COMPUTATIONAL MODELLING OF PHOTONIC CRYSTALS

T. N. Langtry
CUDOS & Department of Mathematical Sciences
University of Technology, Sydney
NSW 2007, Australia
e-mail: Tim.Langtry@uts.edu.au

L. C. Botten
CUDOS & Department of Mathematical Sciences
University of Technology, Sydney
NSW 2007, Australia
e-mail: Lindsay.Botten@.uts.edu.au

A. A. Asatryan
CUDOS & Department of Mathematical Sciences
University of Technology, Sydney
NSW 2007, Australia
e-mail: Ara.Asatrian@uts.edu.au

M. A. Byrne
CUDOS & Department of Mathematical Sciences
University of Technology, Sydney
NSW 2007, Australia

A. Bourgeois
Magistère de Physique Fondamentale d'Orsay,
Université de Paris-Sud,
Bâtiment 470, 91 405 ORSAY Cedex, France

R. C. McPhedran
CUDOS & School of Physics
University of Sydney
NSW 2006, Australia
e-mail: r.mcphedran@physics.usyd.edu.au

Abstract

Photonic crystals (PC) are a novel class of complex materials which have properties that make them the optical analogues of semiconductors. Accordingly, PCs will likely be key building blocks for future micro-optical and communication technology, specifically because of their ability to tailor the propagation of light on the scale of optical wavelengths with minimal diffraction losses. This paper, which describes the computational modelling of PC based devices being undertaken within the ARC Centre for Ultrahigh-bandwidth Devices and Optical Systems (CUDOS), commences with an outline of multipole and Bloch mode theory, outlines the implementation of the computational algorithms using MPI and Open MP on parallel computer systems, and presents briefly the results of the modelling of some PC based devices.

Key words and phrases

Photonic crystals, Monte Carlo, MPI, OpenMP

1 INTRODUCTION

Photonic crystals (PC) (Joannopoulos, Meade and Winn, 1995; Soukoulis, 1995) are a novel type of complex material in which the refractive index varies periodically with position. Through interference action known as Bragg reflection, PCs can inhibit the propagation of electromagnetic waves for various ranges of some frequencies, giving rise to "band gaps" in their spectrum. It is these band gaps that make them the optical analogues of semiconductors which have underpinned the electronics revolution. Because PCs have the capacity to tailor the propagation of light on the scale of optical wavelengths with minimal loss, it is likely that they will be a key building block for future micro-optical and communication technology. It is anticipated that they will provide optical functionality analogous to that available in electronic chips and facilitate information processing of unprecedented speed (Parker and Charlton, 2000).

Of substantial, contemporary interest is the design of PC based devices that will allow the miniaturization of photonic components, and the use of PCs to modify the radiation dynamics of sources embedded within them. With respect to the former, various devices and interconnections such as bends, junctions and filters have been modelled and fabricated (Lin, Chow, Hietala, Villeneuve and Joannopoulos, 1998; Mekis, Chen, Kurland, Fan, Villeneuve and Joannopoulos, 1996). The exploitation of their technological potential however requires a thorough understanding of the coupling and guiding mechanisms. Coupling problems, which are both challenging and general, involve interfacing various PC media and waveguide devices to one another or to some external medium. While most investigations to date have been largely computational, we have recently developed a new approach (Botten, Asatryan, Langtry, White, de Sterke and McPhedran, 2003), utilising analytic techniques based on Bloch modes and elaborated in §2.2. The method exploits the underlying physics and provides substantial analytic and computational advantages.

The crystals described in this paper consist of infinitely long cylindrical inclusions embedded in a background matrix. Sources, where appropriate, are modelled by line sources parallel to the cylinder axes. Such structures represent idealised 2D models of actual devices consisting of finite length cylinders bounded by slabs of a uniform material. For sufficiently long cylinders it is hoped that the current 2D models will provide useful insight into the behaviour of the 3D crystals.

The potential to exert control over the radiation dynamics through the ability to vary the density of states available to electromagnetic radiation was seen as a prime motivation for the advent of PCs (Yablonovitch, 1987; John, 1987). A key quantity that determines the dynamics of a fluorescent source embedded in a photonic crystal is the spatially resolved or *local density of states* (LDOS). In three dimensions, the LDOS provides the spectral distribution of modes to which a fluorescent source couples: if, for a particular frequency, the LDOS is large, then this indicates that the light emission at that frequency is enhanced by the structure; correspondingly, a small value of the LDOS indicates suppression of light emission at that frequency. In recent years, we have developed a novel theoretical method (Asatryan, Busch, McPhedran, Botten, de Sterke and Nicorovici, 2001) based on multipole expansions which is elaborated in §2.1 and (Botten, Nicorovici, Asatryan, McPhedran, de Sterke and Robinson, 2000). These theories work well with both regular (periodic) and irregular or finite clusters, yielding exceptional computational speed and accuracy, and considerable analytic

Figure 1: Device designs which may be seen as either finite clusters of scatterers or extended devices (when one or more regions is assumed to be semi-infinite in extent). (a) A directional coupler exhibiting disorder in the coupling region (M_2). (b) Waveguide coupling to free space.

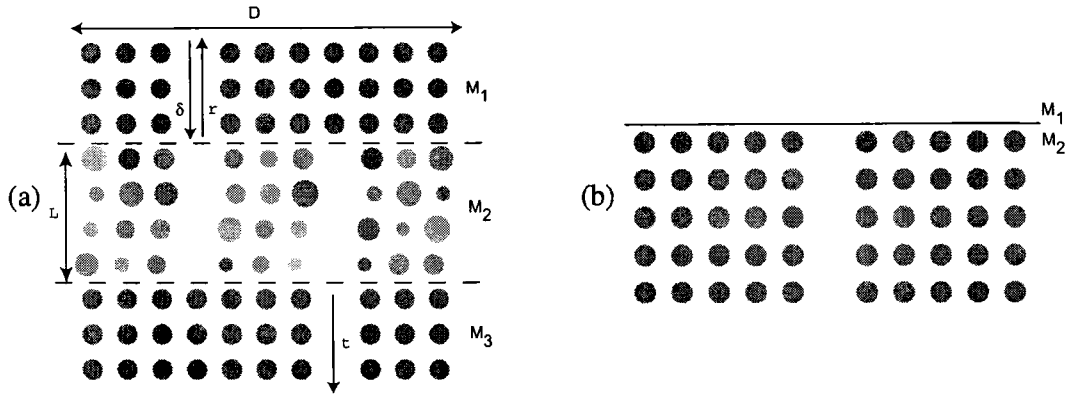

insight, albeit with some restrictions on the range of geometries that can be considered.

Despite the substantial research in this area, the strong similarities that exist between photonic phenomena and their electronic counterparts suggests that this is a rich vein of potential new opportunities that should be actively "mined". One such is the transfer of the concept of electron conductance, one of the main properties of nano-scale wires, to photonics, which we deal with briefly in §2.3.

Another key aspect of the theoretical and computational programs being undertaken within CUDOS, the ARC Centre for Ultrahigh-bandwidth Devices for Optical Systems, is the modelling of the performance tolerance of PC structures to fabrication defects. It is here that the modelling programs exploit parallel computation. Accordingly, this paper commences with an outline the multipole theory for a cluster, indicates its extension to the modelling of linear, periodic structures, makes use of this in the determination of Bloch modes, and exploits these tools in the modelling of linear PC devices (Botten *et al.*, 2003). From there, we outline the design of our simulation codes that are implemented on the APAC and ac3 systems, discuss the computational performance that is achievable, and present briefly some results from our simulations.

2 THEORETICAL FORMULATION

In this section, we outline theoretical models for the diffraction of light by (a) a finite cluster of scatterers and for (b) extended devices. Fig. 1, when considered as an ensemble of cylinders, exemplifies a simple 2D cluster of finite size which is modelled using a multipole expansion method (§2.1). The analysis of an extended device, such as that shown in Fig. 1, in which the front and rear channels (or waveguides) seen in regions M_1 and M_3 are arbitrarily long, requires a more elaborate analysis involving the basis of natural (Bloch) modes of the various structural elements. Section §2.2 treats the structure in Fig. 1, known as a *folded directional coupler*, and finally, §2.3 outlines the theoretical treatment for the quantification

of photon conductance.

