

Photonic Crystals: From Theory to Practice

Chapter 1

by

Steven G. Johnson

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 18, 2001

© Massachusetts Institute of Technology, 2001. All Rights Reserved.

Chapter 1

Photonic Crystals and Other Obscure Topics

1.1 In the Beginning, There Was Maxwell

The thing that hath been, it is that which shall be; and that which is done is that which shall be done: and there is no new thing under the sun.

Ecclesiastes 1:9.

In 1864, James Clerk Maxwell presented his treatise, “A dynamic theory of the electromagnetic field,” and revealed the equations underlying all phenomena in electricity and magnetism [1]. This remarkable work laid bare the deep connections between the electricity of Franklin, the magnetic lodestones of the ancient mariners, and the light that had filled the first eyes gazing into the heavens. And yet Maxwell’s equations, encompassing as they are, were only the beginning.

Of course, it is not enough to understand electricity and magnetism—one must also comprehend the stuff upon and within which those forces operate. Maxwell himself knew the behavior of certain kinds of materials, such as dielectrics (of which glass is a famous example), and many other important systems would emerge over the years: from the semiconductors in our computers to the black holes that warp the universe itself along with its light. (On the other hand, one material of Maxwell’s analyses turned out *not to exist*: the “luminiferous ether” was vanquished by relativity’s assertion that light propagates without a fixed medium.) Indeed, Maxwell’s electromagnetism has figured prominently in virtually every subsequent advance of physics, from quantum mechanics to relativity to the “standard model” of particle theory. Forgive us our myopia, however, if we allow ourselves to forget those extraordinary discoveries for a moment—what, then, remains to be found in Maxwell’s equations themselves?

An equation is only useful to the extent that you can solve it, and even in simple circumstances the solution of Maxwell's equations is not easy (as any student in an electromagnetism course will unhappily inform you). Nevertheless, the phenomena that can emerge in such basic problems were largely solved and understood by clever Maxwell and his contemporaries. Since then, it often seemed that the only remaining task was the "engineering" challenge of solving and utilizing those equations in ever more realistic and messier systems. Even the glassy fiber-optic cables that today span our oceans for telecommunications would surprise Maxwell only in their uses and in their perfection of manufacture, not in their underlying principles. The story of *condensed-matter physics*, however, has been that even simple, well-known equations can have unexpected consequences, especially in complicated circumstances.

1.2 Periodic Surprises

Milo tried very hard to understand all the things he'd been told, and all the things he'd seen, and, as he spoke, one curious thing still bothered him. "Why is it," he said quietly, "that quite often even the things which are correct just don't seem to be right?"

Norton Juster, *The Phantom Tollbooth*

Imagine shining a light through a large block of glass with a single bubble of air in it. When the light strikes the bubble, some of it will reflect and some of it will continue forward at a slightly different angle (be *refracted*). As you look through the glass, this scattered light allows you to see the bubble, perhaps with an attractive sparkling caused by all of the reflections and refractions. Picture now a second bubble in the glass, just like the first but at a different place. As before, the light will reflect and refract, this time from both bubbles, sparkling in a more intricate pattern than before. All of this is exactly predicted by Maxwell's equations, although solving for the precise details would probably require

the help of a computer. Next, suppose that we fill the glass with *millions* of bubbles, all identical and arranged in a perfectly periodic lattice extending in all directions. Surely, this is not much different from before, although the sparkles will form a far more complicated mess? (“Ugh,” says the engineer.) In fact, if you are *very* careful, you may see something quite different: *nothing at all*. What would Maxwell say?

James would be gratified to learn that he has not been contradicted—everything is predicted by his equations, as always. The engineer, working out the integrals carefully on her computer, would discover to her amazement that all of those fiendish reflections and refractions *cancel* one another before reaching our eyes, and the light passes through the crystal unimpeded. This is no accident, but is a general consequence of the periodicity of the structure—actually, the explanation is closely related to one of the great mysteries of physics from a century ago, in which the cancellation was of electrons rather than of light.

