Adaptive curvilinear coordinates in a plane-wave solution of Maxwell's equations in photonic crystals

G. J. Pearce,* T. D. Hedley, and D. M. Bird
Department of Physics, University of Bath, Bath BA2 7AY, United Kingdom
(Received 20 December 2004; published 19 May 2005)

A method is described to compute the modes propagating at a given frequency in dielectric systems that are periodic in two dimensions and uniform in the third dimension, using a plane-wave basis expressed in a system of generalized curvilinear coordinates. The coordinates are adapted to the structure under consideration by increasing the effective plane-wave cutoff in the vicinity of the interfaces between dielectrics, where the electromagnetic fields vary most rapidly. The favorable efficiency and convergence properties of the method are shown by comparison with the conventional plane-wave formulation of Maxwell's equations. Although the method is developed to study propagation in photonic crystal fibers, it is also applicable more generally to plane-wave modal solutions of structured dielectrics.

DOI: 10.1103/PhysRevB.71.195108 PACS number(s): 42.70.Qs, 78.20.Bh, 02.70.Hm, 71.15.Dx

I. INTRODUCTION

Photonic crystals, materials whose structure varies on the scale of the wavelength of light, have been an area of great interest in recent years.[1,2] By means of a periodic variation in structure, it is possible in such materials to introduce "photonic band gaps" that forbid propagation of light over a range of frequencies or wave vectors, providing control of the optical properties of the material. Photonic crystal fibers (PCFs) are a class of photonic crystals that are invariant in one spatial dimension but structured in two dimensions.[3,4] Similar to conventional optical fibers, PCFs are usually fabricated from glasses such as silica but the microstructure, consisting of air holes running along the length of the fiber, leads to a wide range of new physical properties. In particular, the ability of PCFs to guide light in an air (rather than solid) core has been demonstrated.[5] These hollow-core PCFs open up a wide range of novel opportunities,[3,4] including enhanced nonlinear effects resulting from their long interaction lengths with gases,[6] and particle guidance.[7] There has been a corresponding increase in the need for computational modeling tools, largely owing to the failure of the theoretical techniques previously used with conventional fibers to describe PCF adequately. Our aim in this paper is to present a method for solving Maxwell's equations to understand the physics of light propagation in PCF. Although the method is formulated specifically for PCF, the ideas behind it are applicable more generally to other types of photonic crystals.

The solution of Maxwell's equations in photonic systems can be approached by a variety of means. Methods generally fall under the broad class of either time-domain or frequency-domain methods, each of which is more ideally suited to different problems. In time-domain methods, such as finite-difference time domain (FDTD),[8] fields are represented on a real-space grid and, by using Maxwell's equations, are evolved over time; both the grid spacing and time interval must be sufficiently fine to ensure that the solutions are physically accurate. This is a flexible and general approach that is well suited to dynamical problems, but it is less suited to determining eigenmodes. Frequency-domain methods instead rely on expanding fields in basis states of a definite frequency. A cutoff is imposed to make the basis set finite (but sufficiently large to give physically accurate results), and the resulting eigenproblem is solved to yield the eigenmodes of the system.

In the case of modeling PCF, and often more generally in photonics, it is the eigenmodes of a given structure that are needed in order to determine the position of band gaps and hence guidance properties. Such optical eigenproblems lend themselves naturally to frequency-domain methods. However, there is generally a choice of variable in using such methods. One method is to fix the wave vector of light $\mathbf{k}$ and compute a set of modes characterized by frequency ω (the "fixed-wave-vector" method). The other possibility is instead to fix ω and, for a d-dimensional calculation, $(d-1)$ components of $\mathbf{k}$, and compute a set of the remaining component (the "fixed-frequency" method). Although the fixed-wave-vector method is in common use,[9] we choose here to develop the fixed-frequency method, as it is a natural approach when studying PCF. In experiments, light enters PCF at a given frequency, and it is of interest which modes (each characterized by the component of $\mathbf{k}$ along the length of the fiber) may propagate. Working at a fixed frequency has the additional advantage of simplifying the inclusion of material dispersion. If the dielectric constant in a material is a function of frequency, we simply choose the dielectric constant relevant to the chosen ω.

Defining the z direction as that along the axis of the fiber, the translational invariance of PCF along its length allows us to write the magnetic field $\mathbf{H}$ as

$$\mathbf{H} = (\mathbf{h}_t + h^z\hat{\mathbf{z}})\exp(i\beta z), \tag{1}$$

where β is the z component of the wave vector (the "propagation constant"), $\mathbf{h}_t(x,y)$ and $h^z(x,y)$ are the transverse and longitudinal components of the magnetic field, respectively, and a time dependence of the form $e^{-i\omega t}$ is implicit throughout. Providing we assume the medium is linear and nonmagnetic $[\varepsilon=\varepsilon_0 n^2(x,y),\ \mu=\mu_0]$, Maxwell's equations for $\mathbf{h}_t$ in this geometry take the form[10]

$$\{\nabla_t^2 + n^2 k_0^2 + \nabla_t \ln n^2 \times \nabla_t \times\}\mathbf{h}_t = \beta^2 \mathbf{h}_t, \tag{2}$$

where ∇_t^2 is the transverse Laplacian operator, $k_0 = \omega/c$ is the free-space wave vector, and $n(x,y)$ is the refractive index. The longitudinal component h^z and the electric field can be derived from $\mathbf{h}_t$.[10] We hold k_0 (and hence frequency) constant, and the eigenvalues of the equation are a set of allowed β^2. Unlike in the fixed-wave-vector method,[9] this eigenproblem is not Hermitian, and therefore a generalized complex eigensolver must be used. This precludes the use of many well-developed variational methods for Hermitian systems, but it brings the advantage that the method, once implemented, can handle materials with complex dielectric constants with no additional effort. The need to develop nonvariational methods also simplifies the process of finding interior eigenvalues, as discussed in detail in Sec. II A.

To solve Eq. (2) computationally, the fields must be expanded in a finite basis set. Issues to consider when choosing the basis set include generality [i.e., placing as few restrictions on $n^2(x,y)$ as possible], the compactness of the representation of the fields (i.e., the number of basis functions needed for accurate results), and the speed and memory efficiency of the computation. A natural choice in studying photonic crystals is to use plane waves as the basis. This choice brings several advantages: it is implicitly periodic; it is general, placing no restrictions on $n^2(x,y)$ other than periodicity; and, crucially, it allows the use of the fast Fourier transform (FFT) to perform rapid calculations in real space in addition to reciprocal space, as is described below in Sec. II B. Although the plane-wave basis is not particularly compact in terms of basis states, the speed increase brought about by using FFTs outweighs this problem. For these reasons, the plane-wave basis is frequently used in photonics,[9,11] and also in other areas of physics such as electronic structure computation.[12]

The implicit periodicity of the plane-wave basis makes it ideal to study propagation in PCF cladding, i.e., the periodic structure that surrounds the central region of the fiber. Knowledge of, for example, the density of photonic states and the location of band gaps in a cladding structure is important to understand the guidance properties of PCF.[13] However, real PCFs have a different microstructure at their center (e.g., an air core[3]) which can be thought of as a defect in the periodic cladding structure. To model the effect of defects in the PCF cladding, it is still possible to use the plane-wave method by employing a supercell approximation. A supercell is a large "unit cell" containing a single defect and many repeated unit cells of cladding; the exponential decay of fields in the cladding region ensures that the interaction between adjacent defect images (introduced by the periodicity of the supercell lattice) is minimal. Of course, this false periodicity introduced by the supercell prevents the computation of radiation losses of a real (finite-sized) fiber. Approaches such as the transfer matrix[14] and multipole[15] methods, which use field expansions localized to defects, are better suited to such confinement loss calculations. However, they do not have the flexibility or speed of the plane-wave method for accurate modal calculations.

