

Balancing accuracy against computation time: 3-D FDTD for nanophotonics device optimization

Geoffrey W. Burr¹ and Ardavan Farjadpour²

¹*IBM Almaden Research Center,
650 Harry Road, San Jose, California 95120*

²*Massachusetts Institute of Technology,
Department of Materials Science & Engineering, Cambridge, MA 02139*

ABSTRACT

The finite-difference time-domain (FDTD) approach is now widely used to simulate the expected performance of photonic crystal, plasmonic, and other nanophotonic devices. Unfortunately, given the computational demands of full 3-D simulations, researchers can seldom bring this modeling tool to bear on more than a few isolated design points. Thus 3-D FDTD—as it stands now—is merely a *verification* rather than a *design optimization* tool.

Over the long term, continuing improvements in available computing power can be expected to bring structures of current interest within general reach. In the meantime, however, many researchers appear to be exploring alternative modeling techniques, trading off flexibility of approach in return for more rapid turnaround on the devices of specific interest to them. In contrast, we are trying to improve the efficiency of 3-D FDTD by reducing its computational expense without sacrificing accuracy. We believe that these two approaches are completely complementary — even with vast amounts of computational power, any real-world system will still require a modular approach to modeling, spanning from the nanometer to the millimeter scale or beyond.

Keywords: photonic crystals; FDTD

1. INTRODUCTION

The advantages of FDTD are numerous [1–3] — its disadvantages can be boiled down to a single statement: in order to be confident you will get the right answer, 3-D FDTD simulations must be large and slow. Large simulations come from the need to discretize at a fine spatial scale, both to follow device features as well as to avoid error (from “numerical dispersion” [1]). The community of researchers using FDTD for microwave modeling has known of these problems for some time. Many of their solutions require near-reinvention of the FDTD formulation, at the cost of simplicity and stability. But some are quite simple, and can be easily applied to or modified for nanophotonics modeling. We have previously shown that by using fairly straightforward sub-cell methods [1–5] and simple corrections for numerical dispersion [1–3, 6], accurate reflectivity coefficients and photonic crystal band-diagrams can be produced at significantly coarser gridding than would otherwise be possible [2, 3].

Slow simulations come from a combination of simulation size (many grid points to be updated every timestep), the small timestep, and the need to run simulations for many optical cycles to attain high frequency resolution. Because the timestep scales with the spatial step, the coarser gridding described in the previous paragraph helps with both of the first two items. However, the frequency resolution of a nanophotonic device’s simulated spectral response will still scale inversely with the time-span of the simulation. And it is this frequency resolution which is of critical importance when using FDTD to study photonic crystal devices such as resonators [7] or waveguides [8].

However, during optimization we may not need the same high-resolution spectrum that one would want for design verification or for publication, as long as the design metric is evaluated accurately leading to the “right” next design iteration. In addition, the *a priori* knowledge that we are searching for an isolated sharp

For further information, contact G. W. Burr at burr@almaden.ibm.com




Figure 1. Dispersion diagram (normalized ω vs. normalized k) for a 2-D square lattice of high-index ($\epsilon=8.9$) cylindrical rods (of radius $r = 0.2a$) in air ($\epsilon=1.0$), for TM polarization only. Insets show the real-space and reciprocal lattices. The frequency is in units of $2\pi c/a$; the spatial frequency is normalized to the symmetry points Γ ($k_x = k_y = 0$), X ($k_x = \pi/a, k_y = 0$), and M ($k_x = k_y = \pi/a$). Two dotted boxes show regions of the dispersion diagram that are discussed in detail by Figures 5 (near the X point), and 6 (near the M point).




Figure 2. Dispersion diagram (TE polarization only) for the 2-D triangular lattice of air holes ($r/a=.48$) in dielectric ($\epsilon=13$), along with the real-space and reciprocal lattices. The unit cell used here is the true primitive unit cell (the rhombus shown in the inset), rather than the larger rectangle which would lead to folded band-structure.

frequency resonance can be used to obtain information about that resonance even in the absence of good frequency resolution. In fact, something like this is already widely done in the analysis of high-Q photonic crystal resonators, where researchers gauge the Q from the power lost by the device rather than from the frequency spectrum [7].

In this paper, we focus on (and try to quantify) the tradeoffs between the accurate determination of resonance frequencies using such shortcuts, and the duration (and size) of the resultant FDTD simulations, for common photonic bandgap structures such as waveguides and resonators. We will then try to use these computationally lighter 3-D FDTD simulations to demonstrate such design optimization on the intrinsic out-of-plane radiation losses in photonic crystal waveguides.

2. NUMERICAL DISPERSION AND SUB-CELL METHODS

The FDTD algorithm is capable of arbitrarily low error in simulating Maxwell's Equations as the grid spacing approaches zero [1]. However, for reasonable choices of grid spacing, there are two main sources of error: numerical dispersion and staircasing error. To demonstrate this, we will use two band-structure diagrams as computed by FDTD and compared to accepted results of the plane-wave method. The TM band-diagram for a 2-D square lattice of dielectric pillars in air is shown in Figure 1. Figure 2 shows the TE band-diagram, real-space lattice, and irreducible Brillouin zone for the triangular lattice of air holes, as calculated by both the plane-wave method [9] and by FDTD.

The reason for the increasing error at high normalized frequencies in the dispersion diagrams calculated by FDTD is that the number of cells per effective wavelength is getting smaller for these frequencies. For instance, even at the uppermost normalized frequency of $\omega a/2\pi c = a/\lambda_0 = 0.8$ from Figure 2, with the grid cell-spacing used for these simulations of 30 cells/ a it would seem that there should be at least 37.5 cells per wavelength. According to Figure 4.4 in Reference [1], this sampling would seem to suffice. However, given the high index of the dielectric material (nearly 3), the number of cells per wavelength *in the material* is much lower (around 13). As a result, numerical dispersion starts to have a noticeable effect already at these frequencies.

