Appl. Phys. B 81, 283–293 (2005)
DOI: 10.1007/s00340-005-1823-4

Applied Physics B
Lasers and Optics

S.G. JOHNSON[1], ✉
M.L. POVINELLI[2]
M. SOLJAČIĆ[1]
A. KARALIS[1]
S. JACOBS[3]
J.D. JOANNOPOULOS[1]

Roughness losses and volume-current methods in photonic-crystal waveguides

[1]Massachusetts Institute of Technology, Cambridge MA 02139, USA
[2]Ginzton Laboratory, Stanford University, Stanford, CA 94305, USA
[3]OmniGuide Communications, One Kendall Square, Building 100 #3, Cambridge MA 02139, USA

Received: 1 February 2005

ABSTRACT We present predicted relative scattering losses from sidewall roughness in a strip waveguide compared to an identical waveguide surrounded by a photonic crystal with a complete or incomplete gap in both 2d and 3d. To do so, we develop a new semi-analytical extension of the classic "volume-current method" (Green's functions with a Born approximation), correcting a longstanding limitation of such methods to low-index contrast systems (the classic method may be off by an order of magnitude in high-contrast systems). The resulting loss predictions show that even incomplete gap structures such as photonic-crystal slabs should, with proper design, be able to reduce losses by a factor of two compared to an identical strip waveguide; however, incautious design can lead to increased losses in the photonic-crystal system, a phenomena that we explain in terms of the band structure of the unperturbed crystal.

PACS 42.25.Fx, 42.70.Qs, 42.79.Gn

1 Introduction

In this paper, we present results on the analysis of scattering losses from surface-roughness disorder for optical waveguides in photonic crystals with both complete and incomplete bandgaps, and in particular for the two- and three-dimensional structures of Fig. 1 compared to an isolated strip waveguide. Moreover, we correct a longstanding limitation of the standard "volume-current method" (Green's function Born approximation) and related techniques that model a perturbation $\Delta\varepsilon$ in the dielectric constant as a current $\mathbf{J} \sim \Delta\varepsilon\mathbf{E}$ proportional to the unperturbed field $\mathbf{E}$ in order to compute losses [1–12]—this technique is in general only accurate for small $\Delta\varepsilon$ and thus for surface roughness with low index contrast, while we show that its use for high-contrast systems (e.g., photonic crystals) [6, 7, 10–12] can introduce errors of an order of magnitude in the loss (depending on the polarization and roughness statistics). Here, we extend the method to compute the correct $\mathbf{J}$ for surface roughness with arbitrary index contrast, and in doing so enable quantitative predictions of roughness losses with modest (personal-computer) resources; the main previous alternative was brute force discretizations of Maxwell's equations [13–15] that become nearly prohibitive in three dimensions because of the high resolution (or many multipole moments [16–18], etc.) needed to capture small imperfections. Given our semi-analytic tool, we present predictions of how the bandgap will affect roughness-induced radiation and reflection loss, and show two significant results: first, with a proper design, one can substantially reduce overall losses (by a factor of two) while keeping reflection loss fixed (following a general theorem from [14]); second, an incautious design can also substantially *increase* losses by introducing new modes or "hidden" resonances via the bandgap. We can now quantitatively distinguish these cases, permitting future optimizations and potential experimental confirmation in structures amenable to rapid fabrication.

Photonic crystals are periodic dielectric structures possessing a bandgap, a range of wavelengths in which light cannot propagate in some or all directions, yielding unprecedented control over optical phenomena [19]. By introducing linear "defects" into a crystal, one can produce a waveguide structure in which light is confined to propagate along the defect by the bandgap; or, in an incomplete gap structure such as a photonic-crystal slab [20], light is confined in some directions by the bandgap and in other directions by index-guiding ("total internal reflection") [21]. Such waveguides, in an ideal system, are lossless, but in reality losses are introduced by material absorption and fabrication disorder such as surface roughness (which causes light to scatter into both radiation and reflection). Surface roughness, in particular, is known to be a primary source of loss in waveguides made from high-index contrast materials, such as those required in photonic-bandgap structures, because the scattering from a small perturbation grows with the index contrast—thus, the study of this loss is critically important, and is also challenging because high index contrasts prohibit most analytical treatments and demand high resolution in brute-force simulations of small perturbations.

It turns out that many classic semi-analytical methods, such as perturbation theory and coupled-mode theories, require modification in the high-index-contrast, strongly periodic systems now prevalent in nanophotonics [14, 22–25], so it comes as no surprise that the volume-current method

---

✉ Fax: +39-0382-987-563, E-mail: stevenj@math.mit.edu

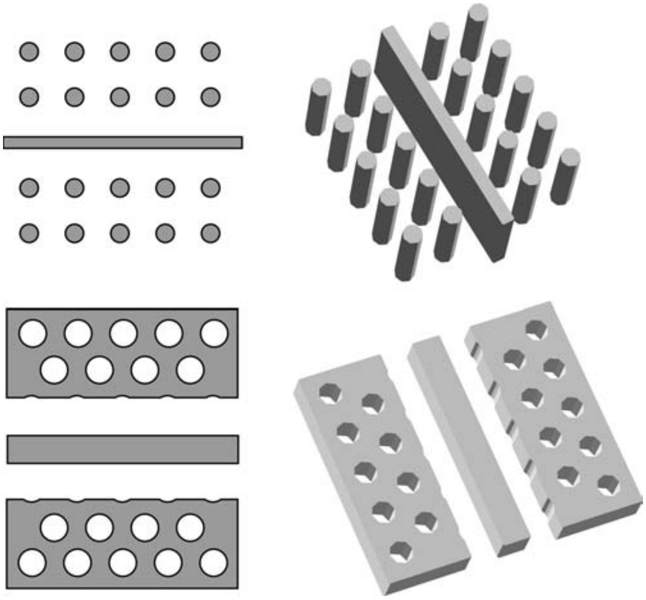

FIGURE 1 Strip waveguides surrounded by 2d-periodic photonic crystals in 2d (*left*) and 3d (*right*) for both a rod crystal (*top*) and a hole crystal (*bottom*). We will compare the losses due to surface roughness for the crystal waveguide versus the isolated strip

is altered as well. Indeed, our corrections here also apply to almost all other perturbative approaches when they are used for small surface "bumps." (There is one exception: 2d systems with TM polarization—**E** out of the plane—have fields everywhere continuous so that the classic methods are correct.) Nevertheless, we show that the basic Green's function strategy is preserved: one finds an effective current **J** to represent a small imperfection, computes the power radiated by this **J** in the *unperturbed* waveguide (a modest calculation that does not require high resolution or sensitivity), and equates this to the scattering loss. For uncorrelated roughness (or, more precisely, disorder whose correlation function only determines the local "bump" shape), the power from multiple scatterers adds incoherently on average, so one obtains the mean loss rate by multiplying the loss from one scatterer by the average scatterer density and amplitude. An alternative semi-analytical picture based on coupled-mode theory was presented in [14], which was well-suited to deriving an analytical theorem stating conditions under which a photonic crystal does not increase reflection loss (contradicting an intuitive picture where the bandgap reflects radiation into reflection). This previous work, however, did not analyze the radiative loss in *incomplete* bandgap systems such as photonic-crystal slabs, and presented quantitative predictions only for a two-dimensional (2d) system in which brute-force simulation was practical. Here, we confirm our earlier theorem by an independent method, analyze situations with incomplete gaps and where the photonic-crystal inadvertently introduces new reflection channels, and present numerical results in both 2d and 3d. (In general, coupled-mode approaches could also be used to predict disorder loss [14, 25, 26], but: (i) they require many modes in order to capture the effect of a continuum of radiation states, and (ii) the coupling matrix elements in many coupled-mode and perturbation theories must also be modifed [14] to use our new **J** for small, high-contrast bumps.) We also discuss the impact of general scaling laws that obtain when the crystal alters the group velocity v_{g} of a mode (e.g., near a band edge [19, 27]): absorption and radiation loss scale as $1/v_{\mathrm{g}}$, and reflection loss as $1/v_{\mathrm{g}}^2$, all other things being equal, as was also pointed out in [12, 28].

In the following, we first review the standard volume-current method and then describe its corrected form for boundary perturbations in high-contrast systems. (Note: $\varepsilon_0 = \mu_0 = 1$ units are used throughout.) Then, we evaluate the shape dependence of the corrected form for typical types of surface roughness. A check of the perturbation current **J** is performed by using it to calculate frequency perturbations (rather than loss), compared to a brute-force computation. Finally, we predict the relative roughness loss introduced by the bandgap in the photonic-crystal structures of Fig. 1, rods in air and holes in dielectric in 2d and 3d, in particular focusing on the case of sidewall roughness [29]. We conclude with general considerations of other loss mechanisms and scaling laws.

2 Corrected volume-current method

The volume-current method can be derived in a number of ways; for example, it is equivalent to the first Born approximation for the Green's function of the perturbed system in terms of the unperturbed one [2]. One can then show that a small shift $\Delta\varepsilon$ in the dielectric constant ε can be represented by a current $\mathbf{J} = -\mathrm{i}\omega\Delta\varepsilon\mathbf{E}$ in the scatterer (which can be replaced by a point source if the scatterer is small) for an unperturbed field **E** at a frequency ω (time dependence $e^{-\mathrm{i}\omega t}$), to first order in $\Delta\varepsilon$. (In a waveguide, **E** is the unperturbed waveguide mode's field at the scatterer.)

