

Large Colloids in Cholesteric Liquid Crystals

K. Stratford^{1,2} · A. Gray¹ · J. S. Lintuvuori^{3,4}

Received: 13 October 2015 / Accepted: 14 October 2015 / Published online: 18 November 2015
© The Author(s) 2015. This article is published with open access at Springerlink.com

Abstract We describe a coarse-grained Landau-de Gennes model of liquid crystals (LCs) including hydrodynamics based on the Beris–Edwards equations. The model is employed to study the impact of large colloids on the long range LC defect structure in the cholesteric LC blue phases. ‘Large’ here means that the particle size is comparable to the cholesteric pitch, the length scale on which the LC order undergoes a helical twist. We investigate the case of a single particle, with either normal or degenerate planar anchoring, placed initially in an equilibrium blue phase LC. It is found that in some cases, well defined steady disclination structure emerges at the particle surface, while in other cases no clear steady state is reached in the simulations, and disclination reorganisation appears to proliferate through the bulk LC. These systems are of potential interest in the context of using LCs to template self-assembly of colloid structure, e.g., for opto-electronic devices. Computationally, we demonstrate a parallel approach using mixed message-passing and threaded model on graphical processing units allows effective and efficient progress for this problem.

Keywords Liquid crystals · Colloids · Simulation

1 Introduction

Liquid crystals (LCs) are both interesting and useful complex fluids. Typically composed of rod-like or disc-like molecules of a few nanometres in size, long range orientational order

✉ K. Stratford
kevin@epcc.ed.ac.uk

¹ EPCC and SUPA School of Physics and Astronomy, The University of Edinburgh, Edinburgh EH9 3JZ, UK

² Present Address: The Pawsey Supercomputing Centre, 26 Dick Perry Avenue, Perth, WA, Australia

³ SUPA School of Physics and Astronomy, The University of Edinburgh, Edinburgh EH9 3JZ, UK

⁴ Present Address: Laboratoire de Physique des Solides, Université Paris-Sud, UMR 8502, 91405 Orsay, France

in LCs means that regular structure occurs at optical wavelengths of several hundreds of nanometres—the source of the useful optical properties seen in everyday electronic devices. In cholesteric LCs, the molecules are chiral, and order is characterised by a helical rotation in one dimension with a length scale referred to as the pitch. As a liquid, hydrodynamics can be important. At the microscopic level, shear flow tends to rotate the rigid LC molecules, while hydrodynamics also influences processes such as phase separation kinetics and is central in phenomena such as permeation flow (where flow is possible while leaving large scale orientational structure unremoved). There is also considerable interest in LC systems into which solid colloidal particles are introduced (such particles may range from nanometres to microns in diameter). These systems are not only worthy of study in themselves, but also have potential as a source of new optical properties or routes to self-assembly of colloidal structure with a template provided by underlying LC order. Again, hydrodynamics is likely to play an important role in the conformational rearrangements required to reach any useful assembly of colloidal structure.

Highly chiral LCs also exhibit equilibrium phases with three-dimensional periodicity which are known as the cholesteric blue phases (BPs, [36]). In these phases, a subtle interplay between different contributions to the energy mean that the LC would like to form so-called double-twist cylinders. However, these cylinders cannot be patched together in three-dimensional space without introducing loci of defects referred to as disclination lines. Two phases in which the disclinations exhibit cubic symmetry have been identified experimentally (BPI and BPII) while a third (BPIII) is thought to be amorphous [16]. Recently, there have been significant advances in the technology of blue phases where polymer-stabilised formulations allow BPs to be used in a practically useful temperature range [8, 19]. This holds out the promise of useful application in, e.g., photonics [5, 6]. More recently, there has also been interest in nanoparticle stabilisation of BPs [18, 38]. The particles used in these experiments are typically a few nanometres in size, and stabilisation is explained in terms of defect elimination (where particles sit on the disclination lines and so reduce defect energy locally). Experiments with larger particles (100 nm or more [10]) suggest this mechanism is less effective; here one might expect the larger particles to have a more wide-ranging influence on the disclination structure. It is this regime which is the subject of this work.

The range of length scales in this problem presents a formidable challenge to computational modelling. Resolution is required at the scale of the helical twist of the LC, the particles themselves, and the large scale LC/particle order. Atomistic or even coarse-grained molecular models, while able to provide intricate detail, are likely to be impracticable for such mesoscale simulations. In this context, a coarse-grained approach which represents LC orientational order is suitable for discrete simulation in fluid dynamics. Such an approach is provided by adopting a combination of the Landau-de Gennes picture [11] and dynamics described by the Beris–Edwards equation [3]. Here, a coarse-grained orientational order is represented on a regular lattice, and discretised dynamics is consistent with an evolution towards thermodynamic equilibrium in the absence of flow. In the presence of flow, full hydrodynamic coupling is possible via advection and rotation of the local LC order by the local flow or shear, and the addition of a stress in the Navier–Stokes equations representing the action of elastic forces on the local flow.

