



Photonic crystals in the optical regime — past, present and future

Thomas F. Krauss*, Richard M. De La Rue

Department of Electronics and Electrical Engineering, The University of Glasgow, G12 8LT, Scotland, UK

Abstract

During the last decade, photonic crystals, also known as photonic microstructures or photonic bandgap structures, have matured from an intellectual curiosity concerning electromagnetic waves to a field with real applications in both the microwave and optical regime. In this review, we shall focus on progress and the prospects for semiconductor structures that mainly involve guided modes interacting with periodic structures, but we also evaluate alternative material systems and fabrication methods, e.g. those based on self-organisation. We shall go from basic concepts, via a discussion of the state of the art, to device applications. Naturally, the discussion of the applications will be more speculative, but we attempt to evaluate the real prospects offered by photonic crystals at optical frequencies while considering practical limitations. In doing so, we identify a variety of areas such as the combination of quantum dot light emitters with photonic crystals that seem particularly promising. We discuss the prospects for enhanced light–matter interactions in photonic crystals and the related material and design issues. Overall, the aim of this review is to introduce the reader to the concepts of photonic crystals, describe the state of the art and attempt to answer the question of what uses these peculiar structures may have. © 1999 Elsevier Science Ltd. All rights reserved.

Keywords: Photonic crystal; Photonic bandgap; LED; Laser; Optoelectronic materials; Semiconductors

1. Introduction	52
2. The past — how did it all start?	53
2.1. The basic concept	53
2.2. Initial realisation	54
3. The present — state-of-the-art	57
3.1. How to make a photonic crystal?	57
3.2. Different methods of characterising photonic crystals	67

* Corresponding author.

E-mail address: t.krauss@elec.gla.ac.uk (T.F. Krauss)

4. The future — applications for photonic crystals	70
4.1. Light emitting diodes (LEDs)	70
4.2. Lasers	78
4.3. Other materials	80
4.4. Photonic VLSI	83
4.5. Surface plasmons	85
4.6. Quasiperiodic structures	86
4.7. Quantum cascade-based light emission	89
5. Conclusions	90
Acknowledgements	92
References	92

1. Introduction

From the beginning of civilisation, man has sought to illuminate the darkness. Fires, gas lights and Edison's light bulb are but a few in a long chain of light sources that have been employed. Considering the high level of technical sophistication that we have achieved, it is surprising that we still rely mainly on such archaic light sources as the incandescent light bulb or on fluorescent strip-lights, particularly since these light sources typically waste 80–95% of the electrical energy input. This is a state of affairs that needs to be improved, so better light sources are required. This need has already been recognised in a recent report compiled by the U.S. Academy of Sciences [1], who cite lighting by LEDs as having a major impact on the economy in the next century, both in terms of lighting and energy saving. In order to make more efficient LEDs, we need to use the concepts of photonic crystals, which brings us to the main topic of this paper.

As far as the material for LEDs is concerned, semiconductor-based light emitters are the obvious choice. Mainly driven by the needs of the information technology industry, semiconductor light-emitting materials have reached a level of sophistication that surpasses that of any other light-emitting material known to man, with *internal* quantum efficiencies above 90% being routinely obtained. Unfortunately, most of this light is trapped inside the material, and only 3–20% can escape, the remainder being eventually re-absorbed and lost as heat (Fig. 1a).

Laser diodes offer one way out. They rely on stimulated emission and are therefore more efficient than LEDs, but are not applicable to every problem. Some applications, such as displays, demand the simpler and cheaper LED, since the high level of coherence of laser light is not desirable, or the generally higher cost and temperature dependence of laser diodes is prohibitive. Therefore, a concept that offers the ability to improve the efficiency of LEDs is a welcome addition to the toolbox.

Photonic crystals represent this concept, and their advantage lies in addressing the extraction problem mentioned above. Moreover, photonic crystals not only offer a simple improvement of the extraction of light from LEDs, they can also alter the fundamental process of light generation in the material itself (Fig. 1b). Perhaps more significantly, they can improve the operating characteristics of laser diodes. Laser




Fig. 1. (a) “Classical” LED: the emission cone (shaded) is very small and only a fraction of the light generated in the material can escape the device, the rest is trapped by total internal reflection. (b) Microstructured emitter: light that would not radiate into the emission cone is suppressed by the photonic crystal, i.e. there are no modes for the emission to radiate into, so all spontaneous emission is channelled into the “useful” modes of the device. In principle, the external efficiency of such a device can reach the internal efficiency of the material of $> 90\%$.

diodes have experienced a tremendous development in recent years, and specific devices have already shown total external (“wallplug”) efficiencies of 50% and above [2–4]. Most types of laser, however, still waste too much light, light that could be re-used, or not be emitted in the first place, if photonic band gap principles were applied.

So, be it a laser or an LED, photonic crystals promise improvement. Their ability to control the flow of light and their capacity to concentrate light and enhance light–matter interaction is certainly inspiring. Examining whether this promise is realistic and what else they have to offer, e.g. in the realm of passive devices, such as filters and waveguides, is the purpose of this paper. In Section 2, we introduce photonic crystals and show how they have emerged. A review of the state of the art is given in Section 3 and a speculative forward-look at what they may be used for in the short and longer term is attempted in Section 4.

2. The past — how did it all start?

In 1946, Purcell [5] was the first to discuss the fact that the presence of a mirror can substantially alter the radiation properties of an electromagnetic dipole. This insight was further developed over the years, and in 1987 led to the concept of a photonic band gap [6,7], a concept that was borrowed from semiconducting crystals in analogy to their electronic band gap.

2.1. The basic concept

The general idea is that photonic crystals do to photons what semiconductor crystals do to electrons, i.e. they create a situation whereby photons in a certain energy range cannot travel through the crystal and are reflected when impinging onto the crystal or are not allowed to propagate at all when generated inside it. The latter point is conceptually very important, because one should not imagine a situation whereby light is emitted from the source, reflected back by the crystal, eventually re-absorbed, re-emitted and so on. The situation is rather that the photon is not emitted in the first

place, because there is no quantum mechanical state into which it can go. For a general introduction to the concepts and underlying theory of photonic crystals, the reader should consult earlier review papers [8–10] and the excellent textbook by Joannopoulos et al. [11].

