JOURNAL OF MODERN OPTICS, 1994, VOL. 41, NO. 2, 173–194

Photonic crystals

E. YABLONOVITCH
Department of Electrical Engineering,
University of California at Los Angeles,
Los Angeles, California 90024-1594, USA

(*Received 11 August 1993 and accepted 26 August 1993*)

Abstract. In this paper, we review the early motivation for photonic crystal research which was derived from the need for a photonic bandgap in quantum optics. This led to a series of experimental and theoretical searches for the elusive photonic bandgap structures: those three-dimensionally periodic dielectric structures which are to photon waves, as semiconductor crystals are to electron waves. We shall describe how the photonic semiconductor can be 'doped', producing tiny electromagnetic cavities. Finally, we shall summarize some of the anticipated implications of photonic band structure for quantum electronics and the prospects for the creation of photonic crystals in the optical domain.

1. Introduction

In this paper we shall pursue the rather appealing analogy [1, 2] between the behaviour of electromagnetic waves in artificial, three-dimensionally periodic dielectric structures, and the rather more familiar behaviour of electron waves in natural crystals.

These artificial two- and three-dimensionally periodic structures we shall call 'photonic crystals'. The familiar nomenclature of real crystals will be carried over to the electromagnetic case. This means that the concepts of reciprocal space, Brillouin zones (BZs), dispersion relations, Bloch wavefunctions, Van Hove singularities, etc., must now be applied to photon waves. It makes sense then to speak of photonic band structure and of a photonic reciprocal space which has a Brillouin zone approximately 1000 times smaller than the Brillouin zone of the electrons. Owing to the periodicity, photons can develop an effective mass, but this is in no way unusual, since it occurs even in one-dimensionally periodic, optically layered structures. We shall frequently leap back and forth between the conventional meaning of a familiar concept such as 'conduction band', and its new meaning in the context of photonic band structure.

Under favourable circumstances, a 'photonic bandgap' can open up: a frequency band in which electromagnetic waves are forbidden, irrespective of propagation direction in space. Inside a photonic bandgap, optical modes, spontaneous emission and zero-point fluctuations are all absent. Because of its promised ability to control spontaneous emission of light in quantum optics, the pursuit of a photonic bandgap has been a major motivation for studying photonic band structure.

2. Motivation

Spontaneous emission of light is a major natural phenomenon, which is of great practical and commercial importance. For example, in semiconductor lasers, spontaneous emission is the major sink for threshold current, which must be

surmounted in order to initiate lasing. In heterojunction bipolar transistors (HBTs), which are non-optical devices, spontaneous emission nevertheless rears its head. In some regions of the transistor current–voltage characteristic, spontaneous optical recombination of electrons and holes determines the HBT current gain. In solar cells, surprisingly, spontaneous emission fundamentally determines the maximum available output voltage. We shall also see that spontaneous emission determines the degree of photon number state squeezing, an important new phenomenon [3] in the quantum optics of semiconductor lasers. Thus the ability to control spontaneous emission of light is expected to have a major impact on technology.

The easiest way to understand the effect of a photonic bandgap on spontaneous emission is to take note of Fermi's golden rule. The downward transition transition rate w between the filled and empty atomic levels is given by

$$w = \frac{2\pi}{\hbar} |V|^2 \rho(E), \tag{1}$$

where $|V|$ is sometimes called the zero-point Rabi matrix element and $\rho(E)$ is the density of final states per unit energy. In spontaneous emission, the density of final states is the density of optical modes available to the photon. If there are no optical modes available, there will be no spontaneous emission.

Before the 1980s, spontaneous emission was often regarded as a natural and inescapable phenomenon, one over which no control was possible. In spectroscopy it gave rise to the term 'natural linewidth'. However, in 1946, an overlooked note by Purcell [4] on nuclear spin levels had already indicated that spontaneous emission could be controlled. In the early 1970s, the interest in this phenomenon was re-awakened by the surface-adsorbed dye molecule fluorescence studies, of Drexhage [5]. Indeed during the mid-1970s, Bykov [6] proposed that one-dimensional periodicity inside a coaxial line could influence spontaneous emission. The modern era of inhibited spontaneous emission dates from the Rydberg atom experiments of Kleppner. A pair of metal plates acts as a waveguide, with a cut-off frequency for one of the two polarizations. Rydberg atoms are atoms in very high-lying principal quantum number states, which can spontaneously emit in the microwave region of wavelengths. Kleppner and coworkers [7] showed that Rydberg atoms in a metallic waveguide could be prevented from undergoing spontaneous decay. There were no modes available below the waveguide cut-off.

There is a problem with metallic waveguides, however. They do not scale well into optical frequencies. At high frequencies, metals become more and more lossy. These dissipative losses allow for virtual modes, even at frequencies which would normally be forbidden. Therefore it makes sense to consider structures made of positive-dielectric-constant materials such as glasses and insulators, rather than metals. These can have very low dissipation, even all the way up to optical frequencies. This is ultimately exemplified by optical fibres which allow the light propagation over many kilometres, with negligible losses. Such positive-dielectric-constant materials can have an almost purely dielectric response with low resistive losses. If arrayed into a three-dimensionally periodic dielectric structure, a photonic bandgap should be possible, employing a purely real reactive dielectric response.

The benefits of such a photonic bandgap for direct gap semiconductors are illustrated in figure 1. On the right-hand side of the figure is a plot of the photon dispersion (frequency against wave-vector). On the left-hand side of figure 1, sharing the same frequency axis, is a plot of the electron dispersion, showing conduction and

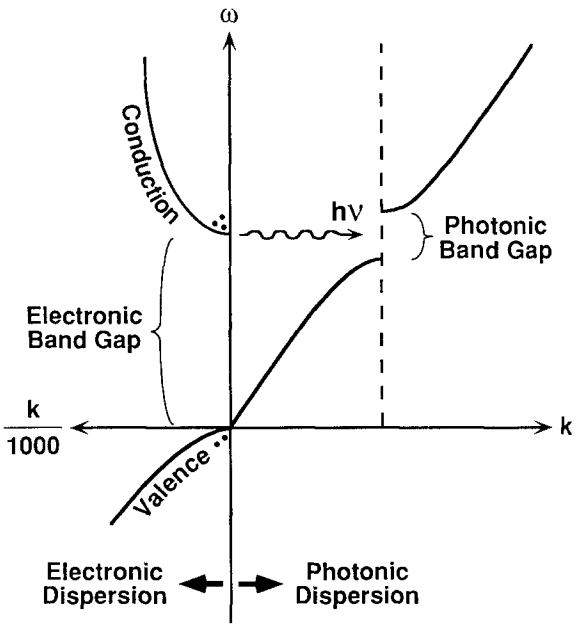

Figure 1. On the right-hand side is the electromagnetic dispersion, with a forbidden gap at the wave-vector of the periodicity. On the left-hand side is the electron wave dispersion typical of a direct-gap semiconductor, the full circles representing electrons and holes. Since the photonic bandgap straddles the electronic band edge, electron–hole recombination into photons is inhibited. The photons have no place to go.

valence bands appropriate to a direct-gap semiconductor. Since atomic spacings are 1000 times shorter than optical wavelengths, the electron wave-vector must be divided by 1000 in order to fit on the same graph with the photon wave-vectors. The full circles in the electron conduction and valence bands are meant to represent electrons and holes respectively. If an electron were to recombine with a hole, they would produce a photon at the electronic band-edge energy. As illustrated in figure 1, if a photonic bandgap straddles the electronic band edge, then the photon produced by electron–hole recombination would have no place to go. The spontaneous radiative recombination of electrons and holes would be inhibited. As can be imagined, this has far-reaching implications for semiconductor photonic devices.

One of the most important applications of inhibited spontaneous emission is likely to be the enhancement of photon-number-state squeezing, which has been playing an increasing role in quantum optics lately. The form of squeezing introduced by Yamamoto *et al.* [3] is particularly appealing, in that the active element producing the squeezing effect is none other than the common resistor. When an electrical current flows, it generally carries the noise associated with the graininess of the electron charge, called shot noise. The corresponding mean square current fluctuations are

$$\langle(\Delta i)^2\rangle = 2ei\Delta f, \tag{2}$$

where i is the average current flow, e is the electronic charge and Δf is the noise bandwidth. While equation (2) applies to many types of random physical processes,

it is far from universal. Equation (2) requires that the passage of electrons in the current flow be a random Poissonian process. As early as 1954, Van der Ziel [8], in an authoritative book called *Noise*, pointed out that good-quality metal film resistors, when carrying a current, generally exhibit much less noise than given by equation (2). Apparently, the flow of electrons in the Fermi sea of a metallic resistor represents a highly correlated process. Far from being a random process, the electrons apparently sense one another, producing shot noise far below equation (2) (so low as to be difficult to measure and to distinguish from thermal or Johnson noise). Sub-Poissonian shot noise has the following meaning. Suppose that the average flow consists of ten electrons per nanosecond. Under random flow, the count in successive nanoseconds could sometimes vary from eight to 12 electrons. With good-quality metal film resistors, the electron count would be ten for each and every nanosecond.