An important feature of plane-wave methods that make use of FFTs is that functions (i.e., n^2 and $\mathbf{h}_t$) are represented on a uniform grid of points in real space (see Sec. II B). In PCF (and other photonic crystals) there are usually large regions of space with no variation in n^2 separated by sharp interfaces. Although the fields vary in the regions of constant n^2, it is at the interfaces where the most rapid variation occurs. The grid points in uniform regions therefore tend to be "wasted," implying a loss of computational efficiency. The aim of our work is a reformulation of the fixed-frequency plane-wave method in generalized curvilinear coordinates that allows a *position-dependent* plane-wave cutoff to be used (or, equivalently, a nonuniform grid spacing in real space) while retaining all of the desirable properties of the plane-wave basis outlined above. We demonstrate that by making an appropriate choice of coordinates adapted to the structure of interest, this provides a considerable increase in computational efficiency.

In Sec. II, we describe the conventional fixed-frequency plane-wave method. The formulation of the method in generalized curvilinear coordinates is presented in Sec. III, and a demonstration of its application to a test system is described in Sec. IV. The method and results are discussed in Sec. V.

II. THE CONVENTIONAL FIXED-FREQUENCY PLANE-WAVE METHOD

A brief outline of the conventional fixed-frequency plane-wave method has been published previously in Ref. 11. Although our key result is the development of the method in curvilinear coordinates, it is useful to set this in the context of the conventional approach. The iterative eigensolver, the use of FFTs, and the need for preconditioning which are described in this section are all common to the formulation in generalized curvilinear coordinates given in Sec. III. We therefore begin with a more detailed description of the method of Ref. 11.

A dielectric function that is periodic in 2D with primitive lattice vectors $\mathbf{R}_1$ and $\mathbf{R}_2$ (defining either a unit cell or supercell) satisfies

$$n^2(\mathbf{x}) = n^2(\mathbf{x} + \mathbf{R}), \tag{3}$$

where $\mathbf{R} = l_1\mathbf{R}_1 + l_2\mathbf{R}_2$, $l_1, l_2 \in \mathbb{Z}$, and $\mathbf{x} = (x,y)$ is the 2D position vector. Defining primitive reciprocal lattice vectors by $\mathbf{G}_\alpha \cdot \mathbf{R}_\beta = 2\pi\delta_{\alpha\beta}$ $(\alpha,\beta = 1,2)$, the dielectric function can be expanded in a plane-wave basis set as

$$n^2(\mathbf{x}) = \sum_{\mathbf{G}} n_{\mathbf{G}}^2 e^{i\mathbf{G}\cdot\mathbf{x}}, \tag{4}$$

where $\mathbf{G} = m_1\mathbf{G}_1 + m_2\mathbf{G}_2$ and $m_1, m_2 \in \mathbb{Z}$. Similarly, with the inclusion of a Bloch wave vector $\mathbf{k}$ in the transverse plane, the transverse magnetic field $\mathbf{h}_t = (h^x, h^y)$ may be written

$$h^i(\mathbf{x}) = \sum_{\mathbf{G}} h_{\mathbf{k},\mathbf{G}}^i e^{i(\mathbf{k}+\mathbf{G})\cdot\mathbf{x}}. \tag{5}$$

Substitution into Eq. (2) yields the vector wave equation in reciprocal-space form

$$\sum_{\mathbf{G}'} \begin{pmatrix} M_{xx} & M_{yx} \\ M_{xy} & M_{yy} \end{pmatrix} \begin{pmatrix} h^x_{\mathbf{k},\mathbf{G}'} \\ h^y_{\mathbf{k},\mathbf{G}'} \end{pmatrix} = \beta^2 \begin{pmatrix} h^x_{\mathbf{k},\mathbf{G}} \\ h^y_{\mathbf{k},\mathbf{G}} \end{pmatrix}, \tag{6}$$

where the elements of M (showing manifestly the non-Hermitian nature of the eigenproblem) are

$$M_{xx} = -|\mathbf{k}+\mathbf{G}'|^2\delta_{\mathbf{G},\mathbf{G}'} + k_0^2 n^2_{\mathbf{G}-\mathbf{G}'} + (G_y - G'_y)(k_y + G'_y)[\ln n^2]_{\mathbf{G}-\mathbf{G}'}, \tag{7a}$$

$$M_{yy} = -|\mathbf{k}+\mathbf{G}'|^2\delta_{\mathbf{G},\mathbf{G}'} + k_0^2 n^2_{\mathbf{G}-\mathbf{G}'} + (G_x - G'_x)(k_x + G'_x)[\ln n^2]_{\mathbf{G}-\mathbf{G}'}, \tag{7b}$$

$$M_{yx} = -(G_y - G'_y)(k_x + G'_x)[\ln n^2]_{\mathbf{G}-\mathbf{G}'}, \tag{7c}$$

$$M_{xy} = -(G_x - G'_x)(k_y + G'_y)[\ln n^2]_{\mathbf{G}-\mathbf{G}'}. \tag{7d}$$

Note that the terms in $|\mathbf{k}+\mathbf{G}'|^2$ arise from the $\boldsymbol{\nabla}^2$ operation in Eq. (2), those in $n^2_{\mathbf{G}-\mathbf{G}'}$ from $n^2k_0^2$, and the remainder from the $\boldsymbol{\nabla}_t \ln n^2 \times \boldsymbol{\nabla}_t \times$ operation. Equation (6) is solved numerically by imposing a finite circular cutoff G_{cut} for the reciprocal-space components of h^x and h^y such that only $\mathbf{G}$ vectors satisfying $|\mathbf{G}| < G_{\text{cut}}$ are included. This creates a matrix eigenproblem $\mathbf{M}\cdot\mathbf{v} = \beta^2\mathbf{v}$, where the vector $\mathbf{v}$ comprises reciprocal-space components of both h^x and h^y.

A. Iterative solution

For sufficiently small plane-wave cutoffs, the matrix $\mathbf{M}$ can be diagonalized directly using standard techniques, requiring $O(N_{\text{PW}}^3)$ operations for an $N_{\text{PW}} \times N_{\text{PW}}$ matrix. N_{PW} is the number of plane waves used in Eq. (5), and is typically of the order of 1000 and 100 000 for single cell and supercell calculations, respectively. Diagonalization is not practical for large supercells because the scaling is so poor. However, when modeling PCF, we generally require only a few eigenvalues of $\mathbf{M}$. In particular, for hollow-core PCF, the eigenvalues of interest have $\beta \approx k_0$, which corresponds to modes that can propagate in air (i.e., in the hollow core) but where a band gap may exist in the cladding. These are *interior* eigenvalues, and consist of only a small fraction of the eigenvalue spectrum: smaller β corresponds to the continuum of unbound modes free to propagate in the air holes of the cladding, and larger β to index-guided modes that are peaked in the glass regions. Although a considerable increase in efficiency can be made by using an iterative eigensolver to locate only these states of interest, many methods currently in use are restricted to finding *extremal* eigenvalues. The use of variational methods for interior eigenvalues generally requires squaring $\mathbf{M}$, which worsens the convergence of iterative solvers.[9] Our method does not use a variational scheme and is ideally suited to determining interior eigenvalues.

The transformed eigensystem $(\mathbf{M}-\sigma\mathbf{I})^{-1}\mathbf{v} = \mu\mathbf{v}$ has identical eigenvectors $\mathbf{v}$ to those of $\mathbf{M}$, and eigenvalues related by $\mu = 1/(\beta^2 - \sigma)$. Thus it is possible to transform the interior eigenvalues β^2 to extremal eigenvalues μ by appropriate choice of σ (e.g., for hollow-core PCF $\sigma = k_0^2$). We use the implicitly restarted Arnoldi method [implemented in ARPACK (Ref. 16)] to locate the eigenvalues; this requires only the result of the matrix-vector multiplication $(\mathbf{M}-\sigma\mathbf{I})^{-1}\mathbf{u}$ for arbitrary vectors $\mathbf{u}$. In this method, repeated multiplication of a random starting vector by $(\mathbf{M}-\sigma\mathbf{I})^{-1}$ builds up a Krylov subspace from which extremal eigenvalues are extracted.[17] Typically, we have found that between 4–10 multiplications per eigenvalue are required.

