

Smoothed particle hydrodynamics techniques for the solution of kinetic theory problems

Part 1: Method

C.V. Chaubal^a, A. Srinivasan^b, Ö. Eğecioğlu^b, L.G. Leal^{a,*}

^a Department of Chemical Engineering, University of California, Santa Barbara, CA 93106, USA

^b Department of Computer Science, University of California, Santa Barbara, CA 93106, USA

Received 13 August 1996; received in revised form 1 November 1996

Abstract

The smoothed particle hydrodynamics (SPH) technique has been applied to a problem in kinetic theory, namely, the dynamics of liquid crystalline polymers (LCPs). It is a Lagrangian solution method developed for fluid flow calculations; its adaption to kinetic theory is outlined. The Lagrangian formulation of the Doi theory for LCPs is first described, and the problem is presented in the general framework of nonparametric density estimation. The implementation of the SPH technique in this specific problem is given, highlighting particular aspects of our implementation of SPH, including the form of the kernel function and use of an adaptive kernel. We then present results which demonstrate convergence and other details of the solution method, and also make comparisons with other solution techniques and discuss other potential applications. © 1997 Elsevier Science B.V.

Keywords: Kinetic theory; Lagrangian technique; Liquid crystalline polymers; Smoothed particle hydrodynamics

1. Introduction

1.1. Kinetic theory of complex fluids

‘Complex fluids’ is a term commonly used to describe a wide class of liquid-like materials, in which the relaxation toward the equilibrium state occurs sufficiently slowly that significant changes in the microstructural configuration, and thus in macroscopic properties, can be induced by flow or the action of other external fields of moderate strength. Polymer solutions and melts, colloidal and fiber suspensions, surfactant systems, and liquid crystals are all

* Corresponding author.

examples of complex fluids. There is much interest in systems of this kind, because it is hoped that by tailoring the microstructure of the fluid, one can produce desired macroscopic behavior. In addition, one would also like to know what specific macroscopic phenomena can be induced by external forces on the system, for example, magnetic fields and flow. For both of these ends, modeling is a crucial element.

A popular theoretical approach to describing the characteristics of complex fluids under flow is the kinetic theory pioneered by Kirkwood [1] and later championed by Bird and co-workers [2]. In this framework, a physical model is proposed which attempts to capture the essential features of the underlying microstructural component of the fluid without being too complicated. The resulting theory has two distinct parts: first is a set of equations, derived by application of classical kinetic theory to an ensemble of microstructural components, which govern the evolution of the microstructural configuration in the flow; second are the macroscopic equations of motion, with an equation to calculate the stress from the microstructural configuration space distribution function. From this coupled set of equations, one can make predictions of the results of various macroscopic experiments. These provide the crucial test of whether the proposed model is valid or not.

The Fokker–Planck formalism is a commonly used description of kinetic theory problems. At the heart of this approach is the Fokker–Planck equation, which governs the evolution of a configuration distribution function. This function gives the probability of finding the microstructural element of the model in a particular configuration. It is a scalar function of a multivariate space, the coordinates of which each represent a single configurational degree of freedom of the microstructural model. The distribution is influenced by energy potentials in the system, as well as dynamic and random Brownian forces. In the Fokker–Planck formalism, this random force tends to smooth out gradients in the distribution function, and leads to a diffusive term in the distribution function equation.

It is generally not possible to obtain exact analytic solutions of the Fokker–Planck equation, and this is especially so in the case of liquid crystalline polymers considered below, where the mean-field Maier-Saupe model for the nematic potential makes it nonlinear. There are two basic approaches that can be used to obtain model predictions. The first is to apply mathematical approximations guided by physical intuition to reduce the complexity of the problem. This includes constructing the asymptotic limit of the equations, or developing so-called closure approximations that allow moments of the distribution to be directly calculated numerically (or in some cases, analytically) without solving for the distribution function itself. This technique is expected to lead to an approximate version of the theory which is manageable enough to provide useful and quantitative results for general cases; however, the danger is that the simplifications may alter or eliminate important physical effects of the model. The second alternative is to retain the full governing equations and attempt a direct numerical solution for the distribution function. Although this will provide accurate results, the computational effort required is often tremendous. This is especially troublesome because the simulation of actual inhomogeneous processing-type flows requires the determination of microstructural behaviour at a large number of material points. A desirable goal is then to find the most efficient numerical technique, in terms of time and memory requirements. With modern computers, one can also take advantage of processor architectures, such as parallel or vectorized computers, so that an efficient scheme is obtained.

1.2. Lagrangian solution techniques for the Fokker–Planck equation

In order to solve the Fokker–Planck equation numerically, it is helpful to take the viewpoint that it describes the convection and diffusion of probability through configuration space. Seen in this manner, one may adopt techniques used to solve conventional flow problems to this different class of problems. The main difference is that the convection and diffusion obey equations of motion that depend on the particular microstructure being described. Another important distinction is that the probability density can vary greatly throughout the configuration space; it can be thought of (in the framework of conventional flow problems) as being analogous to the density of a highly compressible fluid. A promising approach would thus seem to be to adopt numerical techniques that have been designed for simulating the flow of highly compressible fluids.

Lagrangian solution techniques are particularly well-suited for convective transport problems where material points undergo major motion with respect to each other and where the density of these points is expected to be highly non-uniform in a way that is a priori unknown. There have been prior attempts at applying Lagrangian techniques to the solution of configuration distribution functions [3,4]. However, as described later, this class of convective transport problems poses particular challenges that require a more sophisticated approach.

1.3. Other approaches

Another way of solving the Fokker–Planck equation is to expand the configuration distribution function into a series of orthogonal basis functions, and then numerically solve for the coefficients of the series. Truncation of the series at a finite cutoff is necessary, the hope being that the surviving terms accurately represent the true dynamics of the system. This approach has been used for the particular system, liquid crystalline polymers, that we study below [5,6]. The main disadvantage of this method as a general tool is that the procedure is difficult to apply to systems for which the governing equations are very complicated. In addition, the method is best suited for systems in which variations in the probability density are smooth. If the probability density varies greatly, many more terms in the series must be retained, and solving for the coefficients becomes very computationally-intensive.

An approach related to the Fokker–Planck formalism is the Langevin formulation. In this description, the motion of an individual microstructural element is described via a stochastic differential equation. Brownian forces are modeled by a random noise term which obeys prescribed statistics. Recent papers [7,8] have presented the Brownian dynamics implementation of the Langevin approach. An apparent disadvantage of this technique is that, due to its stochastic nature, a very large number of microstructural elements must be tracked in order to obtain statistically accurate answers, thus increasing the computational time and effort.

1.4. Smoothed particle hydrodynamics

Smoothed particle hydrodynamics (SPH) is a Lagrangian simulation technique first developed for astrophysical problems [9]. It has certain features which address the problems encountered by others. Its main advantages are its robustness, simplicity, and its relative accuracy for the

amount of computational effort. The current work represents the first time, to our knowledge, that SPH has been applied to kinetic theory. Its successful application for our particular model system indicates that it may be a useful tool for studying more general kinetic theory problems.

We consider a specific problem in Section 2, namely, a solution for nematic rigid-rod polymers. We first give an overview of the physical problem that we are simulating, and then present the governing equations. In Section 3 we present the general Lagrangian formulation of the problem and discuss various solution strategies. In Section 4, we describe the smoothed particle hydrodynamics technique. We discuss certain implementation issues in Section 5, and present some test results in Section 6.

2. Example problem: kinetic theory of nematic rigid-rod polymers

2.1. Problem overview

Liquid crystals are a sub-class of complex fluids in which there are a multiplicity of thermodynamics phases, each possessing different degrees of molecular ordering. Systems in the *nematic* liquid crystal phase are characterized by liquid-like translational disorder, but crystal-like rotational order. In general, local regions in the system possess a preferred direction of orientation, often denoted by a unit vector that is called the *director*, and the molecules in the region tend to align themselves in this direction. This phenomenon generally arises as a result of excluded volume and electrostatic interactions between the molecules which give rise to collective ordering. The degree to which molecules are mutually aligned is usually characterized by an *order parameter*, which is a scalar measure of this alignment. Nematic liquid crystalline polymer (LCP) systems consist of macromolecules which display this type of behavior.

The most successful kinetic theory model describing the behavior of LCPs is due to Doi [10]. The microstructural element of this model is a rigid axisymmetric rod of infinite aspect ratio; the forces acting on the rods are hydrodynamic, Brownian, and intermolecular forces that give rise to a nematic liquid crystalline phase. These latter forces are approximated by a mean-field potential, which substitutes multi-body interactions by an overall uniform field based on the average behavior of a single molecule. The theory is further simplified by assuming that the nematic potential depends only on the local orientation distribution, but is independent of spatial gradients in the orientation. Hence, the theory strictly applies only to homogeneous systems, in which the director and order parameter are the same throughout the sample, and the velocity field is also uniform. Such situations are encountered approximately, in viscometric flows, and the theory has found success in predicting rheological properties of certain model systems in shear flow; namely, those in which the polymer molecules actually behave as rigid rods in solutions.

