

A PRECONDITIONED NULLSPACE METHOD FOR LIQUID CRYSTAL DIRECTOR MODELING*

ALISON RAMAGE[†] AND EUGENE C. GARTLAND, JR.[‡]

Abstract. We present a preconditioned nullspace method for the numerical solution of large sparse linear systems that arise from discretizations of continuum models for the orientational properties of liquid crystals. The approach effectively deals with pointwise unit-vector constraints, which are prevalent in such models. The indefinite, saddle-point nature of such problems, which can arise from either or both of two sources (pointwise unit-vector constraints, coupled electric fields), is illustrated. Both analytical and numerical results are given for a model problem.

Key words. nullspace method, liquid crystals, saddle-point problems, unit-vector constraints

AMS subject classifications. 65F08, 65F10, 65F50, 65H10, 65N22

DOI. 10.1137/120870219

1. Introduction. Many continuum models for the orientational properties of liquid crystals involve one or more state variables that are vector fields of unit length. The pointwise unit-vector constraints associated with discretizations of such models give rise to indefinite linear systems of saddle-point form, when these constraints are imposed via Lagrange multipliers. In problems such as these, indefiniteness also frequently manifests itself due to another influence (coupling with applied electric fields), and this leads to a *double* saddle-point structure. We are interested in the efficient numerical solution by iterative methods of large sparse linear systems of algebraic equations associated with such problems. We begin by presenting some background on these materials and models.

1.1. Liquid crystals background. Liquid crystals are materials that are capable of existing in certain mesophases “in-between” isotropic liquids and solid crystals. They enjoy properties of both. They are fluid, yet their molecules (on average) possess long-range orientational order. A variety of liquid crystal phases (characterized by different ordering symmetries) are possible, and a number of organic compounds, mixtures, and assemblies exhibit one or more of these phases. Liquid crystals are important in technology because they can be used to control light: the orientational properties in liquid crystal films and cells can be manipulated by electric fields to give desired characteristics for the reflection and transmission of light waves. These effects have led to a number of important technologies, most in the area of display devices. Standard references include [4, 9, 27, 29].

In the area of liquid crystals, numerical modeling is used for several purposes, including optimizing the performance of devices, analyzing experiments, and validating theory. In the past, one-dimensional (1-D) and two-dimensional (2-D) modeling

*Submitted to the journal’s Computational Methods in Science and Engineering section March 15, 2012; accepted for publication (in revised form) October 19, 2012; published electronically February 12, 2013.

<http://www.siam.org/journals/sisc/35-1/87021.html>

[†]Department of Mathematics and Statistics, University of Strathclyde, Glasgow G1 1XH, Scotland (a.ramage@strath.ac.uk).

[‡]Department of Mathematical Sciences, Kent State University, Kent, OH 44242 (gartland@math.kent.edu). This author’s research was supported in part by National Science Foundation grants DMS-0608670 and DMS-1211597.

(one or two space dimensions) has often been adequate. This has been due to the simplicity of many of the geometries (thin films, cylindrical fibers, spherical droplets, etc.) and the fact that many three-dimensional (3-D) problems are symmetry reducible. Since many devices exist to control light, and since many experiments rely on microscopy of some form, the numerical modeling often involves two stages: first, characterizing the molecular orientational properties of the medium, and then second, determining the resulting optical behavior. Several modern devices now require bona fide 3-D modeling—this includes, for example, certain bistable devices and photorefractive cells [7, 18, 20]. Such problems pose greater challenges for numerical modeling, for characterizing both the orientational state as well as the optical properties. We are interested here in addressing the first issue, that is, in developing efficient numerical methods for computing the equilibrium orientational state of a liquid crystal in a realistic device or experiment with a bona fide 3-D structure. In a typical numerical exploration, one would need to solve such problems repeatedly in a parameter-continuation mode.

Materials in the liquid crystal phase are *anisotropic* with respect to essentially all of their physical properties—that is, conductivities, susceptibilities, and the like depend on direction, orientation—and this includes their optical properties. To characterize their optical behavior, one requires the average orientations of the molecular axes throughout the sample, which determine the dielectric tensor at each point. The most commonly used continuum models for equilibrium orientational properties of liquid crystals utilize as state variables one or more *unit-vector fields*. For the *uniaxial nematic* phase (the simplest and most common liquid crystal phase), this is the liquid crystal *director* (unit vector denoting the average orientation of the molecules in a fluid element at a point), traditionally denoted by \mathbf{n} .

1.2. Free energy models. For uniaxial nematic systems in equilibrium, the relevant thermodynamic potential, expressed as an integral functional of the state variables, is the *free energy* \mathcal{F} . For any given parameter values, the globally stable phase of the system is the stationary point of \mathcal{F} of least free energy. Several free-energy models are used; the oldest and still most common is the *Oseen–Frank elastic model*; see [27, Chap. 2] or [29, Chap. 3]—which has the general form

$$(1.1a) \quad \mathcal{F}[\mathbf{n}] = \int_{\Omega} W(\mathbf{n}, \nabla \mathbf{n}), \quad |\mathbf{n}| = 1,$$

with the *free-energy density* W given by

$$(1.1b) \quad 2W = K_1(\operatorname{div} \mathbf{n})^2 + K_2(\mathbf{n} \cdot \operatorname{curl} \mathbf{n})^2 + K_3|\mathbf{n} \times \operatorname{curl} \mathbf{n}|^2 \\ + (K_2 + K_4) [\operatorname{tr}(\nabla \mathbf{n})^2 - (\operatorname{div} \mathbf{n})^2].$$

Here Ω is the region occupied by the liquid crystal, \mathbf{n} is the liquid crystal director field (unit-length vector field on Ω), and K_1 , K_2 , K_3 , and K_4 are temperature-dependent material constants, which must obey certain inequalities to guarantee that W is positive definite; see [27, section 2.2.1] or [29, section 3.4]. A commonly encountered approximation to the above is the *equal-elastic-constant model*:

$$(1.2) \quad K_1 = K_2 = K_3 = K > 0, \quad K_4 = 0, \quad W = \frac{K}{2} |\nabla \mathbf{n}|^2, \quad |\nabla \mathbf{n}|^2 = \sum_{i,j=1}^3 \left(\frac{\partial n_i}{\partial x_j} \right)^2.$$

The model (1.1) is a phenomenological model for the distortional elasticity, constructed to be quadratic in $\nabla \mathbf{n}$ and to satisfy certain symmetry and invariance properties. It simply attaches energy costs to departures of the director field from local

parallel alignment (which is the preferred state of the system and is due to the nature of the intermolecular forces in such materials).

One of the main difficulties in dealing with such a model numerically is the *unit-vector constraint* on \mathbf{n} , which must be satisfied at each point in Ω . In the simplest cases, this can be managed by representing \mathbf{n} in terms of orientation angles, for example,

$$(1.3) \quad \mathbf{n} = \cos \theta \cos \phi \mathbf{e}_x + \cos \theta \sin \phi \mathbf{e}_y + \sin \theta \mathbf{e}_z,$$

which recasts the problem as an *unconstrained* problem for the scalar fields θ and ϕ . In particular, the equal-elastic-constant model (1.2) becomes

$$\mathcal{F}[\mathbf{n}] = \frac{K}{2} \int_{\Omega} |\nabla \mathbf{n}|^2 \quad \leftrightarrow \quad \mathcal{F}[\theta, \phi] = \frac{K}{2} \int_{\Omega} [|\nabla \theta|^2 + \cos^2 \theta |\nabla \phi|^2].$$

Difficulties with this approach result from the facts that first, such a representation degenerates when $\cos \theta = 0$ (and ϕ is undefined), and second, the angles suffer jumps across $\theta \pm 2\pi$, $\phi \pm 2\pi$. The formulation in terms of θ and ϕ is generally more cumbersome analytically (than the \mathbf{n} formulation), especially for the full free-energy density (1.1b); however, it can be used effectively in sufficiently simple settings. When an angle representation can't be employed, it is common to enforce the constraint $|\mathbf{n}| = 1$ via Lagrange multipliers. This provides the cleanest formulation and is robust with respect to the orientation of \mathbf{n} ; however, it changes a potential minimization problem into a *saddle-point problem*.

Several other liquid crystal models involve unit-length vector fields and constraints. These include certain alternate representations of Oseen–Frank in terms of dyadic tensors [2, 10, 21], generalizations of Oseen–Frank for nematics with non-uniform degree of orientational order [29, Chap. 6], several models of *smectic* liquid crystals (the second most common and important liquid crystal phase after the nematic) [27, 28], and others. All of these models acquire an additional indefinite, saddle-point structure when they are coupled with an electric field, as discussed in the next section. Unit-vector constraints arise in other areas as well, including the modeling of ferromagnetic materials—see, for example, [13, 19].

1.3. Coupled electric fields. Most devices and many experiments involve the interaction between a liquid crystal material and an applied electric field. The electric fields are usually created by sandwiching a liquid crystal film between electrodes to which a voltage is applied. In experiments, the electrodes are typically metals connected to controlled power supplies, while in portable devices, there are usually transparent conducting layers connected to small batteries. This is a coupled interaction, with the electric field influencing the orientations of the liquid crystal molecules and the molecular orientational properties in turn influencing the local electric fields through their effect on the dielectric tensor. Thus equilibrium orientation and electric field must be computed in a coupled, self-consistent way. Additional terms must be added to the free-energy densities, and this can be done for any liquid crystal model. We again illustrate the ideas in the context of the Oseen–Frank model.

