand (50) to read

$$\begin{aligned} \boldsymbol{\alpha}^{(n+1/2,*)} &= -(\mathbf{N}\mathbf{K}^T)^{(n+1/2,*)} \mathbf{L}\mathbf{X} \\ &+ \sqrt{\frac{2k_B T}{\Delta t}} (\mathbf{N}\mathbf{K}^T \tilde{\mathbf{M}}^{-1})^{(n+1/2,*)} \tilde{\mathbf{M}}^{1/2} \boldsymbol{\eta} \\ &+ \sqrt{\frac{2k_B T}{\Delta t}} (\mathbf{N}\mathbf{K}^T \tilde{\mathbf{M}}^{-1})^{(n+1/2,*)} \\ &\times (\tilde{\mathbf{M}}^{(n+1/2,*)} - \tilde{\mathbf{M}}) (\tilde{\mathbf{M}}^{-1/2})^T \boldsymbol{\eta}. \end{aligned} \quad (\text{C8})$$

We will Taylor expand the second and third line in succession. Beginning with the second line, we have

$$\begin{aligned} &\sqrt{\frac{2k_B T}{\Delta t}} \left[(\mathbf{N}\mathbf{K}^T \tilde{\mathbf{M}}^{-1})^{(n+1/2,*)} \tilde{\mathbf{M}}^{1/2} \boldsymbol{\eta} \right]_i \\ &= \sqrt{\frac{2k_B T}{\Delta t}} (\mathbf{N}\mathbf{K}^T \tilde{\mathbf{M}}^{-1/2} \boldsymbol{\eta})_i \\ &+ \sqrt{\frac{2k_B T}{\Delta t}} \partial_j (\mathbf{N}\mathbf{K}^T \tilde{\mathbf{M}}^{-1})_{ih} \\ &\times (\tilde{\boldsymbol{\tau}}_j^{(n+1/2,*)} - \tilde{\boldsymbol{\tau}}_j) \tilde{\mathbf{M}}_{hp}^{1/2} \boldsymbol{\eta}_p + \mathcal{O}(\Delta t). \end{aligned}$$

We recognize the first term from (32) 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 (46) into the second line, we get the average drift

$$\begin{aligned} \text{DriftUB}_i &= \sqrt{\frac{2k_B T}{\Delta t}} \partial_j (\mathbf{N}\mathbf{K}^T \tilde{\mathbf{M}}^{-1})_{ih} \\ &\times \sqrt{\frac{k_B T \Delta t}{2}} \tilde{\mathbf{C}}_{ja} \mathbf{K}_{ab}^{-1} \tilde{\mathbf{M}}_{bc}^{1/2} \mathbb{E}[\boldsymbol{\eta}_c \boldsymbol{\eta}_p] \tilde{\mathbf{M}}_{hp}^{1/2} \\ &= (k_B T) \partial_j (\mathbf{N}\mathbf{K}^T \tilde{\mathbf{M}}^{-1})_{ih} (\tilde{\mathbf{M}}_{hb} \mathbf{K}_{ba}^{-T} \tilde{\mathbf{C}}_{aj}^T), \end{aligned}$$

where in the last equality, we used $\tilde{\mathbf{M}} = \tilde{\mathbf{M}}^{1/2} (\tilde{\mathbf{M}}^{1/2})^T$. Rearranging this gives $\tilde{\mathbf{M}} (\tilde{\mathbf{M}}^{-1/2})^T = \tilde{\mathbf{M}}^{1/2}$, which allows us to expand the last line in (C8) as

$$\begin{aligned} \text{DriftUMD}_i &= \mathbb{E} \left[\sqrt{\frac{2k_B T}{\Delta t}} (\mathbf{N}\mathbf{K}^T \tilde{\mathbf{M}}^{-1})_{ih} \partial_j (\tilde{\mathbf{M}}_{hp}) (\tilde{\boldsymbol{\tau}}_j^{(n+1/2,*)} - \tilde{\boldsymbol{\tau}}_j) \tilde{\mathbf{M}}_{qp}^{-1/2} \boldsymbol{\eta}_q \right] + \mathcal{O}(\Delta t) \\ &= (k_B T) (\mathbf{N}\mathbf{K}^T \tilde{\mathbf{M}}^{-1})_{ih} \partial_j (\tilde{\mathbf{M}}_{hp}) \tilde{\mathbf{C}}_{ja} \mathbf{K}_{ab}^{-1} \tilde{\mathbf{M}}_{bc}^{1/2} \mathbb{E}[\boldsymbol{\eta}_c \boldsymbol{\eta}_q] \tilde{\mathbf{M}}_{qp}^{-1/2} \\ &= (k_B T) (\mathbf{N}\mathbf{K}^T \tilde{\mathbf{M}}^{-1})_{ih} \partial_j (\tilde{\mathbf{M}}_{hp}) \tilde{\mathbf{C}}_{ja} \mathbf{K}_{ap}^{-1} \\ &= (k_B T) (\mathbf{N}\mathbf{K}^T \tilde{\mathbf{M}}^{-1})_{ih} \partial_j (\tilde{\mathbf{M}}_{hp} \tilde{\mathbf{C}}_{ja} \mathbf{K}_{ap}^{-1}). \end{aligned} \quad (\text{C9})$$

The last equality makes use of the identity

$$\mathbf{N}_{ik} \mathbf{K}_{kp}^T \partial_j (\tilde{\mathbf{C}}_{jr} \mathbf{K}_{rp}^{-1}) = 0, \quad (\text{C10})$$

which is proven later.

We can now add the two drift terms to obtain the total drift

$$\begin{aligned} \text{Drift}_i &= \text{DriftUB}_i + \text{DriftUMD}_i \\ &= k_B T \partial_j \left(\left(\mathbf{N} \mathbf{K}^T \tilde{\mathbf{M}}^{-1} \right) \left(\tilde{\mathbf{M}} \mathbf{K}^{-T} \tilde{\mathbf{C}}^T \right) \right)_{ij} \\ &= k_B T \partial_j \left(\mathbf{N} \left(\tilde{\mathbf{C}} \mathbf{K}^{-1} \mathbf{K} \right)^T \right)_{ij} = k_B T \partial_j \left(\mathbf{N} \tilde{\mathbf{C}}^T \right)_{ij} \\ &= k_B T \partial_j \left(\mathbf{N}_{ik} \right) \tilde{\mathbf{C}}_{kj}^T. \end{aligned}$$

In the last line above, we used (19) in the second equality, thereby establishing that the one-solve scheme produces the drift (40), as desired.

a. Proof of (C10)

We now prove that (C10) is indeed zero by substituting the definitions of \mathbf{K} and \mathbf{K}^{-1} to get

$$\begin{aligned} \mathbf{N}_{ik} \mathbf{K}_{kp}^T \partial_j (\tilde{\mathbf{C}}_{jr} \mathbf{K}_{rp}^{-1}) &= \mathbf{N}_{ik} \tilde{\mathbf{C}}_{ky}^T \mathcal{X}_{yp}^T \partial_j \left(\tilde{\mathbf{C}}_{jr} \tilde{\mathbf{C}}_{rb}^T \right) \mathcal{X}_{bp}^{-1} \\ &= \mathbf{N}_{ik} \tilde{\mathbf{C}}_{ky}^T \partial_j \left(\tilde{\mathbf{C}}_{jr} \tilde{\mathbf{C}}_{ry}^T \right). \end{aligned}$$

Now, the matrix $(\tilde{\mathbf{C}}_{jr} \tilde{\mathbf{C}}_{rb}^T)$ 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

$$\frac{\partial}{\partial \tilde{\boldsymbol{\tau}}_j} \left(\tilde{\mathbf{C}}_{jr} \tilde{\mathbf{C}}_{ry}^T \right) = -4 \tilde{\boldsymbol{\tau}}_y \quad (\text{C11})$$

for $b \leq 3N$ and zero otherwise (for the parts of $\tilde{\boldsymbol{\tau}}$ associated with constant motions). Multiplying by $\tilde{\mathbf{C}}^T$ and using (15) gives (C10).