2.2. *Initial realisation*

2.2.1. *3-D photonic crystals in the microwave regime*

After a period of speculation, some controversy [12,13] and a large theoretical effort, it was finally shown conclusively that such photonic band gap materials were indeed possible [14,15]. The first experimental demonstration appeared in 1991 [16], using an array of holes drilled into high refractive index material. The artificial crystal thus created (later termed “Yablonovite” after its inventor) showed a stop-band for the transmission of microwave radiation that extended from 13 GHz to about 15 GHz, irrespective of the direction of propagation, so these structures exhibited a truly “forbidden” zone, a band gap. Another feature reminiscent of semiconductors is the possibility of “doping” these structures. By including a defect, i.e. locally adding or removing material, it was possible to demonstrate localised states in the band gap [16]. The semiconductor analogy extends as far as the type of dopant, with material addition leading to a “donor”-type defect, i.e. a state near the “conduction band” of the crystal band structure and material removal resulting in an “acceptor” state near the “valence band”. An immediate application was found by placing an antenna onto the top of an (undoped) crystal [17]. With no radiation being allowed to enter the crystal side of the antenna, the emission into the air side was greatly enhanced, thus clearly showing that photonic crystals can alter the radiation process itself.

2.2.2. *3-D photonic crystals in the optical regime*

Although the microwave regime is a very useful testing ground and offers a variety of applications for photonic crystals, there has always been a drive towards truly “photonic” applications, i.e. structures that operate in the visible or near IR regime of the electromagnetic spectrum. One approach is to take fabrication methods that were first successfully demonstrated on samples in the microwave regime and then scale them down, usually employing some sort of stacking technique. These “layer-by-layer” photonic crystals have now matured to the point of reaching into the far-infrared, e.g. a structure displaying a stop-band between 1.35 and 1.95 μm was demonstrated recently [18].

Another approach that has also been successful at making structures for the 1.5 μm window is to use micro-fabrication techniques and refine them to the extent that structures can be made at the desired wavelength. Impressive results have been achieved on that front [19,20] but have unfortunately failed to show the expected sharp spectral response.

A third approach is to let nature do the ordering, e.g. in the form of opal, a self-organised structure consisting of closely packed silica spheres. This is possibly the most convincing of the different approaches towards fully 3-D periodic structures, although the results shown to-date indicate the promise of the approach rather than

providing a fully convincing demonstration. The main problems are the lack of “pure” material with sufficient uniformity and without unwanted defects, the relatively low intrinsic refractive index contrast and the fact that opal is optically passive. In order to increase the index contrast, and to achieve optical gain, the silica spheres that comprise opal can be doped, e.g. with rare-earth ions, or the voids between the individual spheres can be back-filled (“loaded”) with suitable material. These techniques and their promise will be discussed in more detail in Section 3.1.6.

2.2.3. 2-D photonic crystals in the optical regime

In contrast to the limited experimental successes with three-dimensional structures, work on 2-D structures has flourished and produced many interesting results. This progress is mainly due to the fact that two-dimensional structures are easier to make, in particular with the aid of the tools developed for the silicon microchip industry. Lower dimensionality also implies easier manipulation of the photonic lattice to increase functionality, such as the deliberate inclusion of defects, as mentioned above, or interfacing with “standard” optical elements, such as waveguides, fibres, light sources and detectors. We shall discuss two-dimensional structures in more detail, because so much work has been done on them.

The main factors that determine the properties of two-dimensional photonic crystals are the refractive index contrast, the fraction of high and low index materials in the lattice and the arrangement of the lattice elements. As a general rule, the higher the refractive index contrast, the stronger the expected effects and the larger the band gap achievable. The refractive index contrast is typically limited to that between GaAs ($n = 3.6$) and air. As to lattice configurations, the most promising ones are the “honeycomb” [21] (Fig. 2) and the “graphite” [22] structure where the widest band gaps occur when small regions of semiconductor are surrounded by large regions of air, i.e. when the semiconductor “area fill-factor” is low, a direct consequence of the high refractive index of the semiconductor. The “hand-waving” reason for this condition is that the maximum effect can be achieved when the optical path length in both materials is roughly equal, in direct analogy to the “ $\lambda/4$ ” criterion of dielectric Bragg mirrors.¹

An additional factor is that of light confinement in the third dimension; there, the structure must either: (a) be large compared to the beam size (i.e. many wavelengths deep) or (b) the beam must be sufficiently confined within a waveguide to experience the full interaction with the periodic lattice. Several “infinitely deep” structures have been successfully demonstrated [23–27] (Fig. 3), although most structures presented to date utilise some sort of waveguide confinement. The waveguide configuration has obvious advantages, such as growth of layered structures by established epitaxial methods and compatibility with other planar optoelectronic elements. There are, however, major problems that need to be addressed when designing the structure.

The naive point of view is as follows (Fig. 3): A high refractive index contrast is typically achieved by etching, so the structure is separated into areas of high-index

¹ The “ $\lambda/4$ ” criterion states that each layer in a Bragg mirror should be $\lambda/4n$ thick for highest reflectivity.




Fig. 2. Example of a typical 2-D photonic lattice, consisting of holes etched into semiconductor. The type of lattice is triangular, it is self-similar upon 60° rotation and the arrows illustrate the two main directions of symmetry. Solid-state physics provides the terminology, with the direction indicated by the darker arrow referred to as $\Gamma - K$ and the lighter arrow pointing along $\Gamma - M$.




Fig. 3. (a) A deeply etched waveguide, designed according to “conventional wisdom” for maximum reflectivity with the high and low index layers $\lambda/4n$ thick. This configuration suffers from diffraction losses, because light is not guided in the low index region and diffracts out of the waveguide plane. (b) If the air gaps, i.e. the etched holes or slots, are small enough, there is very little loss. Strong interaction between the periodic structure and the guided mode can then take place as long as the period of the structure is an integer multiple of half-wavelengths.

material, typically semiconductor (with the in-built waveguide structure) and air. Light is then only guided in the semiconductor and not while travelling through the air, so diffraction loss and scattering into the third dimension (i.e. out of the plane of the waveguide) is inevitable. This loss can be minimised by increasing the fill-fraction of semiconductor, for example by designing a lattice of pillars with narrow air gaps or a semiconductor honeycomb structure with small air holes.

The principle of this approach is that light can “hop” across the narrow gaps without suffering excessive loss, whilst still experiencing the full refractive index contrast between semiconductor and air. This solution compromises the highest achievable band gap and is in direct contrast to the “ $\lambda/4$ ” condition given above, but it

is the approach that has so far shown the most convincing PBG effects in structures with low-contrast claddings, e.g. semiconductor heterostructure waveguides. These structures have allowed the observation of a variety of PBG effects, including high transmission, reflection and very sharp filtering characteristics [29,30]. Also, keeping the substrate has the great advantage of allowing current injection. From an optical point of view, however, it would be better to remove the semiconductor substrate. Thin films of semiconductor on glass or in air, for example, support more strongly confined modes that can be inherently non-lossy [31,32] and thereby more suited to guiding and confining light.

3. The present — state-of-the-art

3.1. *How to make a photonic crystal?*

There is a variety of ways of making photonic crystals, many of which have been borrowed from the silicon microelectronics industry.