Yamamoto *et al.* put this property to good use by driving a high-quantum-efficiency laser diode with such a resistor. Suppose that the laser diode quantum efficiency into the cavity mode were 100%. Then for each electron passing through the resistor there would be one photon into the laser cavity mode. A correlated stream of photons are produced whose statistical properties are unprecedented ever since Einstein's interpretation of the photoelectric effect. If the photons are used for optical communication, then a receiver would detect exactly ten photoelectrons each nanosecond. If 11 photons were detected, it would be no mere random fluctuation but would represent an intentional signal. Thus information in an optical communications signal could be encoded at the level of individual photons. The term photon-number-state squeezing is associated with the fixed photon number per unit time interval. Expressed differently, the bit-error rate in optical communication can be diminished by squeezing.

There is a limitation to the squeezing, however. The quantum efficiency into the lasing mode is not 100%. The 4π sterad outside the cavity mode can capture a significant amount of random spontaneous emission. If unwanted electromagnetic modes capture 50% of the excitation, then the maximum noise reduction in squeezing would be only 3 dB. Therefore it is necessary to minimize the spontaneous recombination of electrons and holes into modes other than the lasing mode. If such random spontaneous events were reduced to 1%, allowing 99% quantum efficiency into the lasing mode, the corresponding noise reduction would be 20 dB, well worth fighting for. Thus we see that control of spontaneous emission is essential for deriving the full benefit from photon-number-state squeezing.

We have motivated the study of photonic band structure for its applications in quantum optics and optical communications. Positive dielectric constants and fully three-dimensional forbidden gaps were emphasized. It is now clear that the generality of artificial multidimensional band structure concepts allows for other types of wave, other materials, and various lower-dimensional geometries, limited only by imagination and need.

3. Search for the photonic bandgap

Having decided to create a photonic bandgap in three dimensions we need to settle on a particular three-dimensionally periodic geometry. For electrons, the three-dimensional crystal structures come from nature. Several hundred years of mineralogy and crystallography have classified the naturally occurring, three-dimensionally periodic lattices. For photonic bandgaps, however, we must create an artificial structure using our imagination.

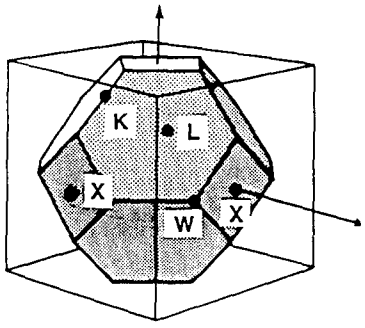

Figure 2. The f.c.c. BZ in reciprocal space.

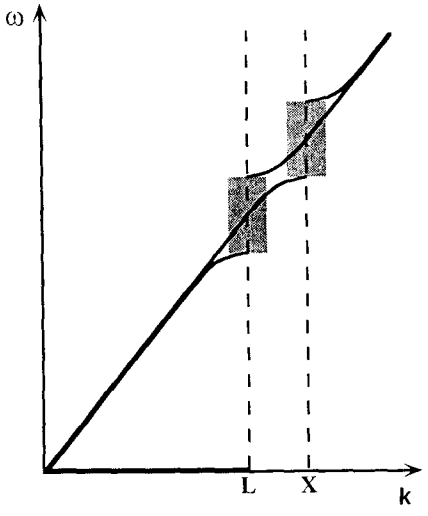

Figure 3. The forbidden gap (shaded) at the L point is centred at a frequency about 14% lower than the X-point forbidden gap. Therefore it is difficult to create a forbidden frequency band overlapping all points along the surface of the BZ.

The f.c.c. lattice appears to be favoured for photonic bandgaps and was suggested independently by Yablonovitch [1] and John [2] in their proposals. Let us consider the f.c.c. BZ as illustrated in figure 2. Various special points on the surface of the BZ are marked. Closest to the centre is the L point oriented towards the body diagonal of the cube. Farthest away is the W point, a vertex where four plane waves are degenerate (which will cause problems later). In the cubic directions are the familiar X points.

Consider a plane wave in the X direction. It will sense the periodicity in the cubic direction, forming a standing wave, opening up a forbidden gap as indicated by the shading in figure 3. Suppose on the other hand that the plane wave is going in the L direction. It will sense the periodicity along the cubic body diagonal, and a gap will form in that direction as well. However, the wave-vector to the L point is about 14% smaller than the wave-vector to the X point. Therefore the gap at L is likely to be centred at a 14% smaller frequency than the gap at X. If the two gaps are not wide

enough, they are unlikely to overlap in frequency. In figure 3 as shown, the two gaps barely overlap. This is the main problem in achieving a photonic bandgap. It is difficult to ensure that a common frequency overlap is assured for all possible directions in reciprocal space.

The lesson from figure 3 is that the BZ should most closely resemble a sphere in order to increase the likelihood of a common frequency overlap in all directions of space. Therefore let us examine two common BZs: the f.c.c. BZ and the b.c.c. BZ. The b.c.c. BZ has pointed vertices which make it difficult to achieve a common frequency overlap in all directions. Likewise most other common BZs deviate even farther from a spherical shape. Among all the common BZs the f.c.c. BZ has the least percentage deviation from a sphere. Therefore until now all photonic bandgaps in three dimensions have been based on the f.c.c. lattice. There has been a report recently of a photonic bandgap in a simple-cubic geometry [9].

The photonic bandgap is different from the idea of a one-dimensional stop band as understood in electrical engineering. Rather, the photonic bandgap should be regarded as a stop band with a common frequency overlap in all 4π sterad of space. The earliest antecedent to photonic band structure, dating [10] back to 1914 and Sir Lawrence Bragg, is the dynamical theory of X-ray diffraction. Nature gives us f.c.c. crystals and X-rays are bona fide electromagnetic waves. As early as 1914, narrow stop bands were known to open up. Therefore, what was missing?

The refractive index contrast for X-rays is tiny, generally 1 part in 10^4. The forbidden X-ray stop bands form extremely narrow rings on the facets of the BZ. As the index contrast is increased, the narrow forbidden rings open up, eventually covering an entire facet of a BZ and ultimately all directions in reciprocal space. We shall see that this requires an index contrast of about 2 to 1 or greater. The high-index contrast is the main new feature of photonic band structure beyond dynamical X-ray diffraction. In addition we shall that electromagnetic wave polarization, which is frequently overlooked for X-rays, will play a major role in photonic band structure.

In approaching this subject, we adopted an empirical viewpoint. We decided to make photonic crystals on the scale of microwaves, and then we tested them using sophisticated coherent microwave instruments. The test set-up, shown in figure 4, is what we would call in optics a Mach–Zender interferometer. It is capable of measuring phase and amplitude in transmission through the microwave scale photonic crystal. In principle, one can determine the frequency against wave-vector dispersion relations from such coherent measurements. Later we used a powerful commercial instrument for this purpose, the HP8510 network analyser. The philosophy of the experiments was to measure the forbidden gap in all possible internal directions in reciprocal space. Accordingly the photonic crystal was rotated and the transmission measurements repeated. Owing to wave-vector matching along the surface of the photonic crystal, some internal angles could not be accessed. To overcome this, large microwave prisms, made out of poly(methyl methacrylate), were placed on either side of the test crystal in figure 4.