A number of authors have pioneered this type of coarse-grained computational approach in simple LC fluids [9, 14, 35]. The somewhat more complex problem of colloidal suspensions in LCs has also been addressed. Here, a computation must allow not only for hydrodynamic interactions between fluid and particles, but also a boundary condition which represents a preferred orientation of the LC order at the solid fluid interface. Simulation studies of particles in simple LCs have addressed topics such as self-assembly [27] and aggregation [2]

in nematic LCs; defect-bonded pairs [17], chain formation [23] and collective motions [21] in cholesteric LCs. Simulation of colloids in bulk BPs [28, 30] have been undertaken, along with work on colloids confined in thin layers [22, 29, 32]. However, these existing studies are mostly limited to the case where the colloidal particles are relatively small compared with the cholesteric pitch (with the notable exception of Ref. [29] in BPs). While this is a natural starting point, it also reflects the computational challenge of addressing the large particle problem. To meet this challenge, we have developed a mixed message passing and threaded implementation of the simulation code to allow sufficiently large systems to be addressed on parallel computers. Specifically, we use a parallel machine consisting of graphical processing units (GPUs).

The work therefore has two aims. First, we concentrate on simulation results for the behaviour of a single colloid in cubic BPI, and in particular the defect structure formed by particles which are relatively large compared to the cholesteric pitch (and hence the BP lattice constant). This regime has not been accessible to simulation before. Second, as the work presents a number of interesting computational issues, we describe briefly some of the details of the implementation, and give an overview of its performance for the LC-colloid problem.

The work is organised as follows. In the following section we provide a brief overview of the salient features of the computational approach. Section 4 examines results of defect structures in BPI for different colloid sizes and surface anchoring conditions, and also comments on the parallel GPU performance of the computation. A summary and some discussion of future directions is given in the final section.

2 Method

The coarse-grained LC order is represented as a symmetric traceless tensor $Q_{\alpha\beta}$ (Greek indices here denote coordinate directions, and repeated indices are assumed to be summed over). The use of a tensor order parameter rather than a vector director—a mean local orientation—allows treatment of disclinations (defect lines) where the order vanishes and the director is undefined, as well as admitting biaxiality. The basis of the model is the Landau de Gennes free energy [11, 36], the density of which is f , with

$$f(Q_{\alpha\beta}) = \frac{1}{2}A_0 \left(1 - \frac{1}{3}\gamma\right) Q_{\alpha\beta}^2 - \frac{1}{3}A_0\gamma Q_{\alpha\beta} Q_{\beta\gamma} Q_{\gamma\alpha} + \frac{1}{4}A_0\gamma \left(Q_{\alpha\beta}^2\right)^2 + \frac{1}{2}K_0(\partial_\beta Q_{\alpha\beta})^2 + \frac{1}{2}K_1(\varepsilon_{\alpha\gamma\delta}\partial_\gamma Q_{\delta\beta} + 2q_0 Q_{\alpha\beta})^2. \quad (1)$$

The constant A_0 sets the energy scale for the free energy; γ is a parameter related to the temperature and controls the proximity to the isotropic/nematic transition; K_0 and K_1 are elastic constants. The first three terms are a bulk contribution, while the second line represents the cost of bend, splay and twist elastic distortions; q_0 is related to the helical pitch $p = 2\pi/q_0$ in a cholesteric; $\varepsilon_{\alpha\gamma\delta}$ is the permutation tensor.

The time evolution of the order parameter $Q_{\alpha\beta}$ is described by the Beris–Edwards equation, which includes the effects of flow:

$$\partial_t Q_{\alpha\beta} + \partial_\gamma(u_\gamma Q_{\alpha\beta}) - S_{\alpha\beta} = -\Gamma H_{\alpha\beta}. \quad (2)$$

Here, the left-hand side contains terms which represent the material derivative for rod-like molecules: the tendency, the divergence of an advective flux involving the velocity field u_α , and a term $S_{\alpha\beta}$ which couples the order parameter to velocity gradients $\partial_\beta u_\alpha$ and represents

the rotation of molecules by shear flow [3]. The right hand side ensures the relaxation of the order parameter toward equilibrium (in the absence of flow) via the molecular field $H_{\alpha\beta}$ —the functional derivative of the free energy with respect to the order parameter. This relaxation takes place on a time scale related to a collective rotational diffusion constant Γ . Hydrodynamics is controlled by the Navier–Stokes equation:

$$\rho \partial_t u_\alpha + \rho u_\beta \partial_\beta u_\alpha = -p_0 \delta_{\alpha\beta} + \eta (\partial_\alpha u_\beta + \partial_\beta u_\alpha) + \partial_\beta P_{\alpha\beta} \tag{3}$$

where an additional stress $P_{\alpha\beta}$ is introduced to provide feedback of the elastic forces exerted by the LC molecules on the fluid itself. $P_{\alpha\beta}$ is again related to the free energy by a rather involved expression omitted here (see, e.g., [7]). The fluid viscosity is η and the isotropic pressure p_0 .