b. Alternate way of obtaining mobility drift

An alternative way of obtaining the drift term (47) is via an RFD, which avoids potentially expensive resistance problems. The alternative expression for \mathbf{U}_{MD} in this case is

$$\begin{aligned} \boldsymbol{\mu}^{(\text{RFD})} &= \mathbf{K}^{-1} \boldsymbol{\eta}^{(\text{RFD})}, \\ \boldsymbol{\tau}^{(\text{RFD})} &= \text{rotate}(\boldsymbol{\tau}, \delta L \boldsymbol{\mu}^{(\text{RFD})}) \approx \boldsymbol{\tau}_p - \delta L \mathbf{C} \mathbf{K}^{-1} \boldsymbol{\eta}^{(\text{RFD})}, \\ \mathbf{U}_{\text{MD}} &= \frac{k_B T}{\delta L} \left(\tilde{\mathbf{M}}(\boldsymbol{\tau}^{(\text{RFD})}) - \tilde{\mathbf{M}}(\boldsymbol{\tau}) \right) \boldsymbol{\eta}^{(\text{RFD})}, \end{aligned} \quad (\text{C12})$$

and it is not hard to show that this generates the drift term (C9) in expectation for small δ . Numerical tests show that using the expression (47), rather than (C12), gives a closer approximation to the equilibrium distribution for larger time step sizes, whereas for smaller time step sizes, the two expressions perform similarly. In this paper, we use (47), but (C12) will be useful for many-fiber systems in which we can apply $\tilde{\mathbf{M}}$ fast, while inverting $\tilde{\mathbf{M}}$ requires an iterative solver.

c. Comparison to scheme of Westwood et al.⁷³

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. (Ref. 73, 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 Ref. 73 defined 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

$$\partial_t \mathbf{X}_i = \mathbf{u} + \boldsymbol{\omega} \times (\mathbf{X}_i - \mathbf{q}) = [\mathbf{K}(\mathbf{X}) \mathbf{U}]_i, \quad (\text{C13})$$

where \mathbf{q} is an arbitrary tracking point for the rigid structure. Equation (E5) in Ref. 73 gives the drift term for rigid bodies as

$$\text{Drift}_i = (k_B T) \partial_j \mathbf{N}_{ij}. \quad (\text{C14})$$

Similar to the midpoint scheme presented here, the gDC scheme captures the thermal drift term by splitting it into two parts,

$$\begin{aligned} \text{Drift}_i &= (k_B T) \partial_j \mathbf{N}_{ij} \\ &= (k_B T) \partial_j \left(\mathbf{N} \mathbf{K}^T \mathbf{M}^{-1/2} \mathbf{M}^{1/2} \mathbf{K}^{-T} \right)_{ij} \\ &= (k_B T) \partial_j \left(\mathbf{N} \mathbf{K}^T \mathbf{M}^{-1/2} \right)_{ik} \left(\mathbf{M}^{1/2} \mathbf{K}^{-T} \right)_{kj} \\ &\quad + k_B T \left(\mathbf{N} \mathbf{K}^T \mathbf{M}^{-1/2} \right)_{ik} \partial_j \left(\mathbf{M}^{1/2} \mathbf{K}^{-T} \right)_{kj}, \end{aligned} \quad (\text{C15})$$

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

$$\begin{aligned} \text{Drift}_i &= (k_B T) \partial_j \mathbf{N}_{ij} \\ &= (k_B T) \partial_j \left(\mathbf{N} \mathbf{K}^T \mathbf{M}^{-1} \mathbf{M}^{1/2} \mathbf{M}^{1/2} \mathbf{K}^{-T} \right)_{ij} \\ &= (k_B T) \partial_j \left(\mathbf{N} \mathbf{K}^T \mathbf{M}^{-1} \right)_{ik} \left(\mathbf{M} \mathbf{K}^{-T} \right)_{kj} \\ &\quad + k_B T \left(\mathbf{N} \mathbf{K}^T \mathbf{M}^{-1} \right)_{ik} \partial_j \left(\mathbf{M} \mathbf{K}^{-T} \right)_{kj}. \end{aligned} \quad (\text{C16})$$

While Eqs. (C15) and (C16) 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 (C15). The second term is calculated efficiently by using an RFD for

$$\mathbf{v}_k = \partial_j \left(\mathbf{M}^{1/2} \mathbf{K}^{-T} \right)_{kj} \quad (\text{C17})$$

since it does not involve N , while the first term is captured using a midpoint time integration scheme of the form

$$\mathbf{X}^{(n+1/2,*)} - \mathbf{X}^{(n)} = \sqrt{\frac{k_B T \Delta t}{2}} \mathbf{K}^{-1} \mathbf{M}^{1/2} \boldsymbol{\eta}^{(n)}, \quad (\text{C18})$$

$$\tilde{\mathbf{V}}^{(n+1/2,*)} = \sqrt{\frac{2k_B T}{\Delta t}} (\mathbf{N}\mathbf{K}^T \mathbf{M}^{-1/2})^{(n+1/2,*)} \boldsymbol{\eta}^{(n)} \quad (\text{C19})$$

so that

$$\begin{aligned} \mathbb{E}[\tilde{\mathbf{V}}^{(n+1/2,*)}] &= (k_B T) \partial_j (\mathbf{N}\mathbf{K}^T \mathbf{M}^{-1/2})_{ik} \mathbb{E}[\boldsymbol{\eta}_k^{(n)} \boldsymbol{\eta}_r^{(n)}] \\ &\quad \times (\mathbf{K}^{-1} \mathbf{M}^{1/2})_{jr} \\ &= (k_B T) \partial_j (\mathbf{N}\mathbf{K}^T \mathbf{M}^{-1/2})_{ik} (\mathbf{M}^{1/2} \mathbf{K}^{-T})_{kj}. \end{aligned}$$

The gDC scheme combines these parts to compute

$$\begin{aligned} \text{Drift}_i &= k_B T (\mathbf{N}\mathbf{K}^T \mathbf{M}^{-1/2})_{ik} \mathbf{v}_k + \mathbb{E}[\tilde{\mathbf{V}}^{(n+1/2,*)}] \\ &= (k_B T) \partial_j \mathbf{N}_{ij}. \end{aligned} \quad (\text{C20})$$

Note that calculating $\tilde{\mathbf{V}}^{(n+1/2,*)}$ according to (C19) requires two applications of $\mathbf{M}^{1/2}$ (one at each time level), and calculating \mathbf{v}_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.

APPENDIX D: SMALL FLUCTUATIONS OF A CURVED FIBER

The goal of this appendix, which is meant to accompany the corresponding calculations for a free filament in Sec. IV, is to study the performance of our spatial and temporal discretization 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,⁷⁷ 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 Appendix D 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 Appendix D 2, we compute the equilibrium variance of each of these modes via Markov Chain Monte Carlo (MCMC) calculations using the equilibrium probability distribution (6). By comparing our spectral discretization to a blob-link one, we verify that our equilibrium distribution (6) 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 (35) around the base state \mathbf{X}_0 , which we then diagonalize in Appendix D 3 a to compute relaxation timescales of the fundamental modes. This aids us in Appendix D 4, where we study the time step size required for our midpoint temporal integrator to accurately sample from the equilibrium probability distribution (6).