ARPACK requires the results of $(\mathbf{M}-\sigma\mathbf{I})^{-1}\mathbf{u}$. We can provide this by noting that each matrix inversion is equivalent to solving the set of linear equations $(\mathbf{M}-\sigma\mathbf{I})\mathbf{w} = \mathbf{u}$ for the vector $\mathbf{w}$ given $\mathbf{u}$. We have tested a wide range of iterative linear solvers, and conclude that the GMRES (Ref. 18) solver gives particularly rapid convergence. To determine $\mathbf{w}$, GMRES requires only matrix-vector multiplications of the form $(\mathbf{M}-\sigma\mathbf{I})\mathbf{y}$ for arbitrary vectors $\mathbf{y}$. We have found that for well-conditioned problems (see Sec. II C), typically 10–40 multiplications may be needed for each $(\mathbf{M}-\sigma\mathbf{I})^{-1}\mathbf{u}$ evaluation. It follows that of the order of 100 matrix-vector multiplication operations are required for each eigenvalue. This compares well with Hermitian variational schemes for exterior eigenvalues, and very favorably with variational schemes for interior eigenvalues.

In summary, providing we are able to perform the operation $(\mathbf{M}-\sigma\mathbf{I})\mathbf{y}$ on any given vector $\mathbf{y}$, we can efficiently locate interior eigenvalues of $\mathbf{M}$ even for non-Hermitian systems. The problem is reduced to that of performing this operation quickly, even for large $\mathbf{M}$; in Sec. II B below we describe how FFTs can be used to do this efficiently without explicitly forming $\mathbf{M}$.

B. Matrix-vector multiplication by FFT

The $O(N_{\text{PW}}^2)$ storage and computational time requirements of storing the matrix $\mathbf{M}$ and carrying out matrix-vector multiplications can be avoided by use of FFTs. The FFT is a discrete Fourier transform, which creates a real-space representation of a function f on a uniform $N_1 \times N_2$ grid from a set of reciprocal-space components F:

$$f(n_1,n_2) = \frac{1}{N_1 N_2} \sum_{m_1=0}^{N_1-1} \sum_{m_2=0}^{N_2-1} F(m_1,m_2) e^{i\mathbf{G}\cdot\mathbf{x}}. \tag{8}$$

Here, $f(n_1,n_2) = f(\mathbf{x} = n_1\mathbf{R}_1/N_1 + n_2\mathbf{R}_2/N_2)$ and $F(m_1,m_2) = F(\mathbf{G} = m_1\mathbf{G}_1 + m_2\mathbf{G}_2)$. In our case f may be n^2, h^x, or h^y. Any geometry of unit cell or supercell is allowed by Eq. (8), as grid sizes N_1 and N_2 corresponding to any basis vectors $\mathbf{R}_1$ and $\mathbf{R}_2$ can be included. In practice, PCFs tend to be based on hexagonal lattices, in which case it is natural to take $N_1 = N_2 = N$. By choosing a circular cutoff for the reciprocal space components of h^i as described previously, we are able to preserve all symmetries exactly; to use the FFT in this case we extend the cutoff to the nearest edge of the FFT grid and set any unfilled elements on the grid to zero. It follows that the linear FFT grid size N and the number of basis states are related by $N^2 \propto N_{\text{PW}}$. The FFT can be evaluated in $O(N^2 \log N)$ time, or equivalently $O(N_{\text{PW}} \log N_{\text{PW}})$ which compares very favorably with the $O(N_{\text{PW}}^2)$ of direct matrix-vector multiplication.

To demonstrate the use of FFTs in performing matrix-vector operations, consider the h^x component of Eq. (6), which is

$$\sum_{\mathbf{G}'} (M_1 + M_2 + M_3) h^x_{\mathbf{k},\mathbf{G}'} + M_4 h^y_{\mathbf{k},\mathbf{G}'}, \tag{9}$$

where

$$M_1 = -|\mathbf{k} + \mathbf{G}'|^2 \delta_{\mathbf{G},\mathbf{G}'}, \tag{10a}$$

$$M_2 = k_0^2 n^2_{\mathbf{G}-\mathbf{G}'}, \tag{10b}$$

$$M_3 = (G_y - G'_y)(k_y + G'_y)[\ln n^2]_{\mathbf{G}-\mathbf{G}'}, \tag{10c}$$

$$M_4 = -(G_y - G'_y)(k_x + G'_x)[\ln n^2]_{\mathbf{G}-\mathbf{G}'}. \tag{10d}$$

The term in M_1 can clearly be carried out in reciprocal space

$$\sum_{\mathbf{G}'} M_1 h^x_{\mathbf{k},\mathbf{G}'} = -|\mathbf{k} + \mathbf{G}|^2 h^x_{\mathbf{k},\mathbf{G}}, \tag{11}$$

but the term in M_2 is best performed as follows:

$$\sum_{\mathbf{G}'} M_2 h^x_{\mathbf{k},\mathbf{G}'} = \sum_{\mathbf{G}'} k_0^2 n^2_{\mathbf{G}-\mathbf{G}'} h^x_{\mathbf{k},\mathbf{G}'} \tag{12a}$$

$$= k_0^2 [n^2(n_1,n_2) h^x_{\mathbf{k}}(n_1,n_2)]_{\mathbf{G}}, \tag{12b}$$

where Eq. (12b) follows from Eq. (12a) by the convolution theorem, and the quantities labeled (n_1,n_2) are stored and multiplied together pointwise on discrete grids of the form of Eq. (8) in real space. It is clear from Eq. (12a) that to obtain an exact representation of the matrix-vector multiplication in Eq. (6) we require components $n^2_{\mathbf{G}}$ at **G**-vector magnitudes up to a maximum of $2G_{\text{cut}}$. In practice, however, we find that convergence is improved by extending the circular cutoffs of both n^2 and $\mathbf{h}_t$ to the edge of the same FFT grid.

The transformation to real space and back to reciprocal space requires 1 FFT for h^x and 1 FFT for h^y. The term in M_3 (and similarly that in M_4) is evaluated as

$$\begin{aligned}\sum_{\mathbf{G}'} M_3 h^x_{\mathbf{k},\mathbf{G}'} &= \sum_{\mathbf{G}'} (G_y - G'_y)(k_y + G'_y)[\ln n^2]_{\mathbf{G}-\mathbf{G}'} h^x_{\mathbf{k},\mathbf{G}'} \\ &= [\{(k_y + G'_y) h^x_{\mathbf{k},\mathbf{G}'}\}(n_1,n_2) \\ &\quad \times \{(G_y - G'_y)[\ln n^2]_{\mathbf{G}-\mathbf{G}'}\}(n_1,n_2)]_{\mathbf{G}},\end{aligned} \tag{13}$$

where the quantities in { } are evaluated in reciprocal space and then transformed to real space, and the quantity in [ ] is transformed to reciprocal space following the pointwise real space multiplication. By combining terms *before* taking FFTs the total number of FFTs required for all the operations in **M** is 5.

The components of both n^2 and $\ln n^2$ are also generated by FFT. For a calculation to be carried out ultimately with an $N \times N$ FFT grid, we sample n^2 on a *fine* real-space grid (usually of size $\approx 8N \times 8N$), then Fourier transform to reciprocal space. Smoothing is then applied by multiplication of the Fourier components by a Gaussian envelope, for the reasons to be discussed in Sec. II D. The resulting smoothed components are transformed back to real space and the logarithm is taken to create the function $\ln n^2$, which is then transformed to reciprocal space. Components of both n^2 and $\ln n^2$ with $|\mathbf{G}| < G_{\max}$ (where $G_{\max}$ corresponds to the smaller $N \times N$ FFT grid) are extracted for use in the mode-solving calculation. This oversampling procedure ensures that the components of n^2 and $\ln n^2$ remain consistent with one another after smoothing (i.e., the same n^2 is described in each case).