For a "bump" in a high-index-contrast boundary with a small volume ΔV and a shift $\Delta\varepsilon$ in the dielectric constant ε, however, one cannot expand in $\Delta\varepsilon$. The problem is not so much that $\Delta\varepsilon$ is not a small parameter, but that the discontinuities in **E** at high-contrast boundaries makes the product $\Delta\varepsilon\mathbf{E}$ ill-defined [14, 22, 28, 30]. (Except for pure TM polarizations in 2d, where the classic **J** is correct as mentioned above and below.) Instead, the most useful viewpoint for our purposes is based on the standard scattering theory for small scatterers at long wavelengths (much larger than the scatterer diameter), for which the quasi-static approximation can be made [31]. In this approximation, which is essentially first-order in ΔV rather than $\Delta\varepsilon$, one takes the local unperturbed electric field **E** to be constant (except for boundary discontinuities) and solves for the induced dipole moment $\mathbf{p} \sim \mathbf{E} \cdot \Delta V$, and then writes $\mathbf{J} \sim \mathbf{p} \cdot \delta(\mathbf{x})$ to replace the scatterer with an oscillating point dipole. Below, we first consider the well-known situation of a scatterer in a homogeneous background dielectric, and then derive two corrections for a $\Delta\varepsilon$ due to a high-contrast surface bump: the relationship $\mathbf{p} = \alpha\mathbf{E}\Delta V$ is in general in terms of a numerically computed polarizability tensor α, and the relationship between **J** and **p** is altered on a surface.

2.1 *Locally uniform media*

For a large $\Delta\varepsilon$, the discontinuous boundary conditions on the interface-normal component of **E** (recall that $D_\perp = \varepsilon E_\perp$ and $\mathbf{E}_\parallel$ are continuous) can no longer be neglected, and the field inside the scatterer must be solved via the static-field boundary-value problem (assuming a constant incident field given by the unperturbed $\mathbf{E}$ at the scatterer location). Then, from the change in the field $\Delta\mathbf{E}$ compared to the unperturbed field $\mathbf{E}$, one finds $\Delta\rho = \nabla \cdot \Delta\mathbf{E}$ and thus the induced dipole moment $\mathbf{p} = \int \mathbf{r}\Delta\rho dV = \alpha\mathbf{E}\Delta V$, where α is the polarizability tensor per unit volume. In particular, there is a classic analytical solution for α in the case of a sphere of ε_1 in a background of ε_2 ($\Delta\varepsilon = \varepsilon_1 - \varepsilon_2$), in which case [31]:

$$\alpha_{\text{sphere}} = \frac{3\Delta\varepsilon}{3\varepsilon_2 + \Delta\varepsilon}. \tag{1}$$

In the limit of $\Delta\varepsilon \to 0$, this gives $\mathbf{p} = \Delta\varepsilon\mathbf{E}\Delta V/\varepsilon_2$, and comparing with above we find that $\mathbf{J}$ is proportional to $\varepsilon_2\mathbf{p}$. The reason for this is that the current $\mathbf{J}$ injects a *free* charge into the system, but $\mathbf{p}$ is the *total* induced dipole moment (including bound charge). Thus, since any free charge is reduced (screened) by a factor of ε_2, the current must be bigger than $\mathbf{p}$ by a factor of ε_2 to induce the same total dipole moment. Therefore, the general relationship between $\mathbf{J}$ and $\mathbf{E}$ in a uniform background ε_2 is:

$$\mathbf{J} = -\mathrm{i}\omega\varepsilon_2\Delta V \cdot (\alpha\mathbf{E}) \cdot \delta(\mathbf{x}) \tag{2}$$

and thus the α_{sphere} solution gives the correct low-contrast limit $\mathbf{J} \to -\mathrm{i}\omega\Delta\varepsilon\Delta V\mathbf{E} \cdot \delta(\mathbf{x})$ from above. In the high contrast limit, however, α goes to a constant ($= 3$) and differs greatly from the low-contrast approximation. In general geometries, one must compute the polarizability α numerically, as described in Sect. 2.3, but this is a modest calculation.

A slightly modified relationship holds in the case of a scatterer with a fixed 2d (xy) cross section with a small area ΔA, such as a cylinder. In this case, we wish to replace the scatterer with a line current $\delta(x)\delta(y)$, where the current at any point along the line is given from the unperturbed field $\mathbf{E}$ at that point. For a field in the xy plane of the cross section, the analysis is the same as above and we get the the same formula (2) except that ΔV is replaced with ΔA, and α is now the 2×2 polarizability tensor per unit area, which in general must be computed numerically. For $\mathbf{E}$ in the uniform-cross-section z direction, however, the problem can be solved analytically because $\mathbf{E}$ is continuous, and one finds: $J_z = -\mathrm{i}\omega\ \Delta A\ \Delta\varepsilon E_z \cdot \delta(x)\delta(y)$—that is, the "standard" volume-current proportional to $\Delta\varepsilon\mathbf{E}$, which is no surprise since the lack of field discontinuities makes the standard treatment valid.

2.2 *Scatterers at interfaces*

For surface roughness, we have not an isolated scatterer but rather a "bump" on an interface between two materials ε_1 and ε_2 with the normal direction denoted x, as depicted in Fig. 2. (Note that in the vicinity of a small bump, any smooth interface can locally be treated as flat.) In this case, the unperturbed $\mathbf{E}$ is no longer approximately constant—rather, the surface parallel and perpendicular components $\mathbf{E}_\parallel$ and $\mathbf{D}_\perp$ are approximately constant in a small region. So, it is natural to express the induced dipole moment $\mathbf{p} = \int \mathbf{r}\Delta\rho\ \mathrm{d}V$ (where $\Delta\rho = \nabla \cdot \Delta\mathbf{E}$ as above) in terms of these quantities:

$$\mathbf{p} = \Delta V(\alpha_\parallel\mathbf{E}_\parallel + \gamma_\perp\mathbf{D}_\perp), \tag{3}$$

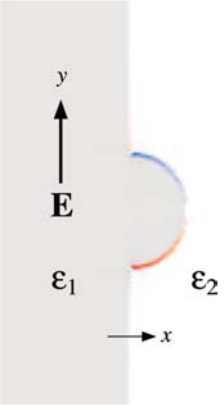

FIGURE 2 A "bump" on the interface between two materials ε_1 and ε_2, with x and y denoting directions perpendicular and parallel to the interface, respectively. An applied electric field $\mathbf{E}$, here in the y direction, will induce a dipole moment via an induced charge density $\Delta\rho$ (blue/red denotes positive/negative $\Delta\rho$) on the surface. We call this a "positive" bump because it juts from ε_1 into ε_2, so $\Delta\varepsilon > 0$ if $\varepsilon_1 > \varepsilon_2$

where $\alpha_\parallel$ and $\gamma_\perp$are polarizability tensors. It is important to note that a "positive" bump with $\Delta\varepsilon = \varepsilon_1 - \varepsilon_2 > 0$ (as in Fig. 2) will in general have *different* polarizability tensors α and γ (in both sign and magnitude) than a "negative" bump (where ε_2 juts into ε_1) with $\Delta\varepsilon = \varepsilon_2 - \varepsilon_1 < 0$, which has important consequences for scattered powers and perturbation theory in the subsequent sections.

Just as in the previous section, $\mathbf{p}$ represents the total dipole moment including all bound charges, whereas $\mathbf{J}$ injects only a free charge, so we must compute the correction factor so that $\mathbf{J}$ produces a total moment $\mathbf{p}$. The correction is found by using the standard method of images [31] to compute the bound charges for a free charge q that approaches the interface, and hence the relationship between free and bound dipole moments for a point dipole that approaches the interface from one side. This gives a current $\mathbf{J}_\parallel = \mathbf{p}_\parallel(\varepsilon_1 + \varepsilon_2)/2$ and $\mathbf{J}_\perp = \mathbf{p}_\perp\varepsilon$, where ε is either ε_1 or ε_2 depending upon which side of the interface $\mathbf{J}$ is on. In systems with sufficient symmetry (e.g., a $y = 0$ mirror plane), the parallel and perpendicular directions are principal axes (so that $\mathbf{p}_\parallel = \alpha_\parallel\mathbf{E}_\parallel$ and $\mathbf{p}_\perp = \gamma_\perp\mathbf{D}_\perp$), and the final expression for $\mathbf{J}$ is then simply:

$$\mathbf{J} = -\mathrm{i}\omega\Delta V\left(\frac{\varepsilon_1 + \varepsilon_2}{2} \cdot \alpha_\parallel\mathbf{E}_\parallel + \varepsilon\gamma_\perp\mathbf{D}_\perp\right) \cdot \delta(\mathbf{x}). \tag{4}$$

Note that we have the convenient property that we can place $\mathbf{J}$ on *either* side of the interface (albeit with different values of $\mathbf{J}_\perp$) and get the *same* total $\mathbf{p}$ and thus the same far-field scattered power. (A related property, important for the perturbation theory of the next section, is that $\mathbf{J} \cdot \mathbf{E}$ is continuous across the interface.) Note also that this gives Eq. (2) when $\varepsilon_1 = \varepsilon_2$.