2.1 Multipole theory for photonic crystal clusters

We solve for the total field $V(\mathbf{r})$ at a point $\mathbf{r}$ in a cluster of N_c cylinders of radii $\{a_l\}$ and refractive indices $\{n_l\}$, due to a line source located at $\mathbf{c}_s$ and radiating at a frequency ω and wavelength λ (related by $\omega = kc$, with $k = 2\pi/\lambda$ and c denoting the speed of light in free space). For simplicity and brevity, we outline only the treatment for $E_{\shortparallel}$ polarised light for which the electric field vector is aligned with the cylinder axes (the $z-$axis). In this case, the only non-trivial component of the Green's function $\mathbf{G}$ is $V \stackrel{\text{def}}{=} G_{zz}$ that satisfies the boundary value problem

$$\nabla^2 V(\mathbf{r};\mathbf{c}_s) + k^2 n^2(\mathbf{r}) V(\mathbf{r};\mathbf{c}_s) = \delta(\mathbf{r}-\mathbf{c}_s), \tag{1}$$

where V and $\nu \cdot \nabla V$ are continuous across all cylinder boundaries. Here, $n(\mathbf{r})$ denotes the refractive index of the cylinders or the background matrix, and ν is the outward normal unit vector to the surface of the cylinders.

In the vicinity of each cylinder l, the exterior field in the background matrix of refractive index $n_b = 1$ is expanded in local coordinates $\mathbf{r}_l = (r_l, \theta_l) = \mathbf{r} - \mathbf{c}_l$, with

$$V(\mathbf{r};\mathbf{c}_s) = \sum_{m=-\infty}^{\infty} \left[A_m^l J_m(kr_l) + B_m^l H_m^{(1)}(kr_l)\right] e^{im\theta_l}, \tag{2}$$

and involves both singular or irregular components characterised by coefficients B_m^l and regular components characterised by A_m^l. This expansion is valid only in the annulus extending from the surface of the cylinder l to the surface of the nearest cylinder or source.

A global field representation, known as the Wijngaard expansion, valid everywhere in the matrix may be derived (for example, see Asatryan, Busch, McPhedran, Botten, de Sterke and Nicorovici, 2001) using Green's theorem:

$$V(\mathbf{r};\mathbf{c}_s) = \frac{1^{\text{ext}}(\mathbf{c}_s)}{4i} H_0^{(1)}(k|\mathbf{r}-\mathbf{c}_s|) + \sum_{q=1}^{N_c} \sum_{m=-\infty}^{\infty} B_m^q H_m^{(1)}(k|\mathbf{r}-\mathbf{c}_q|) e^{im\arg(\mathbf{r}-\mathbf{c}_q)}. \tag{3}$$

Here, the first term is due to the inhomogeneous source term in (1) (with $1^{\text{ext}}(\mathbf{c}_s)$ acting as a source term switch), while the second is associated with scattered radiation sourced by each of the cylinders q.

Equating the global (3) and the local (2) forms of the field expansions in the vicinity of cylinder l, and applying Graf's addition theorem to the global form in order to change the field phase origin to the centre of cylinder l, we derive the Rayleigh multipole field identity

$$A_m^l = K_m^l + \sum_{q=1, q\neq l}^{N_c} \sum_{p=-\infty}^{\infty} S_{mp}^{lq} B_p^q. \tag{4}$$

Equation (4) indicates that the regular part of the local field in the vicinity of cylinder l is due to a combination of (a) the exterior source (which, when expanded in a local expansion, is

characterized by coefficients K_m^l), and (b) sources on all other cylinders ($q \neq l$), the contributions of which to the multipole term of order $m - p$ at cylinder l is given by S_{mp}^{lq}. The particular forms of the S_{mp}^{lq} and K_m^l and details of the derivation may be found in (Botten *et al.*, 2000). The field identity is most conveniently cast in matrix form

$$\boldsymbol{A} = \boldsymbol{S}\boldsymbol{B} + \boldsymbol{K} \tag{5}$$

and typifies the form arising in all problems solved using the Rayleigh multipole method. Here, $\boldsymbol{A} = \left[\boldsymbol{A}^l\right]$, $\boldsymbol{B} = \left[\boldsymbol{B}^l\right]$ and $\boldsymbol{K} = \left[\boldsymbol{K}^l\right]$ denote partitioned vectors and $\boldsymbol{S} = \left[\boldsymbol{S}^{lq}\right]$ denotes a partitioned matrix.

The coefficients $\boldsymbol{A}$ and $\boldsymbol{B}$ are linked by the boundary conditions, viz. the continuity of the tangential components of the electric and magnetic fields across cylinder-matrix interfaces. Commencing with a local expansion of the field interior to cylinder l, we have

$$V(\mathbf{r};\mathbf{c}_s) = \sum_{m=-\infty}^{\infty} Q_m^l H_m^{(1)}(kn_l r_l)e^{im\theta_l} + \sum_{m=-\infty}^{\infty} C_m^l J_m(kn_l r_l)e^{im\theta_l}, \tag{6}$$

in which the first series denotes a possible, interior field source $1_l^{\text{int}}(\mathbf{c}_s)H_0^{(1)}(kn_l|\mathbf{r}-\mathbf{c}_s|)/(4i)$, analogous to the exterior source in (3). The boundary conditions then take the form

$$\boldsymbol{B} = \tilde{\boldsymbol{R}}\boldsymbol{A} + \tilde{\boldsymbol{T}}\boldsymbol{Q}, \tag{7}$$

$$\boldsymbol{C} = \tilde{\boldsymbol{T}}'\boldsymbol{A} + \tilde{\boldsymbol{R}}'\boldsymbol{Q}, \tag{8}$$

where the matrices $\tilde{\boldsymbol{R}}$, $\tilde{\boldsymbol{T}}$, $\tilde{\boldsymbol{R}}'$ and $\tilde{\boldsymbol{T}}'$ are block diagonal matrices containing diagonal blocks of cylindrical harmonic reflection and transmission coefficients. In (7) and (8), the regular exterior field $\boldsymbol{A}$ and the irregular interior field $\boldsymbol{Q}$ are regarded as incident or source fields, generating the irregular exterior $\boldsymbol{B}$ and regular interior $\boldsymbol{C}$ fields.

The Rayleigh identity, in the form of a system of linear equations in the source coefficients $\boldsymbol{B}$, is then deduced from (5) and (7) yielding the solution

$$\boldsymbol{B} = \left(\boldsymbol{I} - \tilde{\boldsymbol{R}}\boldsymbol{S}\right)^{-1}\left(\tilde{\boldsymbol{R}}\boldsymbol{K} + \tilde{\boldsymbol{T}}\boldsymbol{Q}\right), \tag{9}$$

in which the right hand side comprises the multipole scattering operator $\left(\boldsymbol{I} - \tilde{\boldsymbol{R}}\boldsymbol{S}\right)^{-1}$ and a source term which is the sum of the reflection ($\tilde{\boldsymbol{R}}$) of a source $\boldsymbol{K}$ exterior to the cylinder or the transmission ($\tilde{\boldsymbol{T}}$) of an interior source $\boldsymbol{Q}$ through cylinder interfaces. With the multipole source coefficients $\boldsymbol{B}^l$ determined from the solution of the system (9), the Green's function can be reconstructed from its exterior (3) and interior (6) Wijngaard expansions. The LDOS ρ at any point $\mathbf{r}$ is then given by

$$\rho(\mathbf{r};\omega) = -\frac{2\omega}{\pi c^2}\Im\left[V(\mathbf{r},\mathbf{r};\omega)\right],$$

where $\Im$ denotes the imaginary part of the expression.

2.2 Bloch mode modelling of extended PC devices

While multipole methods are effective tools for compact structures, they are not an efficient tool for handling extended structures such as that in Fig. 1(a). The motivation for the development of the Bloch mode tools was to develop an efficient design suite that could handle the wide range of coupling and guiding problems that arise in the design and development of PC devices. While the majority of previous investigations rely almost exclusively on computational techniques, our tools use a combined analytic and computational approach that separates propagation in straight "waveguide" sections from the scattering of modes that occurs at interfaces. The latter is characterized by matrix generalisations of the Fresnel reflection and transmission coefficients that are used widely in thin film optics.