To correct for numerical dispersion, we can apply the simple correction suggested in Reference [1], adjusting the permittivity and permeability of free space, ϵ_0 and μ_0 , to force a particular frequency and angle to propagate at exactly the speed of light, c , within the FDTD grid. As suggested in Reference [6], we apply a different adjustment to each axis to allow for non-square/non-cubic grid-spacing. However, in contrast to the procedure

detailed in Reference [6], we choose to solve an index-adjusted version of Equation 4.5 from Reference [1],

$$\left[\frac{n}{c \Delta t} \sin \left(\frac{\pi c \Delta t}{\lambda_0} \right) \right]^2 = \left[\frac{1}{\Delta x} \sin \left(\frac{\pi n \alpha_x \Delta x}{\lambda_0} \right) \right]^2 + \left[\frac{1}{\Delta y} \sin \left(\frac{\pi n \alpha_y \Delta y}{\lambda_0} \right) \right]^2 + \left[\frac{1}{\Delta z} \sin \left(\frac{\pi n \alpha_z \Delta z}{\lambda_0} \right) \right]^2, \quad (1)$$

for the values of $\alpha_{x,y,z}$ in a two-step process. Here for readability, we have written out ω in terms of the vacuum wavelength λ_0 . First, we remove the effects of different grid-spacings by solving Equation 1 for the three cartesian directions, resulting in

$$\alpha_x = \frac{\lambda_0}{\pi n \Delta x} \text{asin} \left[\frac{n \Delta x}{c \Delta t} \sin \left(\frac{\pi c \Delta t}{\lambda_0} \right) \right] \quad (2)$$

and the equivalent for y and z . At this point, the numerical dispersion is corrected such that the speed of light is identically c along any of the 3 coordinate axes, and is faster at all intermediate angles. It is then straightforward to follow the iterative Newton's method [1] to compute the effective phase velocity at any of these intermediate angles, and to then use this incorrect speed of light to further adjust $\alpha_{x,y,z}$ as in Equation 4.69 of Reference [1]. Thus the numerical dispersion can be exactly compensated out at this particular angle (and wavelength). Alternatively, an angle can be picked which minimizes the numerical dispersion over a particular range of nearby angles, or which minimizes the worst-case numerical dispersion.

The improvements due to the compensation of numerical dispersion in the accuracy with which photonic crystal band-diagrams can be computed are not particularly large [3]. This is inherent in the variation of the speed of light within the grid at frequencies and angles that are not close to the set-point that was chosen for compensation in this simple approach.

In contrast with the subtle and imperfect improvements offered by compensating numerical dispersion, we have found that the correction of staircasing errors can offer a significant and noticeable improvement [3]. For instance, in Figure 3 we plot the first two bands of the triangular lattice dispersion diagram at the K point (spatial frequency $k_x = 4\pi/3a, k_y = 0$), as a function of the size of the cylindrical air hole. Four plots are shown: parts (a) and (b) are for the regular FDTD algorithm (e.g., *with* staircasing error) for both moderate (20 cells/ a) and coarse gridding (12 cells/ a); parts (c) and (d) show the same two bands at the same grid-spacings, but when a sub-cell method is used to compute an effective dielectric constant for every E-field component near a material interface (e.g., two such values per cell in these TE simulations). Such a sub-cell method can be thought of as simply translating effective-medium theory for the Yee lattice.

For photonic crystal structures, sub-cell methods have two important benefits: first, the absolute determination of resonance frequencies becomes more accurate, as can be seen by comparing part (a) of Figure 3 with part (c). This arises because staircasing often causes the amount of high- and low-index material to differ from the intended values (an effect which tends to produce oscillation in the error signal as a function of the number of grid-cells). Second, even when there is an error in computing the frequency, as can be seen in one of the two degenerate modes in part (d), we only need a few FDTD simulations to accurately measure trends or derivatives of interest that involve a change of material boundaries. Here, where we may be interested in the change in the resonance frequency of the 2^{nd} band as a function of hole size, we can accurately compute this derivative after only 2–3 extremely coarse FDTD simulations (Figure 3(d)). Without sub-cell methods, at such coarse levels of gridding (part (b)) we should not try to estimate the trend against cylinder radius without at least 15–20 FDTD simulations (at which point, one ought to consider fewer simulations with more grid-points in each). These considerations are somewhat trivial for 2-D band structures, where each FDTD simulation might take at most a minute and even faster alternatives are readily available [9]. However, the tradeoff shown here does not change for large 3-D FDTD simulations. Thus the reduction in the memory size, run-time, and number of required simulations that is offered by such sub-cell methods can enable simulation studies that would otherwise be simply intractable.

The particular sub-cell method used here assigns an effective dielectric constant to each field component based on a volume (area) integration of the grid-cell-sized region around that field component, as shown in




Figure 3. Lowest three bands of the triangular lattice band-diagram from Figure 2 (bands 2 and 3 are degenerate) at the $\mathbf{K}$ point, as a function of the hole radius, r/a . The results of standard FDTD, with no correction for staircasing, are shown for both (a) moderate (20 cells/a) and (b) coarse (12 cells/a) gridding; parts (c) and (d) show the FDTD results when using a sub-cell method taken from the microwave FDTD literature [4] to compute an effective dielectric constant for every field component located near a material interface. Insets show the effective dielectric constant for $r/a \sim 0.40$.