The introduction of a solid surface with normal anchoring is viewed in the same picture as the addition of a free energy (per unit area) associated with the surface. This is

$$f_s(Q_{\alpha\beta}) = \frac{1}{2} W_1 \left(Q_{\alpha\beta} - Q_{\alpha\beta}^0 \right)^2, \tag{4}$$

where $Q_{\alpha\beta}^0$ is the preferred order parameter tensor to the solid surface, and W_1 is a constant determining the strength of the anchoring. Equation 4 describes homoetropic or normal anchoring, where the preferred surface order is such that the local director is normal to the local surface. A somewhat more complicated expression is relevant to degenerate planar anchoring [13], where the preferred surface order favours a director in the tangent plane at the surface, which we write:

$$f_s(Q_{\alpha\beta}) = \frac{1}{2} W_1 \left(\tilde{Q}_{\alpha\beta} - \tilde{Q}_{\alpha\beta}^\perp \right) \left(\tilde{Q}_{\alpha\beta} - \tilde{Q}_{\alpha\beta}^\perp \right) + \frac{1}{2} W_2 \left(\tilde{Q}_{\alpha\beta} \tilde{Q}_{\alpha\beta} - S_0^2 \right)^2. \tag{5}$$

In this case, $\tilde{Q}_{\alpha\beta} = Q_{\alpha\beta} - \frac{1}{3} S_0 \delta_{\alpha\beta}$ and $\tilde{Q}_{\alpha\beta}^\perp$ is its projection on the local solid surface; S_0 is a characteristic magnitude of order.

The surface boundary condition for a colloid with unit outward normal at the surface n_α gives, for the homoetropic case:

$$\begin{aligned} \frac{1}{2} K_0 (n_\alpha \partial_\gamma Q_{\beta\gamma} + n_\beta \partial_\gamma Q_{\alpha\gamma}) + K_1 n_\gamma \partial_\gamma Q_{\alpha\beta} - \frac{1}{2} K_1 n_\gamma (\partial_\alpha Q_{\gamma\beta} + \partial_\beta Q_{\gamma\alpha}) \\ - K_1 q_0 n_\gamma (\epsilon_{\alpha\gamma\sigma} Q_{\sigma\beta} + \epsilon_{\beta\gamma\sigma} Q_{\sigma\alpha}) - W_1 (Q_{\alpha\beta} - Q_{\alpha\beta}^0) = 0. \end{aligned} \tag{6}$$

In practice, this provides a means to compute gradients of the order parameter $\partial_\gamma Q_{\alpha\beta}$ at fluid sites next to the boundary from the anchoring properties associated with a colloidal particle. Away from the boundary, gradients are computed using a standard finite difference stencil using only fluid information. In this work we consider only the case $W_1 = W_2 = W$ which, along with the assumption of a single elastic constant $K_0 = K_1 = K$, allows us to introduce the relevant dimensionless parameter for anchoring strength at the surface of a spherical particle of radius a : $w = Wa/K$. This measures the strength of the surface anchoring energy Wa^2 compared with the LC fluid elastic energy Ka .

2.1 Implementation

The numerical solution of the problem involves a number of coupled components. A uniform Cartesian discretisation is adopted, and relevant physical quantities stored at each discrete grid or lattice position. As $Q_{\alpha\beta}$ is a symmetric, traceless, tensor it is stored as five independent elements at each lattice position. The gradient tensor $\partial_\gamma Q_{\alpha\beta}$ is computed using a seven-point stencil in three dimensions allowing the free energy (Eq. 1) to be evaluated. The time evolution of the Beris–Edwards equation (2) is computed using a forward-in-time differencing scheme

together with knowledge of the current order parameter and its derivatives, and the velocity field and velocity gradients. The velocity field is provided by the Navier–Stokes equation (3), which here is solved using a standard lattice Boltzmann (LB; see, e.g., [33]) technique involving a D3Q19 model and a multiple relaxation time approach [1, 12]. The additional stress $P_{\alpha\beta}$ required in the Navier–Stokes equation is again computed from the order parameter and its derivatives, and the divergence of this quantity provides a body force density applied to the LB fluid at each lattice site.

Hydrodynamic interactions between the fluid and the colloidal particles are represented via the standard method of bounce-back on links [25] in the LB picture, while the LC surface anchoring condition (Eq. 6) is handled in the finite difference picture [20]. The anchoring properties enter into the gradient calculation near the surface. In practice, Eq. 6 (in the case of both normal and tangential anchoring) represents five equations in the five independent elements of $Q_{\alpha\beta}$ with an additional equation in Q_{zz} . At a flat interface, Eq. 6 then gives rise to a small linear system of six equations in six unknowns with a constraint of tracelessness implemented as a final stage. At interior edges and corners, the linear system is 12×12 or 18×18 , respectively. A no-normal flux condition is implemented at solid-fluid boundaries for the advection term in the Beris–Edwards equation, while the velocity gradient terms entering $S_{\alpha\beta}$ are approximated by assuming a solid body rotation of the colloid at the surface. The net elastic force on the colloid from the fluid is computed by integrating the stress $P_{\alpha\beta}$ over the discrete surface, giving rise to a body force on the particle.

2.2 Parallel Implementation

A common approach to parallel computation is domain decomposition and message passing. For lattice based algorithms this is relatively straightforward. Moving colloidal particles, while somewhat more complicated, also admit efficient implementations (e.g., [31]). One advantage of such a simple domain decomposition implementation is the requirement only for nearest-neighbour communication in both lattice and colloid quantities. However, retaining nearest-neighbour communication places a constraint on the size of the particle relative to the underlying parallel sub-domain (roughly, that the diameter should not exceed the sub-domain extent). When larger colloidal particles are required, this constraint limits opportunity for strong parallel scaling and the associated improvement in time-to-solution. (As an example, a good strong-scaling implementation may be efficient with sub-domain sizes as small as 8^3 lattice sites per MPI task, effectively limiting particle radii to a few lattice units.) A natural way around this constraint is to move to a mixed threaded/message passing approach to allow larger sub-domains. This is also in line with the current hardware trend towards many-core architecture.