Early the question arose of what material should the photonic crystal be made? The larger the refractive index contrast, the easier it would be to find a photonic bandgap. In optics, however, the largest practical index contrast is that of the common semiconductors, silicon and GaAs, with a refractive index $n=3{\cdot}6$. If that index was inadequate, then photonic crystals would probably never fulfil the goal of being useful in optics. Therefore we decided to restrict the microwave refractive index to 3·6, and the microwave dielectric constant to $n^2=12$. A commercial

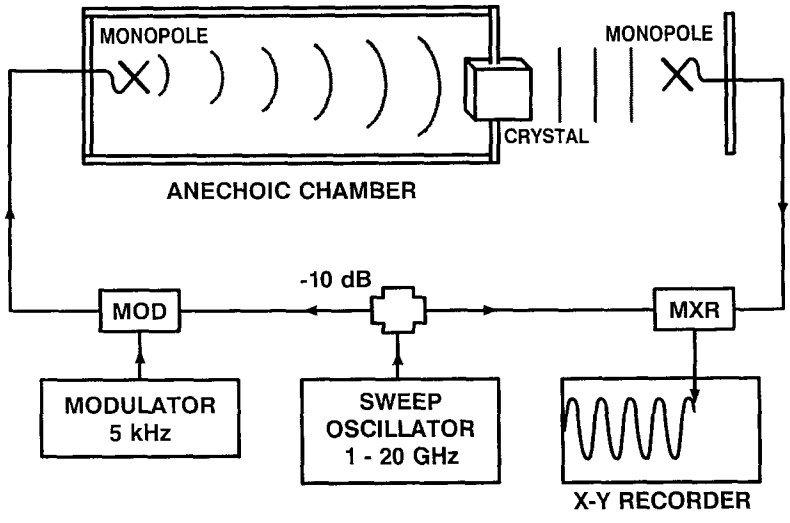

Figure 4. A homodyne detection system for measuring phase and amplitude in transmission through the photonic crystal under test. A sweep oscillator feeds a 10 dB splitter. Part of the signal is modulated (MOD) and then propagated as a plane wave through the test crystal. The other part of the signal is used as local oscillator for the mixer (MXR) to measure the amplitude change and phase shift in the crystal. Between the mixer and the x–y recorder is a lock-in amplifier (not shown).

microwave material, Emerson & Cumming Stycast 12, was particularly suited to the task since it was machinable with carbide tool bits. Any photonic band structure that was found in this material could simply be scaled down in size and would have the identical dispersion relations at optical frequencies and optical wavelengths.

With regard to the geometry of the photonic crystal, there are a universe of possibilities. So far, the only restriction we have made is toward f.c.c. lattices. It turns out that a crystal, with a f.c.c. BZ in reciprocal space, as shown in figure 2, is composed of f.c.c. Wigner–Seitz (W–S) unit cells in real space as shown in fig. 5. The problem of creating an arbitrary f.c.c. dielectric structure reduces to the problem of filling the f.c.c. W–S real-space unit cell with an arbitrary spatial distribution of dielectric material. Real space is then filled by repeated translation and close packing of the W–S unit cells. The decision before us is what to put inside the f.c.c. Wigner–Seitz cells. There are an infinite number of possible f.c.c. lattices since anything can be put inside the fundamental repeating unit. The problem before us is: what do we put inside the f.c.c. W–S unit cell in figure 5? In X-ray language, we have to find a 'form factor' for the W–S unit cell which would produce a crystal with a photonic bandgap.

This question provoked strenuous difficulties and false starts over a period of several years before finally being solved. In the first years of this research, we were unaware of how difficult the search for a photonic bandgap would be. A number of f.c.c. crystal structures were proposed, each representing a different choice for filling the rhombic dodecahedron f.c.c. W–S cells in real space. For example the very first suggestion [1], was to make a three-dimensional 'chequerboard', in which cubes were inscribed inside the f.c.c. W–S real space cells in figure 5. Later on, the experiments [11] adopted spherical 'atoms' centred inside the f.c.c. W–S cell

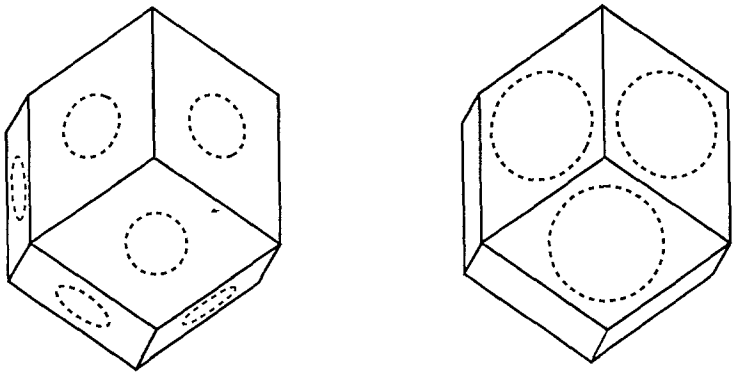

Figure 5. The W–S real-space unit cell of the f.c.c. lattice is a rhombic dodecahedron. (*a*) Slightly oversized spherical voids are inscribed into the unit cell, breaking through the faces, as illustrated by the broken circles. (*b*) W–S cell structure possessing a photonic bandgap. Cylindrical holes are drilled through the top three facets of the rhombic dodecahedron and exit through the bottom three facets. The resulting atoms are roughly cylindrical and have a preferred axis in the vertical direction.

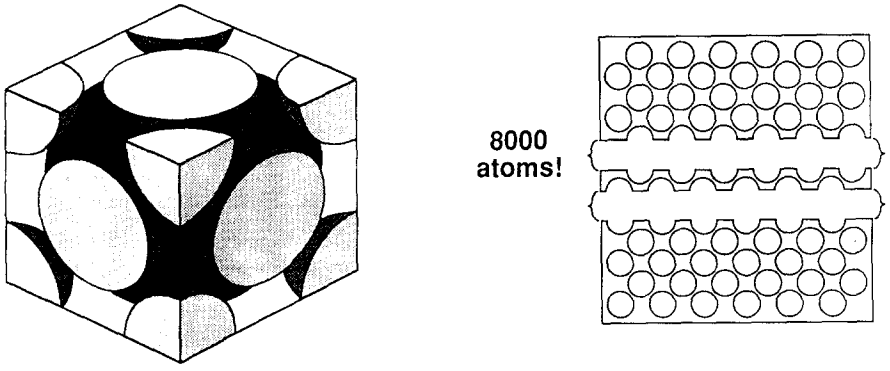

Figure 6. Construction of f.c.c. crystals consisting of spherical voids. Hemispherical holes are drilled on both faces of a dielectric sheet. When the sheets are stacked, the hemispheres meet, producing a f.c.c. crystal.

composed of precision Al_2O_3 spheres, $n \approx 3{\cdot}06$, each about 6 mm in diameter. This structure was tested at a number of filling ratios from close packing to very dilute. Nevertheless, it always failed to produce a photonic bandgap!

Then we tested the inverse structure in which spherical voids were inscribed inside the f.c.c. W–S real space cell. These could be easily fabricated by drilling hemispheres onto the opposite faces of a dielectric sheet with a spherical drill bit as shown in figure 6. When the sheets were stacked so that the hemispheres faced one another, the result was a f.c.c. array of spherical voids inside a dielectric block. These were also tested over a wide range of filling ratios by progressively increasing the diameter of the hemispheres. These also failed to produce a photonic bandgap!

The typical failure mode is illustrated in figure 7. As expected, the 'conduction band' at the L point falls at a low frequency, while the 'valence band' at the W point falls at a high frequency. The overlap of the bands at L and W results in a band structure which is best described as 'semimetallic'.

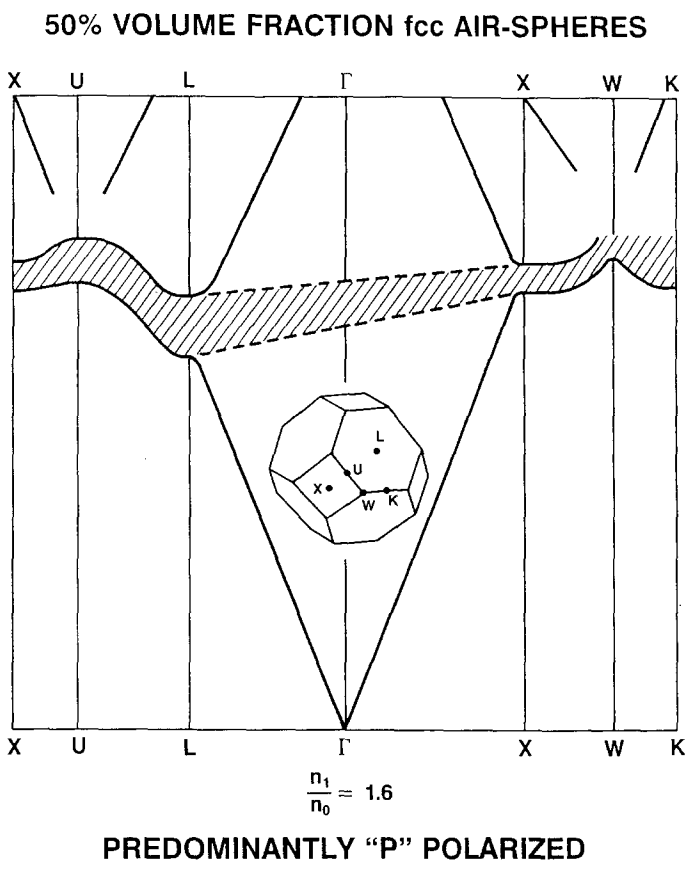

Figure 7. Typical semimetallic band structure for a photonic crystal with no photonic bandgap. An overlap exists between the conduction band at L and the valence band at W.