As an example, consider the propagation problem of Fig. 1(a) in which three distinct media are jointed together. Each medium may be homogeneous (e.g., free space or dielectric) or a periodically modulated structure such as a photonic crystal, possibly penetrated by waveguides as shown. Each medium is characterized by its Bloch modes, generated (Botten *et al.*, 2001 via a transfer matrix technique based on supercell methods that exploit diffraction grating theories. Initially, we consider the region M_2 to comprise L identical layers.

Consider now a uniform structure, such as those of the guides in regions M_1 and M_3 of Fig. 1(a), in which all layers are identical and each has a cylinder removed in the same place. The calculation of Bloch modes characterising a region M_s follows from the solution of an eigenvalue problem derived by considering an infinite region comprising only the layers of region M_s. This infinite structure is unaltered by a translation of a single layer, and hence the electromagnetic fields one layer apart must be related by a simple multiplication factor $\mu = \exp(ik_{0y}d)$ where d is the interlayer spacing and k_{0y} is the relevant component of the Bloch vector. That is, $V(x, y + d) = \mu V(x, y)$. At the boundary of each layer j (that is, midway between the line of centres of the cylinders), we express the fields in expansions of upward and downward propagating plane waves

$$V(x,y) = \sum_{p=-\infty}^{\infty} \left[f_p^{(j)-} e^{-i\chi_p(y-y_j)} + f_p^{(j)+} e^{+i\chi_p(y-y_j)}\right] e^{i\alpha_p x}. \tag{10}$$

We use a supercell method with lateral period D (Fig. 1(a)) chosen to be sufficiently large so that there is negligible cross-talk between channels when the PC structure is operated in a band gap. In this case, the direction sines and cosines of the plane wave terms in (10) are given by $\alpha_p = \alpha_0 + 2\pi/D$ and $\chi_p = \sqrt{k^2 - \alpha_p^2}$. Defining $\boldsymbol{f}_j^{\mp} = \left[f_p^{(j)\mp}\right]$ (vectors of plane wave coefficients), the scattering action of a single layer is

$$\begin{aligned} \boldsymbol{f}_2^- &= \hat{\boldsymbol{T}}\boldsymbol{f}_1^- + \hat{\boldsymbol{R}}\boldsymbol{f}_2^+, \\ \boldsymbol{f}_1^+ &= \hat{\boldsymbol{R}}\boldsymbol{f}_1^- + \hat{\boldsymbol{T}}\boldsymbol{f}_2^+. \end{aligned} \tag{11}$$

In (11), $\hat{\boldsymbol{R}}$ and $\hat{\boldsymbol{T}}$ denote reflection and transmission scattering matrices that characterise the diffraction properties of the cylinder grating layer. These are calculated using a variant of the earlier multipole method (Botten *et al.*, 2000). In this case, the periodicity of a single layer relates the multipole coefficients in each supercell by a quasi-periodicity condition, reducing the solution of the problem to the calculation of multipole source coefficients in only the

primary supercell. The structure of the previous analysis is preserved, with the Rayleigh identity for each of the N_c cylinders in the primary supercell, the analogue of (4), being

$$A_m^l = \sum_{q=1}^{N_c} \sum_{p=-\infty}^{\infty} S_{mp}^{lq} B_p^q + K_m^l, \tag{12}$$

where

$$\begin{array}{rcl} S_{mp}^{ll} & = & \sum_{j\neq 0} H_{m-p}^{(1)}(k|j|D) e^{-i(m-p)\arg j} e^{i\alpha_0 jD}, \\ S_{mp}^{lq} & = & \sum_j H_{m-p}^{(1)}(k(|jD\hat{\boldsymbol{x}} + \boldsymbol{c}_l - \boldsymbol{c}_q|) e^{-i(m-p)\arg(jD\hat{\boldsymbol{x}}+\boldsymbol{c}_l-\boldsymbol{c}_q)} e^{i\alpha_0 jD} \end{array} \tag{13}$$

denote lattice sums (Botten *et al.*, 2000) and $\hat{\boldsymbol{x}}$ is a unit vector in the direction of the supercell axis. As before, the term K_m^l in (12) denotes cylindrical harmonic coefficients of the exterior driving field, due here to incident plane wave trains from above and below. Solving, as before, for the source coefficient B_m^l, we can then generate the outgoing plane wave fields that are radiated by the grating. This in turn leads to the plane wave reflection ($\boldsymbol{R}_s$) and transmission ($\boldsymbol{T}_s$) scattering matrices for medium M_s which, for this up-down symmetric structure, are defined by the reflected ($\hat{\boldsymbol{r}} = \hat{\boldsymbol{R}}_s\hat{\boldsymbol{\delta}}$) and transmitted ($\hat{\boldsymbol{t}} = \hat{\boldsymbol{T}}\hat{\boldsymbol{\delta}}$) fields due to the incident plane wave field $\hat{\boldsymbol{\delta}}$. Here,

$$\hat{\boldsymbol{R}}_s = \frac{2}{D}\boldsymbol{\chi}^{-1/2}\boldsymbol{K}_+ \left(\boldsymbol{I} - \tilde{\boldsymbol{R}}\boldsymbol{S}\right)^{-1} \tilde{\boldsymbol{R}}\boldsymbol{J}_+\boldsymbol{\chi}^{-1/2}, \tag{14}$$

$$\hat{\boldsymbol{T}}_s = \boldsymbol{I} + \frac{2}{D}\boldsymbol{\chi}^{-1/2}\boldsymbol{K}_- \left(\boldsymbol{I} - \tilde{\boldsymbol{R}}\boldsymbol{S}\right)^{-1} \tilde{\boldsymbol{R}}\boldsymbol{J}_+\boldsymbol{\chi}^{-1/2}, \tag{15}$$

where $\boldsymbol{\chi}^{-1/2} = \mathrm{diag}(\chi_p)$. In (14) and (15), the multipole scattering operator $\left(\boldsymbol{I} - \tilde{\boldsymbol{R}}\boldsymbol{S}\right)^{-1}$ is again evident. The terms $\boldsymbol{J}_\pm$ and $\boldsymbol{K}_\pm$ are transformations that perform changes of basis, respectively from plane waves to cylindrical harmonics and *vice versa*.

Solving (11) for $\boldsymbol{f}_2^-$ and $\boldsymbol{f}_2^+$, we obtain an interlayer translation operator $\boldsymbol{\mathcal{T}}$ from which follows the eigenvalue equation that determines the Bloch modes. In particular,

$$\begin{bmatrix} \boldsymbol{f}_2^- \\ \boldsymbol{f}_2^+ \end{bmatrix} = \boldsymbol{\mathcal{T}} \begin{bmatrix} \boldsymbol{f}_1^- \\ \boldsymbol{f}_1^+ \end{bmatrix} = \mu \begin{bmatrix} \boldsymbol{f}_1^- \\ \boldsymbol{f}_1^+ \end{bmatrix} \tag{16}$$

where

$$\boldsymbol{\mathcal{T}} = \begin{bmatrix} \hat{\boldsymbol{T}}_s - \hat{\boldsymbol{R}}_s\hat{\boldsymbol{T}}_s^{-1}\hat{\boldsymbol{R}}_s & \hat{\boldsymbol{R}}_s\hat{\boldsymbol{T}}_s^{-1} \\ -\hat{\boldsymbol{T}}_s^{-1}\hat{\boldsymbol{R}}_s & \hat{\boldsymbol{T}}_s^{-1} \end{bmatrix}. \tag{17}$$

The matrices in (17) have infinite dimension and must be truncated to permit the computational solution. While (16) formally defines the modes, the problem involving $\boldsymbol{\mathcal{T}}$ is not well conditioned, particularly for matrices of large dimension, and thus an alternative method (Smith, Botten, McPhedran and Nicorovici, 2002) is needed to compute the modes.