A popular architecture is the GPU cluster using a mixed MPI and GPU threaded model, for which a number of implementations of LB hydrodynamics have been described in the literature [4, 24, 26, 34, 37]. Recently, we [15] described an implementation for colloidal suspensions which employed the GPU for lattice-based operations, but retained a small number of colloid-based operations on the CPU. We extend that approach here to the case of the LC-colloidal suspension, but restrict ourselves here to comments on the LC sector of the computation, and that involving the interaction between the LC and the colloidal particle.

CUDA kernels associated with the Beris–Edwards dynamics are written by assigning lattice sites to different threads in the case where there are no colloidal particles. In the presence of particles, a number of additional changes are required for an effective GPU implementation. First, while colloidal dynamics are retained on the CPU [15], particle information (position, radius, anchoring properties and so on) must be transferred to the GPU. This is

a relatively small overhead, and can be restricted to those time steps when the particle has changed discrete shape. Second, the calculation of the order parameter gradients must be written in such a way as to prevent undue pressure on registers on the GPU. The gradient computation is performed by inverting the (constant) matrix associated with each of these three different linear algebra problems arising from different geometries (flat surface, inside edge or corner) ahead of time on the CPU, and storing the inverse in constant memory on the GPU. The solution at different points around the particle can then be computed by a kernel which only constructs the necessary right hand side and multiplies by the appropriate inverse. Third, the elastic force on the particle is computed by integrating the stress over the surface of the particle. This involves a reduction operation over lattice sites at the surface of a given particle to obtain the net force on the particle. This is achieved partly in shared memory on the GPU, an partly on the CPU for particles spanning more than one GPU thread block, and finally via MPI for particles spanning sub-domains.

3 Results

3.1 Disclination Structure Around a Single Colloid

We now present a series of results which examine the defect structure formed at a particle surface when a single colloid is placed in a pre-equilibrated BPI disclination network. The LC free energy constants are chosen to be consistent with BPI (see, e.g., [32]): $A = 0.01$, $\gamma \approx 3.086$, and $K \approx 0.000706$. For these values the pitch $p = 32\sqrt{2}$ lattice units and the BPI lattice constant is 32 units providing a well-resolved disclination structure. (We note that a typical blue phase unit cell size in experiment is in the range 100–500 nm [36].) The fluid is initialised to rest, and the order parameter is initialised using a high chirality approximation appropriate for BPI [36]. The order parameter is allowed to relax to a numerically steady state with constant free energy before the colloid is inserted. This process requires 10,000–100,000 LB time steps depending on the system size.

Particle radii of $a = 2.3, 9.71, 15.25, 21.77$ and 30.25 lattice units are considered. This provides a range of a/p between 0.05–0.67 which lets us approach the “large” particle regime where the size of the particle is comparable to the pitch length. (Compared with the experimental unit cell size given above, the largest particle here has a radius corresponding to about 500 nm, which is also experimentally reasonable.) In each case the simulation is performed in a cubic box with side of length L with periodic boundary conditions. In all cases it is arranged that L/a be at least 10 to provide scope for significant rearrangement of the disclination network at the particle surface without interference between periodic images. The cubic cell L is always a whole number of BPI lattice constants; the smallest system is $L = 128$ and the largest system is $L = 320$ to accommodate $a = 30.25$.

In a selection of cases, a number of different initial positions of the colloidal particle relative to the equilibrium disclination structure were tried. It was found that the most favourable position for obtaining a steady state was that predicted by a simple geometrical argument which computes the total energy saving associated with the intersection between an ideal BPI disclination line structure and a sphere. This is consistent with the simple “defect elimination” argument set out above and also observed in previous simulation studies [29]. Simulations were run typically for 4 million LB time steps or until the free energy was observed to be numerically stable, whichever is the sooner. In some cases, particularly for the larger particles with larger surface anchoring energies, the simulation was observed to be unstable in its late

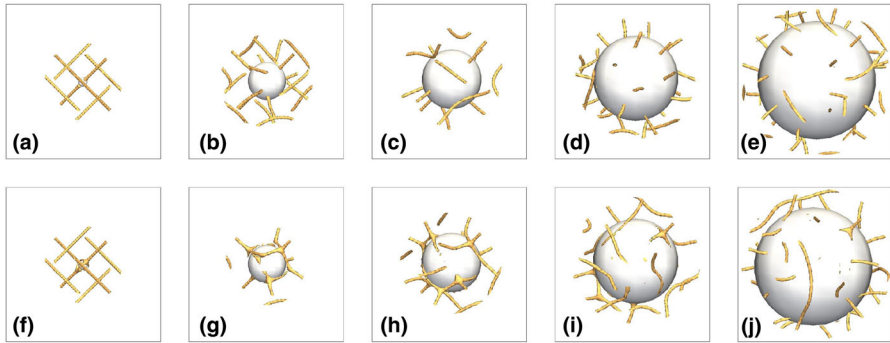


Fig. 1 Examples of disclination structures surrounding a normal anchoring colloidal particle with dimensionless anchoring strength $w = 0.2$ (top row) and $w = 2.0$ (bottom row). **a–j** show a particle of size $a = 2.3$, 9.71, 15.25, 21.77, and 30.25, respectively. The scale is the same in each case and the visible section has $l = 80$ lattice units. Only a limited spherical section of the disclination structure extending 10–12 lattice units beyond the particle radius is shown for clarity. For the smallest particle size (**a**, **f**) the equilibrium cubic BPI structure can still be seen clearly