1. Setup for small fluctuations

For the base state \mathbf{X}_0 , we introduce a curved configuration with $L = 2 \mu\text{m}$ and

$$\boldsymbol{\tau}_0(s) = \frac{1}{\sqrt{2}} \left(\cos(s^3(s-L)^3), \sin(s^3(s-L)^3), 1 \right), \quad (\text{D1})$$

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 $\Delta\mathbf{X} = \mathbf{X} - \mathbf{X}_0$,

$$\mathcal{E}_P = \frac{P}{2} \Delta\mathbf{X}^T \tilde{\mathbf{W}} \Delta\mathbf{X}, \quad P = (1.6 \times 10^4) \frac{k_B T}{L^3}. \quad (\text{D2})$$

This choice of P ensures that the dimensionless discrete L^2 norm $\Delta\mathbf{X}^T \tilde{\mathbf{W}} \Delta\mathbf{X} / L^3$ 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 μm 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,

$$\mathcal{E}_{\text{bend}} = \frac{1}{2} \kappa \Delta\mathbf{X}^T (\mathbf{D}^2)^T \tilde{\mathbf{W}} \mathbf{D}^2 \Delta\mathbf{X}. \quad (\text{D3})$$

In this case, the Ito Langevin equation (35) can be viewed as an equation for $\Delta\mathbf{X}$ rather than \mathbf{X} , and the matrix \mathbf{L} becomes

$$\mathbf{L} = P \tilde{\mathbf{W}} + \kappa (\mathbf{D}^2)^T \tilde{\mathbf{W}} \mathbf{D}^2, \quad (\text{D4})$$

which is the original definition (8) modified to account for the penalty force. Note that the first term in the energy (the penalty term) is independent of κ , while the bending term is proportional to κ . Figure 17 shows \mathbf{X}_0 along with some samples from the Gibbs-Boltzmann distribution (6) 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

$$\mathbf{C}_\alpha = \mathbb{E}[\boldsymbol{\alpha}\boldsymbol{\alpha}^T] = (k_B T) (\mathbf{K}_0^T \mathbf{L} \mathbf{K}_0)^\dagger. \quad (\text{D5})$$

Now, pre- and post-multiplying by \mathbf{K}_0 , we get the expected covariance for $\Delta\mathbf{X}$ as

$$\mathbf{C} = \mathbb{E}[\Delta\mathbf{X}\Delta\mathbf{X}^T] = (k_B T) \mathbf{K}_0 (\mathbf{K}_0^T \mathbf{L} \mathbf{K}_0)^\dagger \mathbf{K}_0^T, \quad (\text{D6})$$

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 and then use the procedure outlined in Appendix D 1 a to project the resulting covariance onto the eigenmodes of \mathbf{C} . This amounts to computing the Rayleigh quotient of the eigenmodes




FIG. 17. 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 (53). The blue line shows \mathbf{X}_0 .

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.

a. Covariance calculations

To compare the covariance from MCMC to (D6), we first use a blob-link discretization to compute the matrix \mathbf{C} exactly as defined in (D6). 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 $\tilde{\mathbf{W}}^{1/2}\mathbf{C}\tilde{\mathbf{W}}^{1/2}$, which is a Hermitian matrix and therefore has eigendecomposition,

$$\tilde{\mathbf{W}}^{1/2}\mathbf{C}\tilde{\mathbf{W}}^{1/2} = \tilde{\mathbf{V}}\mathbf{\Lambda}\tilde{\mathbf{V}}^T, \quad \tilde{\mathbf{V}}^T\tilde{\mathbf{V}} = \mathbf{I}. \quad (\text{D7})$$

The matrix $\tilde{\mathbf{W}}$ is used to compute the L^2 inner product, $\|\mathbf{X}\|_{L^2}^2 \approx \mathbf{X}^T\tilde{\mathbf{W}}\mathbf{X}$. In the case of the blob-link method, $\tilde{\mathbf{W}} = \Delta s \text{diag}(1/2, 1, \dots, 1, 1/2)$. Now, if we use (D7), we observe that

$$\mathbf{C} = \tilde{\mathbf{W}}^{-1/2}\tilde{\mathbf{V}}\mathbf{\Lambda}(\tilde{\mathbf{W}}^{-1/2}\tilde{\mathbf{V}})^T.$$

Thus, if we define $\mathbf{V} = \tilde{\mathbf{W}}^{-1/2}\tilde{\mathbf{V}}$, we have (using the orthonormality of $\tilde{\mathbf{V}}$ in \mathbb{R}^n) the desirable property that

$$\mathbf{C} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T, \quad \mathbf{V}^T\tilde{\mathbf{W}}\mathbf{V} = \mathbf{I}, \quad (\text{D8})$$

which means that 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 $\mathbf{V}^T\tilde{\mathbf{W}}\Delta\mathbf{X}$ since

$$\begin{aligned} \mathbb{E}[\Delta\mathbf{X}\Delta\mathbf{X}^T] &= \mathbf{C}, \\ \mathbb{E}[\mathbf{V}^T\tilde{\mathbf{W}}\Delta\mathbf{X}(\mathbf{V}^T\tilde{\mathbf{W}}\Delta\mathbf{X})^T] &= \mathbf{V}^T\tilde{\mathbf{W}}\mathbf{C}\tilde{\mathbf{W}}\mathbf{V} = \mathbf{\Lambda}, \end{aligned}$$

using (D8) in the last equality.

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 (53), but with $\Delta\mathbf{X} = \mathbf{X} - \mathbf{X}_0$ taking the place of \mathbf{X} in (53) and (54).

a. 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 $\tilde{\mathbf{X}}$ that is generated by rotating the tangent vectors by an oriented angle,

$$\hat{\mathbf{\Omega}}^\perp = 0.1\sqrt{k_B T L}^{-1/2}\boldsymbol{\eta}, \quad (\text{D9})$$

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 $L^{-1/2}$ 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 (D9) 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, although these were not used in this work.

b. Results

Figure 17 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 (D6), 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. 18(a), where we normalize by the corresponding eigenvalue and do the projection in L^2 as discussed in Appendix D 1 a. 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