3.1.1. *Lithography and etching*

3.1.1.1. *Lithography.* As stated above, some of the impressive progress has been made with planar structures that were patterned by lithographic means. Considering the small size of the lattice (periodicity between 200 and 700 nm, with sub-100 nm control of feature size desirable), standard photolithography cannot be employed.² The most popular alternative is electron-beam lithography, which has been exploited by a variety of groups to generate patterns for use in the visible [33] to near-IR [34,35]. Several workers have also succeeded in using multiple exposure holography [36] to generate the triangular pattern. Multiple exposure holography is, in fact, a technique that has been used for some time [37,38] to produce two-dimensionally periodic structures.

3.1.1.2. *Masking.* In several ways, the most suitable structure for a 2-D photonic crystal is one consisting of a matrix of small holes etched deeply into the semiconductor. This requires deep and anisotropic etching with high resolution, so generally, conventional wet etching techniques cannot be used. Anodic etching, i.e. wet etching with an applied electric field, is an exception and is discussed separately in Section 3.1.2.

One of the most important elements contributing to the success of the etching process is the choice of an appropriate mask layer. The most straightforward solution is to use the resist itself, after development. Since the resist erodes at a considerable rate during the dry etch process, however, the obtainable etch depth is limited,

² With optical lithography advancing to shorter and shorter wavelengths, however, it can be expected that by the time devices based on photonic crystals are entering the marketplace in large volume, they will be made by photolithographic means.

particularly for high-resolution patterns that demand thin resists. Therefore, in order to allow deeper etching, a dielectric pattern-transfer layer is usually added. This layer (typically SiO_2 or Si_3N_4) is patterned by plasma etching using fluorine chemistry (we found that CHF_3 is superior to CF_4 or C_2F_6 because of lower resist erosion). The thickness chosen for this dielectric mask is a compromise and is typically around 100–200 nm; the thicker the mask, the longer it withstands the plasma and the deeper one can etch the semiconductor. The thinner the mask, the thinner a resist one can use to pattern it and thus achieve higher resolution. In our work, we found an erosion rate for the dielectric mask of around 15–20 nm/min.

Better masking is provided by a layer of oxidised $\text{Al}_x\text{Ga}_{1-x}\text{As}$ (' AlO_x '), an approach that was first demonstrated by Cheng et al. [20]. We have observed an erosion rate of only 3 nm/min for a layer of oxidised 80% x -fraction AlGaAs , which is a considerable improvement over externally deposited dielectrics. The oxidation is carried out at 400–450°C in a wet steam atmosphere, a process that has been widely studied, mainly for current aperturing in vertical cavity surface-emitting lasers (VCSELs) [39,40]. Despite this obvious improvement in erosion rate, AlO_x masks have not been widely used, in part because they have to be designed into the epitaxial material structure. Finally, one can add a thin layer of metal to the process, and pattern it either by lift-off or by ion milling [19] and then use the patterned metal layer to transfer the structure into the dielectric. This approach gives better masking, but raises the question whether the improvement is worth the additional effort, considering that each processing step increases the cost and failure rate.

3.1.1.3. Dry etching. The standard etch chemistry for $\text{GaAs}/\text{AlGaAs}$ semiconductors is chlorine-based, using either SiCl_4 or Cl_2 . Chemically assisted ion-beam etching (CAIBE) is arguably the most suitable dry etching technique available, with aspect ratios $> 20:1$ achieved [41] (120 nm holes etched $\sim 2.5 \mu\text{m}$ deep). In CAIBE, a collimated high-energy beam impinges on the sample, which results in better directionality and much reduced RIE lag.³

Our standard process for $\text{GaAs}/\text{AlGaAs}$ -based fabrication uses e-beam patterning of PMMA (200 nm thick, 350 K molecular weight), transfer into a ~ 150 nm layer of SiO_2 and etching using SiCl_4 in a standard RIE reactor. Typically, we obtain an aspect ratio of 6–8:1 (e.g. 100 nm holes etched 600–800 nm deep, Fig. 4), which has been sufficient to demonstrate a range of PBG effects in the waveguide geometry, with moderate losses [28,29,42].

InP-based photonic crystals can be etched in either CH_4/H_2 [43] or in chlorine-containing plasmas. Although vertical sidewalls and non-selective etching have both been achieved with CH_4/H_2 , the process is plagued by the buildup of polymer, which tends to clog up small features and, in our case, has prohibited the etching of lattice periods smaller than $0.5 \mu\text{m}$. This problem can be overcome by the use of chlorine chemistry, a solution that is not straightforward, however, because InCl_3 , one of the

³ RIE lag describes the fact that small features experience a smaller etch depth than large features, which is related to the lateral motion of the impinging ions.




Fig. 4. Two examples for 2-D photonic lattices, written by e-beam lithography and etched into a GaAs/AlGaAs waveguide structure using RIE and dielectric masks. The lattice periods are 290 nm (right), respectively. Left micrograph courtesy of C. Smith, University of Glasgow.

main reaction products, is not volatile below 130°C — and so a heated stage is required for the process. Hot etching of InP has been demonstrated successfully in both CAIBE [44] and with our standard RIE reactor, the main problem consists of finding the (reactor dependent) balance between the chemical (temperature, reactive gas flow) and physical (DC bias or acceleration voltage) etching components.

3.1.2. Electrochemistry

An alternative method for producing 2-D lattices involves electrochemistry, namely anodic etching and growth. Impressive aspect ratios have already been obtained, both in silicon (Fig. 5) and alumina. Initial work by Grüning et al. [23] produced 2-D crystals where the elusive goal of an “infinitely high” array of holes was closely approximated. The term “infinitely high” is to be understood in the context of the wavelength; the structures were up to 340 μm deep and exhibited band gaps around 40 μm wavelength. The etched holes had a diameter of 6.2 μm , which constitutes an aspect ratio of $> 50:1$ and compares favourably with the ratios $< 10:1$ that are typically achieved by dry-etching. Due to the generally “macroscopic” (1–50 μm) feature size of the electrochemically etched holes, the material thus produced is termed “macroporous” silicon, in contrast to its “microporous” relative (Section 4.3.1). Macroporous silicon with feature sizes in the sub-micron range has been demonstrated [27,45], although most PBG-related experimental work is aimed at mid- to far-IR wavelengths [25,26].

Much smaller structures that could display photonic band gaps in the visible regime have already been produced by the related technique of anodic growth of alumina pores. The pores were grown on a surface of aluminium embossed with regular arrays of concave dimples that acted as seeds for the growth. Impressive aspect ratios for




Fig. 5. Micrograph of a 2-D photonic crystal fabricated by anodic etching (“macroporous silicon”). The structure depicts a hexagonal or “graphite” lattice as highlighted in the inset. Courtesy of A. Chelnokov and J.-M. Lourtioz, Université Paris-Sud, France.

holes, with diameters as small as 70 nm and periodicities of 100 nm, were reported by Masuda et al. [46].