As in the previous section, the treatment is slightly modified for the case of a scatterer with a fixed 2d (xy) cross section of area ΔA. Again, for fields in the xy plane, one obtains the same Eq. (4), except with ΔA instead of ΔV and polarizabilities α and γ per unit area instead of volume. Just as before, for $\mathbf{E}$ in the z direction, the field is everywhere continuous and one analytically obtains the "standard" volume current $J_z = -\mathrm{i}\omega\Delta A\Delta\varepsilon E_z \cdot \delta(x)\delta(y)$.

2.3 *Computation of polarizability*

In order to compute the polarizability tensors α (or γ), one need only solve the static-field boundary-value problem for a given scatterer geometry and a fixed "external field" $\mathbf{E}$ (or $\mathbf{D}$) at infinity, with three computations (external fields in three directions) to obtain all of the components of the tensor. Very efficient methods, accurate to 5 decimal places with modest computations, have been devised for this problem in the literature [32]. We used, however, a less efficient method accurate to $\sim 2\%$ that had the advantage of employing a freely available software package—in particular, we solve the vectorial Maxwell's equations in a planewave basis by a conjugate-gradient eigensolver approach [33]. This method imposes periodic boundary conditions, so one must choose a computational cell large enough that the boundaries are far from the scatterer. It solves for the eigenmodes at a finite frequency ω and wavevector $\mathbf{k}$, but we obtained the $\omega = 0$ static-field solution by taking the limit as $\mathbf{k}, \omega \to 0$, where the direction of $\mathbf{k}$ and the symmetry of the solution determine the polarization of the "static" external field. Given the difference $\Delta\mathbf{E}$ in the solutions for the cases with the scatterer and without the scatterer, the induced dipole moment is computed by $\mathbf{p} = \int \mathbf{r}\,\nabla \cdot \Delta\mathbf{E}$, whence $\alpha = \mathbf{p}/(\mathbf{E}\,\Delta V)$.

For a scatterer diameter of $2a$, we used a $40a \times 40a$ supercell for scatterers with 2d cross sections and a $10a \times 10a \times 10a$ supercell for scatterers with 3d cross sections, with resolutions of 40 and 10 pixels/a, respectively. Compared to the analytical solutions for the polarizability of a dielectric sphere and cylinder, and to a published solution for a dielectric square [32], we obtained polarizabilities accurate to within 2% for dielectric contrasts up to 16 : 1, an accuracy sufficient for our purposes.

We analyzed both "positive" (ε_1 juts into ε_2) and "negative" (*vice versa*) bumps. The polarizabilities in these two cases are not equal, but *are* related analytically: if we multiply both ε_1 and ε_2 by any constant C, then α is unchanged and γ $(= \alpha/\varepsilon)$ changes to γ/C. Thus, we need only consider the polarizabilities as a function of $\tau = \varepsilon_1/\varepsilon_2$, and a "negative" bump corresponds to $\tau < 1$.

The results for four basic scatterer shapes are shown as a function of τ in Figs. 3 and 4: 2d ("ridge") bumps formed by a semicylinder of radius a or half a square of side $2a$, and 3d bumps formed by a hemisphere of radius a or half a cube of side $2a$. These have sufficient symmetry that the polarizabilities are diagonal and we need only consider α_{yy} and γ_{xx}. (One computed $\Delta\rho$ for α_{yy} with a positive semicylindrical bump is shown in Fig. 2.) We found that the cylinder/sphere and square/cube have polarizabilities that differ by at most 10% and 15%, respectively, which gives some indication that the precise shape of the bump (for a fixed area or volume and height/width ratio) does not have a large impact upon the roughness loss. The solid lines are a least-squares fit to functions of the following form:

$$\alpha(\tau) = \frac{2(\tau-1)}{\tau+1}\left[1 + \frac{\tau-1}{\frac{\tau}{\alpha_\infty/2-1} - \frac{1}{\alpha_0/2-1}}\right], \tag{5}$$

$$\gamma(\tau) = \frac{\tau-1}{\tau}\left[1 + \frac{\tau-1}{\frac{\tau}{\alpha'_\infty/2-1} - \frac{1}{\alpha'_0/2-1}}\right], \tag{6}$$

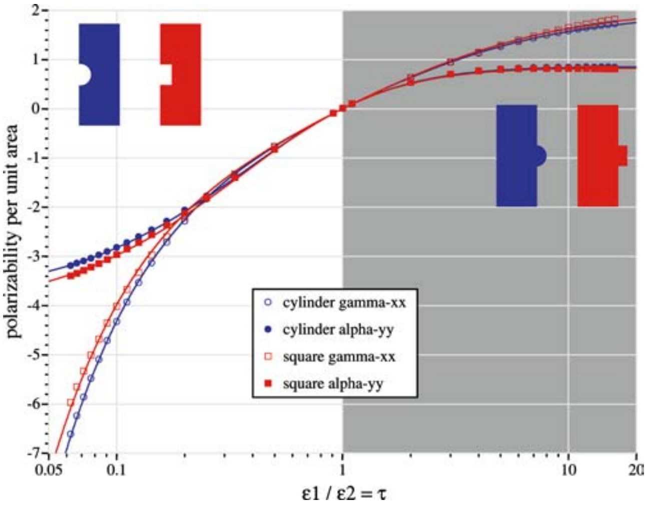

FIGURE 3 Polarizabilities per unit area for "2d bumps" of semicylindrical (blue circle symbols) and half-square (red square symbols) shapes; α_{xx} (filled symbols) is the interface-parallel component and γ_{yy} (hollow symbols) is the interface-normal component. Solid lines are least-squares fits to Eqs. (5)–(6). $\tau > 1$ (shaded) corresponds to "positive bumps" (lower-right inset) and $\tau < 1$ to "negative" bumps (upper-left inset)

where α_0 and α_∞ are fit parameters that (for $\alpha(\tau)$) physically correspond to the polarizabilities when ε_2 or ε_1 are equivalent to perfect metals, respectively (i.e., the $\tau \to 0, \infty$ limits of α are $-\alpha_0, \alpha_\infty$). This form was chosen so that: (i) it has the correct limit for $\tau \to 1$, in which case one must get the standard volume current $\mathbf{J} \sim \Delta\varepsilon\mathbf{E}$; (ii) it fits the known analytical formulas exactly for a cylinder, sphere, and ellipsoid; and (iii) it has finite polarizability α for the "metal" limits. (The divergence in γ for $\tau \to 0$ is a physical result of the fact that we define it in terms of $\mathbf{D}_\perp$, which goes to 0 inside a metal.) The fit parameters and maximum fit errors (for $\tau \in [1/16, 16]$) are: cylindrical α_{yy} has $\alpha_\infty = 0.8510$, $\alpha_0 = 3.882$, error $< 2\%$; cylindrical γ_{xx} has $\alpha'_\infty = 3.905$, $\alpha'_0 = 0.7669$, error $< 3\%$; square α_{yy} has $\alpha_\infty = 0.8214$, $\alpha_0 = 4.162$, error $< 3\%$;

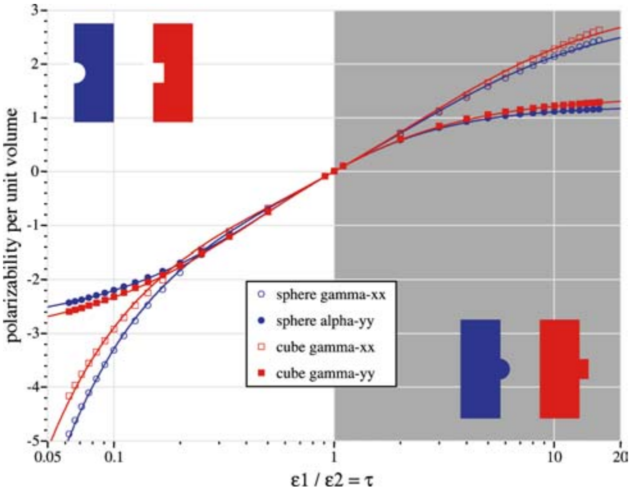

FIGURE 4 Polarizabilities per unit volume for "3d bumps" of hemispherical (blue circle symbols) and half-cube (red square symbols) shapes; α_{xx} (filled symbols) is the interface-parallel component and γ_{yy} (hollow symbols) is the interface-normal component. Solid lines are least-squared fits to Eqs. (5)–(6). $\tau > 1$ (shaded) corresponds to "positive bumps" (lower-right inset) and $\tau < 1$ to "negative" bumps (upper-left inset)

square γ_{xx} has $\alpha'_\infty = 4.082$, $\alpha'_0 = 0.6804$, error < 5%; sphere $\alpha_{yy} = \alpha_{zz}$ has $\alpha_\infty = 1.220$, $\alpha_0 = 2.856$, error < 1%; sphere γ_{xx} has $\alpha'_\infty = 5.860$, $\alpha'_0 = 0.5366$, error < 3%; cube $\alpha_{yy} = \alpha_{zz}$ has $\alpha_\infty = 1.398$, $\alpha_0 = 3.115$, error < 1%; cube γ_{xx} has $\alpha'_\infty = 6.361$, $\alpha'_0 = 0.4376$, error < 5%. These errors are comparable to the likely numerical error in the polarizabilities, and in any case are small enough that the fits alone are suitable for roughness-loss calculations (since they are swamped by experimental uncertainty in the roughness shape, etc.).