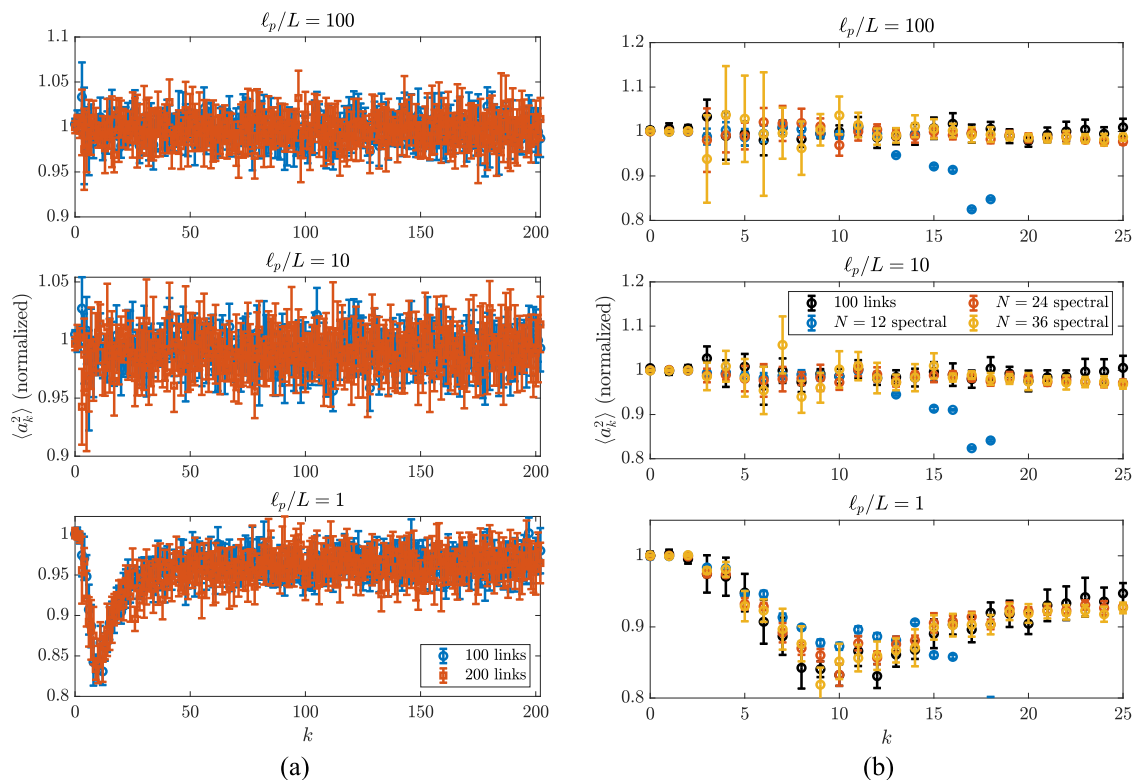


FIG. 18. Covariance of linearized modes for various ℓ_p/L using blob-link and spectral discretizations, estimated using MCMC sampling. We always project the variance onto the same set of L^2 orthonormal modes given by the eigenvectors of the linearized covariance (D6), as discussed in Appendix D 1 a. (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.

mode is exactly that predicted by the theory (D6). 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 sufficiently large for nonlinear effects to matter.

c. 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 nodes,

$$\tilde{\boldsymbol{\Omega}}_{\{p\}} = \frac{0.48}{N} \sqrt{\frac{L}{\ell_p}} \boldsymbol{\eta}_{\{p\}}, \quad (\text{D10})$$

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 $\tilde{\boldsymbol{\Omega}}_p$ and the middle of the fiber is held in place. Note that multiplication of $\boldsymbol{\eta}$ by $L^{-1/2}$ [cf. (D9)] is not necessary since in the spectral discretization, the number of Chebyshev nodes is intended to be relatively small.

The second part of the proposal is to update the fiber midpoint. To do this, we propose a new midpoint via

$$\tilde{\mathbf{X}}_{\text{MP}} = \mathbf{X}_{\text{MP}} + (7.5 \times 10^{-3}) L \boldsymbol{\eta}, \quad (\text{D11})$$

where $\boldsymbol{\eta}$ is a vector of three i.i.d. standard normal numbers. Now, note 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 $\tilde{\boldsymbol{\Omega}}_p$ and $\tilde{\mathbf{X}}_{\text{MP}}$ into a single proposal,

$$\tilde{\mathbf{X}} = \mathcal{X} \begin{pmatrix} \text{rotate}(\boldsymbol{\tau}, \tilde{\boldsymbol{\Omega}}) \\ \tilde{\mathbf{X}}_{\text{MP}} \end{pmatrix}, \quad (\text{D12})$$

and compute the energy via $\frac{1}{2} \Delta \tilde{\mathbf{X}}^T L \Delta \tilde{\mathbf{X}}$; the total acceptance ratio is 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 (D6) in Fig. 18(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).

3. Dynamics

Our task now is to see if our temporal integrators can be used to sample from the Gibbs–Boltzmann distribution (6) 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 main body of this paper to non-dimensionalize the time step sizes.

a. 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 (35) 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 (35) around $\mathbf{X} = \mathbf{X}_0$ is

$$\frac{d\Delta\mathbf{X}}{dt} = -\widehat{\mathbf{N}}_0\mathbf{L}\Delta\mathbf{X} + \sqrt{2k_B T}\widehat{\mathbf{N}}_0^{1/2}\mathcal{W}, \quad (\text{D13})$$

where $\widehat{\mathbf{N}}_0 = \mathbf{K}_0\mathbf{N}_0\mathbf{K}_0^T$, with N defined in (27). 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 (D13) diagonalizes when we consider the first $2N + 3$ generalized eigenvectors of $\widehat{\mathbf{N}}_0^\dagger$ and \mathbf{L} , which satisfy

$$\widehat{\mathbf{N}}_0^\dagger\mathbf{V} = \mathbf{L}\mathbf{V}\mathbf{\Lambda}, \quad (\text{D14})$$

where the $2N + 3$ eigenvectors, which make up the $(3N + 3) \times (2N + 3)$ matrix \mathbf{V} , are normalized such that

$$\mathbf{V}^T\mathbf{L}\mathbf{V} = \mathbf{I} \rightarrow \mathbf{V}^T\widehat{\mathbf{N}}_0^\dagger\mathbf{V} = \mathbf{\Lambda}. \quad (\text{D15})$$

These equations imply that

$$\widehat{\mathbf{N}}_0^\dagger\mathbf{V} = \mathbf{L}\mathbf{V}\mathbf{\Lambda}\mathbf{V}^T\mathbf{L}\mathbf{V} \rightarrow \widehat{\mathbf{N}}_0^\dagger = \mathbf{L}\mathbf{V}\mathbf{\Lambda}\mathbf{V}^T\mathbf{L} \rightarrow \widehat{\mathbf{N}}_0^{-1/2} = \mathbf{L}\mathbf{V}\mathbf{\Lambda}^{1/2}, \quad (\text{D16})$$

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 (D13) and using the definition of the eigenvalues (D14), we obtain the SDE governing the evolution of the eigenmodes $\widehat{\mathbf{X}} = \mathbf{V}^{-1}\Delta\mathbf{X}$,

$$\begin{aligned} \mathbf{V}\frac{d\widehat{\mathbf{X}}}{dt} &= -\widehat{\mathbf{N}}_0\mathbf{L}\mathbf{V}\widehat{\mathbf{X}} + \sqrt{2k_B T}\widehat{\mathbf{N}}_0\mathbf{L}\mathbf{V}\mathbf{\Lambda}^{1/2}\mathcal{W}, \\ \frac{d\widehat{\mathbf{X}}}{dt} &= -\mathbf{\Lambda}^{-1}\widehat{\mathbf{X}} + \sqrt{2k_B T}\mathbf{\Lambda}^{-1/2}\mathcal{W}. \end{aligned} \quad (\text{D17})$$

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. 19, 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 $P = 0$. After the first few modes, the timescales begin to decay with the expected k^{-4} scaling⁷⁸ 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 six 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

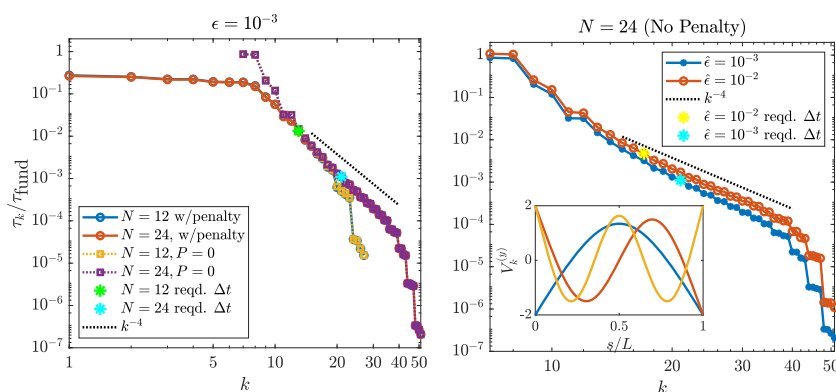


FIG. 19. Timescales in the linearized filament problem. In each plot, we give the timescales (the $2N + 3$ nonzero diagonal entries of $\mathbf{\Lambda}$ in (D14)) made dimensionless by τ_{fund} in (55) (the largest timescale for free fibers) and then mark with a star the time step size required for the midpoint temporal integrator to have an error less than 5% in the equilibrium variance when $\ell_p/L = 10$ (see Appendix D 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 Ref. 77, Fig. 2(a)]. The eigenvalue of the blue mode is the timescale τ_{fund} .