The version of (1.1) for a coupled electric field (at constant potential) has the general form

$$\mathcal{F} = \int_{\Omega} \left[W(\mathbf{n}, \nabla \mathbf{n}) - \frac{1}{2} \mathbf{D} \cdot \mathbf{E} \right], \quad \mathbf{D} = \varepsilon \mathbf{E}.$$

Here W is the distortional elastic free-energy density (as before), \mathbf{D} is the electric displacement, \mathbf{E} is the local electric field, and $\boldsymbol{\varepsilon}$ is the dielectric tensor (a real, symmetric, positive definite tensor field). In a uniaxial medium (with a single optic axis \mathbf{n}), $\boldsymbol{\varepsilon}$ depends on the orientation of the director and takes the form

$$\boldsymbol{\varepsilon}(\mathbf{n}) = \varepsilon_0(\varepsilon_{\perp}\mathbf{I} + \varepsilon_a\mathbf{n} \otimes \mathbf{n}) \quad \Rightarrow \quad \mathbf{D} \cdot \mathbf{E} = \varepsilon_0[\varepsilon_{\perp}|\mathbf{E}|^2 + \varepsilon_a(\mathbf{E} \cdot \mathbf{n})^2].$$

Here ε_0 is the free-space dielectric constant, \mathbf{I} is the identity tensor, $\varepsilon_a = \varepsilon_{\parallel} - \varepsilon_{\perp}$ is the dielectric anisotropy (with ε_{\parallel} and ε_{\perp} the relative dielectric permittivities parallel to and perpendicular to the local director, respectively), and $\mathbf{n} \otimes \mathbf{n}$ denotes the tensor (outer) product ($[\mathbf{n} \otimes \mathbf{n}] = [n_i n_j]$ in Cartesian form). The parameters ε_0 , ε_{\parallel} , and ε_{\perp} are positive, while ε_a can be either positive (the more common case) or negative. In an equilibrium setting, the electric field can be represented by

$$\mathbf{E} = -\nabla U, \quad U = \text{electrostatic potential.}$$

Using the equal-elastic-constant approximation (1.2), we would have

$$\mathcal{F}[\mathbf{n}, U] = \int_{\Omega} W(\mathbf{n}, \nabla \mathbf{n}, \nabla U), \quad 2W = K|\nabla \mathbf{n}|^2 - \varepsilon_0[\varepsilon_{\perp}|\nabla U|^2 + \varepsilon_a(\nabla U \cdot \mathbf{n})^2].$$

It can be seen that minimizing this expression with respect to \mathbf{n} encourages the directors to align *parallel* to \mathbf{E} when $\varepsilon_a > 0$ (“positive dielectric anisotropy”) and *perpendicular* to \mathbf{E} when $\varepsilon_a < 0$ (“negative dielectric anisotropy”). The intrinsic saddle-point nature of the electric-field coupling is also clear: equilibria are minimizing with respect to \mathbf{n} but maximizing with respect to the scalar field U . These then are the computational obstacles that need to be addressed: equilibrium calculations involving director fields in *three space dimensions*, with *unit-vector constraints*, and *coupled electric fields*, in some combination.

The outline of the paper is as follows. In section 2, a model problem and its properties and discretization are described. Section 3 focuses on the Newton iteration used for solving the resulting nonlinear equations and highlights some properties of the associated Hessian matrix. The nullspace method for the inner Hessian system is discussed in section 4. Block preconditioning for the iterative solver is analyzed in section 5. Finally, the overall efficiency of the methods is demonstrated using numerical examples in section 6.

2. Model problem. In this paper, we will illustrate our ideas in the context of a model problem in one space dimension. While this particular problem can be solved quite simply, it embodies almost all of the features of the real target applications, is amenable to some analysis, and provides a satisfactory proof of concept for our approach. In this section, we introduce this problem and discuss the behavior of its solutions and their parameter dependence. We also derive a dimensionless form for the model free energy and an associated finite element discretization, which are used in what follows.

2.1. Twisted nematic device. A simple model problem that captures several of the essential features of interest is the *twisted nematic device* (or *twisted nematic cell*); see [27, section 3.7] or [29, section 4.5]. The basic system involves a thin liquid crystal film sandwiched between conducting and polarizing layers. The device switches between an “off state” (with no applied voltage), in which incident light is reflected, and an “on state” (applied voltage above a critical threshold), in which incident light

is absorbed. The actual optics of the effect are subtle and involve rotation of the polarization of the light through the film. A large number of liquid-crystal-based display devices operate on some variant of this simple idea.

This system can be modeled as one dimensional, with director and potential fields and free energy (in the equal-constant approximation, for simplicity) given, respectively, by

$$\mathbf{n} = u(z)\mathbf{e}_x + v(z)\mathbf{e}_y + w(z)\mathbf{e}_z, \quad U = U(z), \quad 0 < z < d, \quad d = \text{film thickness}$$

$$\mathcal{F}[u, v, w, U] = \frac{1}{2} \int_0^d \left[K(u_z^2 + v_z^2 + w_z^2) - \varepsilon_0(\varepsilon_\perp + \varepsilon_a w^2)U_z^2 \right] dz, \quad u_z = \frac{du}{dz}, \text{ etc.}$$

With λ the Lagrange-multiplier field for the constraint on \mathbf{n} , the first-order necessary conditions give the coupled, constrained equilibrium equations (in strong form)

$$(2.1a) \quad \begin{aligned} -Ku_{zz} &= \lambda u, & -Kv_{zz} &= \lambda v, & -Kw_{zz} - \varepsilon_0 \varepsilon_a U_z^2 w &= \lambda w \\ \frac{d}{dz} [(\varepsilon_\perp + \varepsilon_a w^2)U_z] &= 0, & u^2 + v^2 + w^2 &= 1, & 0 < z < d, \end{aligned}$$

with boundary conditions

$$(2.1b) \quad \begin{aligned} u(0) &= 1, & v(0) &= 0, & w(0) &= 0, & U(0) &= 0 \\ u(d) &= 0, & v(d) &= 1, & w(d) &= 0, & U(d) &= V, \end{aligned}$$

where V is the applied voltage. Nonlinearity manifests itself through the Lagrange-multiplier field $\lambda = \lambda(z)$ and the terms involving U_z^2 and w^2 , which also provide the coupling between \mathbf{E} and \mathbf{n} . In the more realistic case of unequal elastic constants, one would have more couplings among the fields.

The phase of the system is the equilibrium state of least free energy. For $V < V_c$ (critical switching voltage), the stable (“off state”) solution is a 90° helical twist of the director field accompanied by a uniform electric field, given exactly by

$$u(z) = \cos \frac{\pi z}{2d}, \quad v(z) = \sin \frac{\pi z}{2d}, \quad w(z) = 0, \quad U(z) = V \frac{z}{d}, \quad \mathbf{E} = -\frac{V}{d}\mathbf{e}_z.$$

It can be verified directly that u, v, w , and U above satisfy (2.1a) and (2.1b) with constant $\lambda = K\pi^2/(4d^2)$. For $V_c < V$, the pure-twist solution above remains an equilibrium solution, although it is now locally unstable. In this range, a tilted solution ($w \neq 0$), for which the electric field is no longer uniform and for which an analytical solution is only available in implicit form, has lower free energy; see [27, section 3.7]. For $V_c \ll V$, the director field essentially aligns with the electric field throughout most of the cell (except for narrow boundary layers at $z = 0$ and $z = d$), resulting in an “on-state” solution resembling $u, v \approx 0, w \approx 1$ or -1 ($\mathbf{n} \approx \mathbf{e}_z$ or $-\mathbf{e}_z$). There is a symmetry breaking, pitchfork bifurcation at $V = V_c$. Typical parameter values are

$$(2.2) \quad \begin{aligned} K &= 10^{-11} \text{ joules per meter}, & \varepsilon_0 &= 8.8542 \times 10^{-12} \text{ farads per meter}, \\ d &= 10 \text{ micrometers}, & \varepsilon_\parallel &= 15, & \varepsilon_\perp &= 5, & \varepsilon_a &= 10; \end{aligned}$$

see [27, Table D.3, p. 330], for which

$$V_c = \frac{\pi}{2} \sqrt{\frac{3K}{\varepsilon_0 \varepsilon_a}} \simeq 0.9143 \text{ volts.}$$

Figure 2.1 contains plots of the off-state and on-state solutions with the parameter values above for $V = 0.5$ volts and $V = 5.0$ volts expressed in terms of the orientation