Figure 4. Three FDTD grid cells are shown in the vicinity of a material interface (shaded region). For the two E-field components E_y and E_x in the upper right cell, the cell-sized regions centered about that field component are shown in dashed and dotted boxes, respectively. When these are used to compute the effective dielectric constants, two different values will result.

Figure 4. The integration scheme chosen here was introduced by Kaneda and co-workers in Reference [4], and is given as

$$\varepsilon_z^{eff} = \left[\frac{1}{\Delta z} \int_{z_0 - \Delta z/2}^{z_0 + \Delta z/2} \frac{\Delta x \Delta y}{\int_{x_0 - \Delta x/2}^{x_0 + \Delta x/2} \int_{y_0 - \Delta y/2}^{y_0 + \Delta y/2} \varepsilon(x, y, z) dx dy} dz \right]^{-1}. \quad (3)$$

The dimension along which the outer integration is performed corresponds to the direction of the field component of interest located at (x_0, y_0, z_0) , which allows this scheme to satisfy the continuity of the tangential $\mathbf{E}$ and normal $\mathbf{D}$ fields. For an interface parallel to the field component, the outer integration is straightforward, and the inside integral results in a simple averaging of the dielectric constant over the volume of the cube. For an interface perpendicular to the field component, the inside integral is straightforward, and the outer integral leads to an averaging of the inverses of the dielectric constants (like resistors in parallel).

These two simple cases have been shown to require these disparate dielectric-averaging schemes in 2-D FDTD [10,11] (and similar conclusions have been drawn elsewhere [12,13]). However, the advantage of Equation 3 is that it provides for the intermediate case of an arbitrarily tilted interface, and is easy to extend to 3-D [4]. In our implementation of Equation 3, most of the work occurs before the first FDTD timestep: we identify cells at interfaces, perform the sub-cell integration with, for example, a 30×30 sub-grid (or $20 \times 20 \times 20$ in 3-D), compute the 1, 2 or 3 effective dielectric constants (for TM, TE, and 3-D respectively), and store the resulting FDTD update coefficients with each cell. During the FDTD time-stepping, these unique update coefficients are then used instead of the default values. As a result, there is no impact on the stability of the underlying FDTD algorithm, and only minor impact on the memory requirements and execution speed.

Other sub-cell approaches involve straightforward averaging of the dielectric constant around each field component [1], averaging of the dielectric constant over the cell volume itself [14,15], and estimation of the dielectric constant using the intercepts between material structures and the cell boundaries [16]. A new, more intricate sub-cell method has recently been published which incorporates the effects of material boundaries on the tensor relationship between $\mathbf{D}$ and $\mathbf{E}$ at each field component position [5]. While this method requires many more stored variables per affected interface cell, it has been shown to produce even more accurate results [5]




Figure 5. Spectral peaks for the square lattice of pillars at $k_x = 0.8\pi$, corresponding to the left-hand dotted box from Figure 1. Successive curves show the effect of windowing the raw temporal data with a Blackman window before performing the Fourier transform, and the effects of various amounts of truncation (as if the FDTD simulations had been stopped earlier). Curves are offset vertically for visualization purposes. The inset displays only a small portion of the spectral range, showing the strong ripple in the raw, un-windowed spectrum.




Figure 6. Normalized frequency vs. degree of truncation, for the square lattice of pillars at $k_x = k_y = 0.95\pi$, corresponding to the right-hand dotted box from Figure 1. As a smaller portion of the same temporal field evolution is used for post-processing, both Fourier analysis of the windowed data (open squares), and then Fourier analysis of windowed and zeropadded data (filled squares) are unable to resolve the close-set pair of spectral peaks. However, the near-degenerate modes can still be accurately identified by filter diagonalization [17,18] even when the spectral resolution is much coarser

than the spectral separation between the peaks (indicated by an $\times$ on the lower horizontal axis). Dotted horizontal lines show the predictions of the plane-wave method, representing a relative frequency error of only $\sim 0.1\%$.

than the Kaneda method [4] of Equation 3. An analysis of the impact of these various sub-cell techniques on the accuracy of FDTD applied to photonic crystal structures is currently underway [3].

3. FREQUENCY RESPONSE

But there remains an unanswered question here: How many FDTD timesteps are really necessary in order to unambiguously identify the spectral peaks ω_1 ? Through the usual properties of the Fourier transform, the width of each spectral bin in the calculated frequency spectrum is inversely proportional to the total duration of the monitored field component's temporal response. This would seem to imply that accurate determination of the resonance frequencies requires inordinately long FDTD simulations. As an additional complication, due to the inherently abrupt truncation of the ongoing temporal response, the raw spectrum tends to have a significant amount of ripple, as shown by the lowest trace in Figure 5. Here the spectral peaks corresponding to the particular k value of $k_x = 0.8\pi$ are shown, corresponding to the dotted box from Figure 1 that is located just to the left of the $\times$ point. This undesirable ripple tends to complicate the identification of spectral peaks, since the derivative of the spectrum will go through zero in many places other than at the peak.

However, it is important to keep in mind that the main interest here is the presence and position of the spectral peaks, not their width. So multiplying the original temporal response data with a windowing function such as the Blackman window, $f(t) = .42 - .5 \cos\left(2\pi \frac{t}{t_{max}}\right) + .08 \cos\left(4\pi \frac{t}{t_{max}}\right)$, can reduce the ripples in the spectrum without significantly affecting the spectral peak position, as shown by the second trace in Figure 5, labelled as “Windowed.” In essence, the underlying delta functions we are trying to locate (representing a lossless mode of essentially infinite Q) are simply being convolved by a symmetric function. As a smaller and smaller portion of the temporal response captured during the FDTD run is used when performing the Fourier Transform (as if we had truncated the run earlier to save precious wall-clock time), the resulting spectral peaks get wider and wider (remaining traces in Figure 5, shown only with Blackman windowing). However, the centroids of these peaks, representing the normalized frequencies we want to extract to place as dots on the dispersion diagram, remain almost completely unchanged.