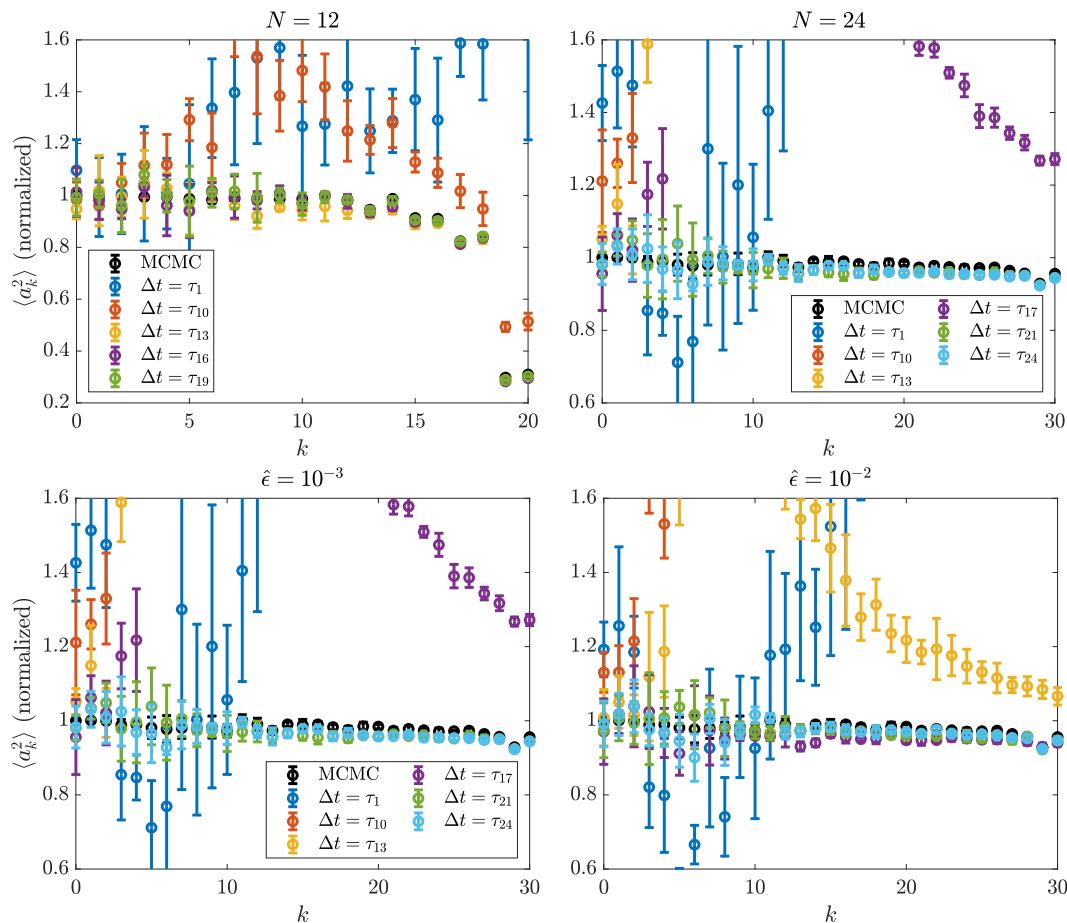


FIG. 20. Convergence of the equilibrium covariance in our midpoint temporal integrator, with Δt expressed in terms of relaxation timescales. In each case, we plot convergence to the MCMC results (see Fig. 18) as a function of the time step size, which is chosen from the modes in Fig. 19. Top: Dependence 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 $\ell_p/L = 10$.

the penalty force contributes a fixed multiple of the identity to 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. 19).

4. Temporal accuracy

We now use the timescale analysis of Appendix D 3 a to understand the performance of our temporal integrator. For each set of parameters, we first compute the relevant timescales using the method of Appendix D 3 a. 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 and then perform the 20 trials a total of five times to generate error bars. The resulting variance of each mode is shown in Figs. 20 and 21. Here, we follow the pattern of Fig. 19 to systematically vary the parameters N (top of Fig. 20), $\hat{\epsilon}$ (bottom), and κ (Fig. 21), while keeping the others constant.

a. Effect of changing N

Beginning with the top of Fig. 20, 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.

b. 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. 20 shows that, for a fixed relative

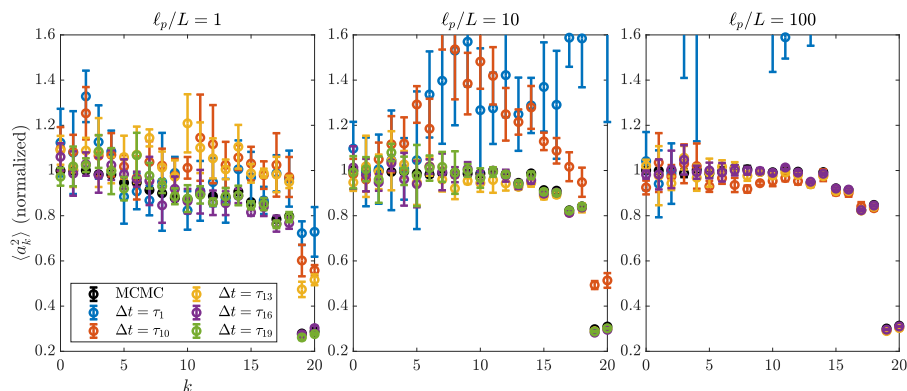


FIG. 21. Convergence of the equilibrium covariance in our midpoint temporal integrator for changing κ , with Δt expressed in terms of relaxation timescales. We fix $N = 12$ and $\hat{\epsilon} = 10^{-3}$ and plot the results for changing κ in terms of ℓ_p/L .

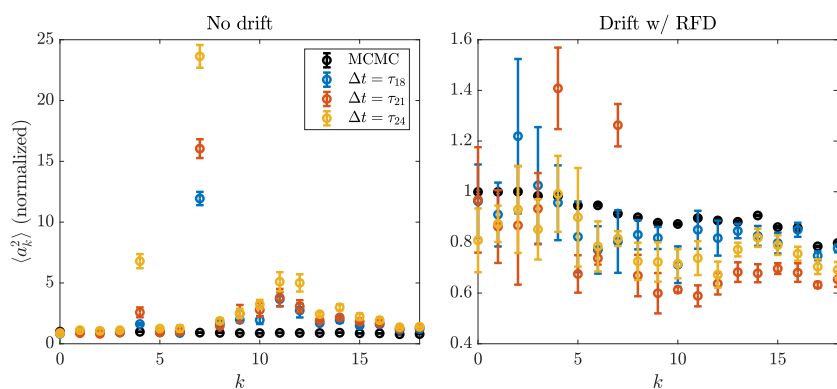


FIG. 22. Results for linearized fluctuations with constant mobility $M = I/(8\pi\mu)$. We fix $N = 12$ and $\ell_p/L = 1$ and show the variance of each of the L^2 -orthonormal modes of the linearized covariance matrix (D6) (as described in Appendix D 1 a) using the simple RFD scheme in Appendix C 1. We show the results without (left) and with (right) the additional RFD drift term in (C6) with $\delta = 10^{-5}$. In all cases, the MCMC results in black are those obtained in Fig. 18(b).

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. 19(b)]. That said, changing the aspect ratio cannot be accommodated by a simple rescaling of time since it scales each mode differently depending on its smoothness.

c. Effect of changing κ

Finally, we consider the variation with ℓ_p/L in Fig. 21, 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).