The high degree of regularity achievable with electrochemical processes surpasses that of structures generated by lithographic means, but constitutes a mixed blessing; the only requirement to initiate the process is some discontinuity on the surface. Once the process has started, the pore size is determined by the electrochemical current, not the size of the seed [46]. Therefore, the growth mechanism has a tendency to accommodate deterministic defects and possibly eliminate them as the process proceeds, producing a defect-free crystal after a certain depth (or height, respectively).

3.1.3. Vertical selective oxidation

The motivation for etching holes or slots into a material is to create a high refractive index contrast. By removing material, however, the surface is left perforated, which makes it difficult to deposit mirrors or contacts afterwards. Also, the device may age prematurely because of chemical reactions at the exposed surface.

An intriguing variant to removing material that addresses these issues is the technique of selective oxidation. A lattice produced by selective oxidation maintains the physical integrity of the material and produces an “all-solid” microstructure. The oxidation technique was first proposed in the early 1990’s [40] as a way of converting AIAs (or AlGaAs with a high Al-content, generally $\text{Al} > 80\%$) into a thermodynamically stable oxide. The process involves placing the sample into a steam environment at 400–450°C and produces an oxide with two desirable qualities: (a) The oxide is a good electrical insulator, which is the reason for its use as a current-blocking layer




Fig. 6. Example for the technique of vertical oxidation, driven by a composition gradient. The oxidation rate in $\text{Al}_x\text{Ga}_{1-x}\text{As}$ is exponentially dependent on the Al-fraction x . By increasing the Al-fraction with depth, an artificial anisotropy is introduced into the oxidation process, leading, ideally, to “oxide pillars”. The sketch illustrates the progression of the oxidation front in four steps with time and the micrograph depicts an oxide pillar that was slightly over-oxidised. In the particular example, the composition was graded from $x = 0.88$ near the surface to $x = 0.95$ near the substrate. The oxidation time was 3 min at 420°C in an open-tube furnace.

in vertical cavity surface-emitting lasers, VCSELs [47], and (b) it has a low refractive index, i.e. between 1.5 and 1.6 [48]. The low refractive index is the most exciting aspect as far as photonic crystals are concerned, because it facilitates the growth of the entire structure, including all electrical and optical functions, by standard epitaxial means with the low-index areas being converted by relatively simple processing afterwards.

The main problem with the technique is that it usually proceeds in the lateral direction, i.e. in a plane parallel to the surface, whereas typical photonic lattices are structured in the vertical direction, i.e. normal to the surface. Our group⁴ has recently explored the possibility of vertical oxidation, with the aim to replace etched holes by “oxide pillars” (Fig. 6). To accomplish substantially vertical oxidation, the Al-content of an AlGaAs layer is graded, with the higher Al-content closer to the substrate. The exponential dependence of the oxidation rate on the Al-content [39] then drives the reaction downwards in an anisotropic fashion. For masking, we have used a thin (50 nm) layer of GaAs, which was either removed locally by wet etching or damaged by exposure to an RIE-generated plasma of CHF_3 . The latter process has the great advantage of leaving a near-perfect surface with the high contrast microstructure embedded in the material underneath.

Our current estimate for the aspect ratio of these pillars is around 1 : 1, which is not as impressive as the ratios of upto 10 : 1 achieved with dry etching techniques. Another problem is the volume shrinkage of up to 13% [39] which can lead to strain, microcracks or structural problems such as delamination of the layer structure if the oxide is too thick. Careful process control is also required, as illustrated by the overshoot seen in Fig. 6.

Overall, the technique of vertical oxidation is not without drawbacks and requires more work, but it eliminates the side-effects of dry etching, such as damage and ion

⁴ In collaboration with Dr. M. Dawson, Strathclyde University, Glasgow, UK.

channelling [49] effects that increase the non-radiative recombination in active microstructures. It also leaves a planar or near-planar surface which opens options for subsequent deposition of mirrors and contacts.

3.1.4. Membranes

The fabrication routine for waveguide-based photonic lattices described above assumes that the light is confined by a “laser-like” semiconductor heterostructure. This approach has a variety of advantages, e.g. simplicity, mechanical stability and, most importantly, it allows current injection. The main drawback, however, is the fact that the light is rather weakly confined by the limited refractive index contrast available with a semiconductor heterostructure, and, given the opportunity, will scatter into the substrate. As mentioned earlier, this scattering loss can be minimised by careful design. It cannot be completely eradicated, however, unless the refractive index contrast between the waveguide core and cladding is increased. Microstructures that consist of a “perforated” waveguide core clad by glass or air, for example, have the ability to support truly lossless modes [31,50] that can propagate through the structure without suffering any loss at all. The purpose of the present section is to present several different ways of making such membranes, i.e. thin layers of semiconductor surrounded by low-index material.

3.1.4.1. Silicon on silica. The easiest approach is to use a material that comes with the desired properties built-in, such as silicon on silica, i.e. quartz. Wafers with thin films of polycrystalline silicon on quartz buffers can be readily obtained and have been used to make high-quality waveguides [51] and microcavities [32].

3.1.4.2. Low-index buffer. The second easiest approach is to use a semiconductor heterostructure and modify it in some way after growth to create a low-index buffer between the waveguide core and the (high-index) substrate. This modification can either be in the form of wet etching or selective oxidation (Fig. 7). Wet etching creates an “airbridge” or suspended membrane and requires a selective wet etch in a suitable




Fig. 7. Semiconductor “membrane”-style waveguide on a partially oxidised buffer. The oxidation was terminated early to study the progression of the oxidation front and to illustrate the process.

material system, e.g. InP/InGaAs [52,53], GaAs/AlGaAs [54] or $\text{Si}_3\text{N}_4/\text{SiO}_2$ [33]. Selective oxidation relies on the steam oxidation process described above and is based on the fact that the oxidation rate of $\text{Al}_x\text{Ga}_{1-x}\text{As}$ depends exponentially on x for $0.8 < x \leq 1.0$. The refractive index of the oxide (“AlOx”) is around 1.6 [48] and thereby provides almost the same strong confinement as the airspace created by wet etching. The oxide buffer, however, creates a structure that is more robust and has better thermal properties because of its “all-solid” nature, although one needs to be aware of strain and possible delamination as already mentioned above.