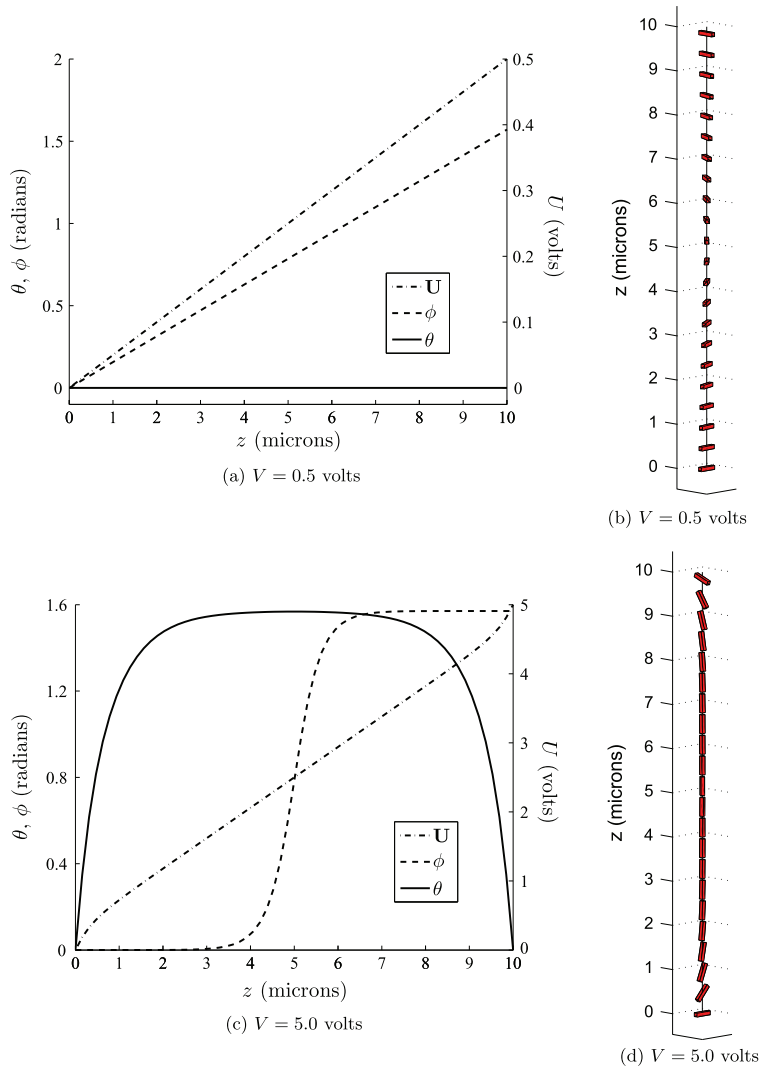


FIG. 2.1. Equilibrium “off-state” and “on state” solutions of (2.1) for $V = 0.5$ volts ((a) and (b)) and $V = 5.0$ volts ((c) and (d)). Parameter values are as given in (2.2). Plots (a) and (c) give the electric potential U and the director field expressed in terms of orientation angles (1.3): $\theta =$ “tilt angle,” $\phi =$ “twist angle.” Plots (b) and (d) render the director field \mathbf{n} as a “peg plot.”

angles θ (“tilt angle”) and ϕ (“twist angle”) from (1.3). Figure 2.2 shows the associated bifurcation diagram, in terms of maximum tilt angle vs applied voltage. As mentioned already, this model problem can be solved quite simply, yet it serves well for illustration and exploration, because it embodies several of the features of interest: pointwise constraints, Lagrange multipliers, and indefiniteness due to coupled electric fields.

2.2. Nondimensionalization. For convenience in terms of both analysis and numerical modeling, we nondimensionalize the problem using

$$\bar{z} := \frac{z}{d}, \quad \bar{U} := \frac{U}{V},$$




FIG. 2.2. Bifurcation diagram for equilibrium solutions of (2.1) with parameters as given in (2.2). Horizontal axis: applied voltage V . Vertical axis: $\theta_{max} = \pm \max_{0 \leq z \leq d} \theta(z) = \theta(d/2) =$ “maximum tilt angle.” Solid line: globally stable solution. Dashed line: locally unstable.

to obtain the nondimensionalized free energy

$$\bar{\mathcal{F}}[\bar{u}, \bar{v}, \bar{w}, \bar{U}] = \frac{1}{2} \int_0^1 [(\bar{u}_{\bar{z}}^2 + \bar{v}_{\bar{z}}^2 + \bar{w}_{\bar{z}}^2) - \alpha^2(\beta + \bar{w}^2)\bar{U}_{\bar{z}}^2] d\bar{z},$$

where $\bar{z} \in [0, 1]$,

$$\alpha^2 := \frac{\varepsilon_0 \varepsilon_a V^2}{K}, \quad \beta := \frac{\varepsilon_{\perp}}{\varepsilon_a}, \quad \bar{\mathcal{F}} := \frac{\mathcal{F}}{K/d},$$

and we have assumed $\varepsilon_a > 0, V \neq 0$. Here α and β are dimensionless parameters, and $\bar{\mathcal{F}}$ is the dimensionless free energy. The main control parameter is α (which is proportional to the magnitude of the applied voltage and is taken to be positive). Its critical switching value is given by

$$(2.3) \quad \alpha_c = \frac{\sqrt{3}\pi}{2} \simeq 2.721.$$

The value of β associated with the parameters in (2.2) is $\beta = \frac{1}{2}$. For the remainder of the paper, we will work with these nondimensionalized equations but for ease of notation will omit the overbars.

2.3. Discretization. Our model problem can be discretized in a variety of ways, starting from the strong or weak forms of the equilibrium equations or from the free-energy functional itself, and utilizing various types of finite elements or finite differences or other discretization methods. If we approximate \mathcal{F} directly via a uniform piecewise-linear finite element scheme with $N = n + 1$ cells using nodal quadrature, for example, we obtain

$$\mathcal{F}[u, v, w, U] \simeq f(u_1, \dots, u_n, v_1, \dots, v_n, w_1, \dots, w_n, U_1, \dots, U_n),$$

where

$$(2.4) \quad f = \frac{\Delta z}{2} \sum_{j=0}^{N-1} \left\{ \left(\frac{u_{j+1} - u_j}{\Delta z} \right)^2 + \left(\frac{v_{j+1} - v_j}{\Delta z} \right)^2 + \left(\frac{w_{j+1} - w_j}{\Delta z} \right)^2 - \alpha^2 \left[\beta + \left(\frac{w_j^2 + w_{j+1}^2}{2} \right) \right] \left(\frac{U_{j+1} - U_j}{\Delta z} \right)^2 \right\},$$

with

$$z_j = j\Delta z, \quad j = 0, \dots, N, \quad \Delta z = \frac{1}{N}, \quad n = N - 1, \quad u_j \simeq u(z_j), \text{ etc.}$$

The Dirichlet boundary conditions

$$u(0) = 1, \quad v(0) = 0, \quad w(0) = 0, \quad U(0) = 0, \quad u(1) = 0, \quad v(1) = 1, \quad w(1) = 0, \quad U(1) = 1$$

are used to eliminate u_0, v_0, w_0, U_0 and u_N, v_N, w_N, U_N . From now on, we will assume that where these values occur in given formulas, they are replaced by their appropriate boundary values.

To identify stationary points of f subject to the constraint that the directors have length one, we use Lagrange multipliers to impose the condition $u_j^2 + v_j^2 + w_j^2 = 1$ at each interior grid point $z_j, j = 1, \dots, n$. That is, we construct the *Lagrangian*

$$L = f + \frac{1}{2}\lambda_1(u_1^2 + v_1^2 + w_1^2 - 1) + \dots + \frac{1}{2}\lambda_n(u_n^2 + v_n^2 + w_n^2 - 1)$$

and find stationary points of L by solving the system of $5n$ nonlinear equations

$$(2.5) \quad \nabla L(\mathbf{x}) = \mathbf{0} \quad \text{for} \quad \mathbf{x} = \begin{bmatrix} \mathbf{n} \\ \boldsymbol{\lambda} \\ \mathbf{U} \end{bmatrix},$$

where

$$\boldsymbol{\lambda} = \begin{bmatrix} \lambda_1 \\ \vdots \\ \lambda_n \end{bmatrix}, \quad \mathbf{U} = \begin{bmatrix} U_1 \\ \vdots \\ U_n \end{bmatrix}, \quad \text{and} \quad \mathbf{n} = \begin{bmatrix} \mathbf{n}_1 \\ \vdots \\ \mathbf{n}_n \end{bmatrix}, \quad \mathbf{n}_j = \begin{bmatrix} u_j \\ v_j \\ w_j \end{bmatrix}, \quad j = 1, \dots, n.$$

3. Newton iteration. The equations in (2.5) can be solved using a Newton iteration, which at step k leads to a linear system of the form

$$(3.1) \quad \nabla^2 L(\mathbf{x}^{(k)}) \delta \mathbf{x}^{(k)} = -\nabla L(\mathbf{x}^{(k)}),$$

which must be solved for the Newton correction $\delta \mathbf{x}^{(k)}$. In the numerical optimization literature, this straightforward approach is sometimes referred to as the *Lagrange Newton method* [12, section 12.4]. The $5n \times 5n$ coefficient matrix and right-hand-side vector above are formed from the Hessian and the gradient of the Lagrangian L . In block form, they have the structure

$$(3.2) \quad \nabla^2 L = \begin{bmatrix} \nabla_{\mathbf{nn}}^2 L & \nabla_{\mathbf{n}\boldsymbol{\lambda}}^2 L & \nabla_{\mathbf{n}\mathbf{U}}^2 L \\ \nabla_{\boldsymbol{\lambda}\mathbf{n}}^2 L & \nabla_{\boldsymbol{\lambda}\boldsymbol{\lambda}}^2 L & \nabla_{\boldsymbol{\lambda}\mathbf{U}}^2 L \\ \nabla_{\mathbf{U}\mathbf{n}}^2 L & \nabla_{\mathbf{U}\boldsymbol{\lambda}}^2 L & \nabla_{\mathbf{U}\mathbf{U}}^2 L \end{bmatrix}, \quad \nabla L = \begin{bmatrix} \nabla_{\mathbf{n}} L \\ \nabla_{\boldsymbol{\lambda}} L \\ \nabla_{\mathbf{U}} L \end{bmatrix}.$$

In a typical numerical exploration, the large sparse linear system (3.1) must be solved repeatedly in parameter-continuation mode. Our main objective is to develop efficient iterative techniques to solve such systems.

3.1. Nondegeneracy condition. The nonlinear system (2.5) has the form of an equality constrained optimization problem:

$$f \text{ stationary, subject to } g_1 = \dots = g_n = 0,$$

where $f = f(\mathbf{n}, \mathbf{U})$ is the discretized free energy, as in (2.4), and

$$g_j(\mathbf{n}) = \frac{1}{2} (|\mathbf{n}_j|^2 - 1) = \frac{1}{2} (u_j^2 + v_j^2 + w_j^2 - 1).$$

Classical theorems concerning the existence of Lagrange multipliers require the *non-degeneracy condition* that $\nabla g_1, \dots, \nabla g_n$ are linearly independent at the critical point: see, for example, [12, section 9.1], [15, section 3.4.1], or [25, section 7.2]. For the simple constraints of our problem, this is equivalent to

$$\mathbf{n}_1, \dots, \mathbf{n}_n \neq \mathbf{0} \iff u_j^2 + v_j^2 + w_j^2 \neq 0, \quad j = 1, \dots, n.$$

This condition, which necessarily holds at any constrained stationary point (by virtue of $g_j = 0 \iff |\mathbf{n}_j| = 1$), guarantees the existence of the Lagrange multipliers. It also has implications, explored in what follows, for iterative solution algorithms.