stages. This may reflect the fact that the parameter set was chosen from experience with small particles; some further testing may be required for the larger ones used here for the first time. However, until shortly before instability occurs the simulation output (free energy, conserved quantities and so on) is well-behaved, and we take the results to be reliable. In some cases, again mainly for the larger particles, no steady state was achieved in a simulation of 4 million LB time steps. These unsteady cases are discussed further below. We note that with the current parameters, including a collective rotational diffusion constant $\Gamma = 0.5$, the physical time scale associated with 1 million LB steps is roughly 0.001 s. This can be estimated by assuming a typical value for A_0 of 10^6 Pa [36], and rotational diffusion constant of (equal to $2q^2/\Gamma$) of 0.1 Pa (see also [32]).

Calculations have been repeated for a range of different dimensionless anchoring strengths $w = Wa/K = 0, 0.2$, and 2.0 . For the case of $w = 0$ (not shown) the presence of the colloidal particle has relatively little impact on the surrounding disclination structure. For $w = 0.2$ and $w = 2.0$ the case for normal anchoring is shown in Fig. 1. These panels show the disclination structure, which is identified by plotting an isosurface of a local scalar order parameter; this is the largest eigenvalue of the diagonalised order parameter tensor $Q_{\alpha\beta}$. Low values of this quantity unambiguously determine vanishing order (locally a defect), and hence the disclination lines. The first result that may be seen clearly is that the size of the particle strongly influences the surrounding disclination structure. For the smallest particle, which is small compared with the BPI unit cell size ($a/p = 0.07$), the particles merely sits on a disclination line, again consistent with the defect elimination argument. The (equilibrium) cubic BPI disclination structure is otherwise unaffected. As the particle becomes larger, more severe rearrangements are observed. These include reconnection of pairs of disclination lines to avoid the particle (e.g., panel c), and more complex topological reconnections giving rise to 4-way junctions (reminiscent of the BPII structure [36]). In general, stronger anchoring appears to induce better defined topological reconnections in the simulations.

The analogous results for planar degenerate anchoring particles are shown in Fig. 2. Broadly, these show a similar tendency to increasing rearrangement for larger particles. There is some suggestion from the visualisations that the disclination lines are forced away from the particle to a greater extent by the planar anchoring boundary condition than in the

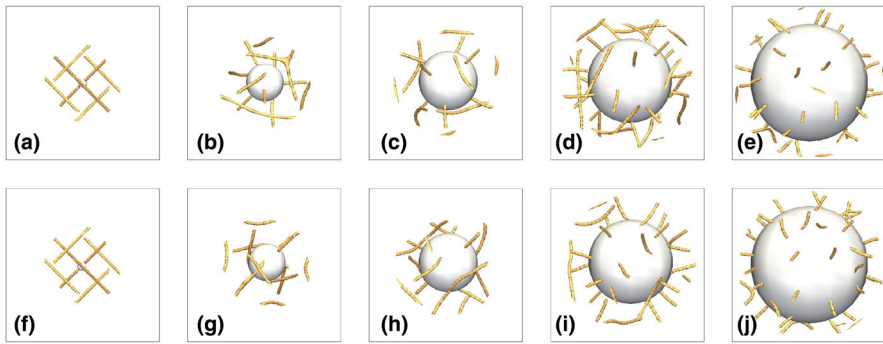


Fig. 2 As for Fig. 1, but with planar anchoring boundary conditions at the particle surface. The *top row* shows dimensionless anchoring strength $w = 0.2$ and the *bottom row* shows $w = 2.0$. The particle sizes are $a = 2.3, 9.71, 15.25, 21.77,$ and 30.25 lattice units in both (a–j), respectively




Fig. 3 A very strong anchoring regime with dimensionless anchoring strength $w = 20.0$ for normal anchoring surfaces. The topological reconnections are more clearly in evidence in this case. The particle sizes are **a–e** $a = 2.3, 9.71, 15.25, 21.77$ and 30.25 lattice units

homeotropic case. This is supported by examining the surface free energy in the simulation (f_s , not shown), which is uniformly lower in the planar case for a given w . It is noticeable for the larger particles in Fig. 2 that the disclinations do not appear to have reformed away from the surface but still intersect the surface.

This last observation may be related to the observation that the free energy in the simulation for larger planar anchoring particles ($a = 21.77$ and $a = 30.25$) does not seem to have converged before the onset of instability. What appears to be happening in the case of planar anchoring is that the topological rearrangements are occurring in discrete steps (again clear from a time series of the surface free energy). For the larger particles it appears these episodic rearrangements are frustrated, or simply destabilise the computation when they do take place. In contrast, the evolution of the free energy appears much smoother in the case of normal anchoring, which allows the simulation to proceed without instability. As an example, Fig. 3 shows the normal anchoring case when $w = 20.0$. (Here, we have also increased the value of Γ to 5.0 to allow the simulation to approach steady state more quickly in LB time steps.) Here, the topological reconnections in the vicinity of the surface are very clearly seen. This shows caging of the particles by the BP disclination lines similar to that observed in simulations of strongly confined BP-colloid composite systems [29]. However, it should be noted that this very strong value is likely to be significantly larger than values currently seen in experiment (typically $w \sim 0.1$, e.g., [38]).