5. Importance of drift terms

To conclude, we look at what happens to the variance when we exclude the drift terms from the overdamped Langevin equation (35). We do this with the RFD scheme of Appendix C 1, which allows us to explicitly exclude the drift terms to see what kind of covariance we obtain in X . To demonstrate that the drift terms are important even when the mobility is position-independent, we use the constant mobility $M = I/(8\pi\mu)$, which we apply to a fiber with $N = 12$ tangent vectors and $\ell_p/L = 1$. In Fig. 22, 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 Appendix D 2 c. In the left panel, we show the results when we simulate the dynamics *without* adding the RFD term (C6). 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 also note that the RFD scheme appears to require *all* 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.

REFERENCES

- ¹H. C. Berg, *E. Coli in Motion* (Springer, 2004).
- ²E. Lauga and T. R. Powers, "The hydrodynamics of swimming microorganisms," *Rep. Prog. Phys.* **72**(9), 096601 (2009).
- ³A. Mogilner and G. Oster, "Cell motility driven by actin polymerization," *Biophys. J.* **71**(6), 3030–3045 (1996).
- ⁴B. Alberts, A. Johnson, J. Lewis, M. Raff, K. Roberts, and P. Walter, *Molecular Biology of the Cell* (Garland Science, 2002).
- ⁵D. Vavylonis, J.-Q. Wu, S. Hao, B. O'Shaughnessy, and T. D. Pollard, "Assembly mechanism of the contractile ring for cytokinesis by fission yeast," *Science* **319**(5859), 97–100 (2008).
- ⁶M. Murrell, P. W. Oakes, M. Lenz, and M. L. Gardel, "Forcing cells into shape: The mechanics of actomyosin contractility," *Nat. Rev. Mol. Cell Biol.* **16**(8), 486–498 (2015).
- ⁷C. P. Brangwynne, F. C. MacKintosh, and D. A. Weitz, "Force fluctuations and polymerization dynamics of intracellular microtubules," *Proc. Natl. Acad. Sci. U. S. A.* **104**(41), 16128–16133 (2007).
- ⁸B. Nöding and S. Köster, "Intermediate filaments in small configuration spaces," *Phys. Rev. Lett.* **108**(8), 088101 (2012).
- ⁹P. Pawelczyk, N. Mücke, H. Herrmann, and N. Willenbacher, "Attractive interactions among intermediate filaments determine network mechanics in vitro," *PLoS One.* **9**(4), e93194 (2014).
- ¹⁰E. E. Charrier and P. A. Janmey, "Mechanical properties of intermediate filament proteins," *Methods Enzymol.* **568**, 35–57 (2016).
- ¹¹F. Gittes, B. Mickey, J. Nettleton, and J. Howard, "Flexural rigidity of microtubules and actin filaments measured from thermal fluctuations in shape," *J. Cell Biol.* **120**(4), 923–934 (1993).
- ¹²M. D. Welch, A. Mallavarapu, J. Rosenblatt, and T. J. Mitchison, "Actin dynamics in vivo," *Curr. Opin. Cell Biol.* **9**(1), 54–61 (1997).
- ¹³T. T. Falzone, M. Lenz, D. R. Kovar, and M. L. Gardel, "Assembly kinetics determine the architecture of α -actinin crosslinked F-actin networks," *Nat. Commun.* **3**(1), 861 (2012).
- ¹⁴S. L. Freedman, C. Suarez, J. D. Winkelman, D. R. Kovar, G. A. Voth, A. R. Dinner, and G. M. Hocky, "Mechanical and kinetic factors drive sorting of F-actin cross-linkers on bundles," *Proc. Natl. Acad. Sci. U. S. A.* **116**(33), 16192–16197 (2019).
- ¹⁵O. Maxian, R. P. Peláez, A. Mogilner, and A. Donev, "Simulations of dynamically cross-linked actin networks: Morphology, rheology, and hydrodynamic interactions," *PLoS Comput. Biol.* **17**(12), e1009240 (2021). Software available at <https://github.com/stochasticHydroTools/SlenderBody>.
- ¹⁶S. Chen, C. P. Broedersz, T. Markovich, and F. C. MacKintosh, "Nonlinear stress relaxation of transiently crosslinked biopolymer networks," *Phys. Rev. E* **104**(3), 034418 (2021).
- ¹⁷S. Chen, T. Markovich, and F. C. MacKintosh, "Motor-free contractility of active biopolymer networks," *arXiv:2204.00222* (2022).
- ¹⁸Y. Mulla, F. C. MacKintosh, and G. H. Koenderink, "Origin of slow stress relaxation in the cytoskeleton," *Phys. Rev. Lett.* **122**(21), 218102 (2019).
- ¹⁹C. Barbier, R. Dendievel, and D. Rodney, "Numerical study of 3D-compressions of entangled materials," *Comput. Mater. Sci.* **45**(3), 593–596 (2009).
- ²⁰T. Kim, W. Hwang, and R. D. Kamm, "Computational analysis of a cross-linked actin-like network," *Exp. Mech.* **49**(1), 91–104 (2009).
- ²¹R. D. Groot, "Mesoscale simulation of semiflexible chains. I. End-point distribution and chain dynamics," *J. Chem. Phys.* **138**(22), 224903 (2013).
- ²²F. Vargas-Lara and J. Douglas, "Fiber network formation in semi-flexible polymer solutions: An exploratory computational study," *Gels* **4**(2), 27 (2018).
- ²³I. Pincus, A. Rodger, and J. R. Prakash, "Dilute polymer solutions under shear flow: Comprehensive qualitative analysis using a bead-spring chain model with a FENE-Fraenkel spring," *J. Rheol.* **67**, 373 (2022); *arXiv:2206.01870*.
- ²⁴C. Cruz, F. Chinesta, and G. Régner, "Review on the Brownian dynamics simulation of bead-rod-spring models encountered in computational rheology," *Arch. Comput. Methods Eng.* **19**(2), 227–259 (2012).
- ²⁵V. Siva Krishna, P. Kumar, B. Sharma, and I. Saha Dalal, "Petascale Brownian dynamics simulations of highly resolved polymer chains with hydrodynamic interactions using modern GPUs," *arXiv:2208.06559* (2022).
- ²⁶F. Balboa Usabiaga, B. Kallemov, B. Delmotte, A. P. S. Bhalla, B. Griffith, and A. Donev, "Hydrodynamics of suspensions of passive and active rigid particles: A rigid multiblob approach," *Commun. Appl. Math. Comput. Sci.* **11**(2), 217–296 (2016). Software available at <https://github.com/stochasticHydroTools/RigidMultiblobsWall>.
- ²⁷D. C. Morse, "Theory of constrained Brownian motion," *Adv. Chem. Phys.* **128**, 65–189 (2003).
- ²⁸H. C. Öttinger, *Stochastic Processes in Polymeric Fluids: Tools and Examples For Developing Simulation Algorithms* (Springer Science & Business Media, 2012).
- ²⁹S. 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>.
- ³⁰J. Rotne and S. Prager, "Variational treatment of hydrodynamic interaction in polymers," *J. Chem. Phys.* **50**(11), 4831–4837 (1969).
- ³¹E. Wajnryb, K. A. Mizerski, P. J. Zuk, and P. Szymczak, "Generalization of the Rotne–Prager–Yamakawa mobility and shear disturbance tensors," *J. Fluid Mech.* **731**, R3 (2013).
- ³²J. E. Butler and E. S. G. Shaqfeh, "Brownian dynamics simulations of a flexible polymer chain which includes continuous resistance and multibody hydrodynamic interactions," *J. Chem. Phys.* **122**(1), 014901 (2005).
- ³³V. A. Beck and E. S. G. Shaqfeh, "Ergodicity breaking and conformational hysteresis in the dynamics of a polymer tethered at a surface stagnation point," *J. Chem. Phys.* **124**(9), 094902 (2006).
- ³⁴J. Wang, E. J. Tozzi, M. D. Graham, and D. J. Klingenberg, "Flipping, scooping, and spinning: Drift of rigid curved nonchiral fibers in simple shear flow," *Phys. Fluids* **24**(12), 123304 (2012).
- ³⁵S. F. Schoeller, A. K. Townsend, T. A. Westwood, and E. E. Keaveny, "Methods for suspensions of passive and active filaments," *J. Comput. Phys.* **424**, 109846 (2021).
- ³⁶O. Maxian, B. Sprinkle, C. S. Peskin, and A. Donev, "Hydrodynamics of a twisting, bending, inextensible fiber in Stokes flow," *Phys. Rev. Fluids* **7**, 074101 (2022). Software available at <https://github.com/stochasticHydroTools/SlenderBody>.
- ³⁷E. Grazi, "What is the diameter of the actin filament?," *FEBS Lett.* **405**(3), 249–252 (1997).
- ³⁸P. A. Janmey, S. Hvidt, J. Käs, D. Lerche, A. Maggs, E. Sackmann, M. Schliwa, and T. P. Stossel, "The mechanical properties of actin gels. Elastic modulus and filament motions," *J. Biol. Chem.* **269**(51), 32503–32513 (1994).
- ³⁹J. L. McGrath, E. A. Osborn, Y. S. Tardy, C. F. Dewey, and J. H. Hartwig, "Regulation of the actin cycle *in vivo* by actin filament severing," *Proc. Natl. Acad. Sci. U. S. A.* **97**(12), 6532–6537 (2000).
- ⁴⁰T. T. Bringley and C. 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," *J. Comput. Phys.* **227**(11), 5397–5425 (2008).
- ⁴¹B. Kallemov, A. P. S. Bhalla, B. Griffith, and A. Donev, "An immersed boundary method for rigid bodies," *Commun. Appl. Math. Comput. Sci.* **11**(1), 79–141 (2016).
- ⁴²M. J. Shelley and T. Ueda, "The nonlocal dynamics of stretching, buckling filaments," in *Advances in Multi-Fluid Flows* (SIAM, 1996), pp. 415–425.
- ⁴³M. J. Shelley and T. Ueda, "The Stokesian hydrodynamics of flexing, stretching filaments," *Physica D* **146**(1–4), 221–245 (2000).
- ⁴⁴J. B. Keller and S. I. Rubinow, "Slender-body theory for slow viscous flow," *J. Fluid Mech.* **75**(4), 705–714 (1976).
- ⁴⁵R. E. Johnson, "An improved slender-body theory for Stokes flow," *J. Fluid Mech.* **99**(2), 411–431 (1980).
- ⁴⁶A.-K. Tornberg and M. J. Shelley, "Simulating the dynamics and interactions of flexible fibers in Stokes flows," *J. Comput. Phys.* **196**(1), 8–40 (2004).
- ⁴⁷E. Nazockdast, A. Rahimian, D. Zorin, and M. Shelley, "A fast platform for simulating semi-flexible fiber suspensions applied to cell mechanics," *J. Comput. Phys.* **329**, 173–209 (2017).
- ⁴⁸O. Maxian and A. Donev, "Slender body theories for rotating filaments," *J. Fluid Mech.* **952**, A5 (2022).