2.2. Governing equations

The Doi theory leads to the Fokker–Planck equation which gives the time evolution of the orientation distribution function $\psi(\mathbf{u}, t)$, where \mathbf{u} is a unit vector representing a given orientation. $\psi(\mathbf{u})$ gives the probability density of rods being oriented in the direction \mathbf{u} . As mentioned

above, it is assumed in the present development that this probability is the same no matter where in the system one looks. The Fokker–Planck equation for the Doi model is often given in the following form:

$$\frac{\partial \psi}{\partial t} = -\frac{\partial}{\partial \mathbf{u}} \cdot [(\boldsymbol{\Omega} \cdot \mathbf{u} + \mathbf{E} \cdot \mathbf{u} - \mathbf{E} : \mathbf{u}\mathbf{u}\mathbf{u})\psi] + \frac{\partial}{\partial \mathbf{u}} \cdot D_r \left[\frac{\partial \psi}{\partial \mathbf{u}} + \psi \frac{\partial}{\partial \mathbf{u}} \left(\frac{U_{\text{MF}}}{kT} \right) \right], \quad (1)$$

where \mathbf{E} is the rate-of-strain tensor, $\boldsymbol{\Omega}$ is the vorticity tensor, and D_r is the rotational diffusivity. U_{MF} represents the nematic mean-field potential, which will be modelled here in the Maier-Saupe form,

$$U_{\text{MF}} = -\frac{3}{2} UkT \mathbf{u}\mathbf{u} : \mathbf{S}, \quad (2)$$

where \mathbf{S} is related to the second moment of the distribution function,

$$\mathbf{S} = \langle \mathbf{u}\mathbf{u} \rangle - \frac{1}{3} \delta,$$

with $\langle \mathbf{u}\mathbf{u} \rangle = \int \mathbf{u}\mathbf{u}\psi(\mathbf{u})d\mathbf{u}$. U is a dimensionless parameter which measures the strength of interaction between the rods. Brownian forces are accounted for via the diffusive term $\partial\psi/\partial\mathbf{u}$, which tends to smooth out gradient in ψ . The scalar order parameter S is defined by:

$$S \equiv \left(\frac{3}{2} \mathbf{S} : \mathbf{S} \right)^{1/2}.$$

3. Lagrangian formulation of equations

The problem to be solved is the time evolution of the orientation distribution function $\psi(\mathbf{u}, t)$ as governed by Eq. (1). As noted earlier, this equation is actually an example of a classical convection diffusion problem, except in this case the dependent variable is the probability density, whose domain is the surface of the unit sphere in configuration space. In order to see this interpretation, it is instructive to recast Eq. (1) into two coupled equations:

$$\frac{\partial \psi}{\partial t} + \nabla \cdot (\dot{\mathbf{u}}\psi) = 0, \quad (3)$$

$$\dot{\mathbf{u}} = \boldsymbol{\Omega} \cdot \mathbf{u} + \mathbf{E} \cdot \mathbf{u} - \mathbf{E} : \mathbf{u}\mathbf{u}\mathbf{u} - D_r \left[\frac{\partial U_{\text{MF}}}{\partial \mathbf{u}} + \frac{1}{\psi} \frac{\partial \psi}{\partial \mathbf{u}} \right], \quad (4)$$

The Eqs. (3) and (4) represent the conservation relation (3) and equation of motion (4) in the configuration space which describes the system. The orientation is specified by the unit vector \mathbf{u} and hence can be specified by a point on the surface of a unit sphere. The probability density function of interest here is the orientation of the polymer rods. Thus one can regard the system as analogous to a compressible fluid on the surface of a sphere, with the probability density used instead of the mass density. Compressibility, in this context, refers to the fact that the (probability) density is not uniform, and the motion of particles in the system is influenced by the probability density at their locations. The equation of motion (4) depends on the overall global behavior of the orientation distribution through the nematic potential. Additionally, since we consider the system to be spatially uniform, the velocity gradient of the fluid, given in terms of $\boldsymbol{\Omega}$ and \mathbf{E} , enters into Eq. (4) as a constant tensor. Other than this distinction, the system (3)

and (4) resembles a typical convection-diffusion problem. Therefore, we can use techniques employed in compressible flow simulations to obtain solutions.

Some difficulties with the traditional Eulerian schemes are well known. For example, let us consider solving Eq. (1) for the time evolution of ψ . In an Eulerian approach, we would maintain a mesh on the domain of our problem, which in this case would be the surface of the unit sphere, and calculate the evolution of the probability density at the mesh points. Gradients of ψ can then be determined via standard difference formulae. However, in the present class of problems, there are typically regions with steep density gradients, and this requires a very high degree of spatial resolution to minimize discretization error. Although one could, in principle, employ an adaptive mesh the positions of regions with steep gradients change with time, which makes it necessary to change the mesh with time, and the management of such adaptive meshes can get extremely complicated. In general, Eulerian methods are difficult to use when different resolution is required at different regions in space, especially when these regions change with time.

In a general Lagrangian framework, the distribution function is discretized in terms of the positions of a finite set of interpolation points on the surface of the unit sphere. Each point is tracked as it is convected by the equation of motion (4). The discrete version of this equation for the rod with orientation \mathbf{u}_i is given by

$$\dot{\mathbf{u}}_i = \boldsymbol{\Omega} \cdot \mathbf{u}_i + \mathbf{E} \cdot \mathbf{u}_i - \mathbf{E} : \mathbf{u}_i \mathbf{u}_i \mathbf{u}_i - D_r \left[\frac{\partial U_{MF}}{\partial \mathbf{u}} + \frac{1}{\psi} \frac{\partial \psi}{\partial \mathbf{u}} \right]_i \quad (5)$$

The positions of all the points collectively represent the distribution function ψ . Because of the term in square brackets, the motion of each individual point depends on the entire distribution. Eq. (5) must be solved in a manner which is consistent with the conservation relation (3). The main advantage of the Lagrangian formulation is that the density of interpolation points automatically increases in regions requiring high resolution (and vice versa).

Up until now, the discussion has been general to any Lagrangian method. The distinction between different techniques is the way in which ψ is represented, and consequently the way terms depending on ψ are calculated. In the next section, we will briefly describe various ways of representing a distribution from a discrete set of points, including the method which is ultimately invoked in SPH.

3.1. Non-parametric estimation of probability density functions

Non-parametric density estimation is concerned with the estimation of a probability density function, given discrete samples from the distribution. No assumption is made about the type of distribution from which the samples are taken, in contrast to parametric estimation in which the density is assumed to be of a given parametric form, and the parameters are estimated. The solution of Eq. (5) involves repeatedly estimating the term in square brackets as time marches forward, essentially a problem of non-parametric density estimations. The conservation relation (3) will in general always be satisfied if the method of density estimation conserves probability.

The problem of non-parametric density estimation has been well studied. Good descriptions of various methods and their analyses can be found in [11,12]. Results of the comparison of some widely used methods is also given in [13,14]. Among the various methods that have been

proposed for general non-parametric density estimation are histograms, kernel methods [15–17], orthogonal series methods [18,19], and the nearest neighbor method [11].

The histogram method is the simplest to implement. The domain is divided into a number of cells. The probability density function within a cell is taken to be proportional to the number of samples from the distribution in the cell, scaled by the size of the cell. In general, this method is very sensitive to the location and sizes of the cells. This causes it to require extremely small cells and hence large sample sizes to give accuracy comparable to other methods.

The nearest neighbor method is based on the idea that in regions of high probability density, a sample point can be expected to be closer to its k th nearest neighbor than in a low density region. So, this method estimates the true probability density at a sample point in terms of the distance from the k th nearest neighboring sample point. One of the drawbacks of this method is that it has discontinuous derivatives and so we may expect sensitivity to local noise. In the context of iterative calculations, the noise will arise from errors in the discretization, as well as round-off due to the finite precision of the arithmetic.

The orthogonal series method [18,19] approximates the probability density by a series expansion, and determines the coefficients from the locations of the sample points. One of the problems with this method is that the series approximation can give a negative probability density estimate in low density regions. This would lead to non-convergence for most iterative calculations. There are also results for some series estimators suggesting that stronger conditions on the true probability density are required for them to give the same optimal convergence rates as kernel estimators [8]. However, when taking into account the computational effort involved, these methods may yet prove competitive.

3.2. The kernel method

The kernel method for non-parametric density estimation is asymptotically more accurate than the histogram method as the number of particles increases, does not suffer from the drawbacks mentioned above for the nearest neighbor and orthogonal series methods, has been well studied theoretically, and so is very popular in applications. We shall discuss the kernel method in greater detail here, since this is the method used in SPH.