The results here raise a number of questions. The first is that of the protocol: this is, at best, questionable from the experimental standpoint. There is no clear mechanism by which one can insert a colloidal particle into an equilibrium fluid with equilibrium disclination structure. This may mean the simulation initial condition is a deeply metastable state which

cannot easily relax. This problem may be particularly acute for particles with $a/p > 1/2$, and for planar anchoring where the simulation results show the near-surface disclination network appears largely unreconstructed (Fig. 2e, j). The second problem is the difficulty in approaching a steady state owing to proliferation of reconstructions in the disclination lines further away from the particle, which to some extent is a finite-size effect. To avoid these potential problems, future simulations will investigate a quench protocol (cf. [32]) in larger systems. This alternative protocol provides a role for hydrodynamics and flow to enter the problem to a greater extent, and could also help to bypass metastable states.

3.2 Comments on Performance

The simulations are performed on Titan at Oak Ridge Leadership Computing Facility, which provides NVIDIA K20X GPU hardware. One GPU has a peak performance of 1.31×10^{12} double precision floating point operations (flops) per second and a maximum main memory bandwidth of 250 GB/s (ratio: 5.24 flops per byte). The host hardware is one 16-core AMD Opteron processor with peak performance of 1.41×10^{11} flops per second and main memory bandwidth of 52 GB/s (ratio: 2.71 flops per byte). Charging policy is 16 host CPU cores plus 14 GPU streaming multiprocessor cores, i.e., 30 cores per node irrespective of actual core usage, meaning GPU use is cost-efficient.

It is instructive to compare the hardware capability with what is needed by the computation. For example, in the Beris–Edwards update used to compute the time evolution expressed by Eq. 2, the simulation requires at each lattice site $Q_{\alpha\beta}$, the gradient $\partial_\gamma Q_{\alpha\beta}$, the velocity field u_α , and the velocity gradient tensor $\partial_\beta u_\alpha$. This represents some 56 double precision reads per site and five double precision writes per site for the updated $Q_{\alpha\beta}$. By counting floating point operations in the code (found to be reliable when compared with profiler output for GPU kernels), the update takes around 1100 flops per lattice site. This gives a computation which involves around two flops per byte, so is largely memory bound for the GPU while roughly balanced on the CPU. This is fairly typical for a lattice based computation, which are often memory bandwidth limited. The situation for the computation of the divergence of the stress is also memory bound. However, here the main performance issue is likely to be the reduction required in the presence of particles to obtain the net body force on each particle. Finally, we note that for a system size $L = 128$, host execution using 16 MPI tasks is around 3–4 times slower than a single GPU. This represents a fair comparison on this system. The best kernel performance on the GPU—the LB collision stage here—achieves around 10 % of peak floating point performance and around 50 % of peak memory bandwidth. The overall efficiency of the code is less than this (and presents a continuing challenge for computations involving these complex systems).

3.2.1 Performance Breakdown

A series of actual performance figures is shown in Fig. 4. Figure 4a shows the breakdown of the computation as a function of the number of MPI tasks (GPUs) for a fixed problem size of $L = 128$. Three cases are considered. The first has no colloid and shows five main contributions to the computational time: the LB hydrodynamics, the order parameter gradient calculation, the calculation of the force via the divergence of $P_{\alpha\beta}$, the order parameter update, and total communication. (Note that there is a communication overhead even for one MPI task which counts the periodic boundary transfers.) The same problem on 8 MPI tasks shows a noticeable increase in the proportion of the time taken in communication overhead as the local domain size becomes smaller. The second case has a single colloid with radius $a = 2.3$,

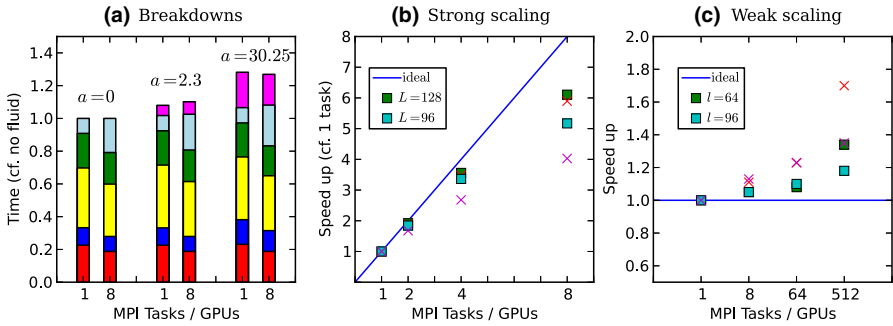


Fig. 4 Performance figures for a number of benchmark problems on Titan. **a** Shows breakdowns of the computation for 1 and 8 MPI tasks (with one GPU per MPI task) and with either no particle ($a = 0$) or one particle of radius a in a system of $L = 128$. The breakdowns show the main computational loads: from the bottom *red* is LB fluid, *dark blue* is computation of order parameter gradients, *yellow* is computation of fluid force from the divergence of $P_{\alpha\beta}$, *green* is the order parameter update, *light blue* is total communication overhead, and *pink* is particle computation. The times are normalised to the time for the fluid only computation ($a = 0$). **b** Shows strong scaling for total system $L = 128$ and $L = 96$; the *squares* are the fluid only problem, and *crosses* are for particles. **c** Shows weak scaling for local domain sizes $l = 64$ and $l = 96$; *squares* are the fluid only simulation and *crosses* with a single particle (Color figure online)

which might be considered as measuring the overhead incurred by having a particle at all. It can be seen that this overhead is significant at 5–10 % and reflects the additional complexity in the code required for colloid operations such as the reduction required to obtain the elastic force on the particle. However, the situation on moving to eight MPI tasks is not significantly different to the case of no particle. The third case shows a single colloid with radius $a = 30.25$, which represents a significant solid volume fraction in this system (around 5 %). Here the overhead for the particle computation is higher, with a noticeable growth in the proportion of the time in the gradient calculation; this reflects the increased number of surface elements requiring boundary condition computations.