- ⁴⁹R. Cortez and M. Nicholas, “Slender body theory for Stokes flows with regularized forces,” *Commun. Appl. Math. Comput. Sci.* **7**(1), 33–62 (2012).
- ⁵⁰R. Cortez, “Regularized Stokeslet segments,” *J. Comput. Phys.* **375**, 783–796 (2018).
- ⁵¹C. S. Peskin, “The immersed boundary method,” *Acta Numer.* **11**, 479–517 (2002).
- ⁵²U. M. Ascher and C. Greif, *A First Course on Numerical Methods* (SIAM, 2011).
- ⁵³O. Maxian, A. Mogilner, and A. Donev, “Integral-based spectral method for inextensible slender fibers in Stokes flow,” *Phys. Rev. Fluids* **6**, 014102 (2021). Software available at <https://github.com/stochasticHydroTools/SlenderBody>.
- ⁵⁴H. Manikantan and D. Saintillan, “Subdiffusive transport of fluctuating elastic filaments in cellular flows,” *Phys. Fluids* **25**(7), 073603 (2013).
- ⁵⁵H. Manikantan and D. Saintillan, “Effect of flexibility on the growth of concentration fluctuations in a suspension of sedimenting fibers: Particle simulations,” *Phys. Fluids* **28**(1), 013303 (2016).
- ⁵⁶K. Liu, J. Lowengrub, and J. Allard, “Efficient simulation of thermally fluctuating biopolymers immersed in fluids on 1-micron, 1-second scales,” *J. Comput. Phys.* **386**, 248–263 (2019).
- ⁵⁷P. M. Ryan and C. W. Wolgemuth, “A finite volume algorithm for the dynamics of filaments, rods, and beams,” *J. Comput. Phys.* **466**, 111375 (2022).
- ⁵⁸C. J. Cyron and W. A. Wall, “Finite-element approach to Brownian dynamics of polymers,” *Phys. Rev. E* **80**(6), 066704 (2009).
- ⁵⁹Y. Lin, X. Wei, J. Qian, K. Y. Sze, and V. B. Shenoy, “A combined finite element-Langevin dynamics (FEM-LD) approach for analyzing the mechanical response of bio-polymer networks,” *J. Mech. Phys. Solids* **62**, 2–18 (2014).
- ⁶⁰C. J. Cyron and W. A. Wall, “Numerical method for the simulation of the Brownian dynamics of rod-like microstructures with three-dimensional nonlinear beam elements,” *Int. J. Numer. Methods Eng.* **90**(8), 955–987 (2012).
- ⁶¹C. J. Cyron, K. W. Müller, A. R. Bausch, and W. A. Wall, “Micromechanical simulations of biopolymer networks with finite elements,” *J. Comput. Phys.* **244**, 236–251 (2013).
- ⁶²Of course, we could also take $d\mu_0$ as the base measure and put additional “entropic” or “metric” factors in $\varepsilon_{\text{bend}}$. Here, however, we take $\varepsilon_{\text{bend}}$ to be the standard elastic bending energy of a semiflexible chain.
- ⁶³J. L. Aurentz and L. N. Trefethen, “Block operators and spectral discretizations,” *SIAM Rev.* **59**(2), 423–446 (2017).
- ⁶⁴T. A. Driscoll and N. Hale, “Rectangular spectral collocation,” *IMA J. Numer. Anal.* **36**(1), 108–132 (2015).
- ⁶⁵D. Li, M. Gurnis, and G. Stadler, “Towards adjoint-based inversion of time-dependent mantle convection with nonlinear viscosity,” *Geophys. J. Int.* **209**(1), 86–105 (2017).
- ⁶⁶A more concrete justification for this is as follows: suppose that the filament satisfied the free fiber boundary conditions $\partial_x^2 \mathbb{X}(0, L) = \partial_x^3 \mathbb{X}(0, L) = \mathbf{0}$. 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 $U = \kappa/2 \mathbf{X}^T \tilde{\mathbf{W}} \mathbf{D}^4 \mathbf{X}$. Taking the derivative, we have the force $\mathbf{F} = -\kappa \tilde{\mathbf{W}} \mathbf{D}^4 \mathbf{X}$. However, we know that force density should be $\mathbf{f} = -\kappa \mathbf{D}^4 \mathbf{X}$, which can only be obtained from \mathbf{F} using (12).
- ⁶⁷This is true even for straight fibers. In continuum, the only nontrivial solutions of $\left(\int_{L/2}^s \Omega(s') ds'\right) \times \boldsymbol{\tau} + \mathbf{U}_{\text{MP}} = \mathbf{0}$ are those with $\Omega = \Omega^{\parallel} \boldsymbol{\tau}$. Our discretization preserves this property. Since in our method $\int \Omega(s) ds$ [and, for non-straight fibers $\int \Omega(s) \times \boldsymbol{\tau}(s) ds$] 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.
- ⁶⁸B. Sprinkle, F. Balboa Usabiaga, N. A. Patankar, and A. 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>.
- ⁶⁹E. Chow and Y. Saad, “Preconditioned Krylov subspace methods for sampling multivariate Gaussian distributions,” *SIAM J. Sci. Comput.* **36**(2), A588–A608 (2014).
- ⁷⁰A. M. Fiore, F. Balboa Usabiaga, A. Donev, and J. W. Swan, “Rapid sampling of stochastic displacements in Brownian dynamics simulations,” *J. Chem. Phys.* **146**(12), 124116 (2017).
- ⁷¹M. Hütter and H. C. Öttinger, “Fluctuation-dissipation theorem, kinetic stochastic integral and efficient simulations,” *J. Chem. Soc., Faraday Trans.* **94**(10), 1403–1405 (1998).
- ⁷²M. Fixman, “Simulation of polymer dynamics. I. General theory,” *J. Chem. Phys.* **69**(4), 1527–1537 (1978).
- ⁷³T. A. Westwood, B. Delmotte, and E. E. Keaveny, “A generalised drift-correcting time integration scheme for Brownian suspensions of rigid particles with arbitrary shape,” *J. Comput. Phys.* **467**, 111437 (2022).
- ⁷⁴S. Delong, B. E. Griffith, E. Vanden-Eijnden, and A. Donev, “Temporal integrators for fluctuating hydrodynamics,” *Phys. Rev. E* **87**(3), 033302 (2013).
- ⁷⁵S. 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>.
- ⁷⁶J. Wilhelm and E. Frey, “Radial distribution function of semiflexible polymers,” *Phys. Rev. Lett.* **77**(12), 2581 (1996).
- ⁷⁷V. Kantsler and R. E. Goldstein, “Fluctuations, dynamics, and the stretch-coil transition of single actin filaments in extensional flows,” *Phys. Rev. Lett.* **108**(3), 038103 (2012).
- ⁷⁸S. L. Poelert and A. A. Zadpoor, “Analytical and numerical methods for capturing the thermal fluctuations of semiflexible polymers,” *Macromol. Theory Simul.* **21**(6), 357–371 (2012).
- ⁷⁹B. Obermayer, W. Möbius, O. Hallatschek, E. Frey, and K. Kroy, “Freely relaxing polymers remember how they were straightened,” *Phys. Rev. E* **79**(2), 021804 (2009).
- ⁸⁰P. Dimitrakopoulos, “Longitudinal relaxation of initially straight flexible and stiff polymers,” *Phys. Rev. Lett.* **93**(21), 217801 (2004).
- ⁸¹O. Maxian, A. Donev, and A. Mogilner, “Interplay between Brownian motion and cross-linking controls bundling dynamics in actin networks,” *Biophys. J.* **121**(7), 1230–1245 (2022). Software available at <https://github.com/stochasticHydroTools/SlenderBody>.
- ⁸²L. Chaudbet, A. R. Chaudhary, H. K. Heris, A. J. Ehrlicher, and A. G. Hendricks, “Dynamic actin cross-linking governs the cytoplasm’s transition to fluid-like behavior,” *Mol. Biol. Cell* **31**(16), 1744–1752 (2020).
- ⁸³M. L. Gardel, J. H. Shin, F. C. MacKintosh, L. Mahadevan, P. Matsudaira, and D. A. Weitz, “Elastic behavior of cross-linked and bundled actin networks,” *Science* **304**(5675), 1301–1305 (2004).
- ⁸⁴K. M. Schmoller, O. Lieleg, and A. R. Bausch, “Structural and viscoelastic properties of actin/filamin networks: Cross-linked versus bundled networks,” *Biophys. J.* **97**(1), 83–89 (2009).
- ⁸⁵O. Lieleg, K. M. Schmoller, C. J. Cyron, Y. Luan, W. A. Wall, and A. R. Bausch, “Structural polymorphism in heterogeneous cytoskeletal networks,” *Soft Matter* **5**(9), 1796–1803 (2009).
- ⁸⁶P. A. Kuhlman, J. Ellis, D. R. Critchley, and C. R. Bagshaw, “The kinetics of the interaction between the actin-binding domain of α -actinin and F-actin,” *FEBS Lett.* **339**(3), 297–301 (1994).
- ⁸⁷D. T. Gillespie, “Stochastic simulation of chemical kinetics,” *Annu. Rev. Phys. Chem.* **58**, 35–55 (2007).
- ⁸⁸D. F. Anderson, “A modified next reaction method for simulating chemical systems with time dependent propensities and delays,” *J. Chem. Phys.* **127**(21), 214107 (2007).
- ⁸⁹S. Le, X. Hu, M. Yao, H. Chen, M. Yu, X. Xu, N. Nakazawa, F. M. Margadant, M. P. Sheetz, and J. Yan, “Mechanotransmission and mechanosensing of human alpha-actinin 1,” *Cell Rep.* **21**(10), 2714–2723 (2017).
- ⁹⁰If we convert the time step sizes from Fig. 21 to the units we use here, we need a time step size of 3×10^{-3} , 1.4×10^{-3} , and 1.2×10^{-3} s to resolve the equilibrium fluctuations for $\ell_p/L = 1, 10$, and 100, respectively. All these are at least an order of magnitude larger than that required to resolve the explicit treatment of CLs.