3.2. Gradient components. The components of ∇L for our model problem are given by

$$\begin{aligned} \frac{\partial L}{\partial u_j} &= \frac{1}{\Delta z} (-u_{j-1} + 2u_j - u_{j+1}) + \lambda_j u_j, \\ \frac{\partial L}{\partial v_j} &= \frac{1}{\Delta z} (-v_{j-1} + 2v_j - v_{j+1}) + \lambda_j v_j, \\ \frac{\partial L}{\partial w_j} &= \frac{1}{\Delta z} (-w_{j-1} + 2w_j - w_{j+1}) + \lambda_j w_j \\ &\quad - \frac{\alpha^2}{2\Delta z} w_j [(U_j - U_{j-1})^2 + (U_{j+1} - U_j)^2], \\ \frac{\partial L}{\partial \lambda_j} &= \frac{1}{2} (u_j^2 + v_j^2 + w_j^2 - 1), \\ \frac{\partial L}{\partial U_j} &= a_{j+\frac{1}{2}} \frac{U_{j+1} - U_j}{\Delta z} - a_{j-\frac{1}{2}} \frac{U_j - U_{j-1}}{\Delta z}, \end{aligned} \tag{3.3}$$

where

$$a_{j-\frac{1}{2}} := \alpha^2 \left(\beta + \frac{w_{j-1}^2 + w_j^2}{2} \right), \quad a_{j+\frac{1}{2}} := \alpha^2 \left(\beta + \frac{w_j^2 + w_{j+1}^2}{2} \right).$$

Formulas (3.3) and (3.4) are valid for $j = 1, \dots, n$. The associated discrete constrained equilibrium equations $\nabla L = \mathbf{0}$ can be seen to be second-order-consistent approximations of the continuous constrained equilibrium equations

$$\begin{aligned} u_{zz} + \lambda u &= 0, \quad v_{zz} + \lambda v = 0, \quad w_{zz} + \alpha^2 U_z^2 w + \lambda w = 0, \\ \frac{d}{dz} [(\beta + w^2)U_z] &= 0, \quad u^2 + v^2 + w^2 = 1, \quad 0 < z < 1, \end{aligned}$$

which correspond to a nondimensionalized version of (2.1) and which have a differently normalized Lagrange-multiplier field than in (3.3) above.

3.3. Hessian components. In this section, we summarize the nature of the individual blocks of the Hessian (3.2), which are obtained by differentiations of the gradient formulas in (3.3). We will write the general block structure as

$$\nabla^2 L = \begin{bmatrix} \nabla_{\mathbf{nn}}^2 L & \nabla_{\mathbf{n}\lambda}^2 L & \nabla_{\mathbf{n}U}^2 L \\ \nabla_{\lambda\mathbf{n}}^2 L & \nabla_{\lambda\lambda}^2 L & \nabla_{\lambda U}^2 L \\ \nabla_{U\mathbf{n}}^2 L & \nabla_{U\lambda}^2 L & \nabla_{UU}^2 L \end{bmatrix} = \begin{bmatrix} A & B & D \\ B^T & O & O \\ D^T & O & -C \end{bmatrix}, \tag{3.5}$$

where O denotes a zero matrix of appropriate dimensions.

The Lagrange multipliers occur linearly in the A submatrix, which has the form

$$A = A_0 + \Lambda, \quad A_0 = \begin{bmatrix} A_{11}^0 & A_{12}^0 & & & \\ A_{21}^0 & \ddots & \ddots & & \\ & \ddots & \ddots & A_{n-1,n}^0 & \\ & & A_{n,n-1}^0 & A_{nn}^0 & \end{bmatrix}, \quad \Lambda = \begin{bmatrix} \Lambda_1 & & & & \\ & \ddots & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & \Lambda_n \end{bmatrix},$$

where

$$(3.6) \quad A_{jj}^0 = \frac{1}{\Delta z} \begin{bmatrix} 2 & & \\ & 2 & \\ & & 2 - \gamma_j \end{bmatrix} \quad \text{with } \gamma_j = \frac{\alpha^2}{2} [(U_j - U_{j-1})^2 + (U_{j+1} - U_j)^2],$$

and

$$A_{j,j-1}^0 = A_{j,j+1}^0 = -\frac{1}{\Delta z} \begin{bmatrix} 1 & & \\ & 1 & \\ & & 1 \end{bmatrix}, \quad \Lambda_j = \begin{bmatrix} \lambda_j & & \\ & \lambda_j & \\ & & \lambda_j \end{bmatrix}.$$

The A_0 matrix reflects the distortional stiffness (analogous to a discrete vector Laplacian, in finite element scaling) and the influence of the electric field on the director (here felt only by w), while the Lagrange multipliers embody the reactions to the constraints.

The (1,2) block of (3.5) is a $3n \times n$ matrix that has the typical structure of a ‘‘corner matrix’’ in an equality constrained problem, with columns made up of the gradients of the constraints (which here have a very simple form). It takes the form

$$(3.7) \quad B = \nabla_{\mathbf{n}\lambda}^2 L = [\nabla_{\mathbf{n}} g_1, \dots, \nabla_{\mathbf{n}} g_n] = \begin{bmatrix} \mathbf{n}_1 & & \\ & \ddots & \\ & & \mathbf{n}_n \end{bmatrix}.$$

We observe that, assuming nondegeneracy of the directors (section 3.1),

$$\text{rank}(B) = \text{rank}(B^T) = \text{rank}(B^T B) = \text{rank}(B B^T) = n.$$

Furthermore, $B^T B$ is a diagonal matrix such that, if the directors exactly satisfy the unit-vector constraints, then $B^T B = I_n$.

The (3,3) block is given by the $n \times n$ symmetric tridiagonal matrix

$$(3.8) \quad C = -\nabla_{\mathbf{U}\mathbf{U}}^2 L = \frac{1}{\Delta z} \begin{bmatrix} a_{1/2} + a_{3/2} & -a_{3/2} & & & \\ -a_{3/2} & \ddots & \ddots & & \\ & \ddots & \ddots & & \\ & & & \ddots & -a_{n-1/2} \\ & & -a_{n-1/2} & a_{n-1/2} + a_{n+1/2} & \end{bmatrix},$$

where $a_{j\pm 1/2}$ are as given in (3.4). For our model, $\beta > 0$, so if there is an applied voltage of any magnitude ($V \neq 0 \Rightarrow \alpha^2 > 0$), then $a_{j\pm 1/2} \geq \alpha^2 \beta > 0$, $j = 1, \dots, n$, and C is diagonally dominant with positive real diagonal entries and therefore positive definite. Note that the degenerate case of zero voltage, $V = 0$, implies $\mathbf{U} = \mathbf{0}$, in which case, these degrees of freedom should be removed from the model. In the case of the off-state solution ($w = 0$), we see from (3.4) that $a_{j\pm 1/2} = \text{const}$, so C is

proportional to a discrete Laplacian, as would also be the case if the medium were isotropic ($\varepsilon_{\parallel} = \varepsilon_{\perp} \Rightarrow \varepsilon_{\mathbf{a}} = 0$).

The (1,3) block is a $3n \times n$ matrix that takes the form

$$D = \begin{bmatrix} D_{11} & D_{12} & & & \\ D_{21} & \ddots & \ddots & & \\ & \ddots & \ddots & & \\ & & & D_{n-1,n} & \\ & & & D_{n,n-1} & D_{nn} \end{bmatrix},$$

where

$$D_{j,j-1} = \frac{\alpha^2}{\Delta z} \begin{bmatrix} 0 \\ 0 \\ w_j(U_j - U_{j-1}) \end{bmatrix}, \quad D_{j,j+1} = -\frac{\alpha^2}{\Delta z} \begin{bmatrix} 0 \\ 0 \\ w_j(U_{j+1} - U_j) \end{bmatrix},$$

and

$$D_{jj} = -D_{j,j-1} - D_{j,j+1} = \frac{\alpha^2}{\Delta z} \begin{bmatrix} 0 \\ 0 \\ w_j[(U_{j+1} - U_j) - (U_j - U_{j-1})] \end{bmatrix}.$$

The matrix D represents the main couplings between the electric field and the director field. The rank of D is at most $n - 1$, and is exactly that when $U_{j+1} - U_j \neq 0$, $j = 0, \dots, n$. This must be the case sufficiently close to the true solution of $\nabla_{\mathbf{U}}L = \mathbf{0}$, which satisfies

$$a_{j+1/2} \frac{U_{j+1} - U_j}{\Delta z} - a_{j-1/2} \frac{U_j - U_{j-1}}{\Delta z} = 0, \quad j = 1, \dots, n, \\ U_0 = 0, \quad U_{n+1} = 1,$$

with $a_{j\pm 1/2} \geq \alpha^2\beta > 0$, $j = 1, \dots, n$, implying $0 = U_0 < U_1 < \dots < U_n < U_{n+1} = 1$.