Let $\mathbf{u} \ominus \mathbf{s}$ denote the distance along the surface of the unit sphere, between points \mathbf{u} and \mathbf{s} . If the kernel method is applied to the problem of estimating the probability density on the surface of a sphere, the distribution function is approximated in the form;

$$\psi_n(\mathbf{u}) = \frac{1}{nA_h} \sum_i K\left(\frac{\mathbf{u} \ominus \mathbf{u}_i}{h}\right).$$

where \mathbf{u}_i are the positions of the n samples, K is a so-called *kernel function* that depends on the distance between sample points \mathbf{u}_i and the points of interest \mathbf{u} . The parameter h is known as the *window width*; it provides the length scale over which the kernel function varies. A_h normalizes the right-hand side to make it a probability density function, and ψ_n is an estimate of ψ using n sample points. We take K to be a symmetric bounded and non-negative function with $\int K(\mathbf{u} \ominus \mathbf{s}) d\mathbf{u} = A_h$ for any given points \mathbf{s} on the surface of the unit sphere. If the derivative of the probability density is also desired, then it is estimated by replacing the kernel by the derivative of the kernel in the above equation, e.g.

$$\frac{\partial \psi_n}{\partial \mathbf{u}} = \frac{1}{nA_h} \sum_i \frac{\partial}{\partial \mathbf{u}} K\left(\frac{\mathbf{u} \ominus \mathbf{u}_i}{h}\right).$$

One needs to consider the accuracy of the above results. There are, of course, a number of error measures. The most widely used is the mean integrated square error (MISE), defined as

$$\text{MISE} = \int (\psi_n(\mathbf{u}) - \psi(\mathbf{u}))^2 d\mathbf{u},$$

where the integration is performed over the entire domain. As the sample size approaches infinity, we need to ensure that the MISE approaches 0. This criterion for convergence of the estimate to the true probability density is satisfied under the following condition [11]:

$$h \rightarrow 0 \quad \text{and} \quad nh^d \rightarrow \infty \quad \text{as} \quad n \rightarrow \infty, \quad (6)$$

where d is the dimension of the domain. On the surface of the sphere, $d = 2$. Eq. (6) follows from the fact that the MISE is approximately bounded by

$$\alpha k_1 h^4 + \beta n^{-1} h^{-d}, \quad (7)$$

where α and β depend on the specific form of the kernel function, and k_1 depends on the true probability density. Intuitively one may consider the first term, the square of the bias, as being due to the non-local nature of the kernel, while the true probability density is a locally determined quantity. The second term, the variance, can be considered as due to a bad sample. If more sample points contribute to the density estimate at a given point, we can expect the error due to the variance to be lower. However, there is a trade off between these two terms. The bias gets smaller as the kernel function approaches a delta function, but this could cause the variance to increase if the number of sample points contributing to the estimate at a given position gets smaller. The convergence criterion states that the window width should get narrower as the sample size increases, but also gives a bound on how fast the window width can decrease. The equation for the bound on the MISE suggests a form for an ideal kernel function, referred to as the Epanechnikov kernel. However, it can be shown that most of the kernel functions that have been considered are close to this kernel in their effectiveness in estimating the probability density [11]. Therefore we may choose the form of the kernel function based on factors other than accuracy, such as the ease of evaluation. Nevertheless, we shall later discuss some constraints which arise because of the details of the application for which the kernel method is currently being used.

Convergence of the probability density estimate does not imply convergence of the estimate for the derivative of the probability density. Generally, convergence of the derivative requires stronger conditions. In the present case, convergence of the derivatives requires a slower rate of decrease of the window width than convergence of the estimate of the density [20,12,21].

4. Smoothed particle hydrodynamics

In the following section we show how the kernel method of non-parametric probability density estimation can be incorporated into a dynamical simulation. We then describe the smoothed particle hydrodynamic (SPH) technique in detail, and discuss how it compares or is related to more conventional methods.

4.1. Precursors of SPH

The basis of all Lagrangian schemes for solving Eq. (1) [or Eqs. (3) and (4)] is a set of interpolation points that move on the surface of the unit sphere according to Eq. (5). The distinction between these schemes is the method of interpolation used to estimate ψ from the spatial distribution of these points.

In mesh-based Lagrangian schemes, the interpolation points are viewed as defining a mesh on the sphere surface, which distorts as the points move, and the density estimation is accomplished via a variant of the nearest neighbor method. Specifically, we observe that for any given interpolation point, the neighbors on the mesh define a polygon. The probability density will be inversely related to the area of this polygon, since in high density regions we expect a large number of points, causing the neighboring points to be closer than in a region of lower density. Gradients of the probability density can be computed using a finite difference scheme on the mesh. The advantage of the Lagrangian mesh (as opposed to an Eulerian mesh) is that points naturally convect to regions with high density, thus providing high resolution in regions where we need a high resolution without the need for active mesh management. This method was utilized by Szeri and Leal to solve a problem related to the current one [4].

There is a difficulty, however, with mesh based Lagrangian schemes. As the configuration of the interpolation points evolves with time, the mesh can get extremely distorted and the grid can become meaningless since the points that are the neighbors of any point on the mesh may move apart, rendering the finite difference calculations inaccurate. Indeed, in simulations performed on the system of interest here using the method of Szeri and Leal mentioned above, this precise problems was noted. The conventional solution is mesh reconnection every few iterations. In regions of high distortion, however, this leads to a high computational overhead since the mesh is distorted very rapidly.

In order to overcome the difficulty with distortion in a mesh based Lagrangian scheme, a number of so-called ‘particle methods’ were invented for the analogous fluid flow problem [22–24]. Most of these methods shared a common feature [9], namely, the use of a combination of an Eulerian mesh, and a set of Lagrangian points, or ‘particles’.

The earliest and best known of these schemes is the Particle In Cell (PIC) method [22]. This method involves an interpolation from a Lagrangian set of particles to an Eulerian grid, followed by certain calculations on the Eulerian grid, and finally an interpolation back to the Lagrangian particles at the end of each iteration. The probability density estimation in the original PIC technique was carried out via a histogram method. Suppose the positions of the particle (or interpolation points) are known at a certain time step. The surface of the sphere is divided into a number of cells. The probability density at the center of each cell is computed by dividing the number of the particles in the cell by the area of the cell, and then normalizing the result with the total number of particles. The value of ψ at the cell center is then interpolated back to the Lagrangian particles. There are a number of ways the interpolation can be done. The simplest is just to assign the value of the probability density at the cell center to each particle in that cell. Gradients of the probability density at the center of each cell can be calculated via finite difference formulae using the values of the probability density at neighboring cells. These values can then be interpolated back to the particles, and then the new positions of the particles can be computed. The major disadvantages of this scheme are the extra memory

required for the Eulerian mesh, and the inaccuracy of the interpolation from the particles to the Eulerian mesh.

Smoothed particle hydrodynamics [9,25] can be considered a generalization of PIC, but involving the more sophisticated kernel method of probability density estimation, as discussed in Section 3.2. This technique avoids the two disadvantages of the PIC method mentioned above, at the cost of possible increase in the computational effort for the same number of particles being tracked. However, we can expect the SPH method to require fewer particles for the same accuracy as PIC, leading to a net savings in computational effort. This method will be discussed further in the next section.

4.2. Description of SPH

SPH associates an interpolating kernel function with each Lagrangian particle (or interpolation point). The averaging or ‘smoothing’ of the influence of each particle, inherent in the use of kernel functions for interpolation, is the source of the name SPH. As discussed earlier, and also shown in Fig. 1, the probability density at any point is computed by summing the contribution of the kernels of all of the interpolation points within a distance of $O(h)$ from that location.

Gradients of the probability density at specific positions are calculated by analytically differentiating the interpolating kernel, and summing the contributions from the same set of Lagrangian points. The conditions under which this estimate of the density gradient approaches the true density gradient were explained in Section 3.2. Since this particle method is meshless, it does not suffer from any problems due to mesh distortion, and thus avoids the computational effort of frequent mesh reconnection encountered in mesh-based Lagrangian techniques. The SPH algorithm is most often discussed in the context of astrodynamical calculations [26,25]. Here, we describe SPH in the context of our particular problem.

As just discussed, the kernel method involves an estimate of the probability density on the surface of the sphere in terms of a sum of contributions from the kernel functions associated with a set of N interpolation points or particles, i.e.