⁹¹The pseudo-inversion in the calculation of N_r is problematic when the fibers are nearly straight, and there are eigenvalues near, but not equal to, zero in the resistance matrix $\mathbf{K}_r^T \tilde{\mathbf{M}}^{-1} \mathbf{K}_r$. 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. 11). For $\ell_p/L = 100$, fibers with semiflexible bending fluctuations behave almost identically to rigid fibers (see Fig. 11), which we can simulate without difficulty.⁸¹

⁹²L. N. Trefethen, *Spectral Methods in MATLAB* (SIAM, 2000), Vol. 10.

⁹³N. Y. Yao, D. J. Becker, C. P. Broedersz, M. Depken, F. C. MacKintosh, M. R. Pollak, and D. A. Weitz, "Nonlinear viscoelasticity of actin transiently cross-linked with mutant α -actinin-4," *J. Mol. Biol.* **411**(5), 1062–1071 (2011).

⁹⁴K. W. Müller, R. F. Bruinsma, O. Lieleg, A. R. Bausch, W. A. Wall, and A. J. Levine, "Rheology of semiflexible bundle networks with transient linkers," *Phys. Rev. Lett.* **112**(23), 238102 (2014).

⁹⁵C. P. Broedersz, M. Depken, N. Y. Yao, M. R. Pollak, D. A. Weitz, and F. C. MacKintosh, "Cross-link-governed dynamics of biopolymer networks," *Phys. Rev. Lett.* **105**(23), 238101 (2010).

⁹⁶N. Y. Yao, C. P. Broedersz, M. Depken, D. J. Becker, M. R. Pollak, F. C. MacKintosh, and D. A. Weitz, "Stress-enhanced gelation: A dynamic nonlinearity of elasticity," *Phys. Rev. Lett.* **110**(1), 018103 (2013).

⁹⁷M. Doi and S. F. Edwards, *The Theory of Polymer Dynamics* (Oxford University Press, 1988), Vol. 73.

⁹⁸J.-Y. Chen, Z. Li, I. Szlufarska, and D. J. Klingenberg, "Rheology and structure of suspensions of spherocylinders via Brownian dynamics simulations," *J. Rheol.* **65**(2), 273–288 (2021).