This would seem to imply that it should not be necessary to have long FDTD runs in order to construct an accurate dispersion diagram. And this would be true if all the peaks were isolated. However, it is straightforward to see from Figure 5 that at some degree of blurring it becomes impossible to distinguish a close-set pair of modes

from a single mode. So the longer the time duration on each FDTD run used to build a dispersion diagram, the smaller the uncertainty about having resolved all the (possible) pairs of spectral peaks. And again, this uncertainty would seem to directly follow the frequency resolution of the Fourier transform, or $\Delta f = 1/N\Delta t$, where N is the number of FDTD timesteps of size Δt .

4. FILTER DIAGONALIZATION

However, it turns out that there are more accurate ways than Fourier analysis to extract resonance frequencies. One such technique, developed and then significantly extended by researchers in the Nuclear Magnetic Resonance community [17,19,20], and introduced to the photonic crystal community by Steven Johnson [18,21], is the filter-diagonalization method (FDM). This method recasts the problem of spectral analysis of a short segment of a time-dependent signal into an eigenvalue problem. Time-response data containing a number of resonances, $C(t) = \sum_k d_k \exp(-i\omega_k t)$, is treated as if it were the autocorrelation function of an unknown dynamical system, $C(t) = (\Phi_0, e^{-it\hat{\Omega}} \Phi_0)$, where $(\cdot, \cdot)$ represents the complex symmetric inner product [17]. By choosing Φ_0 appropriately, the unknown operator $\hat{\Omega}$ can be completely described in terms of the original data $C(t)$, and can thus be diagonalized [17]. The advantage of this technique over Fourier analysis is that, rather than simply measuring the correlation of the signal against multiple sinusoids $e^{-i\omega_k t}$ in multiple fixed-width spectral bins, a smaller number of spectral bins are getting precisely tuned in “size” and position to identify the complex frequencies where there is maximal correlation (e.g., the desired resonances). This allows the FDM method to distinguish pairs of peaks with much shorter time signals than Fourier analysis would require. The advantage of filter diagonalization over older harmonic inversion techniques (such as those based on Prony’s method) is that since FDM can concentrate on a small spectral range, the method remains stable and efficient and the matrices involved small even when an enormous number of resonances may exist outside of this range of interest [17,19].

The FDM method can be straightforwardly implemented in a Fourier basis simply involves following the checklist given on page 6761 of Reference [17]. Alternatively, an even easier approach to implementing the FDM method is that one can now simply download Steven Johnson’s “harminv” program from the web [18]. In addition to the checklist from Reference [17], the harminv program refines the eigenvalues by repeated iteration, converging smoothly to an accepted set of resonances.

The advantages of the filter diagonalization method are demonstrated in Figure 6. Here we plot the resonance frequencies for a close-set pair of peaks from Figure 1 as a function of the number of timesteps used in post-processing. The k point here, indicated in Figure 1 as a dotted box just to the right of the **M** point, corresponds to $k_x = k_y = 0.95\pi$. The same temporal evolutions from five monitor points in a single FDTD run are used for all points—the only thing that changes is the fraction of the data-set used and the analysis method. The horizontal axis is plotted both as a function of the number of FDTD timesteps, as well as in terms of the width of the associated normalized frequency bin. The two horizontal dotted lines show the two near-degenerate frequencies as determined by the plane-wave method [9] using fairly high resolution (128 plane-waves with the meshing value set to estimate the local dielectric constant over 7×7 sub-samples). While the error between frequencies computed by FDTD (squares, diamonds) and the plane-wave estimate may seem large on this scale, one should note that this is a relative error of only 0.1%.

At the right edge of Figure 6, both Fourier analysis and filter diagonalization are able to resolve the nearly degenerate modes. The open squares show the results after performing an FFT directly on the truncated data-set; the filled squares correspond to zero-padding the data-set up to a uniform length of 131072 timesteps. Since the peak-finding algorithm used here identifies a double peak in the spectrum by the presence of two significant local maxima, the extra data-points in the FFT after zero-padding helps distinguish the two peaks down to a truncation of 53,000 timesteps (labelled as point ‘B’). The spectrum for this degree of data-set truncation is shown in the inset of Figure 6. It turns out that even after the smaller peak no longer has a distinguishable local maxima, one can qualitatively identify double peaks by counting the number of zero-crossings in the imaginary part of the spectrum. However, when the truncation is so severe that the minimum frequency bin has become comparable to the frequency spacing between the degenerate modes (the point marked by an ‘X’ on the horizontal axis), then even this approach would fail.

However, Figure 6 shows convincingly that by analyzing the same time-response data with the filter diagonalization method, one can produce an accurate estimate for the two degenerate modes down to a truncation of

11,000 timesteps (nearly $5\times$ shorter than the cutoff of our peak-finding algorithm, and more than $2\times$ smaller than the inverse of the mode spacing being resolved). In addition, the FDM method continues to identify the presence of a double peak (although the frequencies are no longer being computed as accurately) down to a truncation of only 4096 timesteps, or more than $20\times$ coarser frequency resolution than the spacing between the nearly-degenerate modes. Thus filter diagonalization is a powerful technique that can significantly reduce FDTD run-times without sacrificing the ability to resolve nearly degenerate modes.