Fig. 1. This sketch demonstrates density estimation using kernels. The solid lines show the kernels at each of the sample points, marked by x 's. The dotted line is the density estimate obtained by summing the contributions of each of the kernels at each point.

$$\psi(\mathbf{u}) = \frac{1}{N} \sum_{i=1}^N W(\mathbf{u} \ominus \mathbf{u}_i; h), \quad (8)$$

where \mathbf{u}_i specifies the position of particle i , N is the total number of particles and $W(d; h)$ is the kernel function. The parameter h is called the smoothing length, and gives a measure of the ‘distance’ away from a particle’s center within which the particle contributes to the probability. The gradient of the probability density function is estimated as

$$\frac{\partial}{\partial \mathbf{u}} \psi(\mathbf{u}) = \sum_{i=1}^N \frac{\partial}{\partial \mathbf{u}} W(\mathbf{u} \ominus \mathbf{u}_i; h). \quad (9)$$

Finally, moments of the distribution are given as follows:

$$\langle \mathbf{u} \mathbf{u} \dots \rangle \equiv \int (\mathbf{u} \mathbf{u} \dots) \psi(\mathbf{u}) d\mathbf{u} = \sum_{i=1}^N \int W(\mathbf{u} \ominus \mathbf{u}_i; h) (\mathbf{u} \mathbf{u} \dots) d\mathbf{u}. \quad (10)$$

Note that the integral in the last term is a function of \mathbf{u}_i only, and the integration may therefore be performed ahead of time; the resulting algebraic expression would simply be evaluated for each term in the summation when needed.

The function W is a special form of the kernel function mentioned in Section 3.2, with $W(r; h) = K(r/h)/A_h$. As noted earlier there is a great amount of flexibility in the choice of this function. But it is important that it satisfy the requirement of a delta function in the limiting case, i.e.

$$\int W(\mathbf{u} \ominus \mathbf{u}'; h) d\mathbf{u} = 1; \quad \lim_{h \rightarrow 0} W(\mathbf{u} \ominus \mathbf{u}'; h) = \delta(\mathbf{u} \ominus \mathbf{u}'). \quad (11)$$

The integration is performed over the entire surface of the sphere.

In practice, it is convenient to choose W such that it has continuous derivatives, as this allows for gradients to be calculated smoothly. Although a Gaussian form for W has been found to give good results for some applications [9], a desirable property for the kernel to possess is that it have compact support, which implies that it becomes identically zero when its argument is sufficiently large. This allows the interpolation procedure to take advantage of the finite cutoff and hence save on calculation time through techniques such as range-searching, as will be explained later in Section 5.

4.3. Comparison of methods

One can study the relationship between PIC, SPH and finite difference methods based on the ideas of density estimation. PIC estimates density based on the histogram. SPH uses the kernel method for this purpose. A Lagrangian finite difference scheme would estimate the density based on the nearest neighbor method. Observing that such an estimate is the true value of the density at some point inside the polygon, and taking a Taylor’s series expansion for the true density about this point, it can be shown that the error in the estimate of the density is bounded by the order of the length of this polygon. If the number of particles is N , then this length is of the order of $1/\sqrt{N}$. If we use a finite difference scheme, with a fixed number of neighbors, to compute the derivatives, then we divide by a quantity that is of the same order of magnitude, leading to a bound on the error that suggests a high error in the estimation of the derivative.

This suggests sensitivity to noise, as noted earlier for the nearest neighbor approach. Increasing the number of neighbors in the calculation leads to an SPH-type computation. It should be noted that the kernel method currently requires a lot more computation than the finite difference method for the same number of points. The SPH method proves advantageous when this is offset by being able to use fewer points for the same accuracy, and the fact that we can dispense with remeshing. Presumably, more efficient algorithms for density estimation will ultimately widen the range of situations in which SPH is preferable.

It is also informative to compare the SPH method as applied to kinetic theory with other methods of simulating complex fluids. The orthogonal basis function expansion technique has been shown to provide good results for this particular problem [5,6]. However, as mentioned earlier, if the distribution function is highly peaked, then more terms in the expansion must be retained, and the efficiency of the method begins to suffer, since the complexity of the problem increases very rapidly as more terms are included. Brownian dynamics can be thought of as a related Lagrangian solution method, since a discrete set of microstructural elements or ‘particles’ are tracked, and they collectively describe the distribution function. These techniques, however, typically require vast numbers of particles, anywhere from 10 000 to 100 000 (e.g., [6,8]). By contrast, SPH requires far fewer points for an equivalent result. Whether this actually results in shorter computations would again depend on the development of efficient algorithms. An example is given in the test studies presented in Section 6.

The major bottleneck for this scheme is that we currently require up to $O(N^2)$ time for calculations involving N particles. This especially shows up in higher dimensions, where more particles are required for a given accuracy. This can be improved to a certain extent by the use of effective range searching algorithms which will be described in Section 5. However, in our application, we observed that in order to perform range searching, the number of operations involved is quadratic in the number of particles. Thus range-searching alone cannot make the algorithm faster than quadratic, and additional approaches need to be tried to speed up the computations. For example, Egecioglu and Srinivasan [27,28] have recently proposed an almost linear algorithm for density estimation, which is conceptually similar to some of the fast algorithms used in N -body calculations. Though the proportionality constant here is fairly large, this method has shown significant improvements in speed in certain situations over the kernel method even with relatively few points, and it is easy to parallelize. However, further studies need to be made to test this algorithm in dynamical situations, in contrast to its present use in density estimation computations. In any case, we expect that further attention to this problem can produce faster methods that will make SPH even more appealing.

5. Implementation issues

In this section, we discuss the details of implementing the SPH technique. Although some of the issues that arise are concerns which are specific to our particular problem, by studying qualitative trends in this case, we are able to offer some generalizations about performance of the method.

5.1. Kernel form

It was mentioned earlier that the form of the kernel function does not have much effect on the accuracy of density estimation. However, for dynamical situations, a more careful choice of the kernel function is needed, due to the requirement of numerical stability, i.e., we wish to avoid a situation where the numerical errors at each iteration keep adding up to produce a large error. If there is an error pushing the system in one direction, we would want the scheme to push the system in the opposite direction at the next iteration in order to bring it closer to the true state. It has, for example [29], been observed in some situations that clustering of particles occurs at distances of the order of the window width. We too observed this initially. This is due to the fact that most of the popular kernels have a continuous derivative which vanishes when the distance to the particle is zero. For example, we originally used the following spline-based kernel:

$$W(x; h) = \begin{cases} 0, & \frac{x}{h} \geq 2, \\ \frac{1}{4} \left(2 - \frac{x}{h}\right)^3, & 1 \leq \frac{x}{h} \leq 2, \\ 1 - \frac{3}{2} \left(\frac{x}{h}\right)^2 + \frac{3}{4} \left(\frac{x}{h}\right)^3, & 0 \leq \frac{x}{h} \leq 1 \end{cases} \quad (12)$$

A perturbation of a nearby point which causes it to be placed at a distance less than $2h/3$ from the sample point leads to a smaller magnitude of the gradient of the kernel. In our calculations, the Brownian force, which keeps particles from coming too close together, depends on the gradient of the density function (Eq. (4)). This force becomes weaker when the particle separation is smaller than this critical distance, due to the peculiarity of this type of kernel, rather than due to the physical situation. In order to avoid this problem in our calculations, we chose to use a kernel that did not have a continuous derivative at the origin, namely

$$W(x; h) = \begin{cases} 0, & \frac{x}{h} \geq 2, \\ 1 - \frac{3}{2h^3} \left(h^2x - \frac{hx^2}{2} + \frac{x^3}{12}\right) & 0 \leq \frac{x}{h} \leq 2 \end{cases} \quad (13)$$

Using this kernel, the magnitude of the derivative of the kernel with respect to distance does not decrease as the sample point is approached. Fig. 2 gives a graph of W vs. x for each of these kernels, as well as their derivatives as a function x . The derivative of the kernel defined in Eq. (13) is not defined at the origin. In our calculations, we defined the derivative as 0 at the origin. This causes a discontinuity in the derivative. However, this does not invalidate the convergence results for the derivative because the error is expressed in terms of integrals involving the derivative of W , and the discontinuity is only on a set of measure 0. Intuitively, one can justify this as follows: as the sample size increases, the number of points contributing significantly to




Fig. 2. Graphs of W and W' vs. x for the cubic spline kernel, Eq. (12) (left), and the kernel defined by Eq. (13) (right).

the estimation of the derivative at any given location keeps increasing, and so an error in the contribution made by a single sample (i.e., the sample whose coordinates are the same as that of the location where we are estimating the derivative of the density) becomes negligible. Clustering disappeared on the use of the kernel defined in Eq. (13).

5.2. Kernel width and convergence

There has been extensive discussion in the literature on choosing the proper window width. However, most of the discussion is in the context of data analysis and involves choosing the optimal window width for a given set of data. The situation is slightly different in iterative calculations.

It will be useful to compare our situation with traditional schemes in this regard. For example, while using finite differences, one might keep increasing discretization until previous results are sufficiently close to the current results. One assumes that these results have converged to the true value since we have convergence results which prove that as the mesh size decreases to zero, the solution approaches the true value. Similarly, in the density estimation case, we do not as much require the optimal window width for the given sample as much as being confident that if we increase the sample size and our results do not change much, then we have indeed converged to the actual solution.

We observe another analogy between a typical finite difference scheme and SPH. If we consider finite difference discretization of partial differential equations, then often we cannot vary the time step independent of the spatial resolution. Results will often converge only if a certain relationship is maintained between the time step size and the spatial resolution. We have an analogous situation with the kernel method. We cannot decrease the window width

arbitrarily. There is a dependence on the number of interpolation points given by Eq. (6). Convergence to the true density is guaranteed only if this relationship is maintained.