C. Preconditioning

The linear problem $(\mathbf{M} - \sigma\mathbf{I})\mathbf{w} = \mathbf{u}$ may be solved by successive application of $(\mathbf{M} - \sigma\mathbf{I})$ to vectors provided by GMRES, but it is generally the case that convergence when using this method is either slow or fails entirely. This behavior can be rectified by preconditioning,[19] which in this case involves multiplying by a preconditioner **P**:

$$\mathbf{P}(\mathbf{M} - \sigma\mathbf{I})\mathbf{w} = \mathbf{P} \cdot \mathbf{u}, \tag{14}$$

where **P** is chosen to be an approximate inverse of $(\mathbf{M} - \sigma\mathbf{I})$.

A good preconditioner has been found to be an exact inverse (determined by LU decomposition) for $|\mathbf{G}|, |\mathbf{G}'| < G^{\text{P}}_{\text{cut}}$ and the Jacobi preconditioner[19] $P_{ij} = \delta_{ij}/(M_{ii} - \sigma)$ for the part of the matrix corresponding to larger **G** vectors. The choice of the preconditioning cutoff $G^{\text{P}}_{\text{cut}}$ determines the size N_P of the exactly inverted square, and minimizing the total time taken by GMRES involves a trade-off between the number of GMRES iterations, the time taken to perform the exact inversion (scaling as N_P^3), and the time required to apply the preconditioner (scaling as N_P^2). We have found in previous work using this method[11,13,20] that for single unit cells a size of $N_P \approx 100$ is optimal, and for supercells $N_P \approx 2000$ is optimal.

D. Analysis

In nonmagnetic materials, Maxwell's equations require that $\mathbf{H}, \hat{\mathbf{n}} \times \mathbf{E}$, and $\hat{\mathbf{n}} \cdot \varepsilon\mathbf{E}$ are conserved across any boundary with normal vector $\hat{\mathbf{n}}$. In PCF, where ε varies only in the transverse plane, conservation of $\hat{\mathbf{n}} \times \mathbf{E}$ implies conservation of $\varepsilon^{-1}\hat{\mathbf{z}} \cdot \nabla_t \times \mathbf{h}_t$,[10] which in turn implies that any discontinuity in the dielectric function inevitably leads to a discontinuous gradient of $\mathbf{h}_t$.

The truncated Fourier representation of a function that is discontinuous in any of its derivatives suffers the appearance of Gibbs' phenomenon.[21] This can cause slow convergence or numerical instability when using the plane-wave method to solve Maxwell's equations, but can be prevented with smoothing; we apply Gaussian smoothing to the dielectric function, but other methods are also in use.[9,22] The choice of smoothing width is crucial. A large smoothing width gives β values which are perturbed from their true values, but facilitates faster convergence with respect to plane-wave cutoff; a narrower smoothing width gives more accurate β values but makes convergence slower. We discuss this effect quantitatively with reference to a test system in Sec. IV.

Because we use FFTs to create a sampled representation of functions in real space, we can consider the effect of

smoothing on plane-wave convergence qualitatively in terms of the sampling theorem. We can expect a reasonable level of convergence to occur only when there is on average at least one real-space sampling point in the smoothed interface region in any given direction normal to the interface, and so either increasing the smoothing width or the density of grid points in real space enhances convergence. This approximate "convergence, in principle" argument suggests that increasing the density of sampling points at the dielectric interfaces will enhance convergence.

However, the FFT creates a *uniformly spaced* grid of real-space sampling points, as shown in Eq. (8). This has been noted as a problem in Ref. 9; it leads to a loss of compactness in the representation of the dielectric function. If there is rapid variation in n^2 only in a localized area of the unit cell (i.e., at the dielectric interfaces), it is reasonable to expect that convergence as the FFT grid size is increased is largely due to the increased number of points in those areas, and it follows that the additional points in regions of constant n^2 are unnecessary. It would therefore be beneficial to increase the density of sampling points *only* in regions of the unit cell where n^2 is most rapidly varying.

Perhaps surprisingly, it is possible to obtain a position-dependent plane-wave cutoff (and, equivalently, a nonuniform real-space grid spacing) while still using the plane-wave method *and* retaining the use of FFTs. This is made possible by reformulating the plane-wave method in arbitrary curvilinear coordinates; our key finding in this paper is that ultimately we can use this method to allow less smoothing and a smaller number of basis states to be used while retaining all of the computational benefits of the plane-wave method.

III. THE FIXED-FREQUENCY PLANE-WAVE METHOD IN GENERALIZED CURVILINEAR COORDINATES

It is sometimes convenient to express Maxwell's equations in generalized curvilinear coordinates (GCCs), for instance when using the finite-difference method,[23] grating theory,[24] and, recently, perturbation-matched coordinates in coupled mode theory.[25,26] The reformulation of the plane-wave method in arbitrary curvilinear coordinates was first proposed in the field of electronic structure in 1992 by Gygi.[27,28] We report here the first application of this method to the modal solution of Maxwell's equations.

In order to understand the GCC method, first consider the conventional formulation of the plane-wave method. A cutoff condition $|\mathbf{G}| < G_{\max}$ is imposed to create a finite set of basis functions, which are connected by the FFT to an equivalent representation on a real-space grid with uniform spacing proportional to $1/G_{\max}$. Suppose instead that the real-space representation of a function was not uniform in the Euclidean coordinate $\mathbf{x}$, but in another coordinate $\boldsymbol{\xi}$, related to $\mathbf{x}$ by a smooth and invertible mapping $\boldsymbol{\xi} \mapsto \mathbf{x}(\boldsymbol{\xi})$. The grid of $\mathbf{x}$ points corresponding to the uniform $\boldsymbol{\xi}$ grid is nonuniform; thus by choice of coordinate transformation we can increase the density of grid points $1/\Delta x$ at chosen locations in $\mathbf{x}$ space while retaining uniformity in $\boldsymbol{\xi}$ space. A function stored on the $\boldsymbol{\xi}$-space grid may be Fourier transformed to a set of coefficients in reciprocal space which now have a position-dependent effective cutoff $G_{\max}^{\text{eff}}(\mathbf{x}) \propto 1/\Delta x(\mathbf{x})$. It is therefore possible to enhance the convergence of plane-wave calculations by choosing to increase the effective cutoff (or, equivalently, the density of real-space grid points) in regions of interest.

A. The plane-wave basis in curvilinear coordinates

Following the method of Gygi,[27,28] we consider the arbitrary curvilinear coordinate $\boldsymbol{\xi}$, a vector of two Cartesian components ξ^1 and ξ^2, and define a $\boldsymbol{\xi}$-space lattice to have primitive lattice vectors $\mathbf{R}_1$ and $\mathbf{R}_2$ identical to those of the periodic structure under consideration. A transformation between curvilinear coordinates and Euclidean coordinates $\mathbf{x}=(x^1,x^2)$ is defined by

$$\mathbf{x}(\boldsymbol{\xi}) = \boldsymbol{\xi} + \sum_{\mathbf{G}} \mathbf{x}_{\mathbf{G}} \exp(i\mathbf{G}\cdot\boldsymbol{\xi}), \tag{15}$$

where the set of vectors $\mathbf{G}_1$ and $\mathbf{G}_2$ are identical to those of the reciprocal lattice in Euclidean space. The coefficients $\mathbf{x}_{\mathbf{G}}$ fully define the transformation, and can be obtained easily by two FFTs of the quantities $(x^p - \xi^p)$. No restriction is placed on the transformation except that it be invertible (i.e., a 1:1 mapping between $\mathbf{x}$ and $\boldsymbol{\xi}$).

It is useful to define two quantities related to the coordinate transformation: the Riemannian metric tensor, defined by

$$g_{pq} = \sum_{k=1}^{2} \frac{\partial x^k}{\partial \xi^p} \frac{\partial x^k}{\partial \xi^q}, \tag{16}$$

and the "geometric vector potential" A_p:

$$A_p = \frac{1}{4} \frac{\partial}{\partial \xi^p} \ln g, \tag{17}$$

where $g = \det g_{pq}$.[28] Both can be calculated simply from the set of coefficients $\mathbf{x}_{\mathbf{G}}$ using FFTs; they are stored numerically at a uniform set of points in $\boldsymbol{\xi}$ space for future use. The elements $[g_{pq}]$ form a position-dependent and symmetric 2×2 matrix. We may also define another position-dependent symmetric matrix $[g^{pq}] = [g_{pq}]^{-1}$ by inversion at any point in space; this is used below in the evaluation of the ∇^2 matrix element.