The empirical search for a photonic bandgap led nowhere until we tested a spherical void structure with oversized voids breaking through the walls of the W–S unit cells as shown in figure 5 (*a*). For the first time, the measurements seemed to indicate a photonic bandgap, and we published [11] the band structure shown in figure 8. There appeared to be a narrow gap, centred at 15 GHz, and forbidden for both possible polarizations. Unbeknownst to us however, figure 8 harboured a serious error. Instead of a gap at the W point, the conduction and valence bands crossed at that point, allowing the bands to touch. This produced a pseudo-gap with zero density of states but *no* frequency width. The error arose owing to the limited size of the crystal. The construction of crystals with about 10^4 atoms required tens of thousands of holes to be drilled. Such a three-dimensional crystal was still only 12 cubic units wide, limiting the wave-vector resolution and restricting the dynamic range in transmission. Under these conditions, it was experimentally difficult to notice a conduction–valence-band degeneracy which occurred at an isolated point in **k** space, such as the W point.

While we were busy with the empirical search, theorists began serious efforts to calculate photonic band structure. The most rapid progress was made, not by

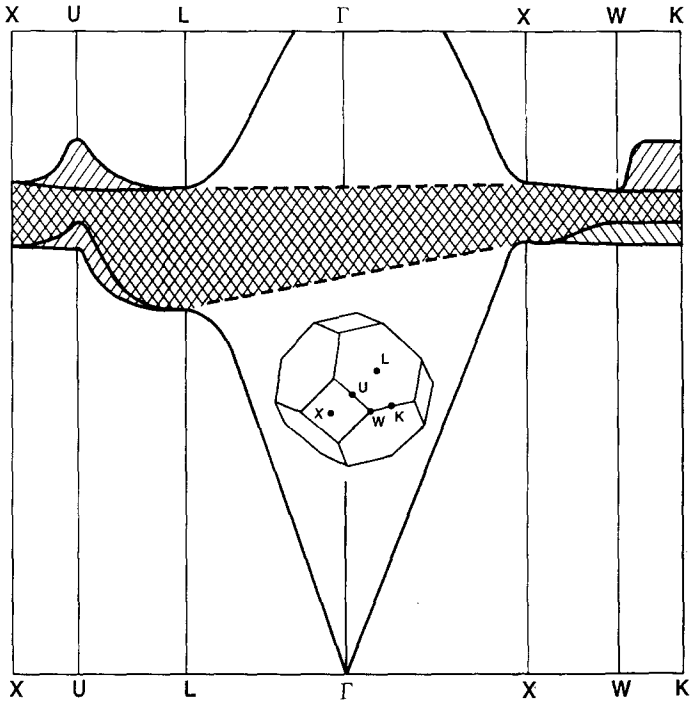

Figure 8. The purported photonic band structure of the spherical void structure shown in figure 5 (*a*). The right-sloping lines represent polarization parallel to the X plane, while the left-sloping lines represent the orthogonal polarization which has a partial component out of the X plane. The shaded region is the reported photonic bandgap. This figure fails to show the crossing of the valence and conduction bands at the W point which was first discovered by theory.

specialists in electromagnetic theory but by electronic band-structure theorists who were accustomed to solving Schrödinger's equation in three-dimensionally periodic potentials. The early calculations [12–15] were unsuccessful, however. As a short cut, they treated the electromagnetic field as a scalar, much as is done for electron waves in Schrödinger's equation. The scalar wave theory of photonic band structure did not agree well with experiment. For example, it predicted photonic bandgaps in the dielectric sphere structure, where none was observed experimentally. The approximation of Maxwell's equations as a scalar wave equation was not working. Finally, incorporating the full vector Maxwell's equations, theory began to agree with experiment. Leung and Liu [16] were probably the first to publish a successful vector wave calculation in photonic band structure, followed by others [17, 18] with substantially similar results. The theorists agreed well with one another, and they agreed well with experiment [11] except at the high-degeneracy points U and particularly W. What the experiment failed to see was the degenerate crossing of valence and conduction bands at those points.

The unexpected pseudo-gap in the f.c.c. crystal triggered concern and a search for a way to overcome the problem. A worried editorial [19] was published in *Nature* but, even before the editorial appeared, the problem had already been solved by the Iowa State group of Ho *et al.* [18]. The degenerate crossing at the W point was very

susceptible to changes in symmetry of the structure. If the symmetry was lowered by filling the W–S unit cell, not by a single spherical atom, but by two atoms positioned along the ⟨111⟩ direction as in diamond structure, then a full photonic bandgap opened up. Their discovery of a photonic bandgap using a diamond 'form factor' is particularly significant since diamond geometry seems to be favoured by Maxwell's equations. A form of diamond structure [20] gives the widest photonic bandgaps requiring the least index contrast, $n \approx 1{\cdot}87$.

More generally, the spherical void symmetry in figure 5 (*a*) can be lowered by distorting the spheres along the ⟨111⟩ direction, lifting the degeneracy at the W point. The W–S unit cell in figure 5 (*b*) has great merit for this purpose. Holes are drilled through the top three facets of the rhombic dodecahedron and exit through the bottom three facets. The beauty of the structure in figure 5 (*b*) is that a stacking of W–S unit cells results in straight holes which pass clear through the entire 'crystal'! The 'atoms' are odd-shaped, roughly cylindrical voids centred in the W–S unit cell, with a preferred axis pointing to the top vertex, ⟨111⟩. An operational illustration of the construction which produces a f.c.c 'crystal' of such W–S unit cells is shown in figure 9.

A slab of material is covered by a mask containing a triangular array of holes. Three drilling operations are conducted through each hole, 35·26° off normal

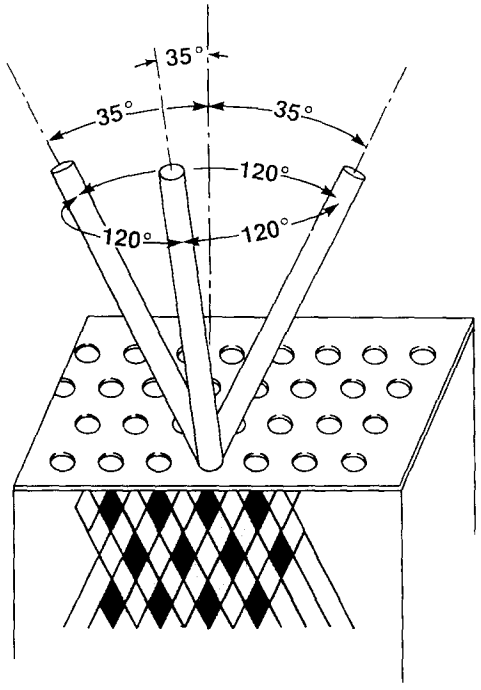

Figure 9. The method of constructing a f.c.c. lattice of the W–S cells as shown in figure 5 (*b*). A slab of material is covered by a mask consisting of a triangular array of holes. Each hole is drilled through three times, at an angle 35·26° away from normal and spread out 120° on the azimuth. The resulting criss-cross of holes below the surface of the slab, suggested by the cross-hatching shown here, produces a full three-dimensionally periodic f.c.c. structure, with unit cells as given in figure 5 (*b*). The drilling can be done by a real drill bit for microwave work, or by reactive ion etching to create a f.c.c. structure at optical wavelengths.

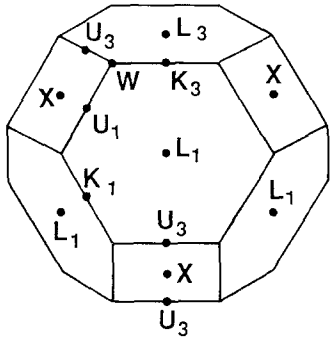

Figure 10. The BZ of a f.c.c. structure incorporating non-spherical atoms, as in figure 5 (*b*). Since the space lattice is not distorted, this is simply the standard f.c.c. BZ lying on a hexagonal face rather than the usual cubic face. Only the L points on the top and bottom hexagons are three-fold symmetry axes. Therefore they are labelled L_3. The L points on the other six hexagons are labelled L_1. The U_3–K_3 points are equivalent since they are a reciprocal-lattice vector apart. Likewise the U_1–K_1 points are equivalent.

incidence and spread out 120° on the azimuth. The resulting criss-cross of holes below the surface of the slab produces a fully three-dimensionally periodic f.c.c. structure, with W–S unit cells given by figure 5 (*b*)! The drilling can be done by a real drill bit for microwave work, or by reactive ion etching to create a f.c.c. structure at optical wavelengths.