Using the above notation, the system for the Newton step (3.1) takes the form

$$(3.9) \quad \begin{bmatrix} A & B & D \\ B^T & O & O \\ D^T & O & -C \end{bmatrix} \begin{bmatrix} \delta \mathbf{n} \\ \delta \boldsymbol{\lambda} \\ \delta \mathbf{U} \end{bmatrix} = - \begin{bmatrix} \nabla_{\mathbf{n}}L \\ \nabla_{\boldsymbol{\lambda}}L \\ \nabla_{\mathbf{U}}L \end{bmatrix}.$$

At regular solution points (away from bifurcation points and turning points), the coefficient matrix is symmetric, nonsingular, and indefinite (that is, it has both positive and negative eigenvalues), and the linear system is in so-called *saddle-point form*. Note that when both pointwise unit-vector constraints and coupled electric fields are present (as above), these problems have a *double* saddle-point structure. For a model with no electric field or for one with no unit-vector constraints (for example, if an angle representation were used for \mathbf{n}), one would obtain instead either of the more common saddle-point forms

$$\begin{bmatrix} A & B \\ B^T & O \end{bmatrix} \quad \text{or} \quad \begin{bmatrix} A & D \\ D^T & -C \end{bmatrix}.$$

A detailed discussion of the nature of the double saddle-point system (3.9), including a derivation of appropriate stability characterizations and an alternative “renormalized Newton method” which takes greater advantage of the special structure of such problems, can be found in [14].

The form of the B matrix is exploited in what follows. Its simple structure is not unique to our setting (one space dimension, equal-elastic-constant free energy, low-order finite elements). More general models, in higher dimensions, discretized by various schemes, would involve B matrices of exactly the same form. All that is needed is that the discretized free energy contain degrees of freedom given by three-component directors ($\mathbf{n}_j \leftrightarrow (u_j, v_j, w_j)$) at nodes (or edges or faces or cell centers). Each such degree of freedom is subject to the constraint $|\mathbf{n}_j|^2 = u_j^2 + v_j^2 + w_j^2 = 1$, which gives rise to the simple structure of B . Higher order finite elements and such would additionally involve derivative degrees of freedom (u_x, u_y , etc.), but these would not be subjected to any a priori constraint. Spectral discretization methods, on the other hand, would incorporate degrees of freedom involving coefficients of polynomial or trigonometric-polynomial expansions, and would need to be handled differently.

4. Nullspace method. One popular method for solving saddle point problems is the *nullspace method* (or *reduced Hessian method*), which can be used to eliminate constraint blocks and reduce the size of the linear system—see, for example, [1, section 6]. In this section, we show that this method is particularly effective when applied to (3.9) with a view to eliminating the unit-vector constraints.

Recall that $B \in \mathbb{R}^{3n \times n}$ has full rank n provided the nondegeneracy condition (section 3.1) holds. We assume that there exists a matrix $Z \in \mathbb{R}^{3n \times 2n}$ with columns that form a basis for the nullspace of B^T —information on how Z can be constructed follows below. Using the notation of (3.9), the solution set of the second block equation ($B^T \delta \mathbf{n} = -\nabla_{\lambda} L$) can be written as

$$\delta \mathbf{n} = \widehat{\delta \mathbf{n}} + Z\mathbf{p}, \quad \mathbf{p} \in \mathbb{R}^{2n},$$

that is, as a linear combination of a particular solution $\widehat{\delta \mathbf{n}}$ and vectors that span the nullspace of B^T . Substituting this into (3.9) leads to the reduced $3n \times 3n$ system of equations

$$(4.1) \quad \begin{bmatrix} \widetilde{A} & \widetilde{D} \\ \widetilde{D}^T & -C \end{bmatrix} \begin{bmatrix} \mathbf{p} \\ \delta \mathbf{U} \end{bmatrix} = - \begin{bmatrix} Z^T (\nabla_{\mathbf{n}} L + A \widehat{\delta \mathbf{n}}) \\ \nabla_{\mathbf{U}} L + D^T \widehat{\delta \mathbf{n}} \end{bmatrix} \equiv \mathcal{H} \tilde{\mathbf{x}} = \tilde{\mathbf{b}},$$

where $\widetilde{A} = Z^T A Z$ and $\widetilde{D} = Z^T D$. The full solution to (3.9) can then be recovered using

$$(4.2) \quad \delta \lambda = -(B^T B)^{-1} B^T (\nabla_{\mathbf{n}} L + A \delta \mathbf{n} + D \delta \mathbf{U}).$$

Note that $(B^T B)^{-1} B^T$ above is the pseudoinverse of B . Note also that $B^T B$ is a diagonal matrix, so the inversion in (4.2) is no more than a vector division. The same observation applies to the construction of $\widehat{\delta \mathbf{n}}$, which can be found cheaply via

$$\widehat{\delta \mathbf{n}} = -B(B^T B)^{-1} \nabla_{\lambda} L.$$

This is not particular to our 1-D model problem: $\delta \lambda$ and $\widehat{\delta \mathbf{n}}$ can be computed directly and cheaply in this way for problems such as these in any number of space dimensions, as the structure of B remains unchanged.

The main perceived difficulty associated with nullspace methods is usually the computation of a well-conditioned null basis Z . For problems of the type we consider here, however, the structure of B^T (where B is as in (3.7)) makes this process

straightforward: one simply needs pairs of linearly independent 3-vectors orthogonal to each \mathbf{n}_j . For example, for a given $\mathbf{n} = [u, v, w]^T$, one could use any of

$$(4.3) \quad \begin{aligned} \mathbf{l} &= \frac{1}{\sqrt{v^2 + w^2}} \begin{bmatrix} 0 \\ -w \\ v \end{bmatrix}, & \mathbf{m} &= \frac{1}{\sqrt{v^2 + w^2}} \begin{bmatrix} v^2 + w^2 \\ -uv \\ -uw \end{bmatrix}; \\ \mathbf{l} &= \frac{1}{\sqrt{u^2 + w^2}} \begin{bmatrix} w \\ 0 \\ -u \end{bmatrix}, & \mathbf{m} &= \frac{1}{\sqrt{u^2 + w^2}} \begin{bmatrix} -vu \\ u^2 + w^2 \\ -vw \end{bmatrix}; \\ \mathbf{l} &= \frac{1}{\sqrt{u^2 + v^2}} \begin{bmatrix} -v \\ u \\ 0 \end{bmatrix}, & \mathbf{m} &= \frac{1}{\sqrt{u^2 + v^2}} \begin{bmatrix} -wu \\ -wv \\ u^2 + v^2 \end{bmatrix}. \end{aligned}$$

Each pair combines with \mathbf{n} to form an orthogonal right-handed triple $\mathbf{l}, \mathbf{m}, \mathbf{n}$ (with \mathbf{l} always normalized and \mathbf{m} normalized if and only if \mathbf{n} is). We use the first, second, or third formula pair depending on which of $|u_j|, |v_j|, |w_j|$ is minimal (which makes the corresponding square root maximal)—at least one of these radicals must be positive, by virtue of the nondegeneracy condition $u_j^2 + v_j^2 + w_j^2 \neq 0$. The columns of the matrix

$$(4.4) \quad Z = \begin{bmatrix} \mathbf{l}_1 & \mathbf{m}_1 & & & \\ & & \mathbf{l}_2 & \mathbf{m}_2 & \\ & & & & \ddots \\ & & & & & \mathbf{l}_n & \mathbf{m}_n \end{bmatrix}$$

therefore form a basis for the nullspace of B^T . Also, if the directors are normalized ($|\mathbf{n}_j| = 1$), then $Z^T Z = I_{2n}$, and $\mathbf{l}_j, \mathbf{m}_j, \mathbf{n}_j$ form an orthonormal right-handed triple.

Note that applying the nullspace method has resulted in a reduction in the number of unknowns from $5n$ to $3n$. We reiterate that the pointwise unit-vector constraints and their Newton corrections behave in exactly the same manner for more realistic problems in higher space dimensions.

5. Iterative solution of the reduced system. Even with the reduction in size of the Hessian due to the nullspace method, for realistic problems, \mathcal{H} will still be large and sparse. An efficient iterative solver for (4.1) is therefore still needed. Because \mathcal{H} is symmetric and indefinite, we use MINRES [24], which is known to be a robust and efficient method for indefinite problems. Typically, the number of MINRES iterations required to converge to a fixed tolerance will grow as the underlying mesh is refined. Since there is also more work involved in carrying out each iteration, the overall cost of solving (4.1) can quickly escalate. This is of course unsatisfactory, as (4.1) must be solved at each outer Newton iteration. Ideally, we would like a solution method for which the number of iterations is independent of the mesh parameter Δz , so that the amount of computational work grows only linearly with the dimension of \mathcal{H} . We will now consider applying a preconditioner to get closer to this goal.

Conceptually, preconditioning involves replacing (4.1) with the systems

$$\mathcal{P}^{-1/2} \mathcal{H} \mathcal{P}^{-1/2} \tilde{\mathbf{y}} = \mathcal{P}^{-1/2} \tilde{\mathbf{b}}, \quad \mathcal{P}^{1/2} \tilde{\mathbf{x}} = \tilde{\mathbf{y}}$$

for some symmetric positive definite *preconditioner* \mathcal{P} , which leaves the solution unchanged and ensures that the modified system matrix $\tilde{\mathcal{H}} = \mathcal{P}^{-1/2} \mathcal{H} \mathcal{P}^{-1/2}$ is symmetric. In practice, the explicit calculation of $\mathcal{P}^{-1/2}$ is not required: all that is needed

within the algorithm is the action of \mathcal{P}^{-1} . Full details can be found in standard texts on iterative methods—see, for example, [16]. Well-known results show that the eigenvalue spectrum of the coefficient matrix influences the convergence rate of MINRES. The preconditioner must be chosen so that the spectrum of $\tilde{\mathcal{H}}$ is more favorable for MINRES convergence than that of \mathcal{H} , while at the same time, the action of \mathcal{P}^{-1} must also be available at a reasonable cost. The art of developing preconditioners is largely about balancing these two issues, namely, their ability to improve convergence and the cost of their use.