5. USE OF *A PRIORI* INFORMATION

In addition to the frequencies (eigenvalues) of the photonic crystal modes, it is often important to know the mode patterns (eigenfunctions) associated with each mode. Such patterns can be obtained with two successive FDTD runs: the first uses only a few monitor points in order to obtain the eigenfrequencies of all the modes. The second then performs a running Fourier transform at every grid point for each frequency of interest. This means that we no longer can rely on filter diagonalization to resolve near-degenerate modes from each other.

Assuming that the first run was performed with sufficient frequency resolution, we have *a priori* knowledge of where the modes that might interfere are located, which we can use to our advantage [22]. For instance, if we are interested in finding the eigenmode at ω_2 while avoiding any contribution from a nearby mode at ω_1 , we can choose an envelope function whose spectrum has a zero at ω_1 . Figure 7 shows the mode patterns (absolute value of $\tilde{E}_z$) for the two nearly degenerate modes of the square lattice near the M point that were the subject of Figure 6. The top row (parts (a), (b), and (c)) shows the lower frequency mode pattern, while the bottom row (parts (d), (e), and (f)) shows the pattern for the higher frequency mode. As can be seen in the inset of Figure 6, the higher frequency mode had a much weaker response in our simulations, simply due to poorer coupling from the particular initial excitation ‘ping’ we happened to choose. Parts (a) and (d) represent the modes as calculated by a long FDTD run (120300 time-steps), so that the two frequencies are clearly resolved by the end of the running Fourier transform. Parts (b) and (e) represent the modes when the running Fourier transform is accumulated over only 12000 time-steps. Since the lower frequency mode is stronger, it dominates the truncated computation of both mode patterns. These were obtained by Fourier transforming at the exact frequencies indicated by filter diagonalization on a previous run, without any Blackman window. (With a Blackman window applied, the higher frequency mode pattern was almost identical to the lower frequency pattern.)


Parts (c) and (f) of Figure 6, in which the two modes are correctly resolved, is obtained with exactly the same reduced number of time-steps. The only difference is that the length of the time-step Δt is slightly adjusted so that the normalized frequency f_1 of the strong mode lies exactly at the center of a frequency bin, as

$$(\Delta t)' = \frac{\lfloor f_1 c \Delta t N \rfloor}{f_1 c N}, \quad (4)$$

where N is the number of timesteps and c the speed of light. Without any spectral windowing, the effective spectral response of each bin is a sinc function (e.g., $\sin(\pi x)/(\pi x)$) whose nulls land at the center of each neighboring spectral bin. This means that a running Fourier transform at the frequency corresponding to the center of the next higher bin will not accumulate any signal from the mode at f_1 , even as it accumulates energy from f_2 (since this bin’s sinc function is non-zero at $f_2 \neq f_1$). This occurs even when this next frequency bin-center is somewhat higher than the actual value of f_2 . As a result, the mode patterns for the weaker mode can be distinguished from the stronger without having to use a large number of time-steps. Such techniques for the Fourier analysis of such photonic crystal modes are possible because the spectral features we are trying to either avoid or detect are essentially delta functions, and because we have *a priori* knowledge of the frequencies involved.

6. INTRINSIC OUT-OF-PLANE RADIATION LOSS FROM PHOTONIC CRYSTAL WAVEGUIDES

In this section, we study the intrinsic out-of-plane radiation loss from the planar photonic crystal shown in Figure 8. Here, instead of the infinitely tall air holes we used earlier, a triangular lattice of air holes is drilled into a finite-thickness slab. A row of holes is left filled to serve as the photonic crystal waveguide. The photonic



mode (b). As described in the text, a simple modification allows the properties of the Fourier transform to “null out” the crosstalk from the strong mode, revealing the weak mode (f).




Figure 8. One example of a planar photonic crystal is a triangular lattice of air holes at lattice spacing a , etched into a thin high-index slab. A photonic crystal waveguide is formed by leaving a row of defects (here, this is a row of missing air holes). Light entering from one of the ridge waveguides at the ends of the photonic crystal waveguide is confined vertically to the slab by total internal reflection (index guiding), and laterally to the defect row by the photonic bandgap effect. Here an air-bridge structure is shown that can offer improved index guiding, with an underlying substrate located some distance below the slab.




Figure 9. Dispersion diagram for the planar photonic crystal triangular lattice of air holes (hole radius $r/a=0.3$, slab thickness $h/a=0.6$) in dielectric ($\epsilon=11.56$), along with the reciprocal lattice. The unit cell used here is the same as the primitive 2-D unit cell in x and y (the rhombus shown in the inset), extruded in z through surrounding air regions and the underlying substrate of identical high-index material (located $3a$ below the slab). PML absorbing boundary conditions truncate the simulation $3a$ above the slab, and $1.5a$ into the substrate. Lower-lying odd modes have been removed (see text).

crystal waveguide is shown connected at both ends to a conventional ridge waveguide, with air both above and below the slab, at least down to an underlying substrate.

The slab modes are shown as calculated by 3-D FDTD in Figure 9. These are, at least for the first three bands, only the “TE-like” modes that are vertically even with respect to the center plane of the slab. The light-line corresponding to the air ($n \sim 1$) cladding is also shown.

Now that we have identified the guided mode gap described by Figure 9, we can engineer a waveguide by removing a row of defects.