2.4 *Errors in* $\mathbf{J} = \Delta\varepsilon\mathbf{E}$

It is instructive to compare the polarizability here and the resulting current $\mathbf{J}$ from Eq. (4) with the simple $\Delta\varepsilon\,\mathbf{E} = 11\mathbf{E}$ formula for a typical high-contrast interface of $\varepsilon = 12{:}1$ (e.g., Si:air). In this case, for a semicylindrical bump, then we get a surface-parallel (y) current proportional to $5.5E_y$ for a positive bump and $19.4E_y$ for a negative bump; for the normal (x) component, if we evaluate $\mathbf{E}$ on the side of the bump, we get a current proportional to $1.64E_x$ for a positive bump $61E_x$ for a negative bump. In other words, depending upon the polarization, the "classical" volume current would be off by a factor of 2–7, leading to an error in the scattered power of 4–50!

Two factors can act to ameliorate this error in common circumstances. First, for a cylindrical or "ridge" bump with $\mathbf{E}$ oriented along the ridge (z), the classic $\Delta\varepsilon\ E_z$ is perfectly correct as described in Sect. 2.2. Moreover, if the bump is very wide and shallow (unlike here, where we assume the bump width is comparable to its height/amplitude), the correct $\mathbf{J}$ becomes $\sim \Delta\varepsilon\ \mathbf{E}$ for *both* surface-parallel polarizations, and only the surface-normal polarization is incorrect in the classic formula [22]. Thus, the magnitude of the error in practice will depend both upon the field polarization at interfaces and on the roughness statistics.

3 Perturbation theory

Before we apply the corrected volume currents from the previous section to compute scattering loss, we first consider their utility in another problem that gives us the opportunity to directly verify our results: the perturbation current $\mathbf{J}$ is also used to compute frequency shifts of eigenmodes in first-order perturbation theory. In particular, if $\mathbf{E}$ is the eigenmode of a structure at a frequency ω (and, in a periodic structure, at a Bloch wavevector $\mathbf{k}$), then it is well known that a small change $\Delta\varepsilon$ in the structure produces a shift $\Delta\omega$ that is approximately (to first-order in $\Delta\varepsilon$) given by [1, 22, 26]:

$$\Delta\omega \cong -\frac{\omega}{2}\frac{\int \mathbf{E}^*\cdot\Delta\varepsilon\mathbf{E}}{\int \varepsilon\,|\mathbf{E}|^2}, \tag{7}$$

where the integrals are over the volume (the unit cell for a periodic structure). (This equation assumes that the unperturbed system is lossless, but can be used to compute e.g., absorption loss: a small imaginary $\Delta\varepsilon$ gives an imaginary $\Delta\omega$, a decay per time, and the decay per distance of a waveguide is $\Delta\omega$ divided by the group velocity.) Another formulation, however, is:

$$\Delta\omega \cong -\frac{\mathrm{i}}{2}\frac{\int \mathbf{E}^*\cdot\mathbf{J}}{\int \varepsilon\,|\mathbf{E}|^2}, \tag{8}$$

where $\mathbf{J}$ is the volume current that appears in the Green's function for the scattering problem. $\mathbf{J} = -\mathrm{i}\omega\Delta\varepsilon\mathbf{E}$ for a small $\Delta\varepsilon$, as discussed above, in which case the two formulations are equivalent. However, the formulation (8) is more general, because it holds even for a boundary perturbation with a large $\Delta\varepsilon$ in a small volume ΔV (assuming $\mathbf{J}$ is computed correctly as in the previous section), in which case the $\Delta\omega$ is correct to first order in ΔV, whereas (7) is ill-defined because it multiplies $\Delta\varepsilon$ by a discontinuous field [22]. A simple proof of this generality follows the same procedure as in [22]: one takes a small ΔV perturbation as the limit of a structure with *smoothed* (continuous) boundaries. In a smoothed system, the fields are continuous, $\Delta\varepsilon$ is small at any point, and (7) and (8) are thus equivalent. Now, we replace $\mathbf{J}$ in the smoothed system by a delta-function point source with the same integral, as in the previous section, which is accurate to first-order in ΔV. Then, we take the limit to the discontinuous system: the limit of (7) is problematic because of the discontinuity, but the limit of (8) is well-defined because $\mathbf{E}^*\cdot\mathbf{J}$ is continuous across the boundary (from Sect. 2.2). Maxwell's equations have a unique solution, so any well-defined limit must be correct, and thus (8) is still valid (up to the first-order-in-ΔV approximation that we made in the smoothed system).

In [22], we derived the correct perturbation theory for a smooth boundary perturbation, one that is locally flat with a shift amplitude Δh, by a method involving anisotropic smoothing of (7). We can now derive the same result via (8), by computing the correct volume current using a procedure like that in the previous section. The quasi-static dipole moment of a flat surface perturbation is easily solved analytically, and one finds that $\mathbf{J} = \Delta h(\Delta\varepsilon\ \mathbf{E}_\parallel - \varepsilon\ \Delta(\varepsilon^{-1})\ \mathbf{D}_\perp)\delta(x)$; when plugged in to (8), this $\mathbf{J}$ gives exactly the surface-integral expression of [22, 30]. Our result is now more general, however, because we can apply it to small "bumps" that are *not* locally flat. (A similar application of $\mathbf{J}$ arises in coupled-mode theory [14, 23].)

Moreover, Eq. (8) gives us a direct method to verify both the $\mathbf{J}$ of Eq. (4) and the numerical polarizabilities computed in Sect. 2.3: we simply compute the exact shift $\Delta\omega$ by a brute-force calculation and compare to the first-order equation, where the latter should give the correct $\Delta\omega$ for small ΔV and must give *exactly* the correct slope $\mathrm{d}\Delta\omega/\mathrm{d}\Delta V$ at $\Delta V = 0$. In particular, we consider a 2d square lattice (lattice constant a) of dielectric squares (side $0.3a$, $\varepsilon = 12$) in air ($\varepsilon = 1$), with a semicylindrical bump on the right side of every square (see insets of Fig. 5). This structure has a mirror plane bisecting the square and the bump, and so we compute two fundamental (lowest-ω) TE-polarized ($\mathbf{E}$ in plane) modes at a Bloch wavevector $k_x = 0.6\pi/a$ (pointing towards the bump): one mode which is even with respect to the mirror plane ($\mathbf{E}$ mainly in the x direction) and one mode which is odd with respect to the mirror plane ($\mathbf{E}$ mainly in the y direction). The two frequencies vs. the area $\Delta A = \pi r^2/2$ of the bump (where negative ΔA indicates a "negative," inward bump) are plotted in Fig. 5 as computed by a planewave method [33] with 512 pixels/a resolution, along with the first-order perturbation prediction from Eqs. (4) and (8) using the α and γ of Fig. 3. (The even mode depends on $\alpha_\parallel$ and the odd mode depends on $\gamma_\perp$.) Several interesting features are apparent in this plot. First, the ω curves have a kink in their slope at $\Delta A = 0$,

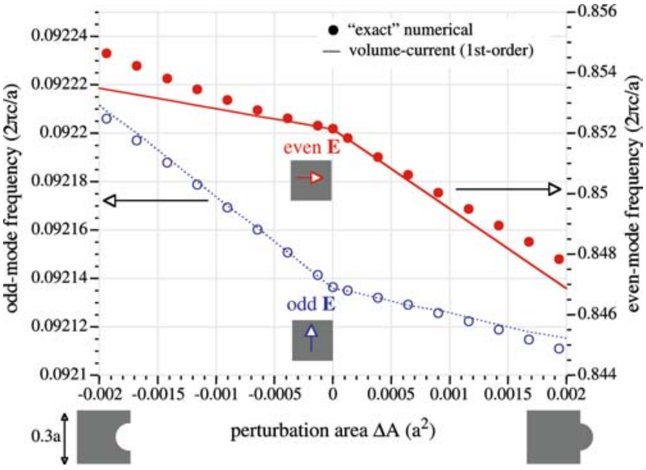

FIGURE 5 Comparison of **J**-based first-order perturbation theory (lines) with exact numerical calculations (dots) for the frequency shifts of two modes in a square lattice of square rods, due to a positive/negative bump in the rods, vs. bump area. The first-order theory correctly predicts the discontinuous slopes at $\Delta A = 0$

which arises precisely from the fact that the polarizabilities α and γ are different for positive and negative bumps. Second, the exact ω curves become tangent to the perturbation-theory curves as they approach $\Delta A = 0$, which verifies that our perturbation theory, volume current, and polarizabilities are all correct. Third, the perturbation theory is accurate for larger values of ΔA with the odd mode—this is simply due to the fact that the odd mode here lies at $1/10$ the frequency of the even mode (the wavelength is 10 times as large relative to ΔA).

4 Roughness losses in photonic crystals

As shown in Fig. 1, the photonic crystals that we consider are a 2d-periodic array of dielectric rods in air and of air holes in dielectric. The rods are in a square lattice with period a, radius $0.2a$, and $\varepsilon = 12$ (similar to Si or GaAs): in 2d, this has a complete gap for the TM polarization (**E** out of the plane), and in 3d there is a gap in the TM-like (odd-symmetry) guided modes when the rods have a finite height $2a$ [20]. The holes are in a triangular lattice with period a, radius $0.35a$, and $\varepsilon = 12$: in 2d, this has a complete gap for the TE polarization (**E** in the plane), and in 3d there is a gap in the TE-like (even-symmetry) guided modes when the dielectric has finite height $0.5a$ [20]. Note that the gap in 3d is not a complete gap, because there are still radiation modes that propagate above and below the slab at all frequencies, which will lead to vertical radiation loss from roughness in a waveguide. In 3d, we will consider both a "membrane" structure where the crystal is suspended in air (as often fabricated for the hole structure [27, 34–37]) and a structure with a low-index ($\varepsilon = 2.25 \sim$ oxide) substrate/superstrate both above and below the slab and having the *same* 2d cross section (i.e., the crystal/waveguide pattern extends into the substrate). The case of an asymmetric substrate that lies only *below* the slab is discussed in Sect. 4.4.