When you flip a wall switch to illuminate the dirty socks on your floor, an electric current—a flow of tiny, charged particles called *electrons*—carries your good intentions to the bulb. Unfortunately, this current requires constant power to maintain, as your monthly electric bill attests; otherwise, the electrons will halt their flow after caroming off an atom or two in the wire (generating heat and light as a side effect). Scientists of the 19th century, however, couldn’t leave well-enough alone: they set out to calculate and measure the rate at which electrons bounced off atoms. What they found was terribly strange. You see, the forces between charged particles are enormous—powerful enough to keep an electron hurtling around a Hydrogen nucleus at a speed of over 1000 miles per second in a circle four billionths of an inch across—so, an electron flying through the metal wire, surrounded by charged nuclei every few billionths of an inch (not to mention all the other electrons), shouldn’t get very far. And yet it was found that, in a pure chunk of a good conductor such as copper wire, an electron apparently meanders blithely past 100 or more nuclei before

bumping into one! The solution to this puzzle came in two pieces (some assembly required). First, in the 20th century, it was reluctantly realized that electrons are actually *waves*: rather than hard little bullets whizzing through space, they are smeared out in a ghostly probability envelope covering an entire volume at once—and, unlike bullets, they can interfere and even cancel if two waves are superimposed. Second, since a conductor like copper forms a periodic lattice of nuclei, a *crystal*, we can apply the lovely* theorem that we alluded to earlier with the bubbles.

This theorem, called alternatively “Bloch’s” [2] or “Floquet’s” [3] theorem depending upon whether you drink beer or wine, says that, in a *periodic* medium, waves can be found that propagate without scattering, and is equally true whether the waves consist of electrons or of light beams (or sound waves, or...). This is why the electrons can drift so far in the metal—they only bounce off of impurities or imperfections—and why a light beam could travel through our bubbly glass without reflecting towards our eyes. And that’s not all: the machinery of the Bloch-Floquet theorem leads us to many other curious phenomena...

1.3 Photonic Crystals

How do you solve a problem like $[\vec{\nabla} \times \frac{1}{\epsilon} \vec{\nabla} \times \vec{H} = \left(\frac{\omega}{c}\right)^2 \vec{H}]$?

How do you hold a moonbeam in your hand?

Oscar Hammerstein II, *The Sound of Music*

By combining Maxwell’s equations with the theorems of solid-state physics, a surprising and simple result has emerged from what seemed like a horrendous problem, that of light bouncing among an infinity of periodic scatterers. But there’s still more to be learned from the humble electron, a behavior of light that Maxwell never guessed although the equa-

* Physicists have strange attractions. (Mathematicians, contrariwise, have “strange attractors.”)

tions were right in front of his eyes: just as there are electrical insulators, which keep the currents in the wires where they belong, one can also build an optical insulator, a *photonic crystal*.

Photonic crystals again arise from the cooperation of *periodic* scatterers—thus, they are called “crystals” because of their periodicity and “photonic” because they act on light.* They can occur when the period (the separation of the bubbles or scatterers) is on the order of the wavelength of the light.† When that’s true, and if your bubbles are the right shape, it is possible that all of those reflections and refractions will cancel not only the light scattered sideways, but the light moving *forward* as well. Then, because the light has to go somewhere (energy is conserved), it has no choice but to go back—it is forbidden from entering the photonic crystal. This happens no matter what direction the light is coming from, in a certain range of wavelengths called the *photonic band gap*.‡ (A similar effect, an *electronic* band gap, is what gives electrical semiconductors their lucrative properties.)

Once you have such a medium, impervious to light, you can manipulate photons in many interesting ways. By carving a tunnel through the material, you have an optical “wire” from which no light can deviate. Even more dramatically, by making a cavity in the center of the crystal, you have an optical “cage” in which a beam of light could be caught and held—although you couldn’t easily admire it, because the very fact that it cannot escape would render it invisible. These kinds of abilities to trap and guide light have many potential applications in optical communications and computing, where one would like to make tiny optical “circuits” to help manage the ever-increasing traffic through the world’s

* Perversely, at about the same time that electrons were reborn as waves, Einstein discovered that the light waves of Maxwell’s equations are “particles” called *photons*.