The data for $L = 128$ are replotted in Fig. 4b which shows the strong scaling, along with data for a system with $L = 96$. These figures suggest that the limit of strong scaling for this problem is a local domain size of $l = 64$ in cases where there are no colloids. When particles are added, the situation deteriorates somewhat: this is attributed to the reduction required for the integration of the stress over the particle surface. The final panel in Fig. 4 shows weak scaling for various local domain sizes. We have found $l = 128$ to be effective for the largest simulations presented in this work.

4 Closing Remarks

We have reported simulations of large colloids in LCs facilitated by a new mixed MPI/GPU implementation of the computation. This has allowed us to study colloids of a size that would have otherwise been impractical. In the context of LCs, this opens a new regime of simulation where the colloids have a similar size to the cholesteric pitch and the related unit cell size in the cubic blue phases. The particle size exerts a strong influence on the topological rearrangements seen around both normal and planar anchoring surfaces. Forthcoming work will involve the role of hydrodynamics which plays little role in these particular computations. We hope this approach opens the way for further innovative simulation work in the future.

Acknowledgments KS acknowledges support from United Kingdom EPSRC Grants EP/I034661 and EP/J007404. AG is supported by the EU CRESTA Project (Seventh Framework Programme ICT-2011.9.13 Grant Agreement 287703) and the ARCHER UK national supercomputing service embedded computational science and engineering programme. JSL acknowledges support from UK EPSRC Grant EP/J007404 and an EU intra-European fellowship 623637 DyCoCoS FP7-PEOPLE-2013-IEF. An award of computer time was provided by the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) programme. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Adhikari, R., Stratford, K., Cates, M.E., Wagner, A.J.: Fluctuating lattice Boltzmann. *Europhys. Lett.* **71**, 473–479 (2005)
- Ario, T., Tanaka, H.: Colloidal aggregation in a nematic liquid crystal: topological arrest of particles by a single-stroke disclination line. *Phys. Rev. Lett.* **97**, 127801 (2006)
- Beris, A.N., Edwards, B.J.: *Thermodynamics of Flowing Systems with Internal Microstructure*. Oxford University Press, Oxford (1994)
- Bernaschi, M., Fatica, M., Melchionna, S., Succi, S., Kaxiras, E.: A flexible high-performance lattice Boltzmann GPU code for the simulations of fluid flow in complex geometries. *Concurr. Comput.* **22**, 1–14 (2010)
- Castles, F., et al.: Blue-phase templated fabrication of three-dimensional nanostructures for photonic applications. *Nat. Mat.* **11**, 599–603 (2012)
- Castles, F., Morris, S.M., Hung, J.M.C., Qasim, M.M., Wright, A.D., Nosheen, S., Choi, S.S., Outram, B.I., Elston, S.J., Burgess, C., Hill, L., Wilkinson, T.D., Coles, H.J.: Stretchable liquid-crystal blue-phase gels. *Nat. Mat.* **13**, 817–821 (2014)
- Cates, M.E., Henrich, O., Marenduzzo, D., Stratford, K.: Lattice Boltzmann simulations of liquid crystalline fluids: active gels and blue phases. *Soft Matter* **5**, 3791–3800 (2009)
- Coles, H.J., Pivnenko, M.N.: Liquid crystal 'blue phases' with a wide temperature range. *Nature* **436**, 997–1000 (2005)
- Denniston, C., Orlandini, E., Yeomans, J.M.: Lattice Boltzmann simulations of liquid crystal hydrodynamics. *Phys. Rev. E* **63**, 056702 (2001)
- Dierking, I., et al.: Stabilising liquid crystalline blue phases. *Soft Matter* **8**, 4355–4362 (2012)
- de Gennes, P.-G., Prost, J.: *The Physics of Liquid Crystals*. Clarendon Press, Oxford (1995)
- d'Humieres, D., Ginzburg, I., Krafczyk, M., Lallemand, P., Luo, L.-S.: Multiple-relaxation-time lattice Boltzmann models in three dimensions. *Phil. Trans. R. Soc. Lond. A* **360**, 437–451 (2002)
- Fournier, J.-B., Galatola, P.: Modelling planar degenerate wetting and anchoring in nematic liquid crystals. *Europhys. Lett.* **72**, 403–409 (2005)
- Fukuda, J.-I.: Effect of hydrodynamic flow on kinetics of nematic-isotropic transition in liquid crystals. *Eur. Phys. J. B* **1**, 173–177 (1998)
- Gray, A., Stratford, K.: Ludwig: multiple GPUs for a fluid lattice Boltzmann application. In: Couturier, R. (ed.) *Designing Scientific Applications on GPUs*. CRC Scientific Computing Series, Boca Raton (2013)
- Henrich, O., Stratford, K., Cates, M.E., Marenduzzo, D.: Structure of blue phase III of cholesteric liquid crystals. *Phys. Rev. Lett.* **106**, 107802 (2011)
- Jampani, V.S.R., Škarabot, M., Ravnik, M., Čopar, S., Žumer, S., Muševič, I.: Colloidal entanglement in highly twisted nematic colloids: Twisted loops, Hopf links, and trefoil knots. *Phys. Rev. E* **84**, 031703 (2011)
- Karatairi, E., et al.: Nanoparticle-induced widening of the temperature range of liquid-crystalline blue phases. *Phys. Rev. E* **81**, 041703 (2010)
- Kikuchi, H., Yokota, M., Hisakado, Y., Yang, H., Kajiyama, T.: Polymer-stabilized liquid crystal blue phases. *Nat. Mat.* **1**, 64–68 (2002)
- Lintuvuori, J.S., Marenduzzo, D., Stratford, K., Cates, M.E.: Colloids in liquid crystals: a lattice Boltzmann study. *J. Mater. Chem.* **20**, 10547–10552 (2010)