4.1 *Waveguide modes*

In the rod structure, we form a waveguide by removing a row of rods and replacing it with a dielectric strip (with $\varepsilon = 12$ and 3d height $2a$) (same material and height as the rods) of width $0.25a$ [21], shown in Fig. 1 (top). In the hole structure, we form a waveguide by first carving out an air trench of width $2.2a$ along a nearest-neighbor direction, and then putting a dielectric strip of width $0.6a$ in its center (with $\varepsilon = 12$ and 3d height $0.5a$), as in Fig. 1 (bot), using the design from [38]; the width of the air trench is a critical parameter to prevent surface states [19] at the crystal boundary from intersecting the waveguide mode and reducing its bandwidth. In both cases, we will then compare the roughness loss rate to the loss rate for the *same* roughness in an identical *isolated* strip waveguide.

The dispersion relations for the rod and hole waveguides in 2d and 3d are shown in Figs. 6–9, as computed by an iterative planewave-expansion method with Bloch-periodic boundaries and a horizontal/vertical supercell [21, 33]. (These plots are for the air-membrane structures, but the presence of a low-index substrate/superstrate changes the curves only slightly [20].) In both cases, the waveguide mode of interest in the bandgap is essentially identical to the "straight-line" dispersion relation of the isolated strip waveguide, "folded" to fit in the Brillouin zone [14, 38]. The reason for this is that the field patterns (shown as insets in Figs. 6–9) are strongly localized in the strip, so that the crystal is only a small perturbation. Because of this, the volume current **J** for roughness on the surface of the strip is identical for the two cases, but its radiation pattern differs markedly due to the bandgap.

However, in the 2d hole-crystal waveguide, there are *two more* guided modes in the bandgap, labelled with red triangles in Fig. 7, that do not correspond to any isolated strip-waveguide mode. Rather, they are localized largely in the air trench, and are analogous to Fabry–Perot resonant modes between the two crystal "mirrors." Because of the presence of these modes, the conditions for the theorem of [14] are violated—the crystal has introduced new reflection/loss channels, which has important consequences for the losses, below. In the 3d hole-crystal waveguide, the 2d Fabry–Perot modes

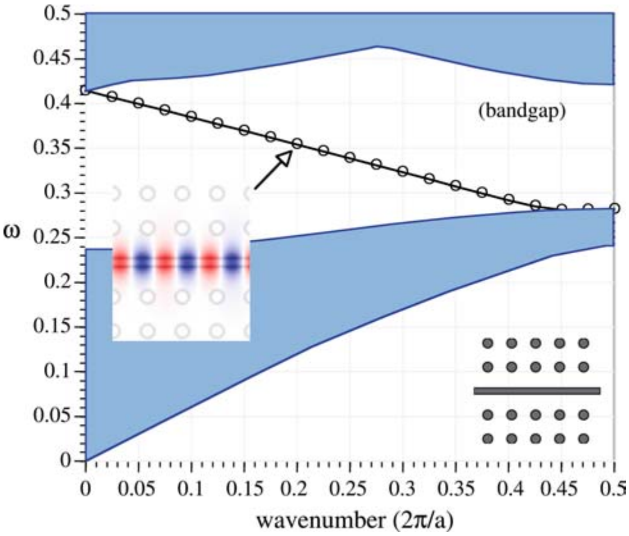

FIGURE 6 TM dispersion relation of 2d strip-waveguide mode (circles), waveguide width $0.25a$, in a dielectric rod square-lattice photonic crystal (inset). Continuum of modes propagating in crystal is shaded in blue, and E_z of guided mode is inset (blue/white/red = positive/zero/negative)

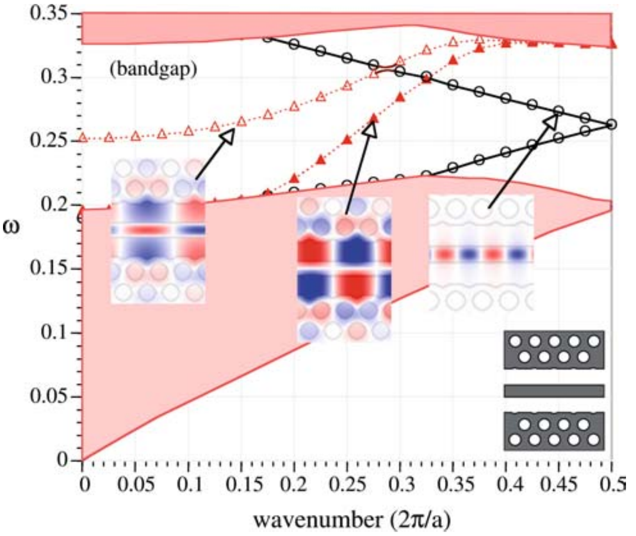

FIGURE 7 TE dispersion relation of 2d strip-waveguide mode (black circles), waveguide width $0.6a$ surrounded by $0.8a$ of air, in a triangular-lattice air-holes-in-dielectric photonic crystal (inset). Continuum of modes propagating in crystal is shaded in pink, and H_z of guided mode is inset (blue/white/red = positive/zero/negative). Two Fabry-Perot like guided modes are shown as pink triangles (hollow/filled triangles = same/opposite parity from strip guided mode)

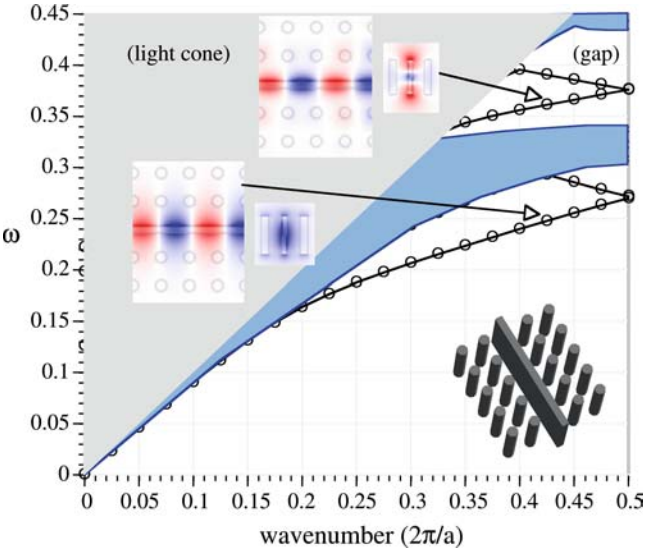

FIGURE 8 TM-like (vertically odd) dispersion relation of 3d strip-waveguide mode (circles), waveguide width $0.25a$, in a photonic-crystal slab of dielectric rods suspended in air (inset). Horizontal/vertical cross sections of E_z of guided modes are inset; note that the guided mode in the gap is higher-order in the vertical direction

lie within the light cone and are therefore not guided, but will still exist as leaky resonances. One possible way to eliminate the Fabry–Perot modes would be to shift the crystal closer to the waveguide from both sides, being careful not to bring the crystal so close that the strip waveguide mode is substantially altered.

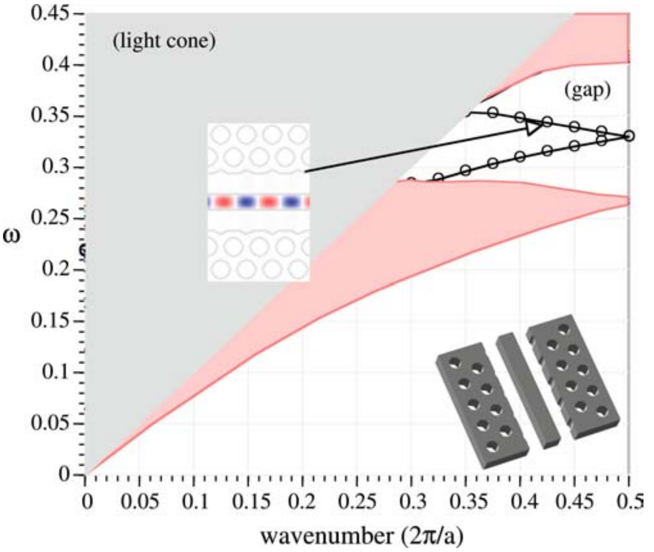

FIGURE 9 TE-like (vertically even) dispersion relation of 3d strip-waveguide mode (circles), waveguide width $0.6a$ surrounded by $0.8a$ of air, in a photonic-crystal slab of holes in a dielectric membrane (inset). Horizontal cross section of guided-mode H_z is inset

4.2 *Computing relative losses*

We compute the *relative* loss rate, the loss rate of the photonic-crystal waveguide divided by the loss rate of the isolated strip waveguide, giving an "apples to apples" measure of the effect of the bandgap. Moreover, in this way we need not specify an average amplitude $\overline{\Delta V}$ of the roughness (which would have to come from experiment)—rather, the roughness amplitude *cancels* in the relative loss rate and we thus obtain a universal, dimensionless, figure of merit for the waveguide. As another advantage, if the field is mostly polarized along one of the principal axes, as is often the case, then the effect of the roughness shape (the polarizability) also cancels. The precise computation is performed as follows.