The solution of Eq. (5) involves time–stepping, with the spatial discretization handled via the SPH approach. Since the time stepping uses a conventional finite difference approach, we discuss only the errors in spatial discretization using the SPH approach, in which the errors in the estimation of the density and its gradient are the major factors.

Fig. 3 is a typical sketch of error vs. window width for different values of the sample size. One can view the convergence results in terms of this figure. We can carry out the following procedure for ensuring convergence of the estimate to the true value. We fix h and increase N . We are assured convergence for high N , though not necessarily to the correct value. Next, we choose a smaller value of h and again increase N until we observe convergence. We continue this process until the results for successive values of h give similar results for high N . This ensures convergence. The alternate is to choose h as a suitable function of N and follow the dashed line. When the results do not change much, we have converged to the true solution. Following this procedure, we can save on calculations. We give further details on this procedure below. We can make h a suitable function of N to satisfy the convergence results for density estimation. For example, an approximate error bound of the form

$$k_1 h^4 + k_2 n^{-1} h^{-2(2r+d+1)}, \quad (14)$$

can be obtained for the error in the r th derivative of the density, where d is the spatial dimension. This yields an optimal window width of the form $h_{\text{opt}} = k/N^{1/9}$ on the surface of a sphere for the first derivative, where $k = 5k_2/(4k_1)$ is a constant. The error bound under the MISE criterion with this variation of h with N is given by




Fig. 3. Sketch of curves for error vs. h . The curves that are lower are for higher values of N . The dotted lines show one possible approach to convergence. We keep h fixed and vary N until the results converge. We repeat this process until the limiting values for successive h 's converge. The other approach is to use the convergence results to simultaneously change h and N . This is shown by the dashed line.

$$\epsilon^2 = \frac{C}{N^{4/9}} \quad (15)$$

where ϵ is the square root of the M.I.S.E., and C is a constant. Provided the true error follows the relation given by the error bound for sufficiently large values of N , then simulations with N_1 and $N_2 = 2^{9/4}N_1$ particles, respectively, will have errors ϵ_1 and ϵ_2 , with $\sqrt{2}\epsilon_2 = \epsilon_1$. This gives $\epsilon_1 \approx 3.4(\epsilon_1 - \epsilon_2)$. Thus we can estimate the error ϵ_1 since the difference in the results for the two simulations is also the difference in their errors. We increase N and stop when our estimate of ϵ is sufficiently small. We also need to choose an appropriate value of k . Although we are guaranteed convergence as $N \rightarrow \infty$, if we choose a sub-optimal value of C , we would require a higher value of N to achieve the desired accuracy than if we had chosen the optimal value. We suggest the following procedure to help solve this problem. Often one knows analytical solutions for some simple cases. We can solve with different values of h with N for those flows and determine a range of k which appears to give a good variation of h with N . Then, in the actual calculations, we can use this value. For example, in the problem we have considered, we know the analytical steady state solution for elongational flow. We tried computations with different values of N and h and noted the error in the scalar order parameter S , which indicates how aligned the polymer molecules are. Using these values of error, we can find a value of k which gives a relationship between h and N such that the error is minimized. We then use this value of k in other calculations where we do not know the exact solution.

5.3. Symmetries of the problem

The domain of our problem is the surface of the unit sphere. Most of the convergence results are for Euclidean domains, though there has also been some work specifically for the sphere [30,31,27,28]. However, we can still use the types of kernels used in Euclidean space since probability density is a local phenomenon, and a sphere is locally Euclidean. In addition, since our phase space is closed upon itself, boundary conditions are not present in this problem. We store the positions of the particles as Cartesian vectors and perform all the calculations in Cartesian coordinates. The distance between the points is measured along the surface of the unit sphere.

There is an additional feature of this particular problem which we can use to our advantage. Because the microstructural elements which comprise this model, i.e., the rigid rods, are fore-aft symmetric, the orientation distribution function $\psi(\mathbf{u})$ is invariant with respect to the transformation $(\mathbf{u}) \rightarrow (-\mathbf{u})$. To make use of this constraint, we create ‘ghost particles’ with antiparallel coordinates for every particle we track. We employ these ghost particles in all density estimations, but we don’t need to keep track of their locations, since this is determined by their corresponding ‘real’ particle. This effectively doubles the resolution of the simulation, without increasing the computational complexity. In general, any additional known symmetries in the distribution function can be utilized by creating other corresponding ghost particles.

5.4. Adaptive rule

Many of the more recent SPH calculations make use of some form of an adaptive kernel. In this modification, each interpolation point has a different kernel width, which is in general

inversely related to the local density. This leads to better accuracy in high density regions. Different adaptive rules with varying degrees of sophistication have been explored [32,33], but there is so far no definitive prescription for how to make the kernel width adaptive.

In our work, we found an adaptive rule of the form $h_i \equiv h_i(\psi_i) = (C_1(\psi_i - C_2))^{-1/2}$ to be effective, where h_i is the kernel width for particle i and $\psi_i \equiv \psi(\mathbf{u}_i)$. The two parameters C_1 and C_2 are chosen such that the following two conditions are met:

$$h(\psi_{\text{iso}}) = H,$$

$$h(\psi_{\text{min}}) = h_{\text{max}}.$$

$\psi_{\text{iso}} \equiv 1/4\pi$ is the value of ψ everywhere for an isotropic distribution (when the orientation probability is uniform). $h_{\text{max}} \equiv \pi/4$ is the highest value that we allow h to obtain. This maximum separation is determined by dictating that the range of influence of a particle's kernel should not extend into the range of influence of that particle's corresponding 'ghost'. Since h is the half-width of the kernel, $h = \pi/4$ means that the kernel's range of influence is one hemisphere (on the unit sphere); the other hemisphere would effectively be covered by the corresponding ghost particle. Finally, ψ_{min} would be the value of ψ at a particle if no particle other than itself contributed to the density at that point; from Eq. (8), this value depends self-consistently on h , as well as the number of points N . This form was chosen because it provides variations of h with ψ that are rapid enough to significantly improve resolution, and yet slow enough to prevent large fluctuations in h which would prevent smooth interpolation. Setting the parameter H has the effect of determining an average value of h ; therefore, increasing or decreasing this value has the effect of biasing the adaptive rule towards larger or smaller kernel widths. In all our simulations, we use H as the parameter which describes the average width of the kernel function.

Generally, it has been suggested that $h(r_i) \propto \psi(r_i)^{-1/d}$ be used [34] in a d dimensional domain. While this is consistent with our formula since we are in a two dimensional domain, it has been suggested by others [35,11] that it may be better to use the first formula in any dimension, since the bias can be shown, in this case, to be of a smaller order than in the case of a fixed width kernel.

5.5. Sharpness of density function

For nematic systems, the existence of a preferred orientation of the microstructure gives rise to a sharp peak in the orientation density (in this case, two peaks, since the microstructure and hence its distribution is fore-aft symmetric), which may be further focussed (or defocussed) by the effect of the flow. The sharpness of the orientation probability density for nematic systems makes it necessary to take a larger sample size than for a uniform distribution. The error from the bias term depends on the second derivative of the probability density. Sharp probability density functions have high bounds on the second derivative. In order to reduce the bias, it is necessary to decrease h . In order to compensate for this in the variance term, it is necessary to increase the sample size. Thus, sharp distributions require more computational effort than more uniform distributions. Using an adaptive kernel mitigates the effect of the sharpness of the density to a certain extent, since it leads to a lower bias as mentioned earlier.

5.6. Moments of the distribution function

Eq. (10) gives the SPH prescription for calculating moments of the distribution function. In that equation, the integral in the last term may be analytically integrated ahead of time. In practice, it is convenient to perform this integration in a reference coordinate system, where the calculations are easier, and then rotate the answer back to the original coordinates. In the reference coordinates, the polar axis is aligned with the particle; this integration then depends only upon the kernel width of the particle. By contrast, the rotation then depends only upon the coordinates of the particle.

Other functionals of the distribution function may be performed in a similar manner. However, it might not be possible to perform the integration analytically, and then some sort of numerical fit would be necessary. Nevertheless, once this fit is determined, it is a simple matter of evaluating this algebraic expression for each term in the summation.