Now consider the set of basis functions given by

$$\chi_{\mathbf{k},\mathbf{G}}(\mathbf{x}) = g^{-1/4}(\mathbf{x}) e^{i(\mathbf{k}+\mathbf{G})\cdot\boldsymbol{\xi}(\mathbf{x})}. \tag{18}$$

Using the property $d^2x = g^{1/2} d^2\xi$ to transform the area element, it can be seen that this set of functions is orthonormal and complete.[28] The magnetic field and dielectric function can be represented as a sum over these basis functions, but note that the coefficients are different from those obtained using the conventional plane-wave method [Eqs. (4) and (5)]:

$$h^i(\mathbf{x}) = g^{-1/4}(\mathbf{x}) \sum_{\mathbf{G}} h^i_{\mathbf{k},\mathbf{G}} \exp[i(\mathbf{k}+\mathbf{G})\cdot\boldsymbol{\xi}(\mathbf{x})], \tag{19}$$

$$n^2(\mathbf{x}) = g^{-1/4}(\mathbf{x}) \sum_{\mathbf{G}} n^2_{\mathbf{G}} \exp[i\mathbf{G}\cdot\boldsymbol{\xi}(\mathbf{x})]. \tag{20}$$

B. Operations in curvilinear coordinates

The reciprocal-space representation of the vector wave equation given as Eq. (6) is modified when curvilinear coordinates are used. Equivalents of the $\nabla_t \mathbf{h}_t$ and $\nabla_t^2 \mathbf{h}_t$ matrix elements are published in Ref. 29, but we find it helpful to derive the gradient operation explicitly in the Appendix. The three required operations [from which all operations in Eq. (6) can be constructed] are

$$[\nabla^2 h^i]_{\mathbf{G}} = -\sum_{p,q=1}^{2} \int d^2\xi \sum_{\mathbf{G}'} h^i_{\mathbf{k},\mathbf{G}'}(k_p + G_p - iA_p) \times g^{pq}(k_q + G'_q + iA_q) e^{i(\mathbf{G}'-\mathbf{G})\cdot\boldsymbol{\xi}}, \tag{21a}$$

$$[n^2 h^i]_{\mathbf{G}} = \int d^2\xi n^2(\boldsymbol{\xi}) \sum_{\mathbf{G}'} h^i_{\mathbf{k},\mathbf{G}'} e^{i(\mathbf{G}'-\mathbf{G})\cdot\boldsymbol{\xi}}, \tag{21b}$$

$$\left[\frac{\partial h^i}{\partial x^j}\right]_{\mathbf{G}} = i\sum_{p=1}^{2} \int d^2\xi \sum_{\mathbf{G}'} h^i_{\mathbf{k},\mathbf{G}'} \frac{\partial \xi^p}{\partial x^j}(k_p + G'_p + iA_p) e^{i(\mathbf{G}'-\mathbf{G})\cdot\boldsymbol{\xi}}. \tag{21c}$$

The application of the Maxwell operator in the conventional formulation is carried out in two parts: the differential operators in reciprocal space, and the "potential" multiplication in real space as described in Sec. II B. In the GCC formulation, all operations make use of both spaces. As an example, consider the ∇^2 operation on the component h^i, rearranged here for clarity:

$$[\nabla^2 h^i]_{\mathbf{G}} = -\sum_{p,q=1}^{2} \int d^2\xi (k_p + G_p - iA_p) g^{pq} e^{-i\mathbf{G}\cdot\boldsymbol{\xi}} T^i_q, \tag{22}$$

where T^i_q is given by

$$T^i_q = \sum_{\mathbf{G}'} (k_p + G'_q) h^i_{\mathbf{k},\mathbf{G}'} e^{i\mathbf{G}'\cdot\boldsymbol{\xi}} + iA_q \sum_{\mathbf{G}'} h^i_{\mathbf{k},\mathbf{G}'} e^{i\mathbf{G}'\cdot\boldsymbol{\xi}}. \tag{23}$$

The sums here are evaluated as inverse FFTs of $(k_q + G'_q) h^i_{\mathbf{k},\mathbf{G}'}$ (one FFT for each of $q=1,2$), and $h^i_{\mathbf{k},\mathbf{G}'}$. These are combined in real space to obtain the two functions T^i_1 and T^i_2 stored on a real-space grid. The sum over q may then be performed, to obtain $U^{ip} = \Sigma_q g^{pq} T^i_q$, and the matrix element becomes

$$[\nabla^2 h^i]_{\mathbf{G}} = -\sum_{p=1}^{2} \int d^2\xi (k_p + G_p - iA_p) U^{ip} e^{-i\mathbf{G}\cdot\boldsymbol{\xi}}, \tag{24}$$

which is evaluated as follows:

$$[\nabla^2 h^i]_{\mathbf{G}} = -\sum_{p=1}^{2} \left\{ (k_p + G_p) \int d^2\xi e^{-i\mathbf{G}\cdot\boldsymbol{\xi}} U^{ip} \right\} + i \int d^2\xi e^{i\mathbf{G}\cdot\boldsymbol{\xi}} \sum_{p=1}^{2} A_p U^{ip}. \tag{25}$$

The first integral is performed by two FFTs, by transforming U^{ip} (one FFT for each of $p=1,2$); the second integral is evaluated by a single FFT of the summed quantity $\Sigma_p A_p U^{ip}$. Thus, by use of a total of 6 FFTs for each of h^x and h^y, the $\nabla^2 \mathbf{h}_t$ operation is evaluated. It is worth noting that the $n^2 h^i$ and ∇h^i terms can be combined before carrying out the FFTs required for $\nabla^2 h^i$ [for instance, adding $n^2 h^i$ in real space to $\Sigma_p A_p U^{ip}$ before taking the transform in the second term of Eq. (25)], and hence no additional FFTs are needed to include these terms.

This increase in number of FFTs (12 vs 5) needed for each matrix-vector operation highlights one of the disadvantages of the GCC formulation: while convergence with respect to the number of plane-wave coefficients may be enhanced, the speed of operations for a given number of plane-wave coefficients is reduced. It is also inevitable that the convergence rate of the iterative linear equation solver will be reduced; that is, a greater number of iterations will be needed. This is a consequence of making the matrix $\mathbf{M}$ more dense—in effect, "spreading out" the matrix elements away from the leading diagonal—which in turn makes it harder to invert. This manifests itself in a decrease in effectiveness of the preconditioner, and larger number of iterations needed by GMRES.

C. Optimization of coordinates

The theory presented is exact for any invertible coordinate transformation. We are therefore free to choose any convenient mapping $\boldsymbol{\xi} \mapsto \boldsymbol{\xi}(\mathbf{x})$ and can expect exact results when converged with respect to the number of plane waves, providing there is no aliasing or truncation error in the set of transformation coefficients $\mathbf{x}_{\mathbf{G}}$ or the metric functions derived from them. Careful choice of the mapping, however, can enhance plane-wave convergence; a logical starting point is to concentrate grid points in the regions of the unit cell where the dielectric function, and hence the magnetic field and its gradient, vary most rapidly.

The original method used by Gygi to determine an appropriate mapping considered the coordinate transformation as a variational parameter, and involved updating the coordinates to minimize an energy functional as a calculation proceeds.[28] Similar methods have become popular in the mathematics community to study systems where large solution variations occur over spatial scales that evolve in time, particularly "blowup" solutions of nonlinear partial differential equations.[30,31] In these cases the mapping itself is evolved as time steps in the simulation proceed.