Given the unperturbed waveguide mode's electric field $\mathbf{E}$ computed in the previous section, along with the polarizabilities α and γ for a given bump shape, we find the volume current $\mathbf{J}$ from Eq. (4). We consider two archetypical bump shapes for sidewall roughness: a semicylindrical "ridge" that runs the height of the slab (giving a line current), and a hemispherical "point" bump (giving a point-dipole current). We also consider the case of surface roughness in the crystal itself, and show that this makes a negligible contribution to the loss. For each shape, using an FDTD calculation described below, we compute the reflected power P_r radiating backwards along the waveguide and the scattered power P_s radiating outward, giving a total loss power $P_\ell = P_\text{r} + P_\text{s}$. This power is averaged over all inequivalent positions of the bumps along the waveguide sidewalls (both horizontally and vertically), and over both "positive" and "negative" bumps. Because uncorrelated roughness adds incoherently, this average power, multiplied by the average $\overline{\Delta V}$ or $\overline{\Delta A}$ would give the mean loss power $\bar{p}_\ell$ per unit length. The power in the waveguide then decays as $\exp(-\bar{p}_\ell z/P_0)$ with distance z, where P_0 is the incident power of the unperturbed mode (with the same field amplitude $\mathbf{E}$ as was used to compute $\mathbf{J}$). However, we instead compute the ratio of the loss rates for the waveguides with and without the crystal, in which case the $\overline{\Delta V}/P_0$ factor cancels.

To compute the P_r and P_s for a given $\mathbf{J}$ (i.e., the unperturbed Green's function), we employ the finite-difference time-domain method (FDTD) with PML absorbing boundary regions [39] and a resolution of 20 pixels/a in 2d and 10 pixels/a in 3d, with an anisotropic sub-pixel averaging technique that increases the effective resolution [33, 40]. (A multipole-based method to compute the Green's function was presented in [41].) Note that this resolution is at least an order of magnitude less than what would be required to resolve a specific bump shape with a realistic experimental amplitude: we only need to resolve the unperturbed crystal, since all information about the bump is contained in $\mathbf{J}$. Our computational cell is ten crystal periods along the waveguide direction and nine periods perpendicular to it; having a sufficiently large cell is important so that the radiated field is clearly distinguishable from the field coupled into the waveguide. The main advantage of FDTD is that we can compute the loss spectrum of the entire gap bandwidth at once: we input a short Gaussian pulse instead of a fixed frequency, and compute the flux spectrum from the Fourier-transformed fields. This involves an approximation, because we determine $\mathbf{J}$ from the eigenmode $\mathbf{E}$ at a particular ω; the mode field pattern for the strip waveguide varies slowly enough across the gap bandwidth, however, that the error in the relative loss rate is typically observed to be under 5%.

4.3 *Discussion of results*

In Fig. 10, we show the relative loss rate for the 2d rod crystal compared to the isolated 2d strip waveguide, both as total loss and divided into reflection and radiation loss. We consider three cases: both the original parameters from Fig. 6 (waveguide width 0.25a), a slightly wider waveguide 0.3a, and the waveguide with the crystal only on one side. In both cases with a full crystal (complete gap), the radiation loss is essentially eliminated within the bandgap, as expected—it is only nonzero because of the finite crystal size. Moreover, as predicted in [14], the reflection loss is the nearly the same with the crystal as without it. With the half crystal, the radiation loss is also reduced, but not to zero—this case is essentially that of an antenna next to a mirror, which can either decrease the radiation (as here) or increase it (as in a case below), depending upon whether the reflection from the mirror interferes constructively or destructively [42]. (Equivalently, the presence of the half-crystal alters the local density of states for the radiation continuum on the other side of the waveguide.) Given equal reflections and reduced or eliminated radiation, the reduction in the total loss rate depends upon the fraction of the loss that scatters into reflection, which depends sensitively upon the waveguide geometry as well as the scatterer type. Here, the bandgap reduces the width-0.25a loss by 20–30%, and the width-0.3a waveguide loss is reduced by 20–50%. The width-0.3a loss for surface roughness was also computed by a brute-force simulation in [14] and the results were ~40% reduction in relative loss around mid-gap, which agrees with our semi-analytical prediction here.

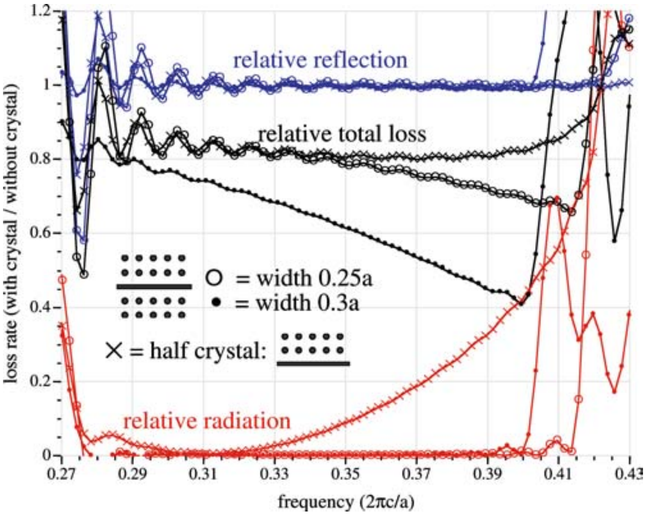

FIGURE 10 Relative loss rates of 2d strip waveguide in rod crystal (inset) compared to isolated strip waveguide, versus ω. Relative reflection (blue), radiation (red), and total (black) loss rates are shown, for three structures: waveguides of width 0.25a (hollow circles), width 0.3a (filled dots), and width 0.25a with the crystal only on one side (×'s)

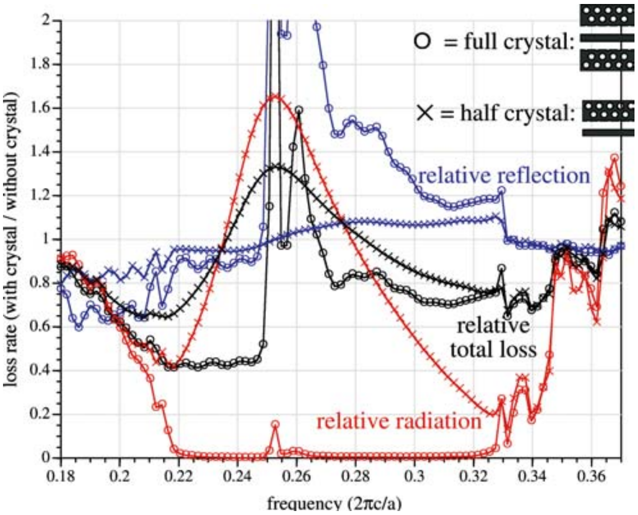

FIGURE 11 Relative loss rates of 2d strip waveguide in hole crystal (inset) compared to isolated strip waveguide, versus ω. Relative reflection (blue), radiation (red), and total (black) loss rates are shown, for two structures: the crystal on both sides of the waveguide (hollow circles) and the crystal only on one side (×'s)

The relative loss for the 2d hole crystal, in Fig. 11, is quite different from the rod case because of the additional Fabry–Perot modes created by the bandgap. We show both a full crystal as in Fig. 7 and also a half-crystal (on one side). In the full crystal, the additional bandgap modes act to *increase* the reflection compared to the isolated strip, and this increase diverges at the zero group-velocity band edge of the upper Fabry–Perot mode. In the half crystal, there are no Fabry–Perot modes, but the *radiation* loss is increased relative to the isolated waveguide, essentially by a constructive reflection off the crystal, as discussed above (this could be changed into a destructive reflection as above by shifting the crystal position). For the full crystal, the radiation is eliminated in the bandgap (other than a small spike where the zero group-velocity mode persists long enough to radiate through the finite number of crystal periods in the simulation), and thus there is a narrow bandwidth where the overall loss rate is reduced by almost 60% relative to the isolated strip.

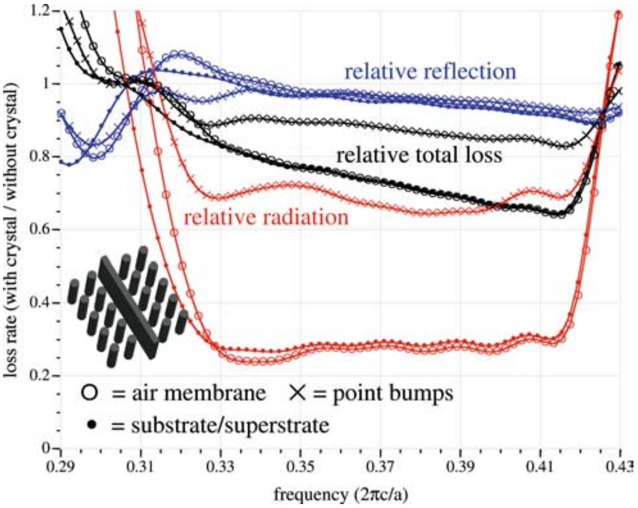

FIGURE 12 Relative loss rates of 3d strip waveguide in rod crystal (inset) for ridge-bump and point-bump roughness compared to isolated strip waveguide. Relative reflection (blue), radiation (red), and total (black) loss rates are shown, for three cases: ridge bumps for a crystal suspended in air (hollow circles) or with an $\varepsilon = 2.25$ substrate/superstrate (filled dots), and point bumps for the crystal suspended in air ($\times$'s)