It can be shown that Bloch modes are paired, respectively corresponding to forward and backward propagation. Only a finite number of modes carry energy through the structure, while the remainder (formally, an unbounded set) carry no energy and are said to be evanescent, decaying exponentially with propagation distance. Modes which carry energy have eigenvalues $|\mu| = 1$ while evanescent modes have either $|\mu| < 1$ or $|\mu| > 1$, respectively

corresponding to forward and backward propagation. For propagating states that carry energy through the structure, their direction of propagation is classified by the direction of energy flow. In this way, we can diagonalise the operator $\mathcal{T}$. The following expression is a simplified form applicable to vertically symmetric structures (such as square or rectangular lattices):

$$\mathcal{T} = \begin{bmatrix} \boldsymbol{F}^- & \boldsymbol{F}^+ \\ \boldsymbol{F}^+ & \boldsymbol{F}^- \end{bmatrix} \begin{bmatrix} \boldsymbol{\Lambda} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{\Lambda}^{-1} \end{bmatrix} \begin{bmatrix} \boldsymbol{F}^- & \boldsymbol{F}^+ \\ \boldsymbol{F}^+ & \boldsymbol{F}^- \end{bmatrix}^{-1}, \tag{18}$$

where the matrices $\boldsymbol{F}^{\mp}$ comprise columns of eigenvector components $\boldsymbol{f}_j^{\mp}$ associated with forward propagating modes with eigenvalues $\boldsymbol{\Lambda} = \text{diag}\,\{\mu_j\}$. In (18), the left and right columns of the partitions respectively refer to the forward and backward propagation.

With the modes defined, the fields at the n^{th} layer boundary in any finite slab of L layers can be expressed as a combination of forward/backward states

$$\begin{bmatrix} \boldsymbol{f}^-(n) \\ \boldsymbol{f}^+(n) \end{bmatrix} = \begin{bmatrix} \boldsymbol{F}^- \\ \boldsymbol{F}^+ \end{bmatrix} \boldsymbol{\Lambda}^n \boldsymbol{c}^- + \begin{bmatrix} \boldsymbol{F}^+ \\ \boldsymbol{F}^- \end{bmatrix} \boldsymbol{\Lambda}^{L-n} \boldsymbol{c}^+, \tag{19}$$

In (19), $\boldsymbol{\Lambda}^n$ and $\boldsymbol{\Lambda}^{L-n}$ characterise propagation of the mode field through the layer. The modes are normalised to carry unit flux and thus $\mathbf{c}^{\mp}$ represent vectors of modal amplitudes. The same applies for semi-infinite media such as M_1 and M_3 of Fig 1(a). In the case of M_1, there in an incident, forward propagating field characterised by the vector $\boldsymbol{\delta}$ of coefficients of mode amplitudes, and a corresponding reflected, backward propagating field $\boldsymbol{r}$. In the case of M_3, the field is represented solely by the forward propagating, transmitted field vector $\boldsymbol{t}$.

Observe that in (19), it is the diagonal matrices $\boldsymbol{\Lambda}^n$ and $\boldsymbol{\Lambda}^{L-n}$ that characterise the propagation of the mode field through the layer. The interfacing of different media (say M_i and M_j) requires application of continuity conditions at the boundaries and is best represented in terms of Bloch mode interface reflection and transmission matrices which are the matrix analogues of the familiar, scalar Fresnel coefficients in optics. To derive their form, we consider a modal field of amplitude $\boldsymbol{c}_-^i$ incident from within a semi-infinite region M_i giving rise to a reflected field $\boldsymbol{c}_i^+$ in M_i and a transmitted field $\boldsymbol{c}_j^-$ in M_j:

$$\begin{pmatrix} \boldsymbol{F}_i^- \\ \boldsymbol{F}_i^+ \end{pmatrix} \boldsymbol{c}_i^- + \begin{pmatrix} \boldsymbol{F}_i^+ \\ \boldsymbol{F}_i^- \end{pmatrix} \boldsymbol{c}_i^+ = \begin{pmatrix} \boldsymbol{F}_j^- \\ \boldsymbol{F}_j^+ \end{pmatrix} \boldsymbol{c}_j^-. \tag{20}$$

Solving these equations, we deduce the generalised Fresnel matrices $\boldsymbol{R}_{ij}$ and $\boldsymbol{T}_{ij}$, respectively defined by $\boldsymbol{c}_i^+ = \boldsymbol{R}_{ij}\boldsymbol{c}_i^-$ and $\boldsymbol{c}_j^- = \boldsymbol{T}_{ij}\boldsymbol{c}_i^-$, are given by

$$\boldsymbol{R}_{ij} = \left(\boldsymbol{F}_i^-\right)^{-1} \left(\boldsymbol{I} - \boldsymbol{R}_j\boldsymbol{R}_i\right)^{-1} \left(\boldsymbol{R}_j - \boldsymbol{R}_i\right) \boldsymbol{F}_i^- \tag{21}$$

$$\boldsymbol{T}_{ij} = \left(\boldsymbol{F}_j^-\right)^{-1} \left(\boldsymbol{I} - \boldsymbol{R}_i\boldsymbol{R}_j\right)^{-1} \left(\boldsymbol{I} - \boldsymbol{R}_i^2\right) \boldsymbol{F}_i^- \tag{22}$$

In (21) and (22), $\boldsymbol{R}_i = \boldsymbol{F}_i^+ \left(\boldsymbol{F}_i^-\right)^{-1}$ denotes the reflection scattering matrix of a semi-infinite crystal of material M_i from free space.

The reflection and transmission matrices $\boldsymbol{R}$ and $\boldsymbol{T}$ of the entire structure of Fig. 1(a), defined by $\boldsymbol{r} = \boldsymbol{R}\boldsymbol{\delta}$ and $\boldsymbol{t} = \boldsymbol{T}\boldsymbol{\delta}$, may now be deduced from the solution of:

$$\boldsymbol{r} = \boldsymbol{R}_{12}\boldsymbol{\delta} + \boldsymbol{T}_{21}\boldsymbol{\Lambda}^L\boldsymbol{c}_2^+ \quad (23) \qquad \boldsymbol{c}_2^+ = \boldsymbol{R}_{23}\boldsymbol{\Lambda}^L\boldsymbol{c}_2^-, \quad (25)$$

$$\boldsymbol{c}_2^- = \boldsymbol{T}_{12}\boldsymbol{\delta} + \boldsymbol{R}_{21}\boldsymbol{\Lambda}^L\boldsymbol{c}_2^+, \quad (24) \qquad \boldsymbol{t} = \boldsymbol{T}_{23}\boldsymbol{\Lambda}^L\boldsymbol{c}_2^-. \quad (26)$$

Eq. (23), expresses the reflected field $\boldsymbol{r}$ as the reflection ($\boldsymbol{R}_{21}$) of the incident field $\boldsymbol{\delta}$ and the transmission ($\boldsymbol{T}_{21}$) of the upward propagating Bloch modes of region M_2 at the M_1-M_2 boundary. Similarly, (24) expresses the downward propagating waves in M_2 at the M_1-M_2 boundary as a transmission of the incident field and a reflection of the upward propagating modes in M_2. Eqns. (25) and (26) fulfil the same roles for the M_2-M_3 boundary, although in the absence of an incident field from M_3. This ultimately leads to the Bloch mode reflection and transmission matrices for the composite structure:

$$\boldsymbol{R} = \boldsymbol{R}_{12} + \boldsymbol{T}_{21}\boldsymbol{\Lambda}^L\boldsymbol{R}_{23}\boldsymbol{\Lambda}^L\left(\boldsymbol{I} - \boldsymbol{R}_{21}\boldsymbol{\Lambda}^L\boldsymbol{R}_{23}\boldsymbol{\Lambda}^L\right)^{-1}\boldsymbol{T}_{12}, \quad (27)$$

$$\boldsymbol{T} = \boldsymbol{T}_{23}\boldsymbol{\Lambda}^L\left(\boldsymbol{I} - \boldsymbol{R}_{21}\boldsymbol{\Lambda}^L\boldsymbol{R}_{23}\boldsymbol{\Lambda}^L\right)^{-1}\boldsymbol{T}_{12}. \quad (28)$$

In both cases, there is a striking resemblance to the scalar Airy formulation for a three layer Fabry-Perot interferometer. We conclude by observing that for more complex structures (with more layers), the treatment above may be extended by recurrence.