3.1.4.3. Epitaxial lift-off/substrate removal. The third possibility is to separate the epi-layer from its original substrate and place it onto another substrate with more desirable properties. A durable joint between the membrane and the new substrate can be made with UV-curable glue or simply by relying on van der Waals forces, which is the basis of the technique known as epitaxial lift-off (ELO). The development of ELO was originally motivated by the need for cheaper substrates in the solar cell industry, since the original substrates can be re-used after the epilayer has been lifted off [55,56]. The technique is also very promising for GaAs-on silicon optoelectronics [57] and the fabrication of efficient light-emitting diodes [58].

Which one of these techniques should be chosen depends on the material system, but also on experience and processing capability. The wet-etched membranes are particularly suitable for the InP/InGaAs/InGaAsP system because of the variety and high selectivity of materials and etches available, while selective oxidation is best suited for the GaAs/AlGaAs system. Substrate removal is an option for many materials that feature semiconductor heterstructures, as long as there is sufficient wet-etch selectivity for different compositions.

3.1.5. Other fabrication methods

The fabrication methods discussed so far have mostly been “borrowed” from the silicon microelectronics industry, i.e. they all involve some kind of lithography followed by transfer into the host material. It is desirable to extend these methods, because lithography usually requires planar substrates whereas the ideal, three-dimensional photonic crystal is periodic in three dimensions. Moreover, some materials, such as light-emitting polymers (Section 4.3.3.), cannot easily be structured in this manner and require other processes.

3.1.5.1. Building up the crystal. The alternative to removing material to form the required microstructure is to build it up. The anodic growth technique mentioned in the context of electrochemical processes (Section 3.1.2) already fits into this category, and there are related techniques such as laser prototyping and stereolithography [59] that have not yet produced truly “photonic” structures, i.e. structures that operate around the $1 \mu\text{m}$ wavelength range, but have great promise for doing so in the near future. A type of photonic crystal that already passes the criterion of size is the one first presented by Kawakami et al. [60,61], a structure that is built-up from a patterned substrate by sputtering alternate layers of SiO_2 and Si to create a truly three-dimensional crystal. The authors report that the sputtering conditions can be

adjusted to yield a constant mark/space ratio pattern as the structure builds up [62]. To underline the feasibility of this approach, they have reported very powerful wavelength selectivity dispersive effects (“superprism phenomena”) [61] that occur because of strong dispersion in certain parts of the photonic bandstructure.

3.1.5.2. Embossing. Light-emitting polymers such as para-phenylene vinylene (PPV) [63] are not easily processed into the shapes and sizes desirable for photonic crystals, so one needs to resort to a different approach; embossing is one of a range of “soft lithography” techniques [64] that offers itself as a solution. The principle of embossing is to bring a negative preform of the desired pattern and the polymer into close contact. By applying suitable pressure and temperature, the pattern is transferred from the preform into the polymer. Embossing is often used for mass production of diffractive optical elements and has a demonstrated resolution in the sub-100 nm regime. Microstructuring PPV by embossing has now been demonstrated [65], which opens up a whole new range of possibilities for light-emitting polymers. One of the problems with this technique is the adhesion of the polymer to the preform once the pressure has been released, an issue that can be addressed by adjusting the chemical composition of the polymer.

3.1.5.3. Fibre-pulling. Stacking hollow glass rods and pulling them into fibres offers yet another method of producing “infinitely deep” photonic crystals. The technique was first demonstrated in 1992 [66] and has since produced several interesting results, such as photonic band-gap behaviour in the visible [24] and the very intriguing photonic crystal fibre. These fibres have a virtually unlimited single-mode bandwidth [67], i.e. they do not (within the useful window determined by absorption losses) allow the propagation of higher order waveguide modes, whatever the wavelength of operation.

3.1.6. Self-organised photonic crystals, notably opal

Most of the fabrication methods discussed so far impose periodicity on a photonic scale onto the otherwise uniform material. This approach has great strengths, particularly the ability to include deterministic defects, but also a variety of weaknesses, for example the lack of a viable fabrication technology for 3-D structures that mimic the type of performance already demonstrated in the microwave regime.

3.1.6.1. Background. A possible solution lies in self-organisation. Self-organisation is a fundamental feature of nature and much of what makes existence possible can be said to rely on self-organisation. Within physical science, the coalescence of solid materials into remarkably perfect single crystals that we refer to as epitaxial growth is something we tend to take for granted, even if our task is to produce the massive single crystals now demanded by, for instance, the silicon industry. Such ‘electronic’ crystals are made up from atomic building blocks which are glued together by relatively strong binding forces — a key point when we turn to the subject of self-organised photonic crystals, such as opal. Opal consists of a regular arrangement of nearly spherical balls of amorphous silica (Fig. 8). Opal occurs naturally as a semi-precious




Fig. 8. Close-up and overview of the structure of opal, with the silicaspheres that make up the material clearly visible. The spheres are approximately 260 nm in diameter. Courtesy of N.P. Johnson and S. Romanov, University of Glasgow.

material renowned for its iridescent behaviour when observed in white light. The synthetic formation of bare opal is also possible, and involves the formation of nearly spherical balls of silica with appropriate size and their arrangement into a multilayer “crystal” with substantial regions of regularity.

Other self-organised structures include the colloidal photonic crystals formed from monodisperse polystyrene particles [68] and the more recent work by Holland, Blanford and Stein in which balls of organic material (latex) are formed into regular arrays and then provide a template for the formation of regular arrays of titania, zirconia or alumina matrices with open-pore structures [69]. Wijnhoven and Vos [70] describe a similar approach based on colloidal silica balls which are removed to form air spheres after infilling with titania. It is worth remarking here that the microcrystalline nature of the material which forms the open-pore structure is such that it exhibits a high level of Rayleigh scattering, so that the out-of-stopband or off-normal incidence transmission may also be negligible, thus masking true photonic band-gap behaviour.

But let us return to synthetic opal. Bogomolov and co-workers [71] have carried out a substantial amount of work on synthetic opal and established a “standard” process, which begins with the colloidal formation of spherical balls of silica [72,73], followed by the sedimentation or centrifuging of the silica spheres into a “crystal” in which the monodispersed spheres pack into a solid assembly. Substantial regions of material typically exhibit a regular face-centred cubic (f.c.c.) arrangement.

3.1.6.2. Loading opal. While properly formed “bare” synthetic opal, like good samples of its naturally occurring counterpart, is clearly a photonic crystal, the maximum refractive contrast available, 1.45 : 1, is insufficient to give full photonic band-gap behaviour. The possibility has been recognised for several years [74,75] that semiconductor material can, using suitable “loading” processes, fully or partially fill the voids inherent in the structure of opal and at the same time provide both a larger

refractive index contrast and a periodic gain medium. Furthermore, by back-filling opal with high-index material, a more favourable reactive index distribution, i.e. one approaching the “ $\lambda/4$ ” criterion (see ¹), can be achieved, where the optical path in the high-index material equals that in the low-index material. Thus, only a small amount of high index material is required, which is typically what the loading process, e.g. with suitable semiconductor material, produces. The technology of the loading process is challenging, but significant progress has been made both using vapour-phase deposition processes [74] and liquid-phase filling processes [75]. Fig. 9 shows an example for the reflected diffraction behaviour of lightly CdS-loaded opal.