† Light, like all waves, has a characteristic *wavelength*: the distance from one wave crest to the next, where each wavelength corresponds to a different color. For example, all of the colors that you can see fall in the range of about 0.4 to 0.6 microns—one micron is 1/1000 millimeters—but the wave crests are too small and move too fast for you to observe directly.

‡ Some people call these systems “photonic band-gap materials” instead of “photonic crystals.”

optical communications networks. Other devices, too, are made possible by this increased control over light: from more-efficient lasers and LED light sources, to opening new regimes for operating optical fibers, to cellular phones that don't pump half of their signal power into your head (as harmless heat, but wasteful).

What's so new about photonic crystals? Aren't they only mirrors, just like the one in which we admire ourselves every morning? Metal mirrors, however, besides operating on entirely different principles, have a major disadvantage: some of the light that strikes them is absorbed rather than reflected, and this absorption becomes intolerably large as you move towards the short wavelengths (infrared and visible) that are most important for lasers and optical communications.* On the other hand, photonic crystals rely only on weak interactions of the light with the material, and can have low absorption at nearly any wavelength.

Actually, a form of photonic crystal can be found in nature, producing the iridescent colors of abalone shells, butterfly wings, and some crystalline minerals. These beautiful materials archetypically comprise alternating microscopic layers of two substances (or two arrangements of the same substance)—they are only periodic in *one* direction, and thus form a *one-dimensional* (1d) photonic crystal, depicted in Fig. 1.1. When light shines on the layers, it is reflected by the photonic band gap.† Because the natural optical contrast between the layers is weak, however, the band gap is narrow—only a *specific color* is reflected, resulting in the brilliant hues we so admire. Moreover, as the angle of the light changes, the periodicity in that direction alters and the band gap shifts wavelengths—that is, different colors are reflected at different angles, shimmering as we tilt our heads.

* Metal mirrors work just fine for *microwave* wavelengths, and Bell Labs tried for many years to use metal tubes as microwave communications “pipes.” How these tubes eventually lost the battle against today's glass fibers makes an interesting story. (For one thing, you couldn't bend the tube.)

† The classical way to understand the 1d band gap is that, if the periodicity is half a wavelength, the cumulative reflections are *in phase*: the reflected wave crests line up, and so are additive.




Figure 1.1: Schematic of a *one-dimensional* photonic crystal: alternating dielectric layers (represented by blue and yellow), with a band gap in one direction. Incident light is reflected for colors of light in the band gap, but it cannot be used to trap light in 3d—the light inevitably escapes “sideways.”

Just such a crystalline mineral attracted the interest of Lord Rayleigh in 1887, and he was (it seems) the first to formulate a precise theory of these phenomena [4]–[5]. In fact, Rayleigh stumbled onto a form of the Bloch-Floquet theorem, deriving quite a general solution for waves in 1d periodic media;* he even correctly showed that *any* such medium exhibits a band gap (although he didn’t call it that). Unfortunately, he failed to pursue this result to its full generality—instead, he later abandoned it for a more direct (but less illuminating) method: performing a heroic calculation to sum the waves reflected and refracted from each interface [6]. No one suspected that the same magical cancellation could happen for light shining from *any* direction in a crystal periodic along all *three* axes. No one, that is, until 100 years later, in 1987, when the first papers were published marrying solid-state physics with Maxwell’s equations and suggesting that a 3d crystal could produce an omnidirectional band gap [10]–[11]. “Suggested,” because an actual *example*

* Rayleigh was guided in his analysis of periodic media by earlier studies of the lunar orbit, in which the moon is periodically tugged by the sun. If the generality of that work had been recognized, we might now call it the “Hill” theorem (since it predates Floquet by six years) [7].