A question of interest in the design of PC devices is the issue of their tolerance to fabrication defects. A statistical analysis of their performance requires that ensembles of realisations of the properties of the exact structure in Fig. 1(a) be analysed. In these, the layers in say M_2 are not identical and so the analysis requires some modification. Instead, we compute matrices $\boldsymbol{R}_2$, $\boldsymbol{R}_2'$, $\boldsymbol{T}_2$, $\boldsymbol{T}_2'$, respectively denoting the free space reflection and transmission matrices for a composite M_2 structure from above and below. Matching fields at the M_1-M_2 and M_2-M_3 interfaces, we derive the aggregate reflection and transmission matrices $\boldsymbol{R}$ and $\boldsymbol{T}$:

$$\boldsymbol{R} = \left(\boldsymbol{F}_1^-\right)^{-1}\left(\boldsymbol{I} - \boldsymbol{R}_{23}\boldsymbol{R}_1\right)^{-1}\left(\boldsymbol{R}_{23} - \boldsymbol{R}_1\right)\boldsymbol{F}_1^-, \quad (29)$$

$$\boldsymbol{T} = \left(\boldsymbol{F}_3^-\right)^{-1}\left(\boldsymbol{I} - \boldsymbol{R}_2\boldsymbol{R}_3\right)^{-1}\boldsymbol{T}_2\left(\boldsymbol{I} - \boldsymbol{R}_1\boldsymbol{R}_{23}\right)^{-1}\left(\boldsymbol{I} - \boldsymbol{R}_1^2\right)\boldsymbol{F}_1^-, \quad (30)$$

$$\boldsymbol{R}_{23} = \boldsymbol{R}_2 + \boldsymbol{T}_2'\boldsymbol{R}_3\left(\boldsymbol{I} - \boldsymbol{R}_2'\boldsymbol{R}_3\right)^{-1}\boldsymbol{T}_2. \quad (31)$$

2.3 Photon Conductance

The model of photon conductance g follows from the definitions for electron conductance developed by Landauer (Landauer, 1957) in which the conductance is directly proportional to the transmittance. We consider the single guide of Fig. 1(b), regarding M_1 as free space and M_2 as either a finite slab bounded by free space M_3, or a semi-infinite medium. The conductance is defined as the total transmission in all propagating output channels caused by unit input in all propagating input channels. Denoting by $\boldsymbol{T} = \left[t_{pq}^{ij}\right]$, the scattering matrix of the structure with incidence in M_i and output in M_j, with p and q respectively denoting output and input channels, it follows that

$$g = \sum_{pq}\left|t_{pq}^{ij}\right|^2. \quad (32)$$

In the case of a slab of $L = 6$ layers, Fig. 2 plots the variation of conductance with increasing width, revealing clear "staircase" behaviour which can be understood easily with a modal

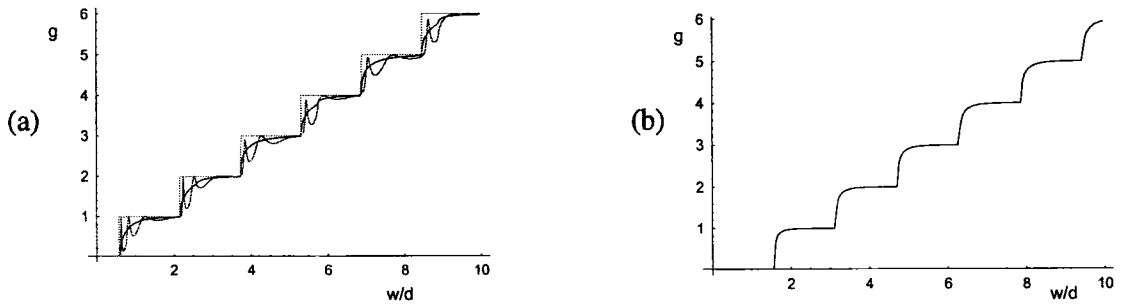

Figure 2: Conductance g versus the guide width w/d, normalised to the lattice constant d. (a) For a PC guide at $\lambda/d = 3.15$. Green curve: the number of propagating modes; Red curve: conductance for a $L = 6$ layer slab; Blue curve: conductance of a semi-infinite PC guide. (b) For a perfectly conducting semi-infinite metal guide for $\lambda/d = 3.15$.

analysis (Botten *et al.*, 2003) similar to that in §2.2. Additional conducting modes, which carry energy through the structure, are introduced with increasing guide width, with the steps in the ladder corresponding to the threshold of each new mode. Oscillations in the steps are easily understood in terms of Fabry–Perot resonances between the front and back interfaces of the slab. The analysis is further enhanced by replacing the slab by a semi-infinite guide, and exploiting the Reciprocity Theorem to write

$$g = \sum_{qp} |t_{qp}^{ji}|^2 \approx \sum_{m} (1 - |\rho_m|^2). \tag{33}$$

The approximate form follows from an asymptotic analysis (Botten *et al.*, 2003) and characterises accurately the staircase behaviour as the sum over all propagating modes m of the transmittance $1 - |\rho_m|^2$ of mode m where $\rho_m \approx \left[\left(\boldsymbol{F}_2^-\right)^{-1} \boldsymbol{R}_2 \boldsymbol{F}_2^-\right]_{mm}$ and medium M_2 denotes the semi-infinite guide. The quantity ρ_m is the effective reflection coefficient experienced by mode m at the interface of the guide with free space. This diminishes rapidly with increasing guide width, allowing the conductance of each mode to saturate rapidly and carry unit energy (Fig. 2(b)).

3 COMPUTATIONAL IMPLEMENTATION

Our initial computational studies of individual finite PC clusters relied on the multipole method of §2.1, encoded in Fortran and executed sequentially. While adequate for small crystals, the code relies on the solution of a dense system of linear equations obtained by truncating the series expansion in (4). However, since the size of the Rayleigh matrix is proportional to the square of the product of N_c and the number of terms in the truncated series, the method requires substantial memory and execution time for large clusters (ie of order $N_c > 200$).

Later studies focussed on the effects of disorder (in the refractive indices, radii or positions of the cylinders) on a cluster. These investigations relied on a Monte Carlo approach in which an ensemble of sample perturbations were applied to a crystal specification, thus generating a

set of realisations, each of which could be treated with the multipole method independently. This approach necessarily placed further demands on execution time. Parallel implementations of the code have been developed in response to these limitations.

Initially, a parallel code for the analysis of a single realisation was developed using OpenMP. The source language used was Fortran90 (MIPSpro 7 Fortran compiler). The code was implemented on the SGI Origin 2400 machine operated by *ac3*. This is a 64-processor shared-memory machine with 32 Gbyte of non-uniform memory access RAM. The cpus are 400MHz R12000 processors, each with a peak speed of about 800 Mflops. The code itself was a direct port of the sequential code, with OpenMP directives used to parallelise the loops responsible for the generation of the Rayleigh matrix, and the reconstruction of the electric field. A full inversion of the Rayleigh matrix was performed by routines imported from the reference version of LAPack (Anderson, Bai, Bischof, Blackford, Demmel, Dongarra, Du Croz, Greeenbaum, Hammarling, McKenney and Sorenson, 1999).

Subsequently, this code was used as the basis for the slave component of a master-slave style program which implemented the Monte Carlo analysis of an ensemble of perturbed PC realisations. The master-slave code was developed using MPI to encapsulate master and slave components as separate processes. The master process was responsible for: creating the perturbed realisations comprising the ensemble by adding randomised perturbations to the parameters of an underlying ideal crystal; farming these out to separate slave processes for analysis; and performing the statistical analysis of the results returned by the slaves. Each slave process could operate either sequentiallly or as a parallel process itself (via OpenMP). In fact, multiple threads were used in each slave to reduce elapsed run time during studies of large crystals. The master process was also responsible for checkpointing, as well as automatic termination after checkpointing and resubmission of the partially completed job to the batch queue.

Early performance data for these codes (on the SGI Origin 2400) was reported in Langtry, Botten, Asatryan and McPhedran (2001). The initial test problems involved clusters of 136 cylinders, and the series expansion in (2) was truncated after 11 terms. The code later successfully handled up to 526 cylinders with series truncation after 11 terms. For larger problems, stack frame buffer overflows occurred. Thus for larger clusters it was decided to port the code to the APAC AlphaserverSC, a distributed system comprising 127 compute nodes, each with 4 Alpha 21264C processors and between 4 and 16 Gbytes of RAM, shared by the four local processors. The architecture of this system is quite different from that of the SGI Origin, and thus the code required some reconfiguration. In particular, since memory can be shared by at most four processors, slaves could employ at most four threads when running with OpenMP. However, system calls used by the master process for resubmitting a partially completed job proved to be incompatible with the combined use of OpenMP and MPI. Consequently the code was recompiled without OpenMP, leaving each slave running on just one processor and communicating with the master process via MPI.