21. Lintuvuori, J.S., Stratford, K., Cates, M.E., Marenduzzo, D.: Self-assembly and nonlinear dynamics of dimeric colloidal rotors in cholesterics. *Phys. Rev. Lett.* **107**, 267802 (2011)
22. Lintuvuori, J.S., Pawsey, A.C., Stratford, K., Cates, M.E., Clegg, P.S.: Colloidal templating at a cholesteric-oil interface: assembly guided by an array of disclination lines. *Phys. Rev. Lett.* **110**, 187801 (2013)
23. MacKay, F.B., Denniston, C.: Modelling defect-bonded chains produced by colloidal particles in a cholesteric liquid crystal. *Europhys. Lett.* **94**, 66003 (2011)
24. Myre, J., Walsh, S.D.C., Lilja, D., Saar, M.O.: Performance analysis of single-phase, multiphase, and multicomponent lattice Boltzmann fluid flow simulations on GPU clusters. *Concurr. Comput.* **23**, 332–350 (2011)
25. Nguyen, N.-Q., Ladd, A.J.C.: Lubrication corrections for lattice Boltzmann simulations of particle suspensions. *Phys. Rev. E* **66**, 046708 (2002)
26. Obrecht, C., Kuznik, F., Tourancheau, B., Roux, J.-J.: Multi-GPU implementation of the lattice Boltzmann method. *Comput. Math. Appl.* **65**, 252–261 (2013). doi:[10.1016/j.camwa.2011.02.020](https://doi.org/10.1016/j.camwa.2011.02.020)
27. Ravnik, M., Škarabot, M., Žumer, S., Tkalec, U., Poberaj, I., Babic, D., Osterman, N., Muševič, I.: Entangled nematic colloidal dimers and wires. *Phys. Rev. Lett.* **99**, 247801 (2007)
28. Ravnik, M., Alexander, G.P., Yeomans, J.M., Žumer, S.: Mesoscopic modelling of colloids in chiral nematics. *Faraday Discuss.* **144**, 159–169 (2010)
29. Ravnik, M., Fukuda, J.-I., Yeomans, J.M., Žumer, S.: Confining blue phase colloids to thin layers. *Soft Matter* **7**, 10144–10150 (2011)
30. Ravnik, M., Alexander, G.P., Yeomans, J.M., Žumer, S.: Three-dimensional colloidal crystals in liquid crystalline blue phase. *Proc. Natl. Acad. Sci. USA* **108**, 5188 (2011)
31. Stratford, K., Pagonabarraga, I.: Parallel simulation of particle suspension's with the lattice Boltzmann method. *Comput. Math. Appl.* **55**, 1585–1593 (2008)
32. Stratford, K., Henrich, O., Lintuvuori, J.S., Cates, M.E., Marenduzzo, D.: Self-assembly of colloid-cholesteric composites provides a possible route to switchable optical materials. *Nat. Commun.* **5**, 3954 (2014)
33. Succi, S.: *The Lattice Boltzmann Equation for Fluid Dynamics and Beyond*. Oxford University Press, Oxford (2001)
34. Tölke, J.: Implementation of a lattice Boltzmann kernel using the Compute Unified Device Architecture developed by NVIDIA. *Comput. Vis. Sci.* **13**, 29–39 (2010). doi:[10.1007/s00791-008-0120-2](https://doi.org/10.1007/s00791-008-0120-2)
35. Tsuji, T., Rey, A.D.: Effect of long range order on sheared liquid crystalline materials Part I: compatibility between tumbling behaviour and fixed anchoring. *J. Non-Newton. Fluid Mech.* **73**, 127–152 (1997)
36. Wright, D.C., Mermin, N.D.: Crystalline liquids: the blue phases. *Rev. Mod. Phys.* **61**, 385–432 (1989)
37. Xian, W., Takayuki, A.: Multi-GPU performance of incompressible flow computation by lattice Boltzmann method on GPU cluster. *Parallel Comput.* **37**, 521–535 (2011)
38. Yoshida, H., et al.: Nanoparticle-stabilized cholesteric blue phases. *Appl. Phys. Express* **2**, 121501 (2009)