(even in theory) was not found until three years later—unlike in 1d, any old periodic structure will not do, and much trial-and-error was required. The first concrete design involved levitating dielectric balls [12], so the engineers asked to build the thing were understandably upset—especially since the balls had to be sub-microns in size for it to work at infrared wavelengths. Nevertheless, after several years, more realistic structures were proposed and even fabricated, using technologies such as those by which tiny circuits are etched onto computer chips.

The challenge of realizing the benefits of photonic crystals in practical systems continues to be a major subject for research. The only such devices currently in production are still based on one-dimensional band gaps like the crystals of Lord Rayleigh (often called *Bragg mirrors*^{*}), and range from dielectric mirrors and reflective coatings to DFB lasers.

Before we continue, we must take a moment to apologize: if it has seemed to you that our explanations of photonic crystals have often failed to explain, then we are in complete agreement. A more substantial understanding can only come after one has had the pleasure of swimming through the actual mathematics of Maxwell's equations and Bloch-Floquet eigensystems, and that would require a work somewhat longer than this chapter. Fortunately, we are able to recommend to the interested reader a marvelous introductory text by an author unrivaled in his eloquence and clarity, and whose signature is coincidentally required on this thesis. That book, cited in Ref. [8], should be accessible to anyone with an undergraduate background in electromagnetism. (A briefer review can be found in the *Nature* article of Ref. [9].) Alternatively, Chapter 2 of this thesis provides a whirlwind tour for the impatient and the mathematically inclined.

* Named after Sir William Henry Bragg, who in 1915 received the Nobel prize in physics for his work on X-ray crystallography—probing periodic *atomic* crystals by bouncing X-rays off them.

1.4 Things to Come

*I'm very well acquainted too with matters mathematical,
I understand equations, both the simple and quadratical,
About binomial theorem I'm teeming with a lot o' news—
With many cheerful facts about the square of the hypotenuse.*

Sir William Schwenk Gilbert, *The Pirates of Penzance*

Although we cannot take credit for the discoveries described in the previous sections, we will endeavor in the subsequent chapters to contribute some marginal advance to the understanding of photonic crystals and related systems. Chapter 2, however, is simply a mathematical introduction to photonic crystals, while Chapter 9 is our conclusion. The remaining chapters are derived from a series of articles published in various scientific journals, meaning that they are either pleasantly self-contained or irretrievably abstruse, depending upon one's perspective. In any case, we shall summarize them here.

1.4.1 Intersecting Optical Pipes

To begin with, we consider an interesting effect that one can achieve with photonic crystals, and more generally with the mathematical and conceptual techniques that they embody. In particular, we show how two optical “wires” can be crossed without interfering, which might be surprising if you think in terms of another wave: sound.

Imagine that you face an intersection of two corridors, deep within the bowels of MIT. On the opposite side of the intersection is one of your fellow graduate students, but you can't reach her because it's noon^{*} and the janitors have roped off the crossing to wax the floor. Instead, you call across to brag that you used the Boston White Pages to pad out your thesis: “Our batty advisor won't even notice!” Meanwhile, speak of the devil, *he* is having his own shouted conversation with a colleague across the same intersection, but in the per-

^{*} In Tech time—2AM, Boston time.

pendicular direction: “No need to feel guilty—he’s young, and has plenty of time to find another trade if we reject his dissertation.” Naturally, neither of you wants the other to overhear his words—is this possible? (No fair talking through tubes, either; both voices must pass through the *same space*.) It turns out that, in certain cases, the answer is yes.

You can see where this is heading: an almost identical situation shows up with light, and poses a substantial problem in designing optical circuits. In order to weave together a group of components, one needs the ability to cross waveguides (optical “corridors”). Doing so in the naïve way, however, would allow the crossed channels to overhear one another—this is called *crosstalk*, and will turn your device into a lemon. (With electronic circuits, this difficulty is circumvented by passing the wires over or under one another, but that solution is less practical in optics for a variety of reasons.)