In spite of non-spherical atoms in figure 5 (*b*), the BZ is identical with the standard f.c.c. BZ shown in textbooks. Nevertheless, we have chosen an unusual perspective from which to view the BZ in figure 10. Instead of having the f.c.c. BZ resting on one of its diamond-shaped facets as is usually done, we have chosen in figure 10 to present it resting on a hexagonal face. Since there is a preferred axis for the atoms, the distinctive L points centred in the top and bottom hexagons are three fold symmetry axes, and are labelled L_3. The L points centred in the other six hexagons are symmetric only under a 360° rotation and are labelled L_1. It is helpful to know that the U_3–K_3 points are equivalent since they are a reciprocal-lattice vector apart. Likewise the U_1–K_1 points are equivalent.

Figure 11 shows the dispersion relations along different meridians for our primary experimental sample of normalized hole diameter $d/a = 0{\cdot}469$ and 78% volume fraction removed (where a is the unit cube length). The oval points represent experimental data with s polarization (perpendicular to the plane of incidence, and parallel to the slab surface), while the triangular points represent p polarization (parallel to the plane of incidence and partially perpendicular to the slab surface). The horizontal abscissa in figure 11 (*b*), L_3-K_3-L_1-U_3-X-U_3-L_3 represents a full meridian from the north pole to the south pole of the BZ. Along this meridian the Bloch wavefunctions separate neatly into s and p polarizations. The s and p-polarized theory curves are the solid broken curves respectively. The dark shaded is the totally forbidden photonic bandgap. The lighter shaded stripes above and below the dark band are forbidden only for s and p polarizations respectively.

At a typical semiconductor refractive index, $n = 3{\cdot}6$; the three-dimensional forbidden gap width is 19% of its centre frequency. Calculations [21] indicate that the gap remains open for refractive indices as low as $n = 2{\cdot}1$ using circular holes. We have also measured the imaginary wave-vector dispersion within the forbidden gap.

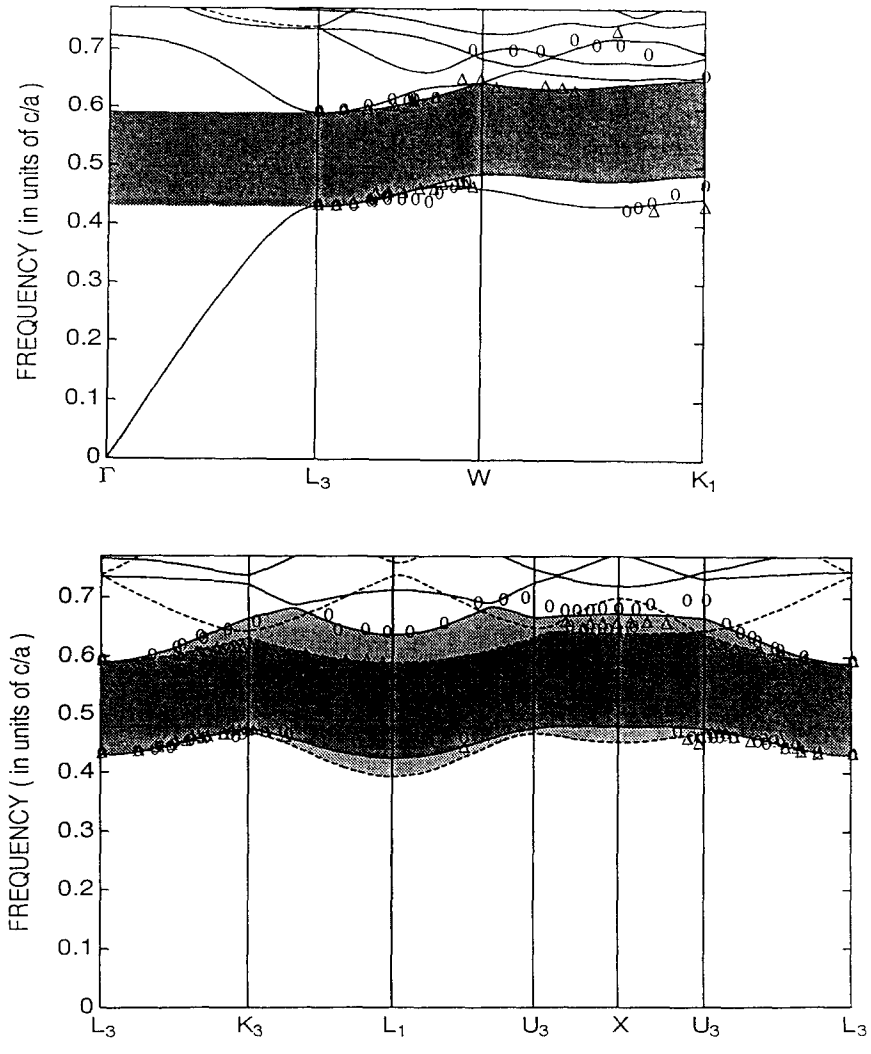

Figure 11. Frequency ω against wave-vector dispersion along the surface of the BZ shown in figure 10, where c/a is the speed of light divided by the f.c.c. cube length: (○), experimental points for s polarization; (△), experimental points for p polarization; (——) calculations for s polarization; (– – –), calculations for p polarization; the dark shaded band is the totally forbidden bandgap; the lighter shaded stripes above and below the dark band are forbidden only for s and p polarizations respectively.

At midgap we find an attenuation of 10 dB per unit cube length a. Therefore the photonic crystal need not be very many layers thick to effectively expel the zero-point electromagnetic field. The construction in figure 9 can be implemented by reactive ion etching as shown in figure 12. In reactive ion etching, the projection of circular mask openings at 35° leaves oval holes in the material, which might not perform as well. Fortunately it was found [21], defying Murphy's law, that the forbidden gap width for oval holes is actually improved, fully 21·7% of its centre frequency.

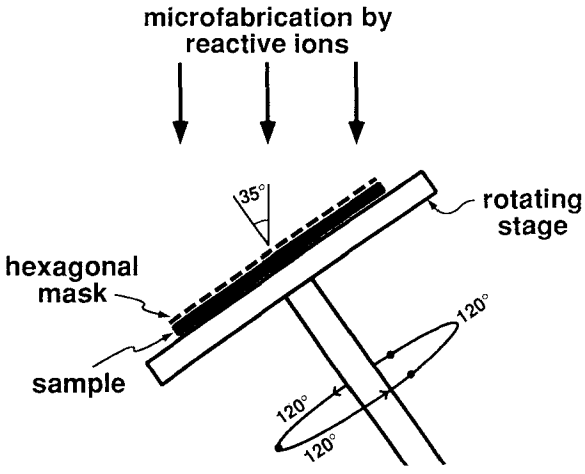

Figure 12. Construction of the non-spherical void photonic crystal of figure 5 (*b*) and figures 9–11 by reactive ion etching.

4. Doping the photonic crystal

The perfect semiconductor crystal is quite elegant and beautiful, but it becomes ever more useful when it is doped. Likewise the perfect photonic crystal can become of even greater value when a defect [22] is introduced.

Lasers, for example, require that the perfect three-dimensional translational symmetry should be broken. Even while spontaneous emission into all 4π sterad should be inhibited, a local electromagnetic mode, linked to a defect, is still necessary to accept the stimulated emission. In one-dimensional distributed feedback lasers [23], a quarter-wavelength defect is introduced, forming effectively a Fabry–Pérot cavity. In three-dimensional photonic band structure a local defect-induced structure resembles a Fabry–Pérot cavity, except that it reflects radiation back upon itself in all 4π spatial directions.

The perfect three-dimensional translational symmetry of a dielectric structure can be lifted in either one of two ways.

(1) Extra dielectric material may be added to one of the unit cells. We find that such a defect behaves very much like a donor atom in a semiconductor. It gives rise to donor modes which have their origin at the bottom of the conduction band.