5.7. Range searching

Since the number of points that contribute significantly to the density estimate at any given point is much fewer than the total sample size, it is more efficient to come up with a search strategy that avoids determining the contributions of points with negligible contribution. We wish to determine the contribution of only those points within a certain cut-off distance. This is the problem of range searching. Several algorithms have been proposed in the computational geometry literature [36] for effective range searching. The more sophisticated algorithms are faster than the cruder methods only when the number of particles is much larger than we consider in our current application. Hence simpler algorithms can be used when the number of particles is reasonably small. In one such scheme, one keeps track of the set of neighbors that contributed in the previous time step; this list is continually updated by considering a particle's neighbor's neighbors as possible new neighbors. The drawback of this scheme is the increase in the amount of memory required. A popular strategy, which is the scheme we employed, is based on projecting the data onto a one dimensional space, sorting this, and then performing range searching. We project the particles from the surface of the sphere to a line segment. In our case, this line is the vector pointing along the preferred orientation, since most points will tend to be clustered around this direction. Let x_i be the coordinate of the projection of point X_i . By the triangle inequality, the distance between X_i and X_j is at least the distance between x_i and x_j . Thus while computing the density at point X_i , with cut-off distance h , we can ignore all points X_j such that $|x_i - x_j| > h$. We can determine the points with $|x_i - x_j| \leq h$ by sorting the projected coordinates of each point, and then searching for those locations that have projected coordinates within a distance h of x_i for each point X_i . It should be noted that in some computations, especially involving gravitational forces, a tree-code similar to the Barnes-Hut scheme has been tried, since it fits in with the rest of the computations.

6. Test results

In this section, we present numerical results for certain test cases. We consider the following cases: (i) uniaxial elongational flow, and (ii) simple shear flow, both with various values for the nematic potential strength. Uniaxial elongation flow is a special case where an exact steady state solution to Eq. (1) can be found analytically, allowing a direct comparison. In this case, the sharpness of the nematic orientation distribution is increased by the aligning tendency of the elongational flow, making this a difficult case for density estimation. Shear flow is normally difficult to compute with mesh based Lagrangian techniques due to mesh distortion. In addition, it proves to be the most difficult flow for mathematical approximations to accurately reproduce [4]. The SPH method, however, is able to provide accurate results with a reasonably small number of particles.

In all the tests reported here, we started with an initial distribution of particles placed randomly over the surface of the unit sphere, using a random-number generator to provide cartesian coordinates. Although a random distribution is not perfectly isotropic, there is no analytic expression for disturbing points isotropically over the unit sphere surface; nevertheless, the initial distribution is sufficiently close to isotropic for test purposes. In addition, the test results, unless otherwise noted, make use of the adaptive scheme described in Section 5, with the parameter H as specified. Although we use a simple first-order time-stepping scheme, other methods could equally well be applied. The nematic interaction parameter U is dimensionless; we can use the following nondimensionalization for the remaining parameters.

For elongational flow,

$$\mathbf{E} = e \begin{bmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 2 \end{bmatrix}; \quad \boldsymbol{\Omega} \equiv \mathbf{0}; \quad Pe \equiv \frac{e}{D_r}.$$

For shear flow,

$$\mathbf{E} = \dot{\gamma} \begin{bmatrix} 0 & \frac{1}{2} & 0 \\ \frac{1}{2} & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}; \quad \boldsymbol{\Omega} = \dot{\gamma} \begin{bmatrix} 0 & -\frac{1}{2} & 0 \\ \frac{1}{2} & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}; \quad Pe \equiv \frac{\dot{\gamma}}{D_r}.$$

Pe is the Peclet number, and gives the ratio of the magnitude of hydrodynamic to Brownian forces.

We first present results for elongational flow, with $U = 7.5$ and $Pe = 2$. Figs. 4 and 5 show the evolution of the scalar order parameter S (as defined in Section 2.2) with time for different values of the window width and two sample sizes of 129 and 289 Lagrangian interpolation points, respectively. The dotted line shows the exact steady state value. The other lines are the results of the numerical computation for different values of window width. From Fig. 4 we can see that even for 129 particles, we obtain results that are accurate to within around 1% of the exact solution. Fig. 5 shows the results with 289 particles. We can observe that, apart from getting more accurate results, the solution obtained is not as sensitive to the window width as with 129 particles. We also observe from Figs. 4 and 5, that more accurate results are obtained




Fig. 4. Variation of the order parameter S vs time for 129 particles in elongational flow. Parameters: $U = 7.5$ $Pe = 2$ $Dt = 0.03$. The solid line is the simulation result with $H = 0.6$. The dashed line is the simulation result with $H = 0.55$. The dashed-dot line is the simulation result with $H = 0.5$. The dotted line shows the exact solution obtained analytically.

for the higher H values. This is not a general trend. We observe this trend in these figures because the H values for the results presented are all smaller than the optimal H value for the particular value of N .

We also wish to consider the variation of the optimal window width with the sample size. Fig. 6 shows the variation of steady state results obtained for the scalar order parameter for different H and N . The dotted line is the exact solution. These curves clearly show that the range of H over which we get good results is wider when we take a larger sample. Using this figure, we also need a regression model to predict the variation of H_{opt} with N . Based upon a prescribed form,




Fig. 5. Variation of the order parameter S vs time for 289 particles in elongational flow. Parameters: $U = 7.5$ $Pe = 2$ $Dt = 0.03$. The solid line is the simulation result with $H = 0.6$. The dashed line is the simulation result with $H = 0.55$. The dashed-dot line is the simulation result with $H = 0.5$. The dotted line shows the exact solution obtained analytically.




Fig. 6. Variation of the order parameter S vs. H in elongation flow. Parameters: $U = 7.5$ $Pe = 2$ $Dt = 0.01$. The dotted line shows the exact solution obtained analytically. Note that the vertical scale is expanded.

$H_{\text{opt}} \approx k/N^c$ we obtained $H_{\text{opt}} \approx 0.734/N^{0.022}$. If we consider convergence results for the derivative of the density given by Eq. (14), we expect $H_{\text{opt}} = k/N^{1/9}$. (We consider the error in the derivative of the density because our equations have terms involving both the density, and its derivative. Since the estimate of the derivative has a larger error, this is the dominating error term). Here, k is higher for sharper distributions. Since H_{opt} does not vary as we would expect based on the asymptotic error bound, we expect that at least one of the terms in our error bound, the bias or the variance, is much looser than the other.

Fig. 7 compares the estimated orientation distribution function to the exact solution for $U = 8$ and $Pe = 0.2$. Because of the symmetry of the flow, the solution is axisymmetric and does not depend on the azimuthal angle. In addition, the distribution has fore-aft symmetry, so ψ on the interval $[\pi/2:\pi]$ is the mirror image of ψ on $[0:\pi/2]$. We can see that even for 289 particles, the computed solution is very close to the true solution. It should be noted that $\theta = 0$ is a degenerate point, due to the peculiarity of the coordinate system; therefore, any difference between the exact value of ψ and its approximated value is greatly exaggerated near this point. For n greater than 513 particles, the estimated solution is nearly indistinguishable from the exact solution on the scale of the graph.

We next consider some results for simple shear flow, with $U = 4$ and $Pe = 1$. Figs. 8 and 9 show the evolution of one of the components of the second moment of the orientation distribution function with time for different values of the window width and sample sizes. In this case, we started with H values which had given good results in elongational flow. Then we performed two sets of numerical experiments. In the first set, we took $H = 0.6$ and $H = 0.6/\sqrt{2} \approx 0.42$. We increased N until we obtained convergence for each of the cases. The error bound, (7), suggests that the error for the smaller H in this example, in the limit of large N , should be half of the error for the larger H . Thus, the results in Figs. 8 and 9 can give us an idea of the accuracy of the above procedure. We can expect the error in the steady state value for $H = 0.42$ to be within twice the difference between the values for this case and for $H = 0.6$. This is approximately an error of 3%. Note that since we consider the situation for large N , the

variance term in the error bound is negligible. Thus, we need to consider this bias term alone. We can see from Eq. (14) that the bias for the error in the derivative of the probability density is of the same form as that for the probability density, itself. It should also be noted that the term involving the bias in the MISE has the square of the bias in it. The bias is thus the square root of that term. Fig. 8 presents the results obtained through this procedure by changing N and H independently. We also present a curve (shown by the dotted line) for $H = 0.42$ and $N = 289$. The bias for this test is the same as for the test with $H = 0.42$ and $N = 1250$. The difference between the two curves is due to the difference in their variance terms. The figure shows that the variance term is small, and that even a relatively small number of particles provides good results. In the next set of experiments, presented in Fig. 9, we started with $H = 0.624$, $N = 129$, and took higher values of N , varying H according to the empirical formula above and observed the convergence trend of the curves. This demonstrates the procedure for obtaining convergence by simultaneously changing H and N described earlier.

6.1. Effectiveness of the adaptive kernel

Next, we show some results which demonstrate the effectiveness of using an adaptive kernel. Fig. 10 shows convergence of a test case using a fixed width kernel with $U = 8$ and $Pe = 21$. As can be seen, even with more than 1000 points, the results have still not converged for this value of $H = 0.15$. By contrast, in Fig. 11, which shows the same case except using the adaptive kernel with $H = 0.354$, the results with 801 points have nearly converged. This example illustrates that the use of an adaptive rule in SPH calculations can greatly improve convergence.