There is a huge amount of literature available on the preconditioning of iterative solvers. Of particular relevance here is the survey paper [1], which highlights some methods appropriate for saddle-point systems, many of which could be applied to (4.1). In the next section, we highlight one example of a simple block preconditioner which is very effective in the current context: the question of how (4.1) should best be preconditioned in general for more complicated liquid crystal models in higher dimensions is an open one which the authors hope to address in future work.

5.1. Ideal block preconditioner. One way of keeping down the cost of evaluating \mathcal{P}^{-1} (or, equivalently, solving a system with \mathcal{P} as coefficient matrix) is to choose a preconditioner that has block-diagonal structure. A natural choice is then a preconditioner of the form

$$\mathcal{P} = \begin{bmatrix} \tilde{A} & O \\ O & S \end{bmatrix},$$

where $S = C + \tilde{D}^T \tilde{A}^{-1} \tilde{D}$ represents the negative Schur complement (see [22]). However, the fact that the matrices \tilde{A} and C have such a simple structure here leads us to suspect that the complication of introducing the Schur complement is not necessary for this problem. Instead we use a simpler block-diagonal preconditioner, setting

$$(5.1) \quad \mathcal{P} = \begin{bmatrix} \tilde{A} & O \\ O & C \end{bmatrix}.$$

We will call this preconditioner the *ideal block preconditioner*, referring to the “ideal” situation where the auxiliary systems with coefficient matrices \tilde{A} and C are solved to high precision. Note that in this section we assume that \tilde{A} is positive definite; circumstances in which this is true are discussed in [14]. The matrix C is always positive definite. Since for our model problem \tilde{A} and C each have very few nonzero diagonals, inverting them is very cheap: in most of the numerical experiments in section 6 we use Matlab’s `backslash` operator. Clearly this strategy may not be practical for more realistic situations with unequal elastic constants, or for more complex two- and three-dimensional problems, as \tilde{A} and C will not have such a simple structure. However, in such cases we conjecture that replacing \tilde{A} in (5.1) by a spectrally equivalent approximation will be a possible way forward (see, for example, [26]). This idea is discussed further in section 6.

Using (5.1), the preconditioned reduced Hessian becomes

$$(5.2) \quad \tilde{\mathcal{H}} = \mathcal{P}^{-1/2} \mathcal{H} \mathcal{P}^{-1/2} = \begin{bmatrix} I & M^T \\ M & -I \end{bmatrix}, \quad M = C^{-1/2} \tilde{D}^T \tilde{A}^{-1/2}.$$

For the off state of the system, D (and hence M) is a zero matrix. The eigenvalues of $\tilde{\mathcal{H}}$ are therefore 1 and -1 , and when the Newton iterate is sufficiently close to the

solution, it can be expected that MINRES will converge extremely rapidly indeed (theoretically, in no more than two iterations, see [22]). For the case when D is not zero, the spectrum of \mathcal{H} is given in the following result. We assume that $\text{rank}(M) = n - 1$, which is true sufficiently close to the true solution when $V \neq 0$, as observed in section 3.3.

LEMMA 5.1. *If $\sigma_1, \dots, \sigma_{n-1}$ are the nonzero singular values of the matrix M , then the $3n$ eigenvalues of \mathcal{H} are*

- (i) 1, with multiplicity $n + 1$,
- (ii) -1 , with multiplicity 1,
- (iii) $\pm\sqrt{1 + \sigma_k^2}$ for $k = 1, \dots, n - 1$.

Proof. Any eigenvalue μ (necessarily real) of $\tilde{\mathcal{H}}$ and corresponding right eigenvector $[\mathbf{X}, \mathbf{Y}]^T$ must satisfy the equations

$$(5.3a) \quad \mathbf{X} + M^T \mathbf{Y} = \mu \mathbf{X},$$

$$(5.3b) \quad M \mathbf{X} - \mathbf{Y} = \mu \mathbf{Y}.$$

Three cases now arise:

(i) $\mu = 1$. In this case, (5.3b) becomes $\mathbf{Y} = \frac{1}{2}M\mathbf{X}$. Then premultiplying (5.3a) by \mathbf{X}^T implies $\mathbf{X}^T M^T M \mathbf{X} = \mathbf{0}$, whence $M\mathbf{X} = \mathbf{0}$ and $\mathbf{Y} = \mathbf{0}$. There are $n + 1$ nontrivial vectors that satisfy $M\mathbf{X} = \mathbf{0}$; so there are $n + 1$ eigenvalues given by $\mu = 1$, with associated right eigenvectors of the form

$$\begin{bmatrix} \mathbf{X}_i \\ \mathbf{0} \end{bmatrix}, \quad i = 1, \dots, n + 1.$$

(ii) $\mu = -1$. In this case, (5.3a) becomes $\mathbf{X} = -\frac{1}{2}M^T \mathbf{Y}$. Then premultiplying (5.3b) by \mathbf{Y}^T implies $\mathbf{Y}^T M M^T \mathbf{Y} = \mathbf{0}$, whence $M^T \mathbf{Y} = \mathbf{0}$ and $\mathbf{X} = \mathbf{0}$. There is only one non-trivial vector which satisfies $M^T \mathbf{Y} = \mathbf{0}$; so there is one eigenvalue given by $\mu = -1$, with right eigenvector of the form

$$\begin{bmatrix} \mathbf{0} \\ \mathbf{Y}_1 \end{bmatrix}.$$

(iii) $\mu \neq 1, -1$. Here (5.3a) implies $\mathbf{X} = \frac{1}{\mu-1}M^T \mathbf{Y}$, so that, from (5.3b), we have $M M^T \mathbf{Y} = (\mu + 1)(\mu - 1)\mathbf{Y}$. Since $M M^T$ has eigenvalues equal to the squares of the singular values of M , we have that all real roots of the quadratic equation

$$\mu^2 - 1 = \sigma_k^2, \quad k = 1, \dots, n - 1$$

give eigenvalues of $\tilde{\mathcal{H}}$. These are the $2(n - 1)$ eigenvalues in part (iii) of the statement of the lemma above. \square

A number of observations follow from this lemma. First, $\tilde{\mathcal{H}}$ has at most $n - 1$ distinct eigenvalues; so in exact arithmetic, MINRES should converge in at most $n - 1$ iterations. Also, in practice the eigenvalues appear to be well clustered, with only a few outliers—see Figure 6.1 for some sample plots. Such an eigenvalue distribution is very favorable for iterative solution via MINRES. Furthermore, the lemma shows that the eigenvalues lie around the origin in the two symmetric intervals

$$[-\beta, -1] \cup [1, \beta], \quad \beta = \sqrt{1 + \sigma_{\max}^2}.$$

It can hence be shown from the standard theory of MINRES convergence that to achieve $\|\mathbf{r}^{(k)}\|_2/\|\mathbf{r}^{(0)}\|_2 \leq \epsilon$ (for some tolerance ϵ and residual vector $\mathbf{r}^{(k)}$), one requires at most

$$(5.4) \quad k \simeq \frac{1}{2} \sqrt{1 + \sigma_{\max}^2} \ln \frac{2}{\epsilon}$$

iterations. That is, the dependence of σ_{\max} , the largest singular value of

$$M = C^{-1/2} \tilde{D}^T \tilde{A}^{-1/2} = C^{-1/2} D^T Z (Z^T A Z)^{-1/2},$$

on the mesh size (in an asymptotic sense) will dictate the growth rate of the number of iterations required for convergence with respect to the grid parameter Δz . For this model problem, σ_{\max} appears in practice to be independent of Δz (see Figure 6.2) and the bound (5.4) is very accurate. This would mean that using the preconditioner (5.1) would lead to an optimal convergence rate for MINRES, in the sense that it is independent of the problem size. This conclusion is supported by the results of the numerical experiments reported in the next section.

6. Numerical experiments. In this section, we illustrate the performance of the ideal block preconditioner in practice via some numerical experiments using Matlab (version R2011b, with the default action of the `backslash` operator for the symmetric solves). For all experiments presented here, we use the nondimensionalized model (2.4), with $\beta = 1/2$, and with associated critical switching value $\alpha_c \simeq 2.721$ (see (2.3)). The initial guess used for \mathbf{n} (as per (1.3)) is taken to be an ad hoc approximate director field with a full, smooth tilt and a 90° twist, while the potential \mathbf{U} is initialized to a linear profile (as would be the true U solution for a homogeneous medium):

$$(6.1a) \quad \mathbf{n}_j = \begin{bmatrix} \cos \theta_j \cos \phi_j \\ \cos \theta_j \sin \phi_j \\ \sin \theta_j \end{bmatrix}, \quad \theta_j = \sin \pi z_j, \quad \phi_j = \frac{\pi z_j}{2}, \quad U_j = z_j, \quad j = 1, \dots, n.$$

The initial guesses for the Lagrange multipliers $\boldsymbol{\lambda}$ are computed from these using

$$(6.1b) \quad \boldsymbol{\lambda} = -B^T \nabla_{\mathbf{n}} f \quad \Leftrightarrow \quad \lambda_j = -\mathbf{n}_j \cdot \nabla_{\mathbf{n}_j} f, \quad j = 1, \dots, n,$$

as $\nabla_{\mathbf{n}} L = \nabla_{\mathbf{n}} f + B \boldsymbol{\lambda} = \mathbf{0}$; for normalized directors,

$$(6.2) \quad \min_{\boldsymbol{\lambda}} \|\nabla_{\mathbf{n}} f + B \boldsymbol{\lambda}\|_2 \quad \Leftrightarrow \quad \boldsymbol{\lambda} = -(B^T B)^{-1} B^T \nabla_{\mathbf{n}} f = -B^T \nabla_{\mathbf{n}} f.$$

The termination criterion for the outer Newton iteration is to stop the iteration at step k if

$$\|\nabla L(\mathbf{x}^{(k)})\|_2 \leq \tau_r \|\nabla L(\mathbf{x}^{(0)})\|_2 + \tau_a,$$

where the relative error tolerance τ_r and the absolute error tolerance τ_a are both 0.0001—see, for example, [8, section 7.2] or [17, section 5.2]. For each inner system (3.9), we used MINRES iteration with a zero initial guess for the Newton corrections, stopping when

$$(6.3) \quad \|\mathbf{r}^{(k)}\|_2 \leq \epsilon \|\mathbf{r}^{(0)}\|_2$$

for a tolerance of $\epsilon = 0.0001$. These tolerances were chosen to ensure accuracy of solutions obtained.