Production runs of the code on both the SGI Origin and the AlphaserverSc systems have provided significant results (Langtry *et al.*, in press). These have in turn suggested further extensions to the code—in particular, the implementation of the Bloch mode method described in §2.2. However, rather than further increasing the complexity of the code, the program was partitioned into a suite of separate programs which can be executed independently.

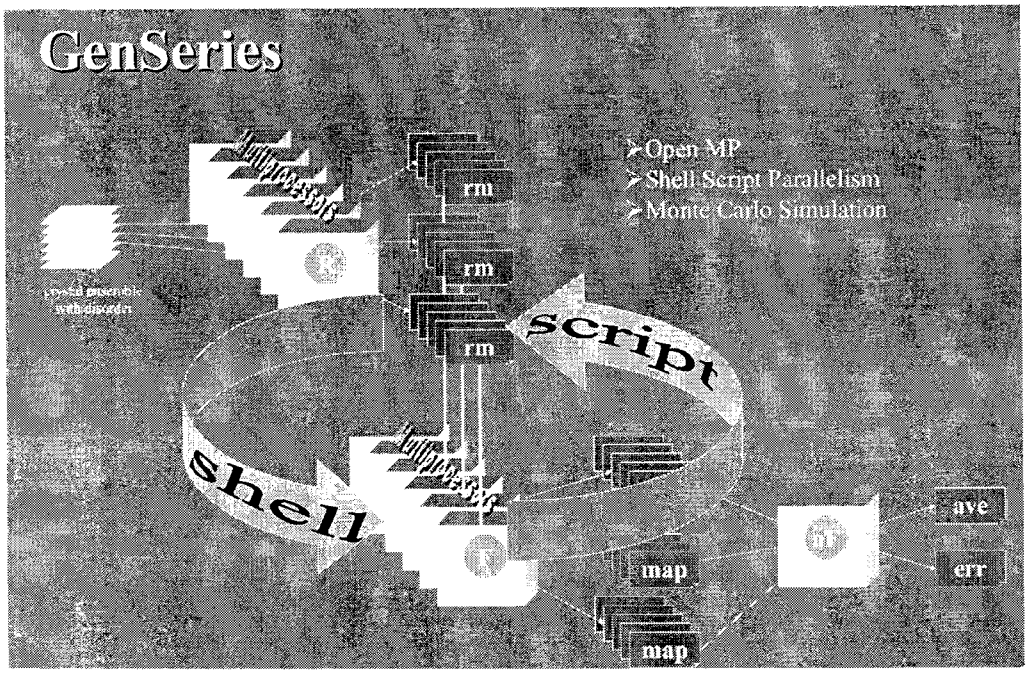

Figure 3: Conceptual design of GenSeries suite: shell scripts govern the processing in parallel of an ensemble of crystal realisations, producing multiple individual map files that are then post-processed by a separate program yielding the statistical analysis.

In its present form, the principal components of this suite (collectively referred to as 'genSeries2'; see Fig. 3) are genCrystalSpec, responsible for generating a crystal realisation with a specified randomisation of one or more defining characteristics of the crystal; genFieldMap, responsible for computing the Rayleigh matrix for a crystal based on its geometry, physical properties and wavelength of radiation, and then computing the property of interest (*eg* local density of states, field intensity) over a two-dimensional grid of specified resolution; and genFieldAveMap, responsible for computing the average of the property over the specified grid along with a statistical measure of the convergence of the Monte Carlo process. Data is communicated between these programs via the use of files, as illustrated in Fig. 4.

The new suite offers a number of advantages over the initial parallel version of the code:

1. The separation of the code into independent program units provides a more loosely coupled system, simplifying the process of extending the functionality of the system.
2. Having the analysis of individual realisations performed by separate instances of the one program allows the task farming aspects of the system to be implemented by shell scripts rather than by MPI code, thus utilising a simpler model of parallel execution.
3. Checkpointing is greatly simplified, particularly for runs analysing large crystals and requiring high resolution in the output field maps.

Site	Version of code	No. cpus	Reported Real time (hours)	Reported user time (hours)	Reported cpu utilisation (%)	Speedup factor
ac3 SGI	1	1	35.24	35.22	100	1
ac3 SGI	1	2	20.11	39.49	197	1.75
ac3 SGI	1	4	23.01	68.52	298	1.53
ac3 SGI	2	4	10.31	40.04	389	
ac3 SGI	3	4	9.20	36.71	399	
APAC-SC	3	1	10.32	10.28	99	1
APAC-SC	3	2	9.02	17.97	199	1.14
APAC-SC	3	4	4.02	15.95	387	2.57
APAC-SC	4	1	1.45	1.73	98	1
APAC-SC	4	2	1.12	2.22	197	1.29
APAC-SC	4	4	0.72	2.62	386	2.01

Table 1: Performance data for versions of the genSeries2 code when executed on the ac3 SGI Origin2400 system and the APAC Compaq Alphaserver-SC system.

4. SMP parallelism can once again be employed via multiple threads and OpenMP, as illustrated in Fig. 5, in order to reduce elapsed run time and take advantage of the customised parallel library routines provided by the proprietary version of LAPack.

Table 1 presents a comparison of the performance of different versions of genSeries2 on the same problem: a study of the intensity of the electromagnetic field due to a line source adjacent to the entry to a straight waveguide embedded in a finite cluster of 252 cylinders. All versions use OpenMP to implement loop parallelism in previously sequential code. In version 1, the calculation of the output field points for an individual realisation was carried out by computing the full inverse of the Rayleigh matrix (9), computing the vector of coefficients $\boldsymbol{B}$ and then using this to compute the field value. Version 1 used the reference (source) version of LAPack to compute the full inverse of the Rayleigh matrix. Version 2 simply replaced this routine by linking in the proprietary version of the same LAPack routine from the vendor-supplied scientific library. In version 3 the full inversion of the Rayleigh matrix is replaced by an LU factorisation of the matrix, again implemented by linking in the appropriate routine from the vendor-supplied scientific library. In version 4 a user-coded routine for performing matrix-vector multiplication used in earlier versions was replaced by the proprietary version of the appropriate BLAS level 2 routine. Program runs were undertaken on the ac3 system (SGI Origin 2400 system with 64 400MHz R12000 processors and 32Gb of RAM) and on the APAC-SC system. The speedup factor quoted is the quotient of the reported real time for the job and the corresponding reported real time for the sequential job described in the first row of each part of the table. For purposes of comparison of performance, individual runs were restricted to four cpus since this is the maximum number available to OpenMP on the APAC-SC system. Note the apparent discrepancy in the third row, where a 4 cpu job appears to have cpu usage corresponding to only 3 cpus. Similar benchmarking jobs are currently being undertaken on the APAC Alphaserver-SC and Linux cluster systems.

The data in Table 1 clearly indicates that the most significant improvements in reported

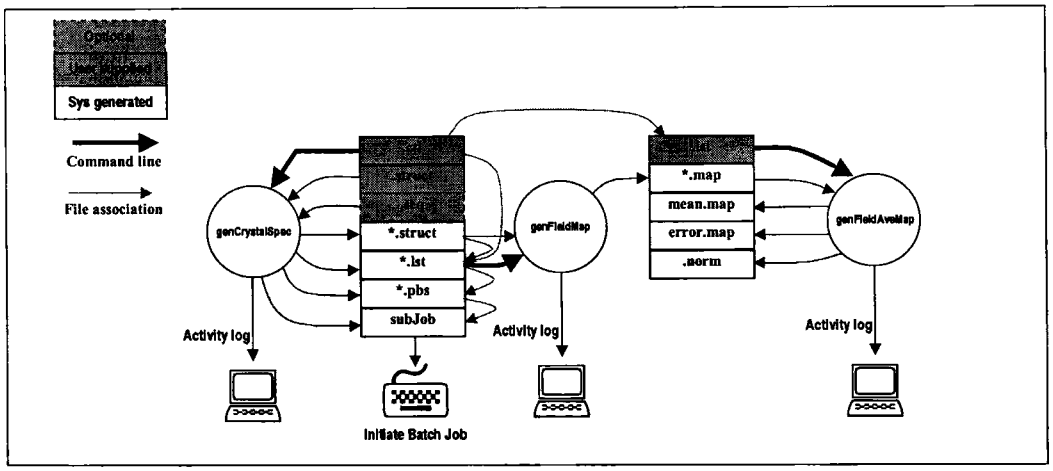

Figure 4: Architectural overview of GenSeries2 suite. Data flows between programs: rectangular elements represent files.