(2) Conversely, translational symmetry can be broken by removing some dielectric material from one of the unit cells. Such defects resemble acceptor atoms in semiconductors. The associated acceptor modes have their origin at the top of the valence band. We shall find that acceptor modes are particularly well suited to act as laser microresonator cavities. Indeed it appears that photonic crystals made of sapphire or other low-loss dielectrics will make the highest-Q single-mode cavities (of modal volume about λ^3) covering electromagnetic frequencies above the useful working range of superconducting metallic cavities. The short-wavelength limit in the ultraviolet is set by the availability of optical materials with refractive index of about 2 or greater, the threshold index [18, 21] for the existence of a photonic bandgap.

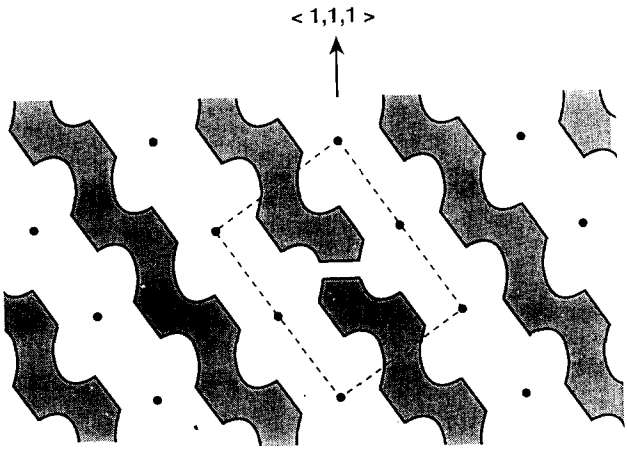

Figure 13. A $\langle 1, \bar{1}, 0 \rangle$ cross-sectional view of our f.c.c. photonic crystal consisting of non-spherical 'air atoms' centred on the full circles. The dielectric material is represented by the shaded area. The broken rectangle is a face-diagonal cross-section of the unit cube. Donor defects consisted of a dielectric sphere centred on an atom. We selected an acceptor defect as shown, centred in the unit cube. It consists of a missing horizontal slice in a single vertical rib.

Figure 13 is a $\langle 1, 1, 0, \rangle$ cross-section of our photonic crystal in figure 5 (*b*) and figures 9–11, cutting through the centre of a unit cube. Shading represents dielectric material. The full circles are centred on the air atoms and the broken rectangle is a face-diagonal cross-section of the unit cube. Since we could design the structure at will, donor defects were chosen to consist of a single dielectric sphere centred in an air atom. Likewise, by breaking one of the interconnecting ribs, it is easy to create acceptor modes. We selected an acceptor defect as shown in figure 13, centred in the unit cube. It comprises a vertical rib which has a missing horizontal slice.

The heart of our experimental apparatus is a photonic crystal embedded in microwave absorbing pads as shown in figure 14. The photonic crystals were eight to ten atomic layers thick in the $\langle 1, 1, 1 \rangle$ direction. Monopole antennae, consisting of 6 mm pins, coupled radiation to the defect mode. The HP 8510 network analyser was set up to measure transmission between the antennas. Figure 15 (*a*) shows the transmission amplitude in the absence of a defect. There is very strong attenuation (about 10^{-5}) between 13 and 16 GHz, marking the valence- and conduction-band edges of the forbidden gap. This is a tribute to both the dynamic range of the network analyser, and the sizeable imaginary wave-vector in the forbidden gap.

A transmission spectrum in the presence of an acceptor defect is shown in figure 15 (*b*). Most of the spectrum is unaffected, except at the electromagnetic frequency labelled Deep acceptor within the forbidden gap. At that precise frequency, radiation 'hops' from the transmitting atenna to the acceptor mode and then to the receiving antenna. The acceptor level frequency, within the forbidden gap, is dependent on the volume of material removed. Figure 16 shows the acceptor-level frequency as a function of defect volume removed from one unit cell. When a relatively large volume of material is removed, the acceptor level is deep as shown in figure 15 (*b*). A smaller amount of material removed results in a shallow acceptor level, nearer the valence band. If the removed material volume falls below a

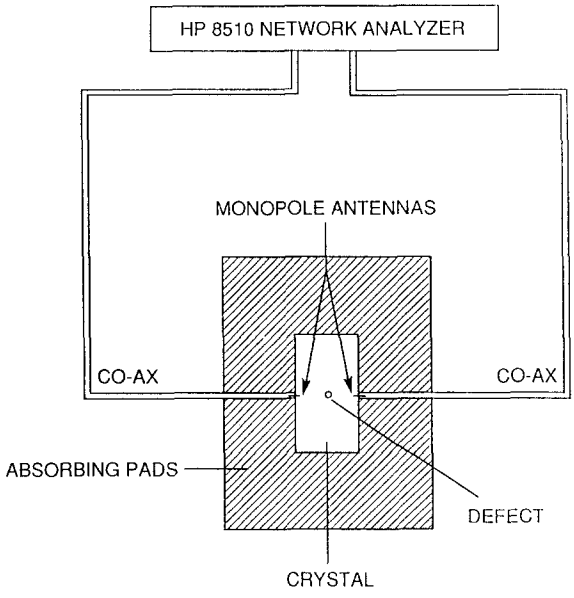

Figure 14. Experimental configuration for the detection of local electromagnetic modes in the vicinity of a lattice defect. Transmission amplitude attenuation from one antenna to the other is measured. At the local mode frequency the signal hops by means of the local mode in the centre of the photonic crystal, producing a local transmission peak. The signal propagates in the $\langle 1, 1, 1 \rangle$ direction through eight to ten atomic layers.

threshold volume, the acceptor level falls within the continuum of levels below the top of the valence band, becoming metastable.

On an expanded frequency scale we can measure the resonator Q of the deep-acceptor mode, which is about 1000, as limited by the loss tangent of the Emerson & Cumming Stycast material of which the photonic crystal was made.

The behaviour of an off-centre donor defect is shown in figure 15 (*c*). In that case the donor volume was only slightly above the required threshold for forming bound-donor modes. Already two shallow-donor modes can be seen in figure 15 (*c*). When the donor is centred in the W–S unit cell, the two modes merge to form a doubly degenerate donor levels as in figure 16. Single donor defects seem to produce multiple donor levels. Figure 16 gives the donor-level frequency as a function of donor volume. As in the case of acceptors, there is a threshold defect volume required for the creation of bound modes below the conduction-band edge. However, the threshold volume for donor defects is almost ten times the acceptor threshold volume. Apparently this is due to the electric field concentration in the dielectric ribs at the top of the valence band. Bloch wavefunctions at the top of the valence band are rather easily disrupted by the missing rib segment.

We have chosen in figure 16 to normalize the defect volume to a natural volume of the physical system, $(\lambda/2n)^3$, which is basically a cubic half-wavelength in the dielectric medium. More specifically, λ is the vacuum wavelength at the midgap frequency, and n is the refractive index of the dielectric medium. Since we are measuring a dielectric volume, it makes sense to normalize to a half-wavelength cube as measured at the dielectric refractive index. On the basis of the reasonable scaling of figure 16, our choice of volume normalization would seem justified.