In the case of eigenmode determination in PCF, only a static mapping is required: we choose to concentrate grid points in regions where the dielectric function varies most rapidly. A simple way to do this, which we have found effective, is to define and then minimize a "fictitious energy" functional

$$E[\mathbf{x}(\boldsymbol{\xi})] = \int -|\nabla n^2[\mathbf{x}(\boldsymbol{\xi})]| + \mu_c \det g^{pq} + \mu_s \mathrm{tr}\, g^{pq} d^2\xi, \tag{26}$$

where the first term represents a "negative" energy favoring points in regions where $|\nabla n^2|$ is largest, and the second two terms, following Gygi,[28] give "compression" and "shear" en-

ergy associated with the distortion that control the maximum effective cutoff. In practice we need use only one of these terms, and choose to set $\mu_s=0$. We minimize the energy E using Powell's direction-set method[32] with a finite number of variable coefficients $\mathbf{x_G}$ [as defined in Eq. (15)].

IV. A TEST SYSTEM

We use as a test system a glass of refractive index $n=1.5$ containing a triangular array of circular air holes ($n=1$) with radius 0.45Λ, where Λ is the lattice pitch (air-filling fraction 73%). This forms a PCF cladding structure that is perfectly periodic. A typical value of $k_0\Lambda=10$ is chosen, together with (for simplicity) Bloch wave vector $\mathbf{k}=\mathbf{0}$. This simple system was chosen such that exact results for comparison were obtainable from an alternative method.

We create a coordinate transformation suitable for the test structure by optimizing the energy functional of Eq. (26) while enforcing full *P*6*mm* symmetry. A circular cutoff is obtained by using the first three stars in reciprocal space of $x^1_{\mathbf{G}}$ [as defined in Eq. (15)]; symmetry constrains components in each star such that, of the total of 25 $x^1_{\mathbf{G}}$ components and 25 $x^2_{\mathbf{G}}$ components, there are only three independent variables.

For the grid generation, we use a Gaussian smoothing of n^2 in Eq. (26) with full-width half maximum (FWHM) of 0.03Λ, and set $\mu_c=0.097$, $\mu_s=0$; these parameters are chosen empirically to obtain a grid with the desired characteristics. (Note that for accurate mode solving we use a much smaller smoothing, but the grid generation and eigenmode determination are performed separately.) The resulting real-space grid is used for all subsequent calculations, and is shown in Fig. 1. Relative to a uniform grid with an equal number of sampling points, the density of points (per unit area) is modified by a factor of between ≈ 13 (at the dielectric interfaces) and ≈ 0.12 (at the center of the holes). Note that conservation of the total number of grid points places an upper limit on the amount of distortion that can be achieved.

Once a suitable grid has been generated, it is necessary to generate the set of coefficients $n^2_{\mathbf{G}}$ of the dielectric function to be used in the mode-solving calculation. However, the use of GCCs introduces a complication in the way in which smoothing is applied to the sharp boundaries of this function. In the conventional representation, smoothing is carried out by multiplication of the reciprocal-space components of the dielectric function by a Gaussian envelope. The convolution theorem ensures that this is equivalent to convolving the dielectric function by a Gaussian in real space, and the FWHM of this function can easily be chosen to give the required amount of smoothing. It is unnecessary to perform the computationally expensive convolution in real space explicitly. In the GCC representation there is no simple method to perform Gaussian smoothing working in reciprocal space only, but we wish to avoid performing a slow real-space convolution. We instead set up a Gaussian function $G[\mathbf{x}(\boldsymbol{\xi})]$ in real space and take the FFT, producing a set of coefficients in reciprocal space which are then multiplied by the unsmoothed coefficients $n^2_{\mathbf{G}}$. This does not give Gaussian smoothing, but it provides control over the smoothing width and allows an an approximate equivalence to Gaussian smoothing to be obtained. We show an interface smoothed with each method in Fig. 2.

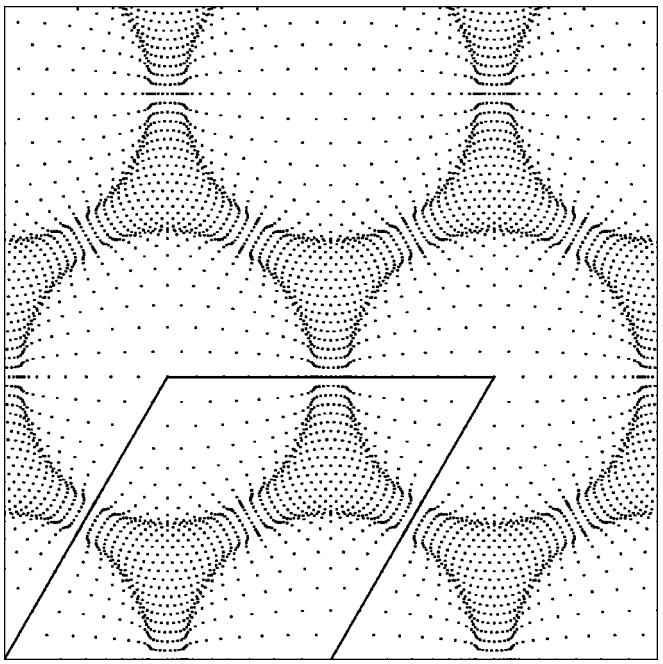

FIG. 1. Grid of real-space sampling points generated for an array of cylindrical air holes of radius 0.45Λ in silica as described in Sec. IV, calculated here with FFT grid size 128×128. To show the grid clearly, every fourth point is plotted to obtain a grid of dimension 32×32. The unit cell is marked.

The GCC method provides greatly enhanced convergence of eigenmodes with respect to the basis set size or, equivalently, the size of the FFT grid used (the equivalence is described in Sec. II B). To demonstrate convergence, we calculate the eight highest-β modes of the test system previously described, and observe the convergence of these with respect

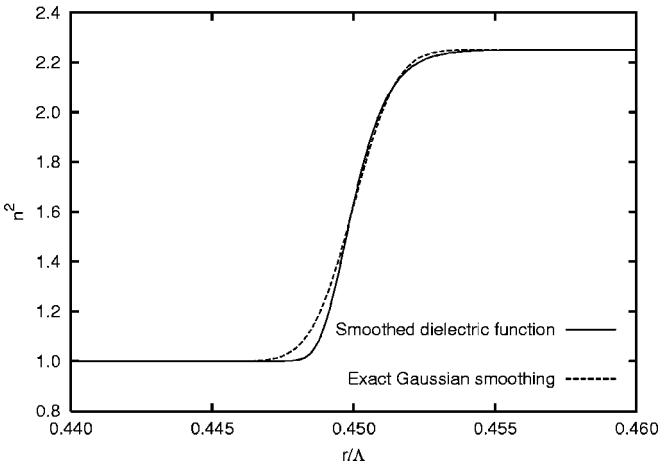

FIG. 2. Smoothed dielectric function of the test system described in Sec. IV plotted in a direction normal to the dielectric interface; the radius r is measured from the center of the hole. Exact Gaussian smoothing of comparable width (FWHM $2.8\times10^{-3}\Lambda$) is shown.