In Chapter 3, we show how, by exploiting general principles of symmetry and resonance, it is possible to design an optical crossing where the two waveguides do not “see” one another—light travels across in one direction or the other, through the same space, without leaking into the perpendicular pipe. Although the inspiration and the ideal realization of this concept resides in full photonic crystals, we show how it can be applied to conventional and hybrid optical systems as well. The application to eavesdropping advisors will, perhaps, be the subject of a future publication.

1.4.2 Photonic-Crystal Slabs

Photonic crystals, in their full three-dimensional glory, are hard to make. Imagine baking ten-layer cake, where each layer is a checkerboard pattern of chocolate and vanilla, and the patterns of adjacent layers must be offset by exactly half a checker. Oh, and the cake is only 1/1000 of an inch high. Wouldn’t it be much easier to bake only a single layer? You can roll out the dough as flat as you need, and even use a cookie-cutter to make the pattern. The optical equivalent of this recipe shortcut is a *photonic-crystal slab*: a




Figure 1.2: Schematic of the photonic-crystal slab chapters, covering first the bulk systems, then waveguides formed by linear defects, then cavities formed by point defects.

dielectric (*e.g.*, Silicon) with only a two-dimensional pattern and a finite thickness, which can be baked via the well-known lithographic techniques used to manufacture computer chips. Slabs are very attractive to the engineers, but they involve a compromise—what confines the light in the vertical direction? Actually, there *is* a confinement mechanism, albeit imperfect: *index guiding*, sometimes misleadingly called *total internal reflection*.

Index guiding is the principle behind today’s optical fibers; simply put, it is the tendency of light to stay in stronger dielectrics (which have a higher *index of refraction*) compared to weaker dielectrics such as air. Just as a bump in the road may throw your car into the ditch, however, any irregularity in the medium can cause light to scatter out, and this ease of loss makes life more difficult with a slab than in a full 3d crystal (where the road is surrounded by impenetrable walls). Fortunately, thanks again to the Bloch-Floquet theorem, at least the *periodic* variation of the slab itself does not cause scattering; if only regularly-spaced bumps did not disturb your driving!

In the next three chapters, as depicted in Fig. 1.2, we explore and explain the unique phenomena that can occur in such hybrid systems, probing both the limitations and the possibilities of photonic-crystal slabs. Chapter 4 lays the foundations by characterizing the




Figure 1.3: A portion of a novel 3d photonic crystal that we propose in Chapter 7. This structure has a large, complete band gap, can be constructed layer-by-layer using traditional lithography, and its behaviour is understandable in terms of much-simpler two-dimensional photonic-crystal structures.

bulk slabs themselves, without defects. Chapter 5 breaks the slab’s periodicity in only one direction by a linear defect, showing how lossless waveguides can be formed. Finally, Chapter 6 breaks the periodicity entirely and attempts to trap light in a cavity. This is not possible, however, due to the lack of a full band gap, and so we propose and investigate various mechanisms to keep the light confined for as long as we can.

1.4.3 A New Photonic Crystal

Despite their difficulty of manufacture, tremendous progress has been made in the fabrication of full three-dimensional photonic crystals. Thus far, however, the menu of practical structures has been limited to a mere handful of possibilities, and so the prospect of new designs with unique and desirable properties holds great attraction in the field. In Chapter 7, we propose a novel photonic crystal that can be (and, in fact, has been) realized even at microscopic lengthscales. Although it may look complex, such a periodic system can easily be analyzed in its entirety with the help of our friends Bloch and Floquet.

Our structure, depicted in Fig. 1.3, has a number of appealing qualities whose combination is unique. First, its photonic band gap covers a wide range and persists even if only weak dielectric materials are employed. Second, it consists of a sequence of patterned layers that can be constructed one at a time via the “cookie-cutter” of modern micro-lithography. Third, each layer has high symmetry (it looks the same from six directions in the plane). Physicists love symmetry for its own sake, but here it has substantial practical advantages as well—it means that complex optical circuits can be constructed by modifying only a *single* layer of the structure. Finally, each layer resembles a common *two-dimensional* photonic-crystal structure, and this allows us to understand its behavior in terms of those vastly simpler systems (for which a substantial body of literature, intuition, and computation is already available).