FIG. 6.1. Eigenvalue plots for ideal block preconditioned Hessian $\tilde{\mathcal{H}}$ (5.2) for $N = 32, 64, 128$ grid cells (from top to bottom) at the first (left) and last (right) Newton iteration with $\alpha = 1.5\alpha_c$.




FIG. 6.2. Variation of $\sigma_{\max}(M)$ (as in (5.2)) with $\alpha > \alpha_c$ for $N = 32$ and $N = 2048$.

Figure 6.1 shows a plot of the eigenvalues of the ideal block preconditioned Hessian $\tilde{\mathcal{H}}$ (5.2) for $N = 32, 64, 128$ grid cells (from top to bottom) at the first (left) and last (right) Newton iteration with $\alpha = 1.5\alpha_c$. Each eigenvalue is represented by a vertical bar at the appropriate point on the x -axis. These plots show that the eigenvalues are very well clustered, with only a few outliers. This type of eigenvalue distribution is also seen for other values of N and α . Figure 6.2 shows a plot of the value of σ_{\max} (the maximum singular value of the matrix M in (5.2)) against α for $\alpha > \alpha_c$ with $N = 32$ and $N = 2048$. It is clear from this plot that σ_{\max} is essentially independent of N , which implies that the number of iterations required for the preconditioned solver to converge (as given by (5.4)) will be independent of Δz . In addition, σ_{\max} tends to a constant value as α increases. Note that the singular values of M are smooth functions of α : the dip shown in Figure 6.2 at around $\alpha = 3.9$ represents a crossover as a different singular value becomes the largest one.

Table 6.1 shows the number of MINRES iterations on the full system (3.9) required to satisfy (6.3) when no preconditioning is used for two representative voltage values ($\alpha = 0.5\alpha_c$ for the off state and $\alpha = 1.5\alpha_c$ for the on state). The column headed d contains the dimension of the matrix in each case. For both states, it appears that doubling N approximately quadruples the number of iterations required,

TABLE 6.1

MINRES iterations required for (3.9) with no preconditioner (matrix dimension $d = 5(N - 1)$).

N	d	Off state ($\alpha = 0.5\alpha_c$)		On state ($\alpha = 1.5\alpha_c$)	
		First step	Last step	First step	Last step
32	155	226	499	291	691
64	315	728	2,004	1,172	3,571
128	635	2,680	8,528	4,106	17,498
256	1,275	10,253	41,666	15,727	85,784
512	2,555	38,809	194,753	57,499	>200,000
1,024	5,115	150,376	>200,000	>200,000	>200,000

TABLE 6.2

MINRES iterations required for (4.1) with no preconditioner (matrix dimension $d = 3(N - 1)$).

N	d	Off state ($\alpha = 0.5\alpha_c$)		On state ($\alpha = 1.5\alpha_c$)	
		First step	Last step	First step	Last step
32	93	59	128	90	172
64	189	187	418	285	557
128	381	660	1,456	1,004	2,002
256	765	2,562	5,455	3,650	7,043
512	1,533	9,983	21,393	13,907	26,504
1,024	3,069	41,267	80,778	55,563	81,821
2,048	6,141	171,385	>200,000	>200,000	>200,000

which is consistent with the general $O(1/(\Delta z)^2)$ conditioning of a discretization of a second-order differential operator. This scaling is independent of the Newton iteration, although the number of iterations required is always higher as the Newton iteration nears convergence. Similar observations apply to the results in Table 6.2, which are the equivalent iteration counts for MINRES applied to the reduced system (4.1). Note, however, that the iteration counts for the reduced system are much smaller; we will therefore focus on this version from now on.

In Table 6.3, the same two problems are solved with a diagonal preconditioner applied. That is, \mathcal{P} is a diagonal matrix formed from the diagonal of \mathcal{H} . It can be seen that in this case, diagonal scaling at best offers no improvement, and in fact makes things worse in some cases. Of course, as mentioned in section 5.1, there are many much more sophisticated possible strategies for preconditioning saddle-point problems such as (4.1) than diagonal scaling. We anticipate that many of these would work well here given that, in the current setting, the matrices \tilde{A} and C are cheap to invert. For example, using an “ideal” constraint preconditioner of the form

$$(6.4) \quad \mathcal{P} = \begin{bmatrix} G & \tilde{D} \\ \tilde{D}^T & -C \end{bmatrix}$$

as described in [11] with $G = \tilde{A}$ always gives convergence of the projected preconditioned conjugate gradient (PPCG) algorithm in [11] in one iteration at each Newton step. Table 6.4 shows the equivalent iteration counts for the ideal block preconditioner (5.1). As the Newton iteration proceeds for the off state and the entries in the off-diagonal blocks in (5.2) tend to zero, convergence is obtained in only one iteration as expected. For the on state, it is clear that we have the desired grid independence as predicted by the theory in section 5. Furthermore, in this case the right-hand side of bound (5.4) evaluates to 5.664 for the first Newton iteration and 6.034 for the last Newton iteration, showing that (5.4) gives a very good prediction of performance in practice.

TABLE 6.3
 MINRES iterations required for (4.1) with diagonal preconditioner.

N	Off state ($\alpha = 0.5\alpha_c$)		On state ($\alpha = 1.5\alpha_c$)	
	First step	Last step	First step	Last step
32	55	129	54	122
64	169	408	167	390
128	573	1,469	565	1,423
256	2,144	5,479	2060	5,301
512	8,254	21,196	8,148	20,804
1,024	33,438	85,154	33,849	80,221
2,048	136,015	>200,000	133,605	>200,000

TABLE 6.4
 MINRES iterations required for (4.1) with ideal block preconditioner (5.1).

N	Off state ($\alpha = 0.5\alpha_c$)		On state ($\alpha = 1.5\alpha_c$)	
	First step	Last step	First step	Last step
32	4	1	5	7
64	4	1	5	7
128	4	1	5	7
256	4	1	5	7
512	4	1	5	7
1,024	4	1	5	7
2,048	4	1	5	7
4,096	4	1	5	7
8,192	4	1	5	7
16,384	4	1	5	7
32,768	4	1	5	7
65,536	4	1	5	7

The above experiments use an ideal version of block preconditioner (5.1), that is, the block systems involving \tilde{A} and C are solved to high precision using the `backslash` operator of MATLAB. As indicated above, for problems other than the one studied here (which has equal elastic constants and low-order finite elements in one space dimension), the structure of \tilde{A} and C may make this impractical. One possibility in such a case would be to solve the block systems iteratively: here we illustrate this idea by using a fixed number of preconditioned conjugate gradient (PCG) iterations [6] in place of the direct solves. The preconditioner is the MATLAB implementation of the HSL_MI20 algebraic multigrid (AMG) preconditioner described in [3]. The number of MINRES iterations required with this approach are given in Table 6.5 for one PCG/AMG iteration per block solve (left) and three PCG/AMG iterations per block solve (right). In all cases, the number of iterations required is almost independent of Δz . Furthermore, with 3 PCG/AMG iterations replacing each block solve, the results look very similar to the ideal case. Although using such approximate solves is not necessary here (so we will not consider them further in this paper), these results suggest that they may be sufficient to retain grid-independent MINRES iteration counts when required for other problems. We note that there are also many other possibilities for an approximate solver (including using AMG without PCG) which we have not considered.

As well as looking at iteration counts, it is important to determine if the preconditioner is expensive to apply. Table 6.6 gives a comparison of the elapsed time (with the `tic` and `toc` functions of MATLAB) to solve the twisted nematic device problem with various methods. The four methods compared are

TABLE 6.5

MINRES iterations required for (4.1) with block preconditioner (5.1) using PCG/AMG for block solves.

N	1 PCG/AMG iteration				3 PCG/AMG iterations			
	Off state ($\alpha = 0.5\alpha_c$)		On state ($\alpha = 1.5\alpha_c$)		Off state ($\alpha = 0.5\alpha_c$)		On state ($\alpha = 1.5\alpha_c$)	
	First	Last	First	Last	First	Last	First	Last
32	6	5	7	9	4	1	5	7
64	6	6	7	9	4	1	5	7
128	7	6	7	9	4	1	5	7
256	7	6	8	9	4	1	5	7
512	7	6	8	9	4	1	5	7
1,024	7	6	8	9	4	2	5	7
2,048	7	6	8	9	4	2	5	7
4,096	7	6	8	9	4	2	5	7
8,192	7	6	8	9	4	2	5	7

TABLE 6.6

Timings using the `tic` and `toc` functions of MATLAB for full solution of the “twisted nematic device” problem (section 2.1) using Newton’s method with initial guesses (6.1), parameter values $\beta = 0.5$ and $\alpha = 1.5\alpha_c$. Direct solution of the inner linear system for the Newton step (3.9) is compared to the solution of the nullspace-method reduced system (4.1) using a direct method, MINRES preconditioned by the ideal block preconditioner (5.1), and PPCG preconditioned by an ideal constraint preconditioner (6.4).