Figure 10 shows the folded slab modes from Figure 9 as the boundaries of two light gray regions of slab




Figure 10. Dispersion diagram for the planar photonic crystal waveguide formed by removing a row of holes along the ΓK direction from the triangular lattice. The unit cell shown in the inset spans $1a$ in the x direction and $14a$ in the y direction, with vertical geometry identical to Figure 9. PML absorbing boundary conditions are placed on the $\pm y$ and $\pm z$ boundaries, with Bloch/Floquet boundary conditions on the $\pm x$ boundaries. When the lattice of holes continues into the PML layers, then slab modes are greatly suppressed, but also require lower Δt values for values of $k_x > \sim 0.8\pi/a$ in order to avoid late-time instability.




Figure 11. Normalized frequency as repeatedly computed at three k_x values from Figure 9, for various choices of the number of surrounding rows of photonic crystal, and of the size of the FDTD cell (in units of cells/ a). These choices are not shown in combination: for the reduced number of photonic crystal rows, the grid spacing is $a/20$; for the coarser grid-spacing, the number of photonic crystal rows is 7. The grid-cells are scaled so that their x and z sizes are identical within the slab, with the y grid step slightly smaller to allow an integer number of steps per $a' = \sqrt{3}a/2$. For most simulations, a smaller z grid-spacing was used for regions of

the 3-D simulation outside $|z| < a$. However, two simulations were run with constant z grid-step of $a/16$ and $a/20$ throughout the unit cell; the resulting frequencies are shown on the plot, but are so close to those from the non-uniform grid simulations that the symbols overlap.

modes, bracketing the mode gap region between normalized frequencies of 0.261 and 0.346. These gray regions represent the presence of modes from throughout the entire Brillouin zone, as bounded by the modes shown in Figures 9 from the periphery of the Brillouin zone. The light line (for air cladding) has also been placed on the dispersion diagram. It is clear that the regions of primary interest are those left unshaded on the figure. If we create a guided mode through defect placement that is in the light gray region, then we can expect it to be lossy by coupling (from our defect mode) to the many slab modes of the photonic crystal which are free to guide light away from our waveguide. If a guided mode runs above the light line, then we can expect it to be lossy by coupling to the many radiation modes which are free to radiate light vertically away from the slab.

To compute the defect modes of the photonic crystal waveguide with FDTD, we use a 3-D unit cell of the waveguide as shown in the inset of Figure 9. Here we take a slice out of the waveguide which is exactly one a long in the guiding direction x , and which encompasses a number of photonic crystal rows in y (here seven on each side) [8, 23]. As before, vertically there is air both above and below the waveguide, and a silicon substrate $3a$ below the slab (not shown in the inset). PML absorbing boundary conditions are used on both the $\pm z$ and $\pm y$ dimensions, and Bloch/Floquet boundary conditions in the $\pm x$ direction [8, 24, 25]. The same type of simulations as before are run, with each FDTD simulation determining the eigenfrequencies of the modes at the k vector being established by the Bloch/Floquet boundary conditions. The difference here is that we need only vary k_x from 0 to π/a .

The guided modes introduced by the defect are shown in Figure 10 as filled circles (connecting lines are present only to guide the eye). The defect mode that is horizontally even, and which has received considerable attention in both theoretical and experimental work [26], is shown in bold with filled circles. Typically the interest in this particular mode is driven by the interest in obtaining strong coupling to and from this photonic crystal waveguide mode using the similarly-shaped fundamental mode of a conventional ridge waveguide. The other two modes in the guided mode gap have different symmetries [8, 27]. While it is clear that part of this mode exists above and part below the light line, the dispersion diagram does not provide any quantitative information




Figure 12. Intrinsic loss due to out-of-plane radiation for the photonic crystal slab waveguide described in Figures 9 and 10, as a function of spatial frequency k_x in units of π/a . Insets show the Blackman-windowed incident field (upper right), the resulting field evolution for a lossy spatial frequency above the light line (center left) and for a low-loss point below the light line (lower left).




Figure 13. Out-of-plane loss as repeatedly computed at three k_x values from Figure 12, for various choices of the number of surrounding rows of photonic crystal, the size of the FDTD cell (in units of cells/ a), and the total number of FDTD timesteps used for the loss measurement. The length of the Blackman window is chosen to be 5 optical cycles shorter (~ 700 timesteps) than the overall length of the simulation.

on the losses. We will take up this topic in the next section.

Before doing so, it is interesting to consider variations on our choice of grid spacings ($a/20$) and the number of surrounding photonic crystal rows (7). To do so, we have picked out three $k_x a/\pi$ values (.25, .60, and .85) and repeatedly calculated the frequency of the primary even mode shown in Figure 10 for different choices of grid spacing and the number of photonic crystal rows (along the y dimension). The results, shown in Figure 11, clearly demonstrate that our original choices were actually quite conservative. It would be prudent to check these choices in combination, but it would seem that a grid-spacing as coarse as 12 cells per a and only 3 surrounding photonic crystal rows (on each side) would be sufficient to deliver low error. Recall that the use of the Kaneda (or another) sub-cell techniques we described before helps greatly in the scaling against grid-spacing; without such techniques, it is quite likely that the errors would be much larger. The relative frequency error is $< 0.1\%$ for all choices ≥ 3 surrounding rows, and $< 0.3\%$ for all grid spacings finer than $a/12$. As described in previous sections, it is likely that the number of timesteps used to compute the eigenfrequencies could be significantly reduced as well. Here, we chose a number of timesteps sufficient to resolve two frequencies near the center of the bandgap that differ by 1% (e.g., to place them two frequency bins apart in the resulting FFT).