1.4.4 So, You Want To Study Photonic Crystals?

Finally, we consider a matter less of physics than of mathematics—given the equations of Maxwell and Bloch-Floquet, it is still a substantial challenge to compute their solutions. Several methods are already known, of course, but in Chapter 8 we attempt to bring together the best developments in numerical analysis, computer science, and physics in order to take a certain class of techniques further towards its ultimate conclusion, and to provide a comprehensive review and comparison of the ideas of ourselves and others. In this way, we hope to furnish a valuable reference for the reader who wishes to engage in these kinds of computations. We offer more than that, however...

In the literature, all too often one sees a proposal for a new computational algorithm, usually with a few sample results, that nevertheless leaves the reader with no way to evaluate or employ the method for herself without investing several months or even years in programming her own implementation on a computer. Moreover, despite the fact that an entire field of scientists and engineers runs the same type of computation on similar types

of problems, to a large extent each group of scientists develops their own software. This duplication of effort is, it seems to us, extraordinarily wasteful and unnecessary. With that in mind, and in the scientific spirit of contributing freely to human knowledge, we have made the results of our work available as free software on the Web—see Ref. [13], <http://ab-initio.mit.edu/mpb/>. Using it, we invite you analyze for yourself the band gaps and the propagation of light in arbitrary three-dimensional photonic crystals.

*And like a silver clarion rung
The accents of that unknown tongue,
Excelsior!*

Henry Wadsworth Longfellow, “Excelsior”

References

- [1] J. Clerk Maxwell, “A dynamical theory of the electromagnetic field,” *Philosophical Transactions of the Royal Society of London* **155**, 459–512 (1865). Abstract: *Proceedings of the Royal Society of London* **13**, 531–536 (1864).
- [2] F. Bloch, “Über die quantenmechanik der electronen in kristallgittern,” *Z. Physik* **52**, 555–600 (1928).
- [3] G. Floquet, “Sur les équations différentielles linéaires à coefficients périodiques,” *Ann. École Norm. Sup.* **12**, 47–88 (1883).
- [4] Lord Rayleigh, “On the maintenance of vibrations by forces of double frequency, and on the propagation of waves through a medium endowed with a periodic structure,” *Philosophical Magazine* **24**, 145–159 (1887).
- [5] Lord Rayleigh, “On the remarkable phenomenon of crystalline reflexion described by Prof. Stokes,” *Philosophical Magazine* **26**, 256–265 (1888).
- [6] Lord Rayleigh, “On the reflection of light from a regularly stratified medium,” *Proceedings of the Royal Society of London* **93**, 565–577 (1917).
- [7] G. W. Hill, “On the part of the motion of the lunar perigee which is a function of the mean motions of the sun and moon,” *Acta Math.* **8**, 1–36 (1886). (This is a reprint of an article published in 1877.)
- [8] J. D. Joannopoulos, R. D. Meade, and J. N. Winn, *Photonic Crystals: Molding the Flow of Light* (Princeton, September 1995).
- [9] J. D. Joannopoulos, P. R. Villeneuve, and S. Fan, “Photonic crystals: putting a new twist on light,” *Nature* **386**, 143–149 (1997).
- [10] E. Yablonovitch, “Inhibited spontaneous emission in solid-state physics and electronics,” *Phys. Rev. Lett.* **58**, 2059–2062 (1987).
- [11] S. John, “Strong localization of photons in certain disordered dielectric superlattices,” *Phys. Rev. Lett.* **58**, 2486–2489 (1987).
- [12] K. M. Ho, C. T. Chan, C. M. Soukoulis, “Existence of a photonic gap in periodic dielectric structures,” *Phys. Rev. Lett.* **65**, 3152–3155 (1990).
- [13] S. G. Johnson and J. D. Joannopoulos, The MIT Photonic-Bands Package home page, <http://ab-initio.mit.edu/mpb/>.