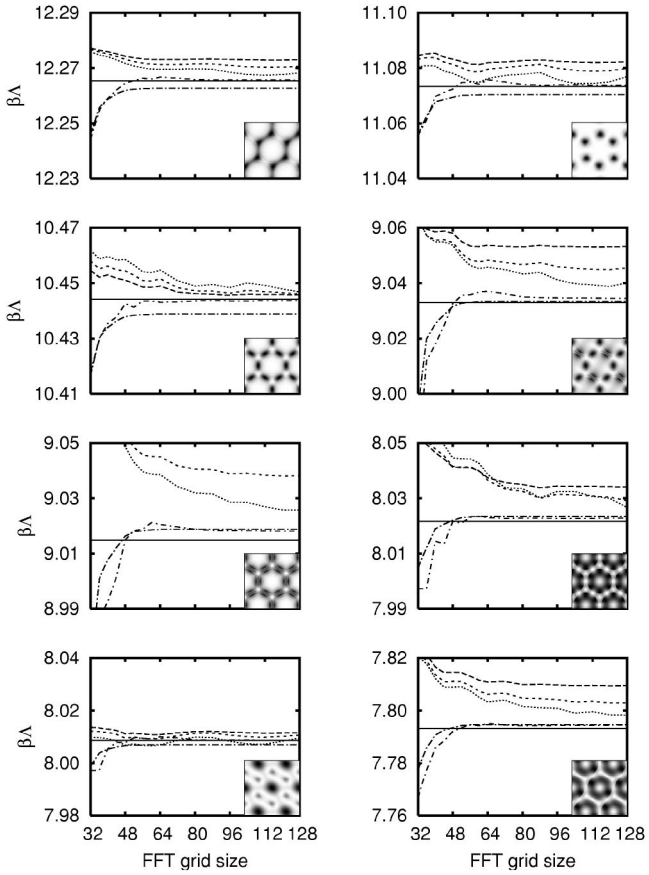

FIG. 3. Convergence of the eight highest-β modes of the test system of Sec. IV as FFT grid size is increased. In each plot the solid horizontal line is the result from a KKR-based method correct to 4 d.p. (Refs. 33,34), the dashed lines show convergence of the conventional method, and the mixed dot/dash lines show the GCC results. The Gaussian smoothing widths for the conventional results are (in order of increasing dash length) $10^{-3}\Lambda$, $6\times10^{-3}\Lambda$, and $10^{-2}\Lambda$. The approximate equivalent widths for the GCC curves are (again in order of increasing dash length) $10^{-3}\Lambda$ and $2.5\times10^{-3}\Lambda$. Inset figures show Poynting vector magnitude (large in black regions). The modes at $\beta\Lambda=12.2654$, 11.0734, 9.0329, 8.0086, and 7.7931 are degenerate pairs; only one of each pair is shown.

to the FFT grid size when different levels of smoothing are applied. In Fig. 3 we plot the calculated value of $\beta\Lambda$ for a range of FFT grid sizes, using three different Gaussian smoothing widths for the conventional calculations and two widths for GCC calculations.

We find that, for all modes examined, the β values of the GCC method converge significantly more quickly with respect to basis size than those produced by the conventional method. (We are not aware of any reason for the GCC β values converging from below, rather than from above as do the conventional plane-wave results.) Comparison of the two least smoothed curves (which are approximately equivalent in smoothing width) makes the enhanced convergence behavior of the GCC method particularly clear. The rapid convergence of the GCC method allows the use of smaller smoothing widths while retaining plane-wave convergence. Consequently, the set of β values from the GCC method are consistently closer than those from the conventional method to the exact (unsmoothed) values, which have been calculated here by a KKR-based method.[33,34] Note that the perturbation in β introduced by the smoothing varies between modes; for example, the most smoothed conventional-method curve for $\beta\Lambda=9.0148$ is not visible on the scale shown: it reaches 9.0529 at FFT size 128. This uncertainty in the accuracy of results is reduced by using smaller smoothing widths throughout.

We can understand qualitatively the reason for the improvement introduced by GCCs in terms of the "convergence in principle" argument of Sec. II D. As the FFT grid size is increased, the density of real-space grid points increases; only when our approximate criterion of having at least one sampling point in the interface region is met can we expect to see convergence of the β values. In the conventional plane-wave case we expect this at an FFT size of roughly 64 for smoothing FWHM $10^{-2}\Lambda$, but we do not expect calculations with smaller smoothing widths to be converged. This can be observed clearly in Fig. 3, where the curves generally do not begin to flatten out until greater FFT grid sizes are used. With curvilinear coordinates, we reach the "converged in principle" criterion at smoothing FWHM $10^{-3}\Lambda$ with an FFT size of roughly 48 because the spacing of sampling points is greatly reduced at the interface. Again, this is visible clearly in Fig. 3: the GCC curves are all generally flat at greater FFT sizes (on the scale shown).

In summary, we find that for a constant FFT grid size a reduction in smoothing width always worsens plane-wave convergence, because the interface regions are narrower and hence less well sampled. By introducing a position-dependent density of sampling points in real space, the curvilinear coordinate method allows convergence to be reached for narrower smoothing widths. Equivalently, the method allows smaller basis sets to be used while keeping the smoothing width constant.

V. DISCUSSION

We have demonstrated that the reformulation of the fixed-frequency plane-wave method in curvilinear coordinates matched to the structure under consideration provides a large enhancement of basis set convergence. For the test system considered, this allows a reduction in FFT grid size from greater than $\approx128\times128$ to $\approx48\times48$ while still obtaining $\beta\Lambda$ values accurate to within 2 d.p. This represents an order of magnitude decrease in the number of basis states needed for a well-converged solution; we expect a similar improvement in other structures.

As with all plane-wave calculations, there is a trade-off between accuracy of results and speed/memory requirements. It is important to consider that, while smoothing perturbs the β values, this is an error that can be controlled. However, calculations that are unconverged with respect to basis set size do not have easily quantifiable errors; Fig. 3 shows that both the magnitude and the rate of convergence varies considerably between modes. This can be seen with reference to, for example, the modes at $\beta\Lambda=9.0148$ and $\beta\Lambda=8.0217$: even at the largest FFT size shown, the conventional plane-wave calculations are clearly unconverged for

some levels of smoothing. However, if convergence were to be gauged by examining, for instance, only the mode at $\beta\Lambda = 8.0086$, the erroneous conclusion that $\beta\Lambda$ values are correct to within 2 d.p. even at FFT size 32 would be reached. Convergence is also strongly dependent on the dielectric system under consideration.[11] We therefore consider it essential to be certain of plane-wave convergence in all calculations of this type, and consider the GCC formulation a useful way to ensure convergence with considerably smaller basis sets than are needed with the conventional method.

Although we have studied a range of grid distortions using the functional of Eq. (26), the choice of parameters μ_s and μ_c in this functional is empirical and cannot be guaranteed to be optimal. In fact, although good results have been obtained with this functional form of E, it is possible that other choices of E or entirely different methods of grid generation may yield coordinate transformations that provide even greater improvements. Grid generation is an area of much current interest in the numerical analysis community, and we intend to make further study of the method of grid generation to effect additional improvements in efficiency.

A further issue that has arisen in the course of our work is the decrease in effectiveness of the preconditioner described in Sec. II C. This can be understood qualitatively by considering the structure of the matrix **M**: the set of Fourier coefficients in the GCC representation provides a more compact description of the dielectric function than the conventional representation, and hence the elements of **M** away from the leading diagonal contain more structural information. The negative effect of using a more compact representation is therefore that the linear problem solved by GMRES becomes less well conditioned. A further direction of our future work will be to improve the method of preconditioning for these problems.

In summary, we have outlined the formulation of the fixed-frequency plane-wave method in curvilinear coordinates for solving Maxwell's equations in PCF, and have demonstrated that this provides a considerable increase in efficiency that is highly desirable in calculations of this type. We anticipate that further research in this area will lead to more accurate and rapid calculation of modal solutions in PCF. We also note that although the ideas used here have been applied to the fixed-frequency plane-wave method, they are equally applicable to the fixed-wave-vector approach and also to fully three-dimensional systems.

ACKNOWLEDGMENTS

This work was supported by the UK Engineering and Physical Sciences Research Council. The authors would like to thank P. J. Roberts for confirmatory calculations.