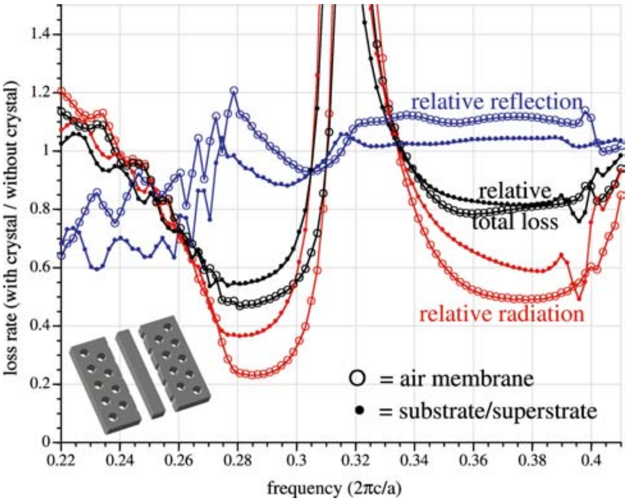

FIGURE 13 Relative loss rates of 3d strip waveguide in hole crystal (inset) for ridge-bump roughness compared to isolated strip waveguide. Relative reflection (blue), radiation (red), and total (black) loss rates are shown, for two structures: suspended in air (hollow circles) and with an $\varepsilon = 2.25$ substrate/superstrate (filled dots)

In the 3d rod crystal, we analyze both the (impractical) air-membrane structure and the $\varepsilon = 2.25$ substrate/superstrate version, both relative to the isolated strip waveguide. Figure 12 shows the relative loss for both "ridge" and "point" roughness. In both structures, the reflection is again nearly identical to that of the isolated strip, which must be the case from [14] since this is a single-mode waveguide with a field pattern close to that of the isolated strip (however, over a short distance, leaky resonance modes can alter the reflection somewhat). Because the bandgap is incomplete, however, it only blocks *some* of the radiation (vertical radiation is not blocked), and the radiation loss is reduced by 60–80% in the gap for ridge bumps. As we see below, however, that argument is a bit too simplistic—the crystal can alter the local density of states above the light line and even increase the radiation loss, so we were somewhat "lucky" with this structure. The total loss, in this case, is therefore reduced by 20–40%. The presence of the substrate/superstrate only increases radiation loss slightly, thanks to the strong field localization. Also shown are the losses for point (hemispherical) bumps for the membrane structure, in which case the crystal only reduces the radiation loss by around 30%. The reason for this is twofold: first, a point dipole (for a point bump) radiates more isotropically than a line dipole (for a ridge bump), so the former has more vertical radiation that is not blocked by the gap; second, a point bump breaks the mirror symmetry of the crystal (except for bumps exactly on the mid-plane) and therefore can couple to the even-symmetry modes that have no bandgap at all [20].

Figure 13 shows the relative loss for "ridge" roughness in the 3d hole crystal. We see a large increase in the radiation loss around mid-gap—this is analogous to what happened in the 2d case, except here the analogue of the 2d Fabry–Perot modes lie above the light line, contributing as leaky resonances to the radiation loss rather than reflection. Thus, we see that in an incomplete gap such as the photonic-crystal slab here, the effect of the crystal on scattering is not necessarily beneficial because the radiation's local density of states may be either decreased (as for the rods) or increased (as here). The reflection, at least, is roughly the same as for the strip because the waveguide is now single-mode, as predicted by [14], where the slight (10%) increase is due to a finite-size effect from the leaky resonance mode. Nevertheless, there is again a narrow bandwidth where the total loss rate is decreased by 50%, a reduction that we are hopeful might be observable in experiment. Moreover, we emphasize that there is no intrinsic reason that hole topology should have Fabry–Perot modes/resonances while a rod topology does not—the frequency (and existence) of such modes can be altered by changing the distance of the crystal from the waveguide, just as the existence of surface states is controlled by the crystal termination. Finally, the substrate increases the radiation loss more substantially than for the rod crystal, above, presumably because the slab here is $1/4$ as thick and thus more easily affected by the substrate.

4.4 *Other crystal loss mechanisms*

Above, we presented the relative losses caused by surface roughness on the waveguide with and without the photonic crystal. However, in general one must also consider whether the crystal itself introduces *new* loss mechanisms, rather than simply altering the density of states for the original roughness. In particular, there are two such new loss mechanisms: surface roughness and disorder within the crystal itself, and losses due to crystal asymmetry that upsets the bandgap confinement. Also, the crystal creates zero group-velocity points, e.g. at the Brillouin-zone edge (more prominent in waveguide modes that "see" the periodicity more strongly), and this exacerbates losses of all kinds.

We studied the effect of surface roughness within the crystal itself by the same volume-current technique, and found that for equal amounts of disorder the contribution of crystal roughness is negligible. In particular, for both the rod and hole structures, the **E** of the waveguide mode is smaller by a factor of 10 or more at surfaces in the crystal (decaying exponentially into the crystal) compared to at the surface of the strip (except near the edges of the bandgap where confinement is lost). Thus, the current **J** is smaller by a factor of ten or more, and consequently the radiated power should be reduced by a factor of $\sim 10^2$ ($\sim |\mathbf{J}|^2$) for the same roughness. Moreover, the reflected power (in a single-mode waveguide) should be reduced by a factor of $\sim 10^4$ because the coupling of **J** back to the waveguide mode is proportional to $|\mathbf{J}^* \cdot \mathbf{E}|^2$. We confirmed these predictions by direct calculation, placing a ridge-bump source at the nearest crystal interface, and the radiation and reflection losses were smaller by a factor of almost exactly 10^2 and 10^4, respectively, compared to a bump on the waveguide surface. Thus, the only way for the crystal roughness to contribute significantly to the total loss in these structures is if the crystal is an order of magnitude more rough than the waveguide. (Of course, the contribution of the crystal will also be larger in a different waveguide structure where the field is not so isolated from the crystal.)

A second concern is that, in an asymmetric crystal structure, e.g., with a substrate below but not above the crystal, the TE-like and TM-like polarizations are hybridized and the bandgap (which is only for one symmetry in the structures here) is destroyed [20]. In this case, the waveguide mode becomes *intrinsically* leaky even for the perfect structure, and this leakage loss can sometimes be substantial [43, 44]. Because this loss is intrinsic, unrelated to the roughness amplitude, however, we cannot include it in the dimensionless measure of relative loss in the preceding section; rather, it must be compared to actual experimental roughness losses of an isolated strip waveguide in order to determine whether the bandgap will be a benefit or a detriment. Ideally, of course, one will fabricate a symmetric structure, either by a suspended air membrane or by having both a low-index substrate and superstrate as in the previous section.

Finally, there is the low group velocity that occurs at band edges in photonic crystals. The scaling of the loss with the group velocity v_{g}, all other things (e.g., field patterns) assumed equal, has been commented on before [12, 28] and its derivation is recapitulated here. One derivation is based on the coupled-mode theory of [14], which shows that the lowest-order loss rate of a initial mode $|i\rangle$ to a final mode $|f\rangle$ is proportional to a matrix element $|\langle i|\hat{\Delta}|f\rangle|^2$ for some operator $\hat{\Delta}$ expressing the perturbation, essentially Fermi's golden rule [45]. Since the fields in this expression are normalized to carry a fixed power, that means that the amplitude $|i\rangle \sim 1/v_{\mathrm{g}}$ and thus the loss goes inversely with the product of the group velocities $1/(v_{\mathrm{g}}^{(\mathrm{i})} v_{\mathrm{g}}^{(\mathrm{f})})$. (The case where $|f\rangle$ is part of a radiation continuum is somewhat more complicated: one integrates over the continuum multiplied by a density of states, and $1/v_{\mathrm{g}}^{(\mathrm{f})}$ divergences are more generally Van Hove singularities [46].) This means that there is a $1/v_{\mathrm{g}}$ divergence whenever the initial *or* a (discrete) final mode approaches zero group velocity, and a divergence as $1/v_{\mathrm{g}}^2$ for reflections (initial = final) when the initial mode's velocity goes to zero. (One example of this is the coupling to the reflected Fabry–Perot mode in Fig. 11.) A similar scaling was argued directly from the Green's functions in [12]. Thus, disorder-induced reflections are likely to be especially problematic in slow-light waveguides (such as coupled-cavity waveguides [47]) [37]. Intrinsic loss mechanisms, such as material absorption, will scale as $1/v_{\mathrm{g}}$ from Eq. (7). Note that these scalings are the losses rates per unit *distance*; the loss rates per *time*, which are more appropriate in many cases where the devices can be shrunk proportional to v_{g} [48], are multiplied by a factor of v_{g} and so are somewhat better.

5 Concluding remarks

In this paper, we have removed the limitations of previous volume-current/Green's-function methods, and indeed of most other perturbative methods, so that they can now be applied to small surface imperfections with arbitrarily high index contrasts. Moreover, we have argued for the utility of a dimensionless figure of merit, comparing the losses of one structure relative to another, that applies universally to all systems with roughness of the given type regardless of the amplitude (as long as it is weak). In applying these approaches to two canonical photonic-crystal systems in 2d and 3d, we have merely scratched the surface of what we believe will be the wide applicability of this method. Indeed, in three dimensions, one hardly has any alternative, because brute-force methods with sufficient resolution to resolve the shapes of small roughness "bumps" become excessively expensive in 3d. One exciting possibility that now lies open is to design a photonic-crystal waveguide specifically to minimize the roughness loss relative to a comparison structure, via direct optimization over the waveguide parameters.