As we write, there has neither been a clear demonstration of laser action nor of strong omnidirectional spontaneous emission inhibition in opal-based materials. These failures may well be surmounted before long but are, we believe, the result of problems in achieving perfect regularity in both the crystal synthesis and the loading process, a sufficient index contrast and a high-quality luminescence from the in-filled material. Our general view on self-organised photonic crystals is that they form a promising and intellectually challenging topic. Possibly, the most obvious application is their use as a periodic gain medium for photopumped laser action, with a broad spectrum available for tuning via an additional element in the laser cavity. Tuning may also be possible by loading opal with a suitable liquid crystal. In the longer run,




Fig. 9. Dependence on the photon energy of the (111)-planed Bragg reflection peak of synthetic opal. The opal has been lightly “doped” with a thin layer of cadmium sulphide (CdS) using a cadmium compound in solution, followed by chemical conversion to CdS. The thickness of the CdS layer is estimated to be in the region of 2–4 nm. The reflection peak for normal incidence ($q = 90^\circ$) is observed at a photon energy of 2.2 eV (i.e. a free-space wavelength of 0.564 μm). The position of this peak is consistent with the Bragg law for a structure consisting of touching silica balls of diameter $D = 260$ nm with an average refractive index of 1.33. The shift in peak position can possibly be explained by taking into account the strong dispersion of CdS near its electronic band gap (2.5 eV) as well as Snell’s law, since the incident ray is refracted as it enters the opal. As a result, the shift from 90 to 70° is smaller than that from 70 to 50° and the average refractive index has increased to about 1.4. The inset confirms that low transmission coincides with high reflection as expected. Courtesy of S.G Romanov, University of Glasgow [75].

electroluminescent structures based on semiconductor-loaded or electroluminescent dye-loaded synthetic opal are a possibility. It may eventually be possible to produce suitably sized opal substrates at low enough cost through mass production — and then apply opal or other self-organised photonic crystals to areas such as display-device technology.

3.2. *Different methods of characterising photonic crystals*

Performing measurements on PBG structures designed for the microwave regime is relatively easy and is usually done with monopole antennas or waveguide horns [17,76]. The same applies to the “large” structures based on opal, macroporous silicon [26] or those made by the fibre-pulling technique [24], structures that can be examined by placing them in the path of a collimated broadband source and measuring the transmission/reflection as a function of wavelength. The characterisation of *waveguide*-based structures in the optical regime, however, is not as obvious and a variety of ingenious approaches have been developed.

3.2.1. *Tuneable laser*

The coherent light of a laser-source has the distinct advantage that it can be focused into a micron-size spot and therefore be coupled into a semiconductor waveguide relatively easily via a single-mode fibre or directly via “end-fire” coupling. These are appealing features, but there is also a variety of drawbacks.

- (a) The tuning range of the laser is limited. Even sources as versatile as dye and Ti : Sapphire lasers tend to have their continuous tuning range limited to 10–20% of the centre wavelength, because of the limited spectrum of the gain medium or the bandwidth of the mirrors. In order to map out a PBG structure, however, a tuning range of 30–40% is often required.
- (b) Coupling to the waveguide can only be performed in a normal or near-normal direction with respect to its end-face. It is often the angular-dependent response of the structure, however, that is of real interest.
- (c) An external source, such as a tuneable laser, can only probe the *external* properties of the lattice and not examine active structures, such as microcavities, where the interest lies in studying the effect of the microstructure on the emission process.

The solution to (a) and (b) is “lithographic tuning” as illustrated in Fig. 10; by stepping the lattice period in a controlled fashion, different parts of the respective stopbands can be assessed with the same, limited wavelength-span of the source. Furthermore, the lattices can be lithographically rotated to probe different symmetry directions, or curved waveguides can be used to direct the light in the desired direction. Using these techniques, we were able to show the first clear experimental signatures of photonic band-gap behaviour in the near-infrared [28].

3.2.2. *In-built photoluminescence*

Using the in-built photoluminescence of the semiconductor material makes for a more versatile measurement, since the source can be excited at any point (Fig. 11).




Fig. 10. Lithographic band-gap tuning illustrated on a 1-D photonic lattice. The structure consists of an array of slots etched into a waveguide structure designed for a third-order stopband in the 800–900 nm wavelength region [30]. The centre element is slightly enlarged to form a defect state in the centre of the stopband. As the lattice spacing is increased, the spectral response shifts to longer wavelength and a different part of the stopband appears in the 820–900 nm wavelength window accessible with our tuneable Ti : Sapphire laser.




Fig. 11. In-built source to probe the photonic lattice. The quantum wells embedded in the structure are photoexcited by a focused laser spot. Some of the luminescence couples to the waveguide mode, travels to the microstructure, interacts with it and finally couples out of the cleaved facet to be collected and spectrally analysed. The distance between the excitation spot and the cleaved facet is limited to between 20 and 100 μm due to reabsorption in the material. Courtesy of D. Labilloy, Ecole Polytechnique, Palaiseau, France.

The waveguides indicated above are not required, but the patterns have to be placed in proximity to a cleaved edge for measurement. The spectral width of the source is even more limited than that of a Ti:Sapphire type tuneable laser, particularly when bulk or quantum well material is used [77]. This limitation is made more severe by the presence of (wavelength-dependent) re-absorption in the unpumped parts of the waveguide.

In a more recent development, self-organised InAs quantum dots with a deliberately broad emission spectrum have shown much improved results [78]. Due to the larger bandwidth and the much reduced re-absorption in the unpumped waveguide, it was possible to map out a large part of the stop-band of an individual lattice (Fig. 12). Due to its versatility, the technique has enabled the simultaneous measurement of transmission, reflection and diffraction, i.e. a complete and quantitative assessment of the interaction between the incoming light and the photonic lattice [29]. Moreover, the PL-technique can be used to excite light inside a structure and examine the *internal* properties of the photonic crystal, e.g. the emission from a photonic crystal microcavity [79].