Fig. 7. Profile of the density distribution function ψ as a function of θ for uniaxial elongational flow. Because the flow is axisymmetric, the solution doesn't depend on the azimuthal angle. Parameters: $U = 8$ $Pe = 0.2$ $Dt = 0.025$. The dotted line is the simulation with $n = 289$, $H = 0.50$. The dashed line is the simulation with $n = 513$, $H = 0.50$. The solid line is the exact solution obtained analytically.




Fig. 8. Convergence of calculations for shear flow by separately reducing the bias and the variance. Parameters: $U = 4.0$ $Pe = 1$ $Dt = 0.05$. The solid line is the simulation result for 1250 particles, with $H = 0.42$. We obtained this curve after simulations which kept H fixed and increased N until the results converged, and thus the error is due to the bias alone. The dashed line is the simulation result for 801 particles, with $H = 0.6$. We obtained this curve too after simulations which kept H fixed and increased N until the results converged. The dotted line shows the results of simulation with 289 particles, $H = 0.42$. Since the curve is close to the solid line, this test demonstrates that we can achieve sufficient accuracy even with a small number of particles.

6.2. High Pe results

The results of our study indicate that good results can be obtained at low Pe using around 800 points. At higher Pe , the number of points required for convergence would in general be expected to be fewer, because the relative importance of the Brownian term, which involves a density estimation Eq. (5), is smaller compared to the terms involving the flow, which is a function only of position of a particle. Fig. 12 shows a simulation with $U = 8$ and $Pe = 100$ for




Fig. 9. Convergence of calculations for shear flow by simultaneously reducing the bias and the variance. Parameters: $U = 4.0$ $Pe = 1$ $Dt = 0.05$. The solid line is the simulation result for 801 particles, with $H = 0.6$. The dashed line is the simulation result for 513 particles, with $H = 0.605$. The dash-dot line is the simulation result for 289 particles, $H = 0.613$. The dotted line shows the results of simulation with 129 particles, $H = 0.624$.




Fig. 10. Convergence of calculations for shear flow using a fixed width kernel with $H = 0.15$. Parameters: $U = 8.0$ $Pe = 21$ $Dt = 0.025$. The solid line is the simulation result for 289 particles. The dotted line is the simulation result for 513 particles. The dashed line is the simulation result for 801 particles. The dashed-dot line shows the results of simulation with 1153 particles.

two values of H and several values of N . The two values of H , 0.354 and $0.354/\sqrt{2} = 0.25$, were chosen because the error of the latter should be half that of the former, for sufficiently large N ; see Eq. (7). On the expanded vertical scale, it can be seen that for $H = 0.354$, $N = 1000$ is sufficient for convergence, while for $H = 0.25$, $N = 1500$ is sufficient. The difference in the value at top of the peak $\langle uu \rangle_{1,1}$ is approximately 2%, indicating that this is the maximum error in the more refined ($H = 0.354$) result. In both cases, however, $N = 500$ is more than sufficient to obtain good results; the difference between $N = 500$ and $N = 1500$ in the terminal value of $\langle uu \rangle_{1,1}$ for $H = 0.25$ is around 3%.




Fig. 11. Convergence of calculations for shear flow using an adaptive kernel with $H = 0.354$. Parameters: $U = 8.0$ $Pe = 21$ $Dt = 0.025$. The solid line is the simulation result for 289 particles. The dotted line is the simulation result for 513 particles. The dashed line is the simulation result for 801 particles. The dashed-dot line shows the results of simulation with 1153 particles.




Fig. 12. Convergence of calculations for high Pe shear flow. Parameters: $U = 8.0$ $Pe = 100$ $Dt = 0.025$. Solid line: $H = 0.354$, $N = 1000$. Dotted line: $H = 0.354$, $N = 500$. Short-dashed line: $H = 0.25$, $N = 500$. Long-dashed line: $H = 0.25$, $N = 1000$. Dot-dashed line: $H = 0.25$, $N = 1500$. Vertical scale expanded on the right side.

7. Discussion

7.1. Comparison with results from other methods

It is illuminating to compare our results with those from Brownian dynamics, which is a currently used method for solving kinetic theory problems that is most closely analogous to SPH, in that a set of discrete particles is tracked in a Lagrangian fashion and the distribution function is represented by their collective positions. In one set of simulation performed on a slightly different LCP model under similar conditions, Brownian Dynamics required 10 000 microstructural elements [6]. Despite the large number of particles, the figures of results that they presented had a noticeable amount of noise. Another simulation [8] performed on the same related system also showed noise in the results; the number of particles used in that study ranged from 2000 to 4000.

In contrast, the results we obtain using SPH show much less numerical noise on the same scale. In addition, even the results with fewer particles are noise-free. In fact, we observe that the qualitative nature of the results does not change when a relatively small number of points is used (see for example Figs. 8 and 9 and Fig. 11). This suggests that using SPH-type calculations with suboptimal numbers of particles might yield results that are still acceptable for making quantitative predictions. This would especially be useful for inhomogeneous flow simulation, in which the distribution function at many material points is calculated. Certainly, the results with a suboptimal number of particles are superior to all mathematical closure-based approximations which have been so far attempted for this particular class of problems. For example, the value

of the stress tensor calculated from simulations with fewer particles would probably not greatly alter the results if they were embedded in a full flow simulation.

The only situation in which having fewer than optimal particles would have a qualitative effect would be near very sensitive boundaries in phase space. In these regions, the dynamics of the system can be profoundly affected by the accuracy of the solution [5]. For these situations, having a converged result is crucial to tracking the correct solution branch, for example, and it would be necessary to have both a sufficiently small H and sufficiently many particles such that the error does not affect the qualitative dynamics. In these applications, the increase in accuracy with computational effort might not be rapid enough, and it would perhaps be best to study the system of interest with a more rigorous, analytic approach.

Hua and Schieber [8] mentioned that in their work the reason for the relatively small upper limit on particle number (compared with Brownian dynamics simulations of other polymer systems) is that this particular calculation is quadratic in the number of particles. This is because the rotational diffusivity D_r , which they employ is a function of ψ and hence must be calculated separately for each particle; in addition, they use a more complicated expression for the nematic potential U_{MF} (the Onsager form) which also depends on ψ , and not simply on $\langle \mathbf{uu} \rangle$. In calculations on simpler systems, Brownian dynamics can be performed in $O(N)$ time, permitting the use of many more particles to obtain desired accuracy. Although we use a simple (constant) rotational diffusivity and the simpler Maier-Saupe nematic potential, our calculations are already quadratic in the number of particles; using the more complicated expression for D_r would not increase the computational complexity of the method, and indeed would not even greatly affect the time per iteration. We would still be able to obtain accurate results for less than 1000 particles. It should be emphasized, however, that the total computational time for each method depends upon the actual set of calculations which are performed.

7.2. Application to other systems

The use of SPH to solve other kinetic theory problems is straightforward. The configuration space of another model would depend on the degrees of freedom of the microstructural element. One would generally have a multi-dimensional vector, the components of which span the configuration space. The equivalent conservation of probability Eq. (3) would be the same except for dimensionality and geometry of configuration space. The equation(s) of motion would depend upon the particular model, but any term which involves the distribution function ψ would be replaced by the equivalent kernel-based density estimation, i.e., the equivalent of (8).

The main issue which arises from the application of SPH to other models is the dimensionality of phase space. Most non-parametric density estimation techniques suffer loss of accuracy with increase in dimensionality. The kernel method is not different from other methods in this respect, and indeed, the similar challenge is posed for related techniques such as Brownian dynamics. More particles would be required to obtain results with the same relative accuracy in any method. At the same time, the flexibility of SPH allows one to adopt different solution strategies. For example, one may employ different kernel functions for each degree of freedom; if the dynamics of any particular direction in the configuration space are unique, the kernel for that direction could be appropriately modified. The algorithm could be made even more adaptive by having extra particles introduced where necessary and removing them when desired.

One may also make use of the more sophisticated algorithms for computing many-body interactions that were discussed in Section 4.3; as the number of particles becomes greater, these algorithms provide proportionally more time savings. In short, there are many variations on the SPH technique, and they have only begun to be explored.

8. Summary

The SPH method appears to be an effective and practical technique for solving Fokker–Planck equations which arise from kinetic theory. It is easy to adapt to a wide variety of systems, and provides robust and accurate solutions. In addition to the Doi model for nematic LCPs, we are also currently undertaking an investigation into a model which attempts to account for flexibility in rigid rod polymers. This is done by incorporating a flexible hinge in the middle of the rod. The kinetic theory equations which result from this model are more complicated, but application of the SPH technique is quite straightforward. We believe that this method could find application in a wide variety of Fokker–Planck type problems, especially complicated ones where other methods would become too cumbersome to implement.

Acknowledgements

We would like to acknowledge Bret Coldren for participation in the MRL-sponsored undergraduate project which gave rise to this work, and Andrew Szeri for providing a related Lagrangian code. CVC's contribution was supported by the MRL Program of the National Science Foundation under Award No. DMR-9123048; AS's contribution was supported in part by NSF/NASA/ARPA Grant No. IRI94-11330 and by a fellowship from the Parsons Foundation.