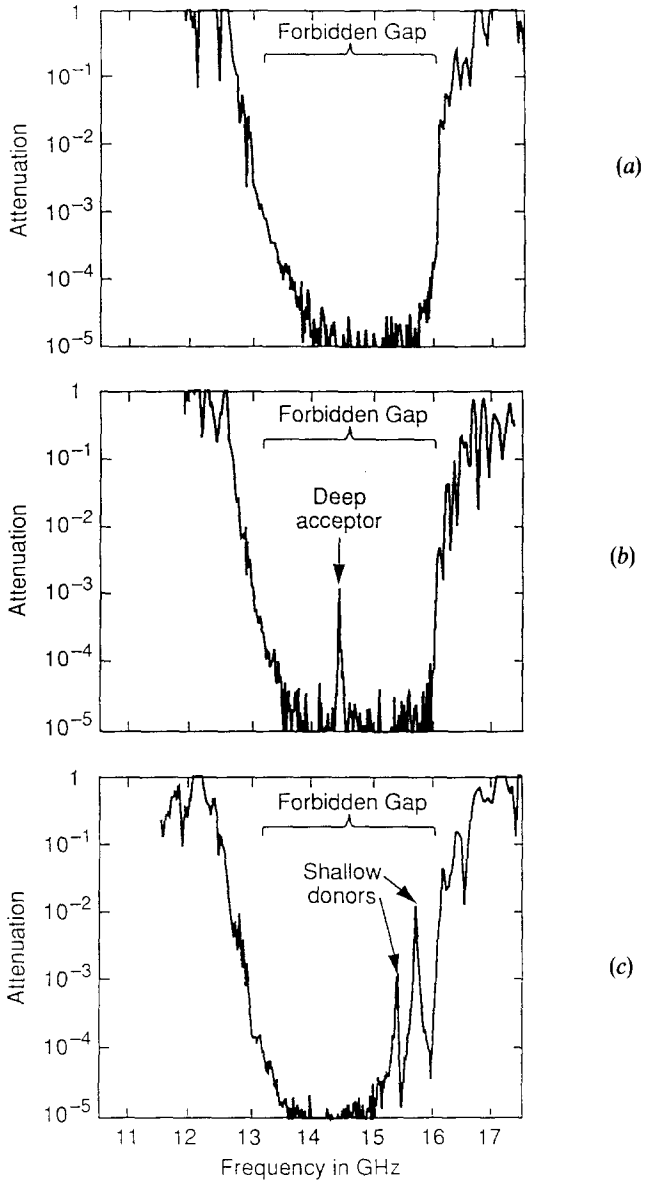

Figure 15. (*a*) Transmission attenuation through a defect-free photonic crystal, as a function of microwave frequency. The forbidden gap falls between 13 and 16 GHz. (*b*) Attenuation through a photonic crystal with a single acceptor in the centre. The large acceptor defect volume shifted its frequency near midgap. The electromagnetic resonator Q was about 1000, limited only by the loss tangent of the dielectric material. (*c*) Attenuation through a photonic crystal with a single donor defect, an uncentred dielectric sphere, leading to two shallow donor modes.

The vertical rib with a missing horizontal slice, as in figure 13, can be readily microfabricated. It should be possible to create it in III–V materials by growing an aluminium-rich epitaxial layer and lithographically patterning it down to a single dot the size of one of the vertical ribs. After regrowth of the original III–IV composition and reactive ion etching of the photonic crystal, HF acid etching, whose [24] selectivity is 10^8 or more, will be used to remove the aluminium-rich horizontal slice from the one rib containing such a layer. The resonant frequency of the microcavity can be controlled by the thickness of the aluminium rich sacrificial layer.

Therefore by doping the photonic crystal, it is possible to create high-Q electromagnetic cavities whose modal volume is less than a half-wavelength cubed. These doped photonic crystals would be similar to metallic cavities, except that they would be usable at higher frequencies where metal cavity walls would become lossy. Using sapphire as a dielectric for example, it should be possible to make a millimetre-wave cavity with $Q \gtrsim 10^9$. The idea is not to compete directly with superconducting cavities, but rather to operate at higher frequencies where the superconductors become lossy. Given the requirement for refractive index greater than 2, doped photonic crystals should work well up to ultraviolet wavelengths where diamond crystals and TiO_2 are still transparent.

5. Applications

The forthcoming availability of single-mode microcavities at optical frequencies will lead to a new situation in quantum electronics. Of course microwave cavities containing a single electromagnetic mode have been known for a long time. At

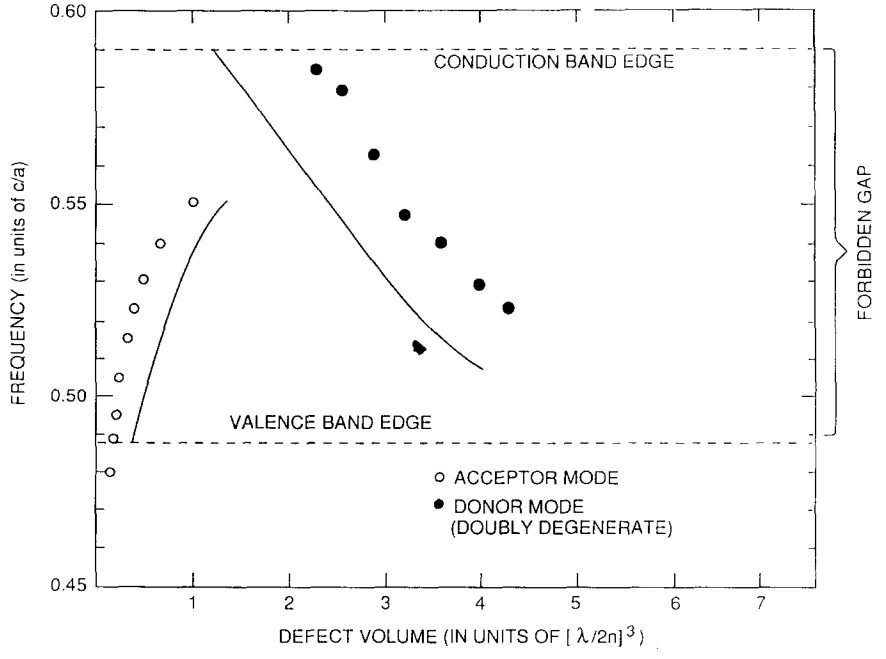

Figure 16. Donor and acceptor-mode frequencies as functions of normalized donor and acceptor defect volume: (●), (○), experimental; ——, calculated. The defect volume is normalized to $(\lambda/2n)^3$, where λ is the mid-gap vacuum wavelength and n is the refractive index. A finite defect volume is required to bind a mode in the forbidden gap.

microwave frequencies, however, spontaneous emission of electromagnetic radiation is a weak and unimportant process. At optical frequencies, spontaneous emission comes into its own. Now we can combine the physics and technology of spontaneous emission with the capability for single-mode microcavities at optical frequencies where spontaneous emission is important. This combination is fundamentally a new regime in quantum electronics.

The major example of this new type of device is the single-mode light-emitting diode (SMLED), which can have many of the favourable coherence properties of lasers, while being a more reliable and threshold-less device. Progress in electromagnetic microcavities, allows all the spontaneous emission of a light-emitting diode (LED) to be funnelled into a single electromagnetic mode.

As the interest in low-threshold semiconductor laser diodes has grown, for example for optical interconnects, its spontaneously luminescent half-brother, the LED, has begun to re-emerge in a new form. In this new form the LED is surrounded by an optical cavity. The idea is for the optical cavity to make available only a single electromagnetic mode for the output spontaneous emission from the semiconductor diode. In fact the figure of merit for such a cavity is β, the fraction of spontaneous emission which is being funnelled into the desired mode. What is new for this application is the prospective ability to make high-β cavities at optical frequencies employing photonic crystals. The three-dimensional character of the cavities ensures that spontaneous emission will not seek out those neglected modes which are found to propagate in a direction away from the optical confinement.

With all the spontaneous emission funnelled into a single optical mode, the SMLED can begin to have many of the coherence and statistical properties normally associated with above-threshold lasing. The essential point is that the spontaneous emission factor β should approach unity. (A closely related concept is that of the 'zero-threshold laser', in which the high spontaneous emission factor produces a very soft and indistinct threshold characteristic in the care of light output against current input of laser diodes.) The idea is to combine the advantages of the LED which is threshhold-less and highly reliable, with those of the semiconductor laser which is coherent and very efficient.

The coherence properties of the SMLED are illustrated in figure 17. In a laser, single-mode emission is the result of gain saturation and mode competition. In the SMLED, there is no gain and therefore no gain saturation, but the output is still a single mode, because only one mode is available for emission. Since a single spatial mode can always be mode converted into a plane wave, the SMLED can be regarded as having spatial coherence.

What about temporal coherence? The spectral linewidth of the SMLED is narrower than the luminescence band of the semiconductor. All the radiation is funnelled into the narrow spectral band determined by the microcavity Q. Thus SMLEDs have both spatial and temporal coherence as represented by the words Directional and Monochromatic in figure 17.