APPENDIX: GRADIENT OPERATOR IN CURVILINEAR COORDINATES

In this appendix, we give the derivation of the gradient operator in reciprocal space in the GCC representation. In GCCs, the plane-wave basis functions are $\chi_{\mathbf{k},\mathbf{G}}(\mathbf{x}) = g^{-1/4}(\mathbf{x}) e^{i(\mathbf{k}+\mathbf{G})\cdot\boldsymbol{\xi}(\mathbf{x})}$, and so the x^j derivative of a field $h^i_{\mathbf{k}}$ at reciprocal lattice vector $\mathbf{G}$ is

$$\left[\frac{\partial h^i}{\partial x^j}\right]_{\mathbf{G}} \equiv \left\langle \chi_{\mathbf{k},\mathbf{G}} \left| \frac{\partial}{\partial x^j} \right| h^i_{\mathbf{k}} \right\rangle = \int d^2x\, g^{-1/4} e^{-i(\mathbf{k}+\mathbf{G})\cdot\boldsymbol{\xi}} \frac{\partial}{\partial x^j} g^{-1/4} \sum_{\mathbf{G}'} h^i_{\mathbf{k},\mathbf{G}'} e^{i(\mathbf{k}+\mathbf{G}')\cdot\boldsymbol{\xi}}. \tag{A1}$$

We now use the chain rule to change variables in the derivative

$$\left[\frac{\partial h^i}{\partial x^j}\right]_{\mathbf{G}} = \int d^2x\, g^{-1/4} e^{-i(\mathbf{k}+\mathbf{G})\cdot\boldsymbol{\xi}} \sum_{p=1}^{2} \frac{\partial \xi^p}{\partial x^j} \frac{\partial}{\partial \xi^p} g^{-1/4} \times \sum_{\mathbf{G}'} h^i_{\mathbf{k},\mathbf{G}'} e^{i(\mathbf{k}+\mathbf{G}')\cdot\boldsymbol{\xi}}, \tag{A2}$$

then evaluate the derivative (where the left-hand side has been omitted)

$$\int d^2x\, g^{-1/4} e^{-i(\mathbf{k}+\mathbf{G})\cdot\boldsymbol{\xi}} \sum_{p=1}^{2} \frac{\partial \xi^p}{\partial x^j} g^{-1/4} \times \left\{ -\frac{1}{4g}\frac{\partial g}{\partial \xi^p} + i(k_p + G'_p) \right\} \sum_{\mathbf{G}'} h^i_{\mathbf{k},\mathbf{G}'} e^{i(\mathbf{k}+\mathbf{G}')\boldsymbol{\xi}}. \tag{A3}$$

Collecting the factors of $g^{-1/4}$ and using $d^2x = g^{1/2} d^2\xi$ provides the change of variables

$$\int d^2\xi\, e^{-i(\mathbf{k}+\mathbf{G})\boldsymbol{\xi}} \sum_{p=1}^{2} \frac{\partial \xi^p}{\partial x^j} \times \left\{ -\frac{1}{4g}\frac{\partial g}{\partial \xi^p} + i(k_p + G'_p) \right\} \sum_{\mathbf{G}'} h^i_{\mathbf{k},\mathbf{G}'} e^{i(\mathbf{k}+\mathbf{G}')\boldsymbol{\xi}}. \tag{A4}$$

The definition of A_p in Eq. (17) can be rearranged as

$$A_p = \frac{1}{4}\frac{\partial}{\partial \xi^p} \ln g = \frac{1}{4g}\frac{\partial g}{\partial \xi^p}, \tag{A5}$$

and hence, by substitution into Eq. (A4),

$$\left[\frac{\partial h^i}{\partial x^j}\right]_{\mathbf{G}} = i \sum_{p=1}^{2} \int d^2\xi \sum_{\mathbf{G}'} h^i_{\mathbf{k},\mathbf{G}'} \frac{\partial \xi^p}{\partial x^j} (k_p + G'_p + iA_p) e^{i(\mathbf{G}'-\mathbf{G})\boldsymbol{\xi}}, \tag{A6}$$

i.e., Eq. (21c). The reciprocal-space form of the Laplacian operator $\nabla^2 h^i$ can be derived using this result by applying Green's first theorem[21] (neglecting the boundary term because of periodic boundary conditions):

$$\langle \chi_{\mathbf{k},\mathbf{G}} | \nabla^2 | h^i_{\mathbf{k}} \rangle = -\int d^2x\, \boldsymbol{\nabla} \chi_{\mathbf{k},\mathbf{G}} \cdot \boldsymbol{\nabla} h^i_{\mathbf{k}}. \tag{A7}$$

*Electronic address: g.j.pearce@bath.ac.uk

[1] J. D. Joannopoulos, R. D. Meade, and J. N. Winn, *Photonic Crystals* (Princeton University Press, Princeton, 1995).
[2] L. Venema, Nature (London) **424**, 809 (2003).
[3] P. S. J. Russell, Science **299**, 358 (2003).
[4] J. C. Knight, Nature (London) **424**, 847 (2003).
[5] R. F. Cregan, B. J. Mangan, J. C. Knight, T. A. Birks, P. S. J. Russell, P. J. Roberts, and D. C. Allan, Science **285**, 1537 (1999).
[6] F. Benabid, G. Bouwmans, J. C. Knight, P. S. J. Russell, and F. Couny, Phys. Rev. Lett. **93**, 123903 (2004).
[7] F. Benabid, J. C. Knight, P. S. J. Russell, and F. F. Couny, Opt. Express **10**, 1195 (2002).
[8] A. Taflove, *Computational Electrodynamics: The Finite-Difference Time-Domain Method* (Artech House, Boston, 1995).
[9] S. G. Johnson and J. D. Joannopoulos, Opt. Express **8**, 173 (2001).
[10] A. W. Snyder and J. D. Love, *Optical Waveguide Theory* (Chapman & Hall, London, 1983).
[11] J. M. Pottage, D. M. Bird, T. D. Hedley, T. A. Birks, J. C. Knight, P. S. J. Russell, and P. J. Roberts, Opt. Express **11**, 2854 (2003).
[12] M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos, Rev. Mod. Phys. **64**, 1045 (1992).
[13] T. A. Birks, D. M. Bird, T. D. Hedley, J. M. Pottage, and P. S. J. Russell, Opt. Express **12**, 69 (2004).
[14] A. Yariv and P. Yeh, *Optical Waves in Crystals* (Wiley, New York, 1984).
[15] J. Lo, R. C. McPhedran, I. M. Bassett, and G. W. Milton, J. Lightwave Technol. **12**, 396 (1994).
[16] URL http://www.caam.rice.edu/software/ARPACK/
[17] D. C. Sorensen, Acta Numerica **11**, 519 (2002).
[18] Y. Saad and M. H. Schultz, SIAM (Soc. Ind. Appl. Math.) J. Sci. Stat. Comput. **7**, 856 (1986).
[19] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. van der Vorst, *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods* (SIAM, Philadelphia, 1994).
[20] F. Luan, A. K. George, T. D. Hedley, G. J. Pearce, D. M. Bird, J. C. Knight, and P. S. J. Russell, Opt. Lett. **29**, 2369 (2004).
[21] K. F. Riley, M. P. Hobson, and S. J. Bence, *Mathematical Methods for Physics and Engineering* (Cambridge University Press, Cambridge, 1998).
[22] R. D. Meade, A. M. Rappe, K. D. Brommer, J. D. Joannopoulos, and O. L. Alherhand, Phys. Rev. B **48**, 8434 (1993).
[23] R. Holland, IEEE Trans. Nucl. Sci. **30**, 4589 (1983).
[24] J. P. Plumey, G. Granet, and J. Chandezon, IEEE Trans. Antennas Propag. **43**, 835 (1995).
[25] M. Skorobogatiy, S. A. Jacobs, S. G. Johnson, and Y. Fink, Opt. Express **10**, 1227 (2002).
[26] M. Skorobogatiy, Phys. Rev. E **70**, 046609 (2004).
[27] F. Gygi, Europhys. Lett. **19**, 617 (1992).
[28] F. Gygi, Phys. Rev. B **48**, 11 692 (1993).
[29] D. R. Hamann, Phys. Rev. B **54**, 1568 (1996).
[30] W. Huang, Y. Ren, and R. D. Russell, SIAM (Soc. Ind. Appl. Math.) J. Numer. Anal. **31**, 709 (1994).
[31] W. Cao, R. Carretero-González, W. Huang, and R. D. Russell, SIAM (Soc. Ind. Appl. Math.) J. Numer. Anal. **41**, 235 (2003).
[32] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in Fortran 77* (Cambridge University Press, Cambridge, 1992), Chap. 10, pp. 406-413.
[33] P. J. Roberts (private communication).
[34] N. A. Nicorovici and R. C. McPhedran, Phys. Rev. E **50**, 3143 (1994).