Appendix A. Convergence of the SPH simulation

We wish to solve for the probability density distribution of orientations. We discretize the system by following the particles. In this section, we first give the general expression for interpolation as used in SPH. We then prove that if we follow the particles and satisfy the conditions for convergence of the derivatives of a probability density function, then our estimate of the distribution function will converge to the true distribution function. We prove this by showing that the differential equation satisfied by the estimate approaches the differential equation satisfied by the true distribution function.

A.1. General interpolation

We can interpolate any function $A(r)$ by

$$A_{in}(r) = \int A(s)W(r-s, h) ds = \int \frac{A(s)}{\psi(s)} \psi(s)W(r-s, h) ds, \quad (A1)$$

where ψ is a probability density function and W is some kernel function. Using the fact that [37]

$$E\left(\frac{1}{N} \sum_{i=1}^N f(x_i)\right) = \int f(x)\psi(x) dx,$$

where E denotes the expected value, f some function, and x_i samples from the probability density function ψ , the interpolated value A_{in} can be estimated as follows.

$$A_E(r) = \sum_i \frac{A(r_i)}{\psi_i} W(r-r_i, h). \quad (A2)$$

Note that A_E is an estimate of the interpolated value A_{in} . The derivative of A too can be estimated by differentiating the kernel in the above equation, as shown below.

$$\nabla A_s(r) = \sum_i \frac{A(r_i)}{\psi_i} \nabla W(r-r_i, h). \quad (A3)$$

A.2. Satisfaction of convergence criteria

The orientation distribution function is defined by

$$\frac{\partial \psi}{\partial t} + \nabla \cdot (\dot{\mathbf{u}}\psi) = 0.$$

Using the SPH technique, the estimate of the density, ψ_E is given by

$$\psi_E(\mathbf{u}) = \frac{1}{nA_h} \sum_i K\left(\frac{\mathbf{u} \ominus \mathbf{u}_i}{h}\right).$$

Using the chain rule, the derivative of the estimate of the density is obtained as follows

$$\frac{\partial \psi_E}{\partial t}(\mathbf{u}) = \frac{-1}{nA(h)} \sum_i \nabla K\left(\frac{\mathbf{u} \ominus \mathbf{u}_i}{h}\right) \cdot \dot{\mathbf{u}} = \frac{-1}{nA_h} \sum_i \nabla \frac{K\left(\frac{\mathbf{u} \ominus \mathbf{u}_i}{h}\right)}{\psi(\mathbf{u}_i)} \cdot \dot{\mathbf{u}}\psi(\mathbf{u}_i).$$

From Eq. (18), we can show that this approaches

$$-\nabla(\dot{\mathbf{u}}\psi(\mathbf{u})),$$

if we satisfy the conditions for convergence of the estimate of the derivative of a probability density function. The above equation is the same as the equation for the true distribution function. Hence, we have completed our proof.

References

- [1] J.G. Kirkwood, *Macromolecules*, Gordon and Breach, New York, 1967.
- [2] R.B. Bird, R.C. Armstrong and Ole Hassager, *Dynamics of Polymeric Liquids*, Vol. 2, Wiley, New York, 1987.
- [3] A.J. Szeri and L.G. Leal, A new computational method for the solution of flow problems of microstructured fluids. Part 1. Theory *J. Fluid. Mech.*, 242 (1992) 549–576.
- [4] A.J. Szeri and L.G. Leal, A new computational method for the solution of flow problems of microstructured fluids. Part 2: inhomogeneous shearing flow of a suspension, *J. Fluid Mech.*, 262 (1994) 171–204.
- [5] R.G. Larson, Arrested tumbling in shear flows of liquid crystalline polymers, *Macromolecules*, 23 (1990) 3983–3992.
- [6] R.G. Larson and H.C. Öttinger, Effect of molecular elasticity on out-of-plane orientations in shearing flows of liquid-crystalline polymers, *Macromolecules*, 24 (1991) 6270–6282.
- [7] R. Keunings, On the Peterlin approximation for finitely extensible dumbbells, *J. Non-Newtonian Fluid Mech.*, submitted.
- [8] C.C. Hua and J.D. Schieber, Application of kinetic theory models in spatiotemporal flows for polymer solutions, liquid crystals and polymer melts using the CONNFESSIT approach, *Chem. Eng. Sci.*, 51 (1995) 1473–1485.
- [9] R.A. Gingold and J.J. Monaghan, Kernel estimate as a basis for general particle methods in hydrodynamics, *J. Comput. Phys.*, 46 (1982) 429–453.
- [10] M. Doi and S.F. Edwards, *The Theory of Polymer Dynamics*, Oxford Press, New York, 1986.
- [11] B.W. Silverman, *Density estimation for statistics and data analysis*, Chapman and Hall, London, 1993.
- [12] E.A. Nadaraya, *Non-parametric estimation of probability densities and regression curves*, Mathematics and Applications (Soviet Series) 1, Kluwer, Boston, 1989.
- [13] J. Hwang, Non-parametric multivariate density estimation: a comparative study, *IEEE Trans. Signal Processing*, 42 (1994) 2795–2810.
- [14] R. Vio, G. Fasano, M. Lazzarin and O. Lessi, Probability density estimation in astronomy, *Astron. Astrophys.*, 289 (1994) 640–648.
- [15] E. Parzen, On estimation of a probability density function and mode, *Ann. Math. Stat.*, 33 (1962) 1065–1076.
- [16] P.J. Bickel and M. Rosenblatt, On some global measures of the deviations of density function estimates, *Ann. Stat.* (1973) 1071–1095.
- [17] G.S. Watson and M.R. Leadbetter, On the estimation of the probability density, I. *Ann. Math. Stat.*, 34 (1963) 480–491.
- [18] S.C. Schwartz, Estimation of probability density by an orthogonal series, *Ann. Math. Stat.*, 38 (1967) 1261–1265.
- [19] R. Kronmal and M. Tarter, The estimation of probability densities and cumulatives by Fourier Series Methods, *Am. Stat. Assoc. J.*, (1968) 925–952.
- [20] M.P. Wand and M.C. Jones, *Kernel Smoothing*, Chapman and Hall, London, 1995.
- [21] P.K. Bhattacharya, Estimation of a probability density function and its derivatives, *Sankhya Ser. A*, 29 (1967) 373–382.
- [22] F.H. Harlow, The particle-in-cell computing method for fluid dynamics, *Meth. Computat. Phys.*, 3 (1964) 319–343.
- [23] B.M. Marder, GAP–A PIC-type fluid code, *Math. Computat.*, 29 (1975) 434–446.
- [24] J.N. Leboeuf, T. Tajima and J.M. Dawson, A magnetohydrodynamic particle code for fluid simulation of plasmas, *J. Computat. Phys.*, 31 (1979) 379–408.
- [25] J.J. Monaghan, Smoothed particle hydrodynamics, *Annu. Rev. Astron. Astrophys.*, 30 (1992) 543–574.
- [26] J.J. Monaghan and J.C. Lattanzio, A refined particle method for astrophysical problems, *Astron. Astrophys.*, 149 (1985) 135–143.
- [27] Ö. Egecioglu and A. Srinivasan, Efficient Non-Parametric Estimation of Probability Density Functions, Tech. Rep. TRCS95-21, University of California at Santa Barbara, 1995.
- [28] Ö. Egecioglu and A. Srinivasan, Numerical non-parametric density estimation, *Communications in Numerical Methods in Engineering*, (in review).

- [29] J.W. Swegle, D.L. Hicks and S.W. Attaway, Smoothed particle hydrodynamics stability analysis, *J. Computat. Phys.*, 116 (1995) 123-134.
- [30] N.I. Fisher, T. Lewis and B.J.J. Embleton, *Statistical Analysis of Spherical Data*, Cambridge University Press, Cambridge, 1993.
- [31] P. Hall, G.S. Watson and J. Cabrera, Kernel density estimation with spherical data, *Biometrika*, 74 (1987) 751–762.
- [32] M. Steinmetz and E. Müller, On the capabilities and limits of smoothed particle hydrodynamics *Astron. Astrophys.*, 268 (1993) 391–410.
- [33] M. Takeda, S.M. Miyama and M. Sekiya, Numerical simulation of viscous flow by smoothed particle hydrodynamics, *Progr. Theoret. Phys.*, 92 (1994) 939–960.
- [34] J.J. Monaghan, Particle methods for hydrodynamics, *Computat. Phys. Rep.*, 3 (1985) 71–124.
- [35] I.S. Abramson, On bandwidth variation in kernel estimates—A square root law, *Ann. Stat.*, 10 (1982) 1217–1223.
- [36] F.P. Preparata and M.I. Shamos, *Computational Geometry*, Springer-Verlag, New York, 1985.
- [37] M.G. Bulmer, *Principles of Statistics*, Dover, New York, 1979.