An alternate technique for computing the intrinsic loss was recently introduced by Kuang and co-workers [28]. This technique is similar to the procedure we used earlier to compute modal patterns, in that once the eigenfrequency of a defect mode of interest is known, we run another nearly identical simulation to focus in on this eigenfrequency. However, rather than ping all temporal frequencies and take a running Fourier transform, instead we preferentially excite just the mode of interest by a continuous-wave excitation of our single spatial dipole as before. By driving the real part of the complex E fields with a sinusoidal carrier, and the imaginary part or bottom layer with a cosine, we can produce an excitation that travels in only one direction across the periodic boundaries [29]. In order to reduce the excitation of modes at other frequencies, we can apodize the temporal waveform with a Blackman window [28]. The envelope of the resulting temporal waveform is shown as the upper right inset in Figure 12; the carrier (one cycle ~ 140 timesteps) oscillates too rapidly here to be seen clearly. At the end of this excitation waveform, a large amount of energy has been injected into the mode of interest, the one specified in spatial frequency k by the boundary conditions and in temporal frequency ω by the incident source. If this mode is lossy, this energy continuously leaks out of the waveguide; if it is not, the injected energy would continue to propagate along the guide.

These two cases are shown in the left-hand insets in Figure 12, where the top inset shows the field evolution at a fairly lossy k -value located above the light line. Even before the excitation pulse is finished, it is apparent that energy is leaking out of the mode at this spatial frequency. In contrast, in the lower inset, the field amplitude

continues to build throughout the excitation. In this case, we could choose to continue the simulation and attempt to measure the slow decay of this field over subsequent time. However, since the lost energy must be leaving the simulation via the absorbing boundary conditions, we can measure the power flow out of the waveguide at any point during this steady-state condition by spatially integrating the Poynting vector just in front of the PMLs. Similarly, we can measure the power flow along the waveguide by integrating the flow across the Floquet boundary conditions. The ratio of the power flow leaving P_{out} to that remaining P_0 can be used to measure both the out-of-plane loss and the intrinsic loss into the slab due to a finite number of rows of surrounding photonic crystal. The loss per centimeter is then $-\log(1 - (P_{out}/P_0))/a$; here we use $a = 420\text{e-}7\text{ cm}$ [28].

In the main plot of Figure 12, we show the out-of-plane loss for the horizontally even mode featured in Figure 10 of the air-bridge, triangular lattice planar photonic crystal shown in Figure 8. This design is identical to the first of the five cases studied in Reference [28], and thus can be directly compared to the appropriate curve in Figure 2(a) of that paper. As expected, the losses are fairly high for the portions of the dispersion curve that lie above the light-line, and are quite low below the light line. The advantages of this technique are that we can measure very low losses (Kuang estimated an accuracy limit of $<.01\text{ cm}^{-1}$ [28]), and that the out-of-plane loss can be isolated from the in-plane loss and effects such as mode coupling.

As we did in the frequency calculations, we repeat these intrinsic loss simulations for various choices of the number of surrounding rows of photonic crystal, for the grid-spacing in cells per lattice constant a , and also as a function of the length of the windowed excitation in time-steps. The results are shown in Figure 13. As before, as few as 4 surrounding rows, 12 grid-cells per a , or 6000 timesteps could be used without significant impact on the computed loss. Of course, the impact of making all three of these choices together would need to be verified. Interestingly, when we initially seeded all these convergence tests with the accurate eigenfrequency obtained with the conservative choices of 7 surrounding rows and 20 cells per a , the errors were somewhat larger at the left edge of each subplot than those shown here. We traced this to the fact that very little energy was being injected into the mode of interest when we combine a slight frequency error and a long time window. As shown in Figure 13, simply using the same conditions (coarse gridding or reduced number of surrounding photonic crystal rows) for both the frequency and loss calculations avoids this problem. Similarly, the use of a shorter time window on the excitation allows any nearby frequency to produce the correct excitation, although this would be exactly the thing to avoid if a second defect mode were located at a nearby frequency. Still, we can reduce the computational load in terms of the number of timesteps so long as no competing defect (or slab) mode is nearby.

7. CONCLUSIONS

It is common to read in photonic crystal papers that the FDTD technique is excessively demanding in computational resources. There is, of course, much truth in this. However, such statements are often associated with assumptions of the necessary grid spacings, grid extent, and the total number of timesteps that are chosen based on rough rules of thumb or approximation, rather than on careful analysis. In some cases, such conservative choices may indeed be prudent. However, we have shown here that by using sub-cell methods, that can permit coarser spatial gridding; filter-diagonalization, that can find resonance frequencies with fewer FDTD timesteps; *a priori* information, to “null out” erroneous contributions at other frequencies; and by maximizing the information obtained with small numerical experiments that involve only the unit cell of the photonic crystal waveguide, we can challenge these assumptions can lead to significant reduction of the computational burden without a significant impact on accuracy.

REFERENCES

1. A. Taflov and S. C. Hagness, *Computational electrodynamics: the finite-difference time-domain method*, Artech House, Boston, 2nd ed., 2000.
2. G. W. Burr, “FDTD as a nanophotonics design optimization tool,” in *PECS-V (International Symposium on Photonic and Electromagnetic Crystals)*, March 2004.
3. A. Farjadpour and G. W. Burr are preparing a manuscript entitled “Accurate FDTD simulation of dielectric nanophotonics structures despite coarse gridding.”
4. N. Kaneda, B. Houshmand, and T. Itoh, “FDTD analysis of dielectric resonators with curved surfaces,” *IEEE Transactions on Microwave Theory and Techniques* **45**(9), pp. 1645–1649, 1997.