Even in the limited example structures and roughness types considered here, however, a number of exciting results emerge. Foremost among them is the prediction that structures such as the strip waveguide in a suspended hole-slab membrane, similar to designs already being fabricated [27, 34–37], should exhibit a measurable reduction in roughness loss compared to an isolated strip with comparable disorder. Moreover, to the extent that realistic roughness resembles the uncorrelated ridge-bumps and point-bumps assumed here, we hope that a quantitative match with the experimental loss spectrum may be obtained. As we have emphasized before [14], a key idea is to employ a waveguide structure that is directly comparable to a conventional waveguide in order to gauge the effect of the bandgap, whereas with a waveguide of completely different mode structure it is more difficult to judge the relative importance of a particular fabrication process and alterations in the waveguide mode versus the intrinsic impact of the bandgap.

Perhaps equally interesting is our quantitative explanation, rooted in a direct consideration of the unperturbed band structure, of why a photonic crystal is not always an improvement even for a similar waveguide mode. Rather, one must consider whether the crystal introduces new guided modes (especially ones with loss-enhancing low group velocity) or, in the case of a crystal with an incomplete bandgap, whether the density of radiating modes is substantially increased (e.g., by leaky resonant modes "hidden" above the light line). These considerations should form a useful guide to future photonic-crystal designs. On the other hand, unless new guided modes are introduced by the crystal, we have confirmed the prediction of [14] that the mere addition of a bandgap does not alter reflection loss for equivalent waveguide modes.

An important topic for future research is that of correlated disorder, in which one cannot simply sum the scattered powers incoherently from different positions. The corrections to the volume current **J** will also be equally necessary in such analyses, however (versus previous methods for correlated roughness in non-periodic waveguides that were limited to low index contrasts [4] or required brute-force computation of many radiation states [26]).

ACKNOWLEDGEMENTS This work was supported in part by the Materials Research Science and Engineering Center program of the National Science Foundation under award DMR-9400334. We are also grateful to S. Hughes for helpful discussions regarding group-velocity scaling.

REFERENCES

1 A.W. Snyder, J.D. Love, *Optical Waveguide Theory* (London, Chapman and Hall, 1983)
2 W.C. Chew, *Waves and Fields in Inhomogeneous Media* (New York, NY, IEEE Press, 1995)
3 M. Kuznetsov, H.A. Haus, IEEE J. Quantum Electron. **19**, 1505 (1983)
4 F.P. Payne, J.P.R. Lacey, Opt. Quantum Elec. **26**, 977 (1994)
5 R.H. Jordan, D.G. Hall, Appl. Phys. Lett. **64**, 3077 (1994)
6 B.E. Little, S.T. Chu, Opt. Lett. **21**, 1390 (1996)
7 K.K. Lee, D.R. Lim, H.-C. Luan, A. Agarwal, J. Foresi, L.C. Kimerling, Appl. Phys. Lett. **77**(11), 1617 (2000)
8 Y. Li, M. Froggatt, T. Erdogan, J. Lightwave Tech. **19**(10), 1580 (2001)
9 B.Z. Steinberg, A. Boag, R. Lisitsin, J. Opt. Soc. Am. A **20**(1), 138 (2003)
10 W. Bogaerts, P. Bienstman, R. Baets, Opt. Lett. **28** (May), 689 (2003)
11 D. Gerace, L.C. Andreani, Opt. Lett. **29**(16), 1897 (2004)
12 S. Hughes, L. Ramunno, J.F. Young, J.E. Sipe, Phys. Rev. Lett. **94**, 033903, (2005)
13 M.A. Kaliteevski, J.M. Martinez, D. Cassagne, J.P. Albert, Phys. Rev. B **66**, 113101 (2002)
14 M.L. Povinelli, S.G. Johnson, E. Lidorikis, J.D. Joannopoulos, M. Soljačić, Appl. Phys. Lett. **84**, 3639 (2004)
15 W.R. Frei, H.T. Johnson, Phys. Rev. B **70**, 165116 (2004)
16 K.-C. Kwan, X. Zhang, Z.-Q. Zhang, C.T. Chan, Appl. Phys. Lett. **82**(25), 4414 (2003)
17 T.N. Langtry, A.A. Asatryan, L.C. Botten, C.M. de Sterke, R.C. McPhedran, P.A. Robinson, Phys. Rev. E **68**(2), 026611, (2003)
18 B.C. Gupta, Z. Ye, J. Appl. Phys. **94**(4), 2173 (2003)
19 J.D. Joannopoulos, R.D. Meade, J.N. Winn, *Photonic Crystals: Molding the Flow of Light*. Princeton Univ. Press, (1995)
20 S.G. Johnson, S. Fan, P.R. Villeneuve, J.D. Joannopoulos, L.A. Kolodziejski, Phys. Rev. B **60**, 5751 (1999)
21 S.G. Johnson, S. Fan, P.R. Villeneuve, J.D. Joannopoulos, Phys. Rev. B **62**, 8212 (2000)
22 S.G. Johnson, M. Ibanescu, M.A. Skorobogatiy, O. Weisberg, J.D. Joannopoulos, Y. Fink, Phys. Rev. E **65**, 066611 (2002)
23 S.G. Johnson, P. Bienstman, M. Skorobogatiy, M. Ibanescu, E. Lidorikis, J.D. Joannopoulos, Phys. Rev. E **66**, 066608 (2002)
24 M. Skorobogatiy, S.A. Jacobs, S.G. Johnson, Y. Fink, Opt. Express **10**(21), 1227 (2002)
25 M. Skorobogatiy, Phys. Rev. E **70**, 046609 (2004)
26 D. Marcuse, *Theory of Dielectric Optical Waveguides*. San Diego: Academic Press, 2nd edn., (1991)
27 M. Notomi, K. Yamada, A. Shinya, J. Takahashi, C. Takahashi, I. Yokohama, Phys. Rev. Lett. **87**, 253902 (2001)
28 S.G. Johnson, M.L. Povinelli, P. Bienstman, M. Skorobogatiy, M. Soljačić, M. Ibanescu, E. Lidorikis, J.D. Joannopoulos, in *Proc. 2003 5th Intl. Conf. on Transparent Optical Networks and 2nd European Symp. on Photonic Crystals*, **1**, 103 (2003)
29 J.H. Jang, W. Zhao, J.W. Bae, D. Selvanathan, S.L. Rommel, I. Adesida, A. Lepore, M. Kwakernaak, J.H. Abeles, Appl. Phys. Lett. **83**(20), 4116 (2003)
30 N.R. Hill, Phys. Rev. B **24**(12), 7112 (1981)
31 J.D. Jackson, *Classical Electrodynamics*. New York: Wiley, 3rd edn., (1998)
32 J. Avelin, R. Sharma, I. Hänninen, A.H. Sihvola, IEEE Trans. Anten. Prop. **49**(3), 451 (2001)
33 S.G. Johnson, J.D. Joannopoulos, Opt. Express **8**(3), 173 (2001)
34 M. Lončar, D. Nedeljković, T. Doll, J. Vučković, A. Scherer, T.P. Pearsall, Appl. Phys. Lett. **77**(13), 1937 (2000)
35 Y. Akahane, T. Asano, B.-S. Song, S. Noda, Appl. Phys. Lett. **83**(8), 1512 (2003)
36 Y. Sugimoto, Y. Tanaka, N. Ikeda, Y. Nakamura, K. Asakawa, Opt. Express **12**(6), 1090 (2004)
37 Y. Tanaka, Y. Sugimoto, N. Ikeda, H. Nakamura, K. Asakawa, K. Inoue, S.G. Johnson, Electron. Lett. **40**(3), 174 (2004)
38 W.T. Lau, S. Fan, Appl. Phys. Lett. **81**, 3915 (2002)
39 A. Taflove, S.C. Hagness, *Computational Electrodynamics: The Finite-Difference Time-Domain Method*. Norwood, MA: Artech, (2000)
40 R.D. Meade, A.M. Rappe, K.D. Brommer, J.D. Joannopoulos, O.L. Alerhand, Phys. Rev. B **48**, 8434 (1993). Erratum: S.G. Johnson, ibid. **55**, 15942 (1997)
41 D.P. Fussell, R.C. McPhedran, C.M. de Sterke, Phys. Rev. E **70**, 066608 (2004)
42 C.A. Balanis, *Antenna Theory: Analysis and Design*. New York: Wiley, 2nd ed., (1996)
43 Y. Tanaka, T. Asano, Y. Akahane, B.-S. Song, S. Noda, Appl. Phys. Lett. **82**(11), 1661 (2003)
44 Y.A. Vlasov, N. Moll, S.J. McNab, J. Appl. Phys. **95**(9), 4538 (2004)
45 C. Cohen-Tannoudji, B. Din, F. Laloë, *Quantum Mechanics*. Paris: Hermann, (1977)
46 N.W. Ashcroft, N.D. Mermin, *Solid State Physics*. Philadelphia: Holt Saunders, (1976)
47 A. Yariv, Y. Xu, R.K. Lee, A. Scherer, Opt. Lett. **24**, 711 (1999)
48 M. Soljačić, S.G. Johnson, S. Fan, M. Ibanescu, E. Ippen, J.D. Joannopoulos, J. Opt. Soc. Am. B **19**, 2052 (2002)