3.2.3. White light measurements

The easiest technique, as far as the light source is concerned, is to use white light, i.e. a broadband light source. Because white light is not coherent, it cannot be focused into a spot small enough to allow launching into narrow waveguides, so the technique is either limited to the case of probing large structures or to techniques that couple into the waveguide via gratings or evanescent fields (e.g. prism-coupling [38,81]). In the white light reflectance measurements presented recently [54,82], the photonic lattice acts as its own grating coupler. As the light impinges on the microstructure, it is generally reflected, unless the wavevector-component of a mode that is allowed to propagate through the structure matches that of the incoming light. In that case, i.e. when the two wavevectors are matched, some of the energy leaks into the waveguide with a corresponding dip in the reflectivity. Due to Snell's law and the high index of




Fig. 12. Transmission through a 2-D semiconductor waveguide based photonic crystal using internal photoluminescence as a light-source, measured by the method shown in Fig. 11. The alternating full and light curves represent data for structures of different lattice constant, which have been pasted together via the normalised lattice constant $u = a/\lambda$. The fact that the different curves join up so well, particularly the fine structure around $u = 0.3$ underlines the reproducibility of the measurement and the accuracy with which these structures can be fabricated using state-of-the-art technology [80]. The lattice period varies between 180 and 330 nm and the wavelength is around 1 μm . Courtesy of D. Labilloy, Ecole Polytechnique, Palaiseau, France.




Fig. 13. Experimentally observed bandstructure for a semiconductor waveguide-based 2-D photonic lattice. The circular and square points are for TE and TM modes, respectively. The points are obtained from white light reflectivity measurements by scanning the angle of incidence; if the \mathbf{k} -vector of the incident light matches that of a lattice mode, a dip in the reflectivity is observed because some of the energy couples to a propagating mode. Courtesy of V. Astratov and M. S. Skolnick, University of Sheffield.

the semiconductor, only modes near the centre of the Brillouin zone can fulfil this matching condition, so only modes near the Γ -point can be addressed (Fig. 13).

4. The future — applications for photonic crystals

4.1. Light emitting diodes (LEDs)

As indicated in the introduction, we see some of the major applications for photonic crystals in novel light sources. Consider, for example, the light emitting diode (LED). The active material of LEDs is still invariably semiconductor, which is the most efficient light-generating material known, with up to 99.7% internal quantum efficiency reported [83]. The problem then lies in the extraction of this light: as a consequence of Snell's law, only the radiation within a small cone (around 16° for GaAs) can escape, everything else being totally internally reflected. A simple calculation shows that the extraction efficiency scales as $1/4n^2$, so the 16° allowed by Snell's law only represent around 2% of the total available solid angle, which is the reason for the low external efficiency of 2–4% observed in standard commercial LEDs.

This discrepancy between external and internal efficiency clearly needs to be addressed. High brightness LEDs already exist, they extract light from more than one facet and achieve external efficiencies around 20%. These devices are beginning to be used in display and lighting (e.g. traffic light) applications because of their superior brightness and lifetime with respect to incandescent lightbulbs [84]. Another

approach is to employ the somewhat unusual technology of placing the active layer on a low index substrate and then roughening it. The resulting surface scatters a large fraction of the light out of the material and leads to a measured efficiency of around 30% [58]. We believe that the most promising approach, however, is based on photonic crystals because both the extraction and the control over the direction of light emission is improved. Many applications require light sources with a more directional output from the light source for coupling into the narrow acceptance cone of a fibre, e.g. applications in short-distance communications, such as fibre to the home, automotive and avionics. Data rates are typically no higher than a few 100 Mbit/s and the price pressure is very high, so LEDs have clear advantages.

The general idea is sketched in Fig. 1. By controlling the light that is emitted by the material, the emission into the unwanted directions is suppressed and that into the desired directions enhanced. Microcavity LEDs (MCLEDs) are the first manifestation of this concept. They cannot genuinely be referred to as photonic crystals, but share the concept of using periodic structures for strong confinement: 23% external efficiency, extracted from a single facet, has already been demonstrated [85,86]. Extending this concept into two and three dimensions should push the limit further and lead to devices with an external efficiency up to 50% or even beyond.

Research is currently being conducted in a number of centres along a variety of avenues that have the common factor of involving a (micro-)cavity of some description. The two main ideas are (a) increased extraction/photon recycling and (b) alteration of the fundamental emission process.

4.1.1. Devices with increased extraction capability

In a microcavity LED, one can generally distinguish between three types of mode: (a) the extracted modes, (b) the modes that are confined to the waveguide formed by the high-index active layer and (c) the modes that leak into the substrate and cannot be collected. The extracted modes constitute the useful output of the device and so should be enhanced.

4.1.1.1. Thin cavities. By reducing the thickness of the optical cavity, the number of allowed modes can be restricted to the bare minimum. The extraction efficiency then increases because if, out of the k modes that the emission process can fill, a single mode is extracted, the extraction efficiency (and hence the external efficiency) scales as $1/k$. This is, in naive terms, the main factor behind the 23% external quantum efficiency already obtained with MCLEDs [85,86]. The other factor involved is photon recycling, which will be discussed below. The realisation of a thin (e.g. one wavelength thick) cavity device involves a λ -thick central layer bounded by a metal mirror on one side and by a GaAs/AlGaAs Bragg stack on the other [86], forming a structure that supports between 4 and 5 modes. Although the cavity spacer is only 1 wavelength thick, which suggests that the cavity can support two modes, the field penetrates substantially into the mirrors and thus increases the effective cavity length. In order to reduce this penetration depth, the substrate index can be reduced by placing the active layer onto a dielectric material, e.g. by using epitaxial lift-off [58] or by using selectively oxidised buffer layers (see also Section 3.1.3). These approaches, however,

make current injection much more difficult and either lateral injection is required [58] or the use of conducting dielectrics, such as Indium Tin Oxide (ITO).

4.1.1.2. Surface roughness/grating couplers. The other way to increase the amount of light extracted is to scatter the waveguide modes out of the material via surface roughness [58] or, in a more controlled manner, via a surface grating (“grating coupler”) thereby converting the waveguide modes into extracted modes. The “ordered” grating method has only recently been proposed for MCLEDs [87] and has yet to prove its superiority over the “random” scattering approach. Furthermore, the extraction angle depends on both the propagation constant of the extracted waveguide mode and the effective period of the grating, which is a function of the propagation direction in the plane⁵ and may therefore limit the maximum extraction. Clever design, however, such as the use of chirped gratings or quasiperiodic structures (Section 4.6), may overcome these difficulties.

4.1.1.3. Photon recycling. A further factor is the apparent increase in the internal efficiency that can be realised by re-absorption and re-emission, or “photon recycling”. Photons, particularly those emitted into the waveguide modes, can have a high probability of reabsorption if the interaction length with the active layer is long enough [86]⁶ and the overlap with the active layer is sufficiently strong [88]. An impressive demonstration of the extent to which the recycling effect can be used was given by Schnitzer and Yablonovitch [83], where each photon was re-emitted for an average 25 times, leading to a total external efficiency of 72%. This type of performance, however, is only possible with material of exceptional quality, the authors quoting 99.7% internal quantum efficiency in that particular example. Also, recycling increases the photon lifetime, thus slowing down the response-time of the device, which may well be a consideration for applications in communications systems.