5. J. Nadobny, D. Sullivan, W. Wlodarczyk, P. Deuffhard, and P. Wust, "A 3-D tensor FDTD-formulation for treatment of sloped interfaces in electrically inhomogeneous media," *IEEE Transactions on Antennas and Propagation* **51**(8), pp. 1760–1770, 2003.
6. J. S. Juntunen and T. D. Tsiboukis, "Reduction of numerical dispersion in FDTD method through artificial anisotropy," *IEEE Transactions on Microwave Theory and Techniques* **48**(4), pp. 582–588, 2000.
7. O. Painter, J. Vuckovic, and A. Scherer, "Defect modes of a two-dimensional photonic crystal in an optically thin dielectric slab," *Journal of the Optical Society of America* **B 16**(2), pp. 275–285, 1999.
8. M. Loncar, T. Doll, J. Vuckovic, and A. Scherer, "Design and fabrication of silicon photonic crystal optical waveguides," *Journal of Lightwave Technology* **18**(10), pp. 1402–1411, 2000.
9. <http://ab-initio.mit.edu/mpb/>.
10. T. Hirono, Y. Shibata, W. W. Lui, S. Seki, and Y. Yoshikuni, "The second-order condition for the dielectric interface orthogonal to the Yee-lattice axis in the FDTD scheme," *IEEE Microwave and Guided Wave Letters* **10**(9), pp. 359–361, 2000.
11. K. P. Hwang and A. C. Cangellaris, "Effective permittivities for second-order accurate FDTD equations at dielectric interfaces," *IEEE Microwave and Wireless Components Letters* **11**(4), pp. 158–160, 2001.
12. R. D. Meade, A. M. Rappe, K. D. Brommer, J. D. Joannopoulos, and O. L. Alerhand, "Accurate theoretical-analysis of photonic band-gap materials," *Physical Review B* **48**(11), pp. 8434–8437, 1993.
13. S. G. Johnson and J. D. Joannopoulos, "Block-iterative frequency-domain methods for maxwell's equations in a planewave basis," *Optics Express* **8**(3), pp. 173–190, 2001.
14. S. Dey and R. Mittra, "A conformal finite-difference time-domain technique for modeling cylindrical dielectric resonators," *IEEE Transactions on Microwave Theory and Techniques* **47**(9), pp. 1737–1739, 1999.
15. O. Hess, C. Hermann, and A. Klaedtke, "Finite-difference time-domain simulations of photonic crystal defect structures," *Physica Status Solidi A* **197**(3), pp. 605–619, 2003.
16. W. H. Yu and R. Mittra, "A conformal finite difference time domain technique for modeling curved dielectric surfaces," *IEEE Microwave and Wireless Components Letters* **11**(1), pp. 25–27, 2001.
17. V. A. Mandelshtam and H. S. Taylor, "Harmonic inversion of time signals and its applications," *Journal of Chemical Physics* **107**(17), pp. 6756–6769, 1997.
18. <http://ab-initio.mit.edu/harminv/>.
19. M. R. Wall and D. Neuhauser, "Extraction, through filter-diagonalization, of general quantum eigenvalues or classical normal-mode frequencies from a small number of residues or a short-time segment of a signal .1. theory and application to a quantum-dynamics model," *Journal of Chemical Physics* **102**(20), pp. 8011–8022, 1995.
20. R. Q. Chen and H. Guo, "Efficient calculation of matrix elements in low storage filter diagonalization," *Journal of Chemical Physics* **111**(2), pp. 464–471, 1999.
21. S. G. Johnson, S. Fan, A. Mekis, and J. D. Joannopoulos, "Multipole-cancellation mechanism for high-Q cavities in the absence of a complete photonic band gap," *Applied Physics Letters* **78**(22), pp. 3388–3390, 2001.
22. G. W. Burr and A. Farjadpour are preparing a manuscript entitled "Accurate FDTD simulation of dielectric nanophotonics structures despite truncated computations."
23. B. D'Urso, O. Painter, J. O'Brien, T. Tombrello, A. Yariv, and A. Scherer, "Modal reflectivity in finite-depth two-dimensional photonic-crystal microcavities," *Journal of the Optical Society of America* **B 15**(3), pp. 1155–1159, 1998.
24. M. Loncar, J. Vuckovic, and A. Scherer, "Methods for controlling positions of guided modes of photonic-crystal waveguides," *Journal of the Optical Society of America* **B 18**(9), pp. 1362–1368, 2001.
25. T. Baba, A. Motegi, T. Iwai, N. Fukaya, Y. Watanabe, and A. Sakai, "Light propagation characteristics of straight single-line-defect waveguides in photonic crystal slabs fabricated into a silicon-on-insulator substrate," *IEEE Journal of Quantum Electronics* **38**(7), pp. 743–752, 2002.
26. For recent work in the photonic crystal field for optics, the reader should consult the proceedings of general conferences such CLEO, LEOS, and Photonics West/Photonics East, as well as the specialized PECS conference (International Symposium on Photonic and Electromagnetic Crystal Structures). In addition, an extensive bibliography and relevant links can be found at www.pbglink.com.
27. S. G. Johnson, P. R. Villeneuve, S. H. Fan, and J. D. Joannopoulos, "Linear waveguides in photonic-crystal slabs," *Physical Review B* **62**(12), pp. 8212–8222, 2000.
28. W. Kuang, C. Kim, A. Stapleton, W. J. Kim, and J. D. O'Brien, "Calculated out-of-plane transmission loss for photonic-crystal slab waveguides," *Optics Letters* **28**(19), pp. 1781–1783, 2003.
29. M. Celuch-Marcysiak and W. K. Gwarek, "Spatially looped algorithms for time-domain analysis of periodic structures," *IEEE Transactions on Microwave Theory and Techniques* **43**(4), pp. 860–865, 1995.