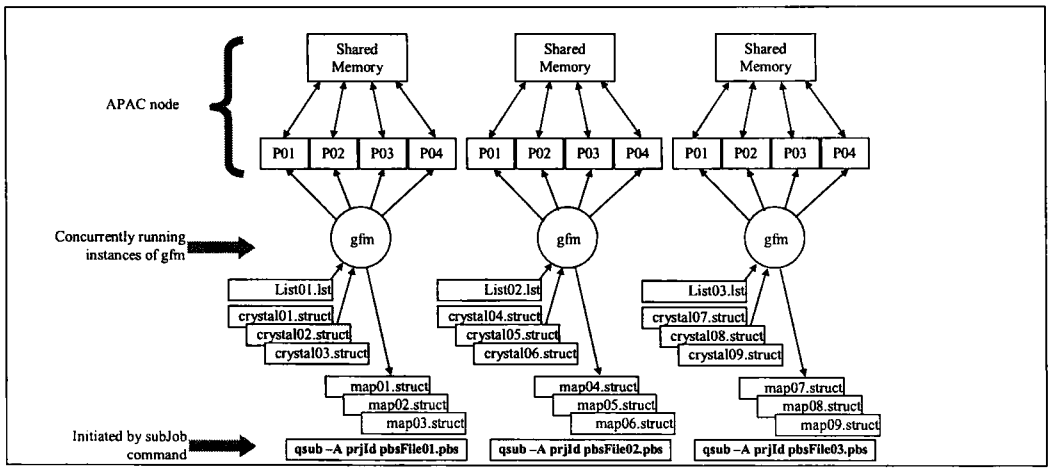

Figure 5: Run-time overview of genFieldMap (gfm) component of the GenSeries2 suite for shared memory processor (SMP) nodes on the APAC AlphaServer.

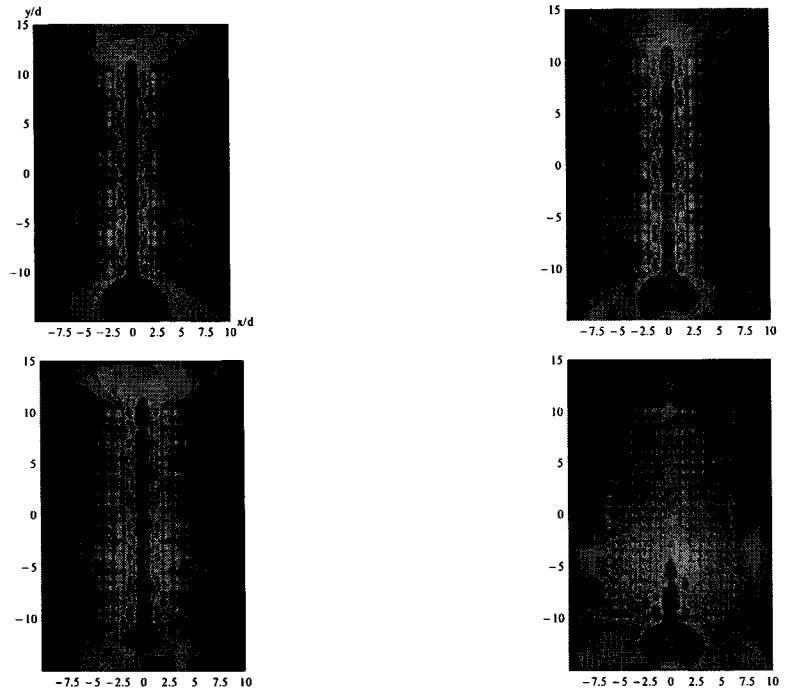

Figure 6: Electric field intensity in a PC waveguide subject to refractive index disorder: (from left to right): 0%, 10%, 20%, 30%.

execution time derive from the replacement of reference version LAPack routines and user-coded routines by vendor-customised proprietary routines. As an indication of the significance of these changes, we note that speedup factors due to increasing the number of cpus available to a given version of the code lie in the range from 1 to 2.57. However, comparing the reported real times for version 3 of the code with 1, 2 and 4 cpus to corresponding times for version 4 yields improvements in these times by factors of 7.21, 8.05 and 5.58, respectively.

4 RESULTS

We begin with an investigation of the effects of disorder in square arrays of dielectric cylinders using the codes described in §3. Our results indicate substantial robustness of the PC bandgap effect and of localised modes when the cylinder positions, radii and refractive indices are perturbed. This is illustrated in Fig. 6 by plots of electric field intensity subject to refractive index disorder in a guide of length $20d$ embedded in a cluster of 252 cylinders of radius $0.3d$ and average index 3.0, with disorder in the cylinder refractive indices given by $n_l = n + \xi$, where ξ is a random variable uniformly distributed in the range $[-Q, Q]$. More succinctly, the level of disorder is defined as Q/n and expressed as a percentage. The guide is excited by a line source parallel to cylinder axes radiating at wavelength $3.3d$.

Fig. 7 plots the normalised local density of states versus wavelength in the middle of the central cylinder of a circular cluster of 317 cylinders of normalised radius $a/d = 0.3$ and

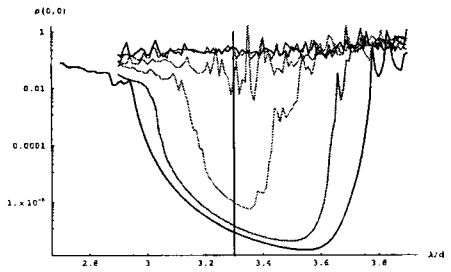

Figure 7: Local density of states versus wavelength in the middle of the central cylinder of a circular cluster subject to refractive index disorder of 0% (bottom curve), 10%, 20%, 30%, 40%, 50% (top).

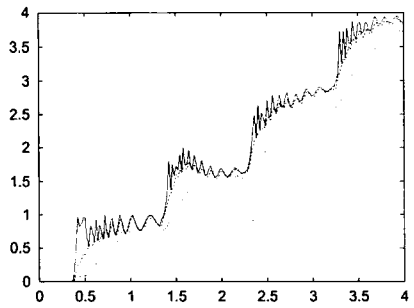

Figure 8: Conductance of a disordered waveguide. As for Fig.2, but with disorder Q=0.0, Q=0.2,Q=0.4,Q=0.8 from top to bottom.

index 3.0. Cylinders are subjected to increasing levels of disorder: 0% (bottom curve), 10%, 20%, 30%, 40%, 50% (top). We note that the band gap is still observable at 20%, but has effectively disappeared at 30%—reflected in the behaviour of the intensity plots in Fig 6.

Fig. 8 displays the conductance of a waveguide subject to disorder in the cylinder refractive indices, given again by $n_l = n + \xi$ where ξ is a random variable uniformly distributed in the range $[-Q, Q]$. It is seen that the ladder behaviour remains observable with disorder of up to 10%, while strong disorder destroys the effect.

Finally, we show in Figs. 9 and 10 field plots for a number of PC devices: in Figs. 9(a) and 9(b) the finite directional coupler (White, Botten, de Sterke and McPhedran, in press) generated using a multipole calculation. Fig. 9(c) displays the transmittance, computed using the modal analysis of §2.2, of such a coupler with an overlap region of $L = 5$ layers. The sharp resonances in the transmission spectrum that are evident may be tuned in various ways to produce ultra-compact devices with very high performance specifications. Fig. 10 shows a field plot for a PC interferometer, also computed using the modal analysis of §2.2.