4.1.2. The Purcell effect. Cavity enhanced light-emitting diodes [89]

One of the key motivations for research into photonic crystals is the promise of spontaneous emission enhancement, i.e. the fact that the spontaneous emission rate of an excited atom can be increased by placing it inside a microcavity. In particular, such a microcavity may be formed by creating a defect in a photonic microstructure. Theoretically, the alteration of the emission rate is usually treated within the framework of Fermi’s golden rule, i.e. the coupling between the atom and the cavity is considered weak. Fermi’s golden rule states that the emission rate is proportional to the product of the mode density (i.e. the spectral density of states) and the matrix element for the atom–field interaction.

To study the effect of the cavity, we first consider the free-space spectral mode density $\rho(\omega)$ and convert it to a volume-normalised mode density $g(\omega) = \rho(\omega)/V$. If the

⁵ The difference between the Γ -M and Γ -K directions in a triangular lattice is a manifestation of this different effective period.

⁶ This is also why in the work De Neve et al. [86], the best external efficiency was obtained for large area devices.

emission process takes place in a material of index n , the effective volume changes accordingly, see Eq. (6). The free-space value for $g(\omega)$ is known from classical mechanics to be proportional to ω^2 . In particular [90],

$$g_f = \frac{\omega^2}{\pi^2 c^3} \tag{1}$$

The cavity restricts the modes to a spectral width $\Delta\omega$ and a volume V_{cav} , which transforms the mode density into

$$g_c = \frac{1}{\Delta\omega V_{\text{cav}}} \tag{2}$$

This expression assumes that (a) the cavity and the atomic emitter are in resonance, and (b) that the linewidth of the cavity resonance exceeds that of the emission linewidth. This is a very important assumption and is illustrated in Fig. 14.

The cavity is usually described by the quality factor Q , which can be expressed as $Q = \omega/\Delta\omega$, so Eq. (2) becomes

$$g_c = \frac{Q}{\omega V_{\text{cav}}} \tag{3}$$




Fig. 14. Difference between a laser and a cavity-enhanced LED, Purcell style. In a semiconductor laser, the resonator typically has a high Q , i.e. a narrow linewidth, and sets the resonance wavelength. The emission at this wavelength is strongly enhanced and the emission at all other wavelengths suppressed. In a cavity-enhanced LED, all wavelengths experience enhanced emission if the material gain bandwidth is smaller than the cavity resonance bandwidth, or, expressed differently, if $Q(\text{material}) > Q(\text{cavity})$. For example, an InGaAs/InP heterostructure operating at room temperature at a wavelength of around $1.6\ \mu\text{m}$ has a material Q of approximately 10, so the cavity Q should ideally not exceed this value. Defining the gain bandwidth of a material as “material Q ” is somewhat unusual, but the analogy is obvious if one interprets the material linewidth $\Delta\lambda$ in the same way as the cavity linewidth via $Q = \lambda_0/\Delta\lambda$.

The mode density alteration caused by the cavity is thus given by g_c/g_f , or

$$\frac{g_c}{g_f} = \frac{\lambda^3 Q}{8\pi V}. \quad (4)$$

Two further considerations are required for the final result.

(a) A real cavity is best described by a Lorentzian line shape function, so Eq. (2) becomes

$$g_c = \frac{2}{\pi} \frac{1}{\Delta\omega V_{\text{cav}}}. \quad (2a)$$

(b) The maximum possible enhancement is achieved when the radiating dipole, in relation to the three orthogonal coordinate axes, is oriented to experience the maximum interaction with the cavity mode. This gives an additional factor of three over the bulk-case.⁷ In total, we obtain the following expression for the cavity enhancement factor, which was first mentioned by Purcell and is therefore often referred to as the Purcell-factor [5]. We shall call it f_p in our discussion,

$$f_p = \frac{g_c}{g_f} \times \frac{2}{\pi} \times 3 = \frac{3}{4\pi^2} \frac{Q\lambda^3}{V} \quad (5)$$

For convenience, we shall write this equation in a slightly different form,

$$f_p = \frac{6}{\pi^2} \frac{Q}{V'} \quad \text{with} \quad V' = \frac{V}{\left(\frac{\lambda}{2\pi}\right)^3} \quad (6)$$

accounting for the fact that the emitter is usually embedded in a material with index n and that the number of modes in the cavity is approximately quantised in integer multiples of the half-wavelength. On closer examination, f_p can be seen to reach values of 100 or more for parameters that have been independently reported in semiconductor microcavities, e.g. Q 's over 1000 [92] and mode volumes as small as 5 cubic half-wavelengths [32]. This shows that there is considerable potential in photonic crystal-based microcavities, because they can produce microcavities with high Q and low volume. It also shows the benefit of two- and three dimensional optical confinement, considering the well-known fact that a maximum value of $f_p = 3$ is possible with a purely one-dimensional structure [93]⁸. Overall, equation (6) tells us that we need to increase Q or reduce V' , or both, if we want to achieve a significant

⁷ This additional factor of 3 cannot be assumed automatically and only describes the *maximum* enhancement possible. A good example for this is given by Gerard et al. [91], where a maximum f_p of 32 was calculated, with a measured value of 5. The difference is due to spatial and orientational overlap, i.e. only $\sim 1/2$ of the dots overlap with the cavity mode and, since the dot emission is randomly polarised, only $\sim 1/3$ of the dots provide gain.

⁸ The enhancement factor of three assumes perfect mirrors. It quickly drops to values below 2 when mirrors with realistically finite reflectivity are put into the equation.

spontaneous emission enhancement. How can we do this, and what is the highest enhancement factor we can expect in practise?

4.1.2.1. High Q cavities. The high- Q cavity seems the most obvious route to take, since high- Q semiconductor-based microcavities already exist, e.g. in vertical cavity surface-emitting lasers (VCSELs). The argument sounds convincing, but there is a catch: very high Q 's are only useful if the line width of the active material is equally narrow, a condition not routinely fulfilled by semiconductor light emitters; Purcell-style cavity enhancement only works if the line width of the emitting material is equal to or smaller than the line width of the cavity, the assumption used in deriving Eq. (2) [14]. Using a high Q cavity and low Q material is typical with lasers, where the emission at one particular wavelength is very strongly enhanced (being predominantly simulated emission) but the emission at all or most other wavelengths in the available gain-spectrum is suppressed (Fig. 14). If the aim is a maximum enhancement primarily of the spontaneous emission, the enhancement must act on the entire emission bandwidth and therefore $Q_{\text{mat}} \leq Q_{\text{cav}}$. As a consequence, standard heterostructure emitters, including quantum