N	Full direct	Reduced direct	Ideal block	Ideal constraint
1,024	9.95e-02	9.70e-02	3.48e-01	3.08e-01
2,048	1.42e-01	1.36e-01	5.32e-01	8.35e-01
4,096	2.91e-01	2.79e-01	1.05e+00	2.73e+00
8,192	6.02e-01	5.90e-01	2.20e+00	9.74e+00
16,384	1.42e+00	1.29e+00	4.69e+00	3.80e+01
32,768	3.36e+00	2.75e+00	9.70e+00	8.25e+02
65,536	9.27e+00	7.41e+00	2.53e+01	—

1. direct solution of the full system (3.1);
2. direct solution of the reduced system (4.1);
3. iterative solution of (4.1) using MINRES with ideal block preconditioning;
4. iterative solution of (4.1) using PPCG with ideal constraint preconditioning.

Note that direct solution refers to use of the `backslash` command in MATLAB (version 2011b). From the first two columns, it can be seen that in terms of direct solution, there is a small advantage in using the nullspace method: for larger problems, the gain of solving a smaller system outweighs the costs incurred in setting up the nullspace method. For this one-dimensional problem, we would not expect the new iterative method proposed here to outperform a direct solve, but the timings in the third column show that using the nullspace method with ideal block preconditioning is competitive. Timings using MINRES with the diagonal preconditioner are not displayed as the method fails to converge in 200,000 iterations for most of the values of N used in Table 6.6. The final column contains timings for the constraint preconditioner (6.4), which performs relatively poorly despite taking only one PPCG iteration per Newton step. As the basic MINRES and PPCG algorithms can be shown to be comparable in terms of computational cost per iteration, the difference in performance can be attributed to the additional costs associated with implementing the more sophisticated constraint preconditioner. The “—” in Table 6.6 indicates that

the calculation required more than 16 GB of memory. We expect, however, that performance of this method may be improved using a more sophisticated implementation (here all inner solves and factorizations were carried out using `backslash`).

7. Summary. We have studied a prototypical liquid crystal director model emphasizing the commonly occurring case of coupled electric field interaction and have elucidated the computational challenges presented by discretizations of such models. General models of this type have a double saddle-point structure due to the presence of nonlinear equality constraints and due to the negative-definite nature of the electric field coupling in the free-energy functional. We have presented a nullspace method which provides a natural and effective way to deal with the pointwise unit-vector constraints that arise in these models (as well as in other areas, such as computational micromagnetics). Although the model problem used for illustration here is in one space dimension, the implementation of this constraint would be exactly the same in two or three dimensions; so the technique should be equally effective there.

To utilize these techniques in numerical modeling and exploration of realistic liquid crystal devices and experiments (via numerical bifurcation and phase analyses), one requires efficiently computable criteria to determine if calculated equilibrium points are stable or unstable. The characterization of local stability of equilibria is complicated by the double saddle-point nature of the problem, and this is addressed in [14]. An important consequence of the analysis in [14] is the realization that the reduced matrix $Z^T A Z$ is *not* positive definite at all constrained equilibrium points of interest. Also discussed in [14] is a simplified outer iteration, which is an alternative to the global Newton scheme and which takes advantage of the simplifications that can be realized when the directors are normalized ($|\mathbf{n}_j| = 1 \forall j$), while maintaining the local quadratic convergence of Newton.

For the types of problems we have considered here, the nullspace-method system of equations is real, symmetric, and indefinite. In the large-scale setting, we contemplate solving this system using a MINRES iteration. The effectiveness of the approach hinges on the ability to precondition this system effectively. We have shown that block-diagonal preconditioning (using the $Z^T A Z$ and C diagonal blocks of the coefficient matrix) gives optimal (mesh-independent) performance for our model problem when $Z^T A Z$ is positive definite—the C block matrix is always positive definite and amenable to fast solution in various ways. The reduced Hessian block $Z^T A Z$ is not positive definite at all points of interest, however, and so preconditioning it is a challenge. There are good ways to precondition the A matrix (whether $Z^T A Z$ is positive definite or indefinite), but there are well-documented difficulties in trying to construct a good preconditioner for $Z^T A Z$ using a good preconditioner for A ; see [1, 5, 23, section 6]. This matter is currently under investigation, but we anticipate that similarly good results may be obtained in fully two- and three-dimensional settings by approximating the block solve with $Z^T A Z$ in a spectrally equivalent way. Our approach will then provide a significant enhancement of scientific computing capabilities for a large class of problems of this type.

REFERENCES

- [1] M. BENZI, G. H. GOLUB, AND J. LIESEN, *Numerical solutions of saddle point problems*, *Acta Numer.*, 14 (2005), pp. 1–137.
- [2] D. W. BERREMAN AND S. MEIBOOM, *Tensor representation of Oseen-Frank strain energy in uniaxial cholesterics*, *Phys. Rev. A*, 30 (1984), pp. 1955–1959.

- [3] J. W. BOYLE, M. D. MIHAJLOVIC, AND J. A. SCOTT, *HSL-MI20: An efficient AMG preconditioner for finite element problems in 3D*, Internat. J. Numer. Methods Engrg., 82 (2010), pp. 64–98.
- [4] S. CHANDRASEKHAR, *Liquid Crystals*, 2nd ed., Cambridge University Press, Cambridge, 1992.
- [5] T. F. COLEMAN AND A. VERMA, *A preconditioned conjugate gradient approach to linear equality constrained minimization*, Comput. Optim. Appl., 20 (2001), pp. 61–72.
- [6] P. CONCUS, G. H. GOLUB, AND D. P. O’LEARY, *A generalized conjugate gradient method for the numerical solution of elliptic partial differential equations*, in Sparse Matrix Computations, J. R. Bunch and D. J. Rose, eds., Academic Press, New York, 1976, pp. 309–332.
- [7] K. R. DALY, G. D’ALESSANDRO, AND M. KACZMAREK, *Regime independent coupled-wave equations in anisotropic photorefractive media*, Appl. Phys. B: Lasers Opt., 95 (2009), pp. 589–596.
- [8] J. E. DENNIS, JR. AND R. B. SCHNABEL, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, Prentice-Hall, Englewood Cliffs, NJ, 1983.
- [9] P. G. DE GENNES AND J. PROST, *The Physics of Liquid Crystals*, 2nd ed., Clarendon Press, Oxford, 1993.
- [10] S. DICKMANN, J. ESCHLER, O. COSSALTER, AND D. A. MLYNSKI, *Simulation of LCDs including elastic anisotropy and inhomogeneous fields*, SID Digest, 24 (1993), pp. 638–641.
- [11] H. S. DOLLAR, N. I. M. GOULD, W. H. A. SCHILDERS, AND A. J. WATHEN, *Implicit-factorization preconditioning and iterative solvers for regularized saddle-point systems*, SIAM J. Matrix Anal. Appl., 28 (2006), pp. 170–189.
- [12] R. FLETCHER, *Practical Methods of Optimization*, 2nd ed., John Wiley & Sons, Chichester, 1987.
- [13] C. J. GARCÍA-CERVERA, *Numerical micromagnetics: A review*, Bol. Soc. Esp. Mat. Apl., 39 (2007), pp. 103–135.
- [14] E. C. GARTLAND, JR. AND A. RAMAGE, *Local Stability and a Renormalized Newton Method for Liquid Crystal Director Modeling*, Tech. report, Department of Mathematics and Statistics University of Strathclyde, submitted.
- [15] P. E. GILL, W. MURRAY, AND M. H. WRIGHT, *Practical Optimization*, Academic Press, London, 1981.
- [16] A. GREENBAUM, *Iterative Methods for Solving Linear Systems*, SIAM, Philadelphia, 1997.
- [17] T. KELLEY, *Iterative Methods for Linear and Nonlinear Equations*, SIAM, Philadelphia, 1995.
- [18] S. KITSON AND A. GEISOW, *Controllable alignment of nematic liquid crystals around microscopic posts: Stabilization of multiple states*, Appl. Phys. Lett., 80 (2002), pp. 3635–3637.
- [19] M. KRUŽÍK AND A. PROHL, *Recent developments in the modeling, analysis, and numerics of ferromagnetism*, SIAM Rev., 48 (2006), pp. 439–483.
- [20] A. MAJUMDAR, C. J. P. NEWTON, J. M. ROBBINS, AND M. ZYSKIN, *Topology and bistability in liquid crystal devices*, Phys. Rev. E, 75 (2007), p. 051703.
- [21] H. MORI, E. C. GARTLAND, JR., J. R. KELLY, AND P. J. BOS, *Multidimensional director modeling using the Q tensor representation in a liquid crystal cell and its application to the π cell with patterned electrodes*, Japan. J. Appl. Phys., 38 (1999), pp. 135–146.
- [22] M. F. MURPHY, G. H. GOLUB, AND A. J. WATHEN, *A note on preconditioning for indefinite linear systems*, SIAM J. Sci. Comput., 21 (2000), pp. 1969–1972.
- [23] S. G. NASH AND A. SOFER, *Preconditioning reduced matrices*, SIAM J. Matrix Anal. Appl., 17 (1996), pp. 47–68.
- [24] C. C. PAIGE AND M. A. SAUNDERS, *Solution of sparse indefinite systems of linear equations*, SIAM J. Numer. Anal., 12 (1975), pp. 617–629.
- [25] A. L. PERESSINI, F. E. SULLIVAN, AND J. J. UHL, JR., *The Mathematics of Nonlinear Programming*, Springer-Verlag, New York, 1988.
- [26] D. SILVESTER AND A. WATHEN, *Fast iterative solution of stabilised Stokes systems part II: Using general block preconditioners*, SIAM J. Numer. Anal., 31 (1994), pp. 1352–1367.
- [27] I. W. STEWART, *The Static and Dynamic Continuum Theory of Liquid Crystals: A Mathematical Introduction*, Taylor and Francis, London, 2004.
- [28] I. W. STEWART, *Dynamic theory for smectic A liquid crystals*, Contin. Mech. Thermodyn., 18 (2007), pp. 343–360.
- [29] E. G. VIRGA, *Variational Theories for Liquid Crystals*, Chapman & Hall, London, 1994.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.