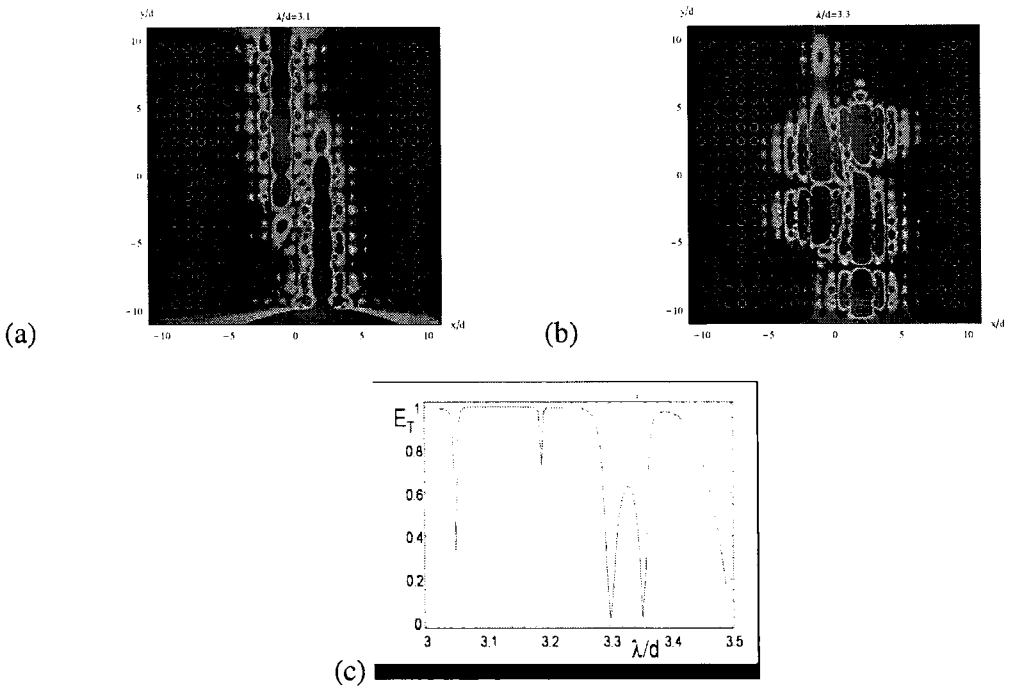

Figure 9: Directional coupler: ideal at (a) $\lambda = 3.1d$ and (b) $\lambda = 3.3d$. The plot in (c) describes transmittance as a function of wavelength.

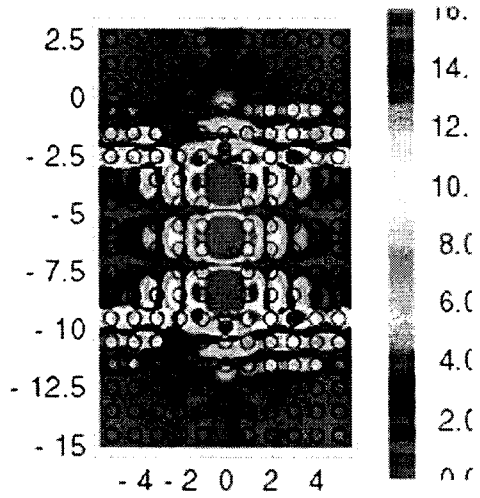

Figure 10: Field plot of photonic crystal interferometer at resonance.

5 CONCLUDING REMARKS

We have described both the theoretical basis and the development of a parallel suite for the analysis of photonic crystals and finite clusters of cylinders. The suite provides a set of tools that may facilitate the analysis and design of certain types of miniaturised PC devices, and we have illustrated the application of these tools to the study of a number of examples. Issues of implementation of the code on different platforms have been considered: of particular interest are the effects on performance of the use of scientific libraries such as LAPack that have been customised for the machine architecture.

Further development of the code is currently under way. We are particularly interested in evaluating its performance, for different classes of problems, on the new cluster systems installed at both ac3 and APAC. We are also in the process of extending the functionality of the code to handle a wider range of design problems, including studies with different types of light sources, such as beams with specified profiles. In future work we hope to extend the code to handle certain types of three-dimensional crystals.

Acknowledgements

This work was produced with the assistance of the Australian Research Council under the ARC Centres of Excellence program. CUDOS (the Centre for Ultrahigh bandwidth Devices for Optical Systems) is an ARC Centre of Excellence. The authors also acknowledge the support of ac3 and APAC for the provision of parallel computing facilities.

References

[1] ANDERSON, E., BAI, Z., BISCHOF, C., BLACKFORD, S., DEMMEL, J., DONGARRA, J., DU CROZ, J., GREEENBAUM, A., HAMMARLING, S., MCKENNEY, A., and SORENSON, D. (1999): *LAPACK User's Guide (3rd edn).* SIAM, Philadelphia.

[2] ASATRYAN, A., BUSCH, K., MCPHEDRAN, R. C., BOTTEN, L. C., DE STERKE, C. M., and NICOROVICI, N. A. (2001): Two-dimensional Green's function and local density of states in photonic crystals consisting of a finite number of cylinders of infinite length. *Phys. Rev. E*, 63(4):046612.

[3] BOTTEN, L. C., ASATRYAN, A. A., LANGTRY, T. N., WHITE, T. P., DE STERKE, C. M., and MCPHEDRAN, R. C. (2003): Semi-analytic treatment for propagation in finite photonic crystal waveguides. *Opt. Lett.*, 28:851–853.

[4] BOTTEN, L. C., NICOROVICI, N. A., ASATRYAN, A. A., MCPHEDRAN, R. C., DE STERKE, C. M., and ROBINSON, P. A. (2000): Formulation for electromagnetic scattering and propagation through grating stacks of metallic and dielectric cylinders for photonic crystal calculations. part i. method. *J. Opt. Soc. Am. A*, 17:2165–2176.

[5] BOTTEN, L. C., NICOROVICI, N. A., MCPHEDRAN, R. C., DE STERKE, C. M., and ASATRYAN, A. A. (2001): Photonic band structure calculations using scattering matrices. *Phys. Rev. E*, 64:046603.

[6] JOANNOPOULOS, J. D., MAEDE, R. D., and WINN, J. N. (1995): *Photonic Crystals*. Princeton University Press, New Jersey.

[7] JOHN, S. (1987): Strong localization of photons in certain disordered dielectric superlattices. *Phys. Rev. Lett.*, 58:2486.

[8] LANDAUER, R. (1957): Spatial variation of currents and fields due to localized scatterers in metallic conduction. *IBM J. Res.. Develop.*, 1:223–231.

[9] LANGTRY, T. N., BOTTEN, L. C., ASATRYAN, A. A., and MCPHEDRAN, R. C. (2001): Local density of states calculations for photonic crystals. In *Proceedings of the 5th International Conference on High Performance Computing in the Asia-Pacific Region.*

[10] LANGTRY, T. N., ASATRYAN, A. A., BOTTEN, L. C., DE STERKE, C. M., MCPHEDRAN, R. C. and ROBINSON, P. A.: Effects of disorder in two-dimensional photonic crystal waveguides. *Phys. Rev. E*, in press.

[11] LIN, S.-Y., CHOW, E., HIETALA, V., VILLENEUVE, P. R., and JOANNOPOULOS, J. D. (1998): Experimental demonstration of guiding and bending of electromagnetic waves in a photonic crystal. *Science*, 282:274–276.

[12] MEKIS, A., CHEN, J. C., KURLAND, I., FAN, S., VILLENEUVE, P. R., and JOANNOPOULOS, J. D. (1996): High transmission through sharp bends in photonic crystal waveguides. *Phys. Rev. Lett.*, 77:3787.

[13] PARKER, G. and CHARLTON, M. (2000): Photonic crystals. *Physics World*, 13(8):29–34.

[14] SMITH, G. H., BOTTEN, L. C., MCPHEDRAN, R. C., and NICOROVICI, N. A. (2001): Cylinder gratings in conical incidence with applications to modes of air-cored photonic crystal fibers. *Phys. Rev. E*, 66(5, art. no. 056604):46612–1 – 46612–4.

[15] SOUKOULIS, C. M., editor (1996): *Photonic Band Gap Materials*, volume 315 of *NATO ASI Series*. Kluwer Academic Publishers, Dordrecht, the Netherlands.

[16] WHITE, T. P., BOTTEN, L. C., DE STERKE, C. M., and MCPHEDRAN, R. C.: Ultra-compact resonant filters in photonic crystals. *Opt. Lett.*, in press.

[17] YABLONOVITCH, E. (1987): Inhibited spontaneous emission in solid state physics and electronics. *Phys. Rev. Lett.*, 58:2059.