What about the modulation speed of SMLEDs in comparison with laser diodes under d.c. modulation? Generally, the modulation speed depends on the carrier lifetime. Since electron–hole pairs in laser diodes experience both spontaneous and stimulated recombination, they have an advantage. However, single-mode cavities concentrate zero-point electric field fluctuations into a smaller volume, creating a stronger matrix element for spontaneous emission. Detailed calculations indicate that spontaneous emission can be speeded up by a factor of about ten owing to this

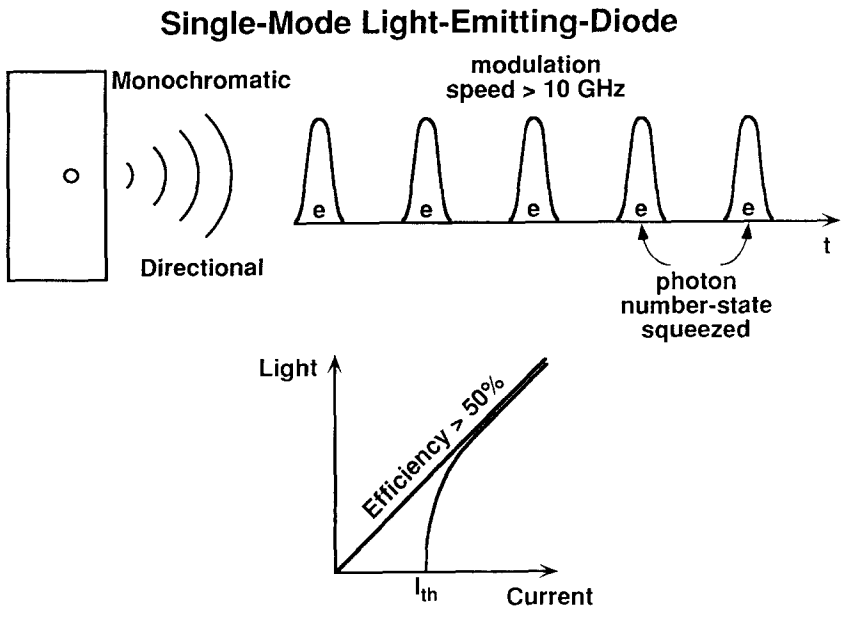

Figure 17. An illustration of the properties of the SMLED, whose cavity is represented by the small open circle inside the rectangular photonic crystal, on the left-hand side. The words Monochromatic and Directional represent the temporal and spatial coherences of the SMLED output as explained in the text. The modulation speed can be greater than 10 GHz, and the differential quantum efficiency can be greater than 50%, competitive with laser diodes, but there is no threshold current for the SMLED as indicated by the L–I curves at the bottom. The regular stream of photoelectrons e are meant to represent photon-number-state squeezing, which can be produced by the SMLED if the spontaneous emission factor β of the cavity is sufficiently high.

cavity quantum electrodynamic (QED) effect. In figure 17 we indicate that a modulation speed greater than 10 GHz should be possible for SMLEDs.

The same cavity QED effects can enhance the spontaneous emission efficiency of SMLEDs since the radiative rate can then compete more successfully with non-radiative rates. External efficiency should exceed 50%, but this can come most easily from intelligent LED design [25] rather than from cavity QED effects.

Shown at the bottom of figure 17 is the curve of light output against current input of SMLEDs and laser diodes. SMLEDs can compete with laser diodes in terms of differential external efficiency, but the SMLEDs can have the advantage of not demanding any threshold current. Lack of threshold behaviour makes the output power and the operating wavelength of a SMLED relatively insensitive to ambient temperature. Combined with the inherent reliability of a LED, this should produce many systems advantages for the SMLED concept.

The final SMLED property illustrated in figure 17 is photon-number-state squeezing, as suggested by the regular sequence of photoelectrons on the horizontal line. Stimulated emission is *not* required for these exotic squeezing effects. The critical variable is absolute quantum efficiency. If the quantum efficiency of the SMLED is high, then these useful correlations will exist in the spontaneous output of the SMLED. This requires, most of all, a high spontaneous emission factor β, our overall figure of merit for microcavities.

There are many other applications for photonic crystals, particularly in the microwave and millimetre-wave regime. They are very imaginative, and they have gone far beyond our initial goals for using photonic crystals in quantum optics.

6. Conclusions

It is worthwhile to summarize the similarities and the differences between photonic band structures and electronic band structure. This is best done by reference to the table.

Electrons are massive and so the underlying dispersion relation for electrons in crystals is parabolic. Photons have no mass; so the underlying dispersion relation is linear. However, as a result of the periodicity the photons develop an effective mass in photonic band structure and this should not be surprising.

Electrons have spin $\frac{1}{2}$, but frequently this is ignored and Schrödinger's equation is treated in a scalar-wave approximation. In electronic band theory the spin $\frac{1}{2}$ is occasionally important, however. In contrast, photons have spin 1, but it is generally never a good approximation to neglect polarization in photonic band-structure calculations.

Finally we come to the accuracy of band theory. It is sometimes believed that band theory is always a good approximation in electronic structure. This is not really true. When there are strong correlations, as in the high-T_c superconductors, band theory is not even a good zeroth-order approximation. Photons are highly non-interacting; so, if anything, band theory makes more sense for photons than for electrons.

The final point to make about photonic crystals is that they are very empty structures, consisting of about 78% empty space, but in a sense they are much emptier than that. They are emptier and quieter than even the vacuum since they contain not even zero-point fluctuation within the forbidden frequency band.

A summary of the differences and similarities between photonic band structure and electronic band structure.

	Electronic band structure	Photonic band structure
Underlying dispersion relation	Parobolic	Linear
Angular momentum	Spin $\frac{1}{2}$ scalar-wave approximation	Spin-1 vector wave character
Accuracy of band theory	Approximate owing to electron–electron interactions	Essentially exact

Acknowledgments

I would like to thank John Gural for drilling all the holes and Tom Gmitter for making the measurements. Thanks are also due to Ming Leung and Bob Meade for their collaborative work. The Iowa State group is thanked for making me aware of their diamond structure results prior to publication.

References

[1] YABLONOVITCH, E., 1987, *Phys. Rev. Lett.*, **58,** 2059.
[2] JOHN, S., 1987, *Phys. Rev. Lett.*, **58,** 2486.
[3] YAMAMOTO, Y., MACHIDA, S., and RICHARDSON, W. H., 1992, *Science*, **255,** 1219.
[4] PURCELL, E. M., 1946, *Phys. Rev.*, **69,** 681.
[5] DREXHAGE, K. H., 1974, *Prog. Optics*, **12,** 165.
[6] BYKOV, V. P., 1975, *Soviet J. quant. Electron.*, **4,** 861.
[7] HULET, R. G., HILFER, E. S., and KLEPPNER, D., 1985, *Phys. Rev. Lett.*, **55,** 2137.
[8] VAN DER ZIEL, A., 1954, *Noise* (New York: Prentice-Hall), pp. 208–209.
[9] SOZUER, H. S., HAUS, J. W., and INGUVA, R., 1992, *Phys. Rev.* B, **45,** 13962.
[10] DARWIN, C. G., 1914, *Phil. Mag.*, **27,** 675.
[11] YABLONOVITCH, E., and GMITTER, T. J., 1989, *Phys. Rev. Lett.*, **63,** 1950.
[12] SATPATHY, S., ZHANG, Z., and SALEHPOUR, M. R., 1990, *Phys. Rev. Lett.*, **64,** 1239.
[13] LEUNG, K. M., and LIU, Y. F., 1990, *Phys. Rev.* B, **41,** 10188.
[14] JOHN, S., and RANGARAJAN, R., 1988, *Phys. Rev.* B, **38,** 10101.
[15] ECONOMU, E. N., and ZDETSIS, A., 1989, *Phys. Rev.* B, **40,** 1334.
[16] LEUNG, K. M., and LIU, Y. F., 1990, *Phys. Rev. Lett.*, **65,** 2646.
[17] ZHANG, Z., and SATPATHY, S., 1990, *Phys. Rev. Lett.*, **65,** 2650.
[18] HO, K. M., CHAN, C. T., and SOUKOULIS, C. M., 1990, *Phys. Rev. Lett.*, **65,** 3152.
[19] MADDOX, J., 1990, *Nature*, **348,** 481.
[20] CHAN, C. T., HO, K. M., and SOUKOULIS, C. M., 1991, *Europhys. Lett.*, **16,** 563.
[21] YABLONOVITCH, E., GMITTER, T. J., and LEUNG, K. M., 1991, *Phys. Rev. Lett.*, **67,** 2295.
[22] YABLONOVITCH, E., GMITTER, T. J., MEADE, R. D., RAPPE, A. M., BROMMER, K. D., and JOANNOPOULOS, J. D., 1991, *Phys. Rev. Lett.*, **67,** 3380.
[23] HAUS, H. A., and SHANK, C. V., 1976, *IEEE Jl quant. Electron.*, **QE–12,** 532; MCCALL, S. L., and PLATZMAN, P. M., 1985, *IEEE Jl quant. Electron.*, **QE–21,** 1899.
[24] YABLONOVITCH, E., GMITTER, T., HARBISON, J. P., and BHAT, R., 1987, *Appl. Phys. Lett.*, **51,** 2222.
[25] SCHNITZER, I., YABLONOVITCH, E., CANEAU, C., and GMITTER, T. J., 1993, *Appl. Phys. Lett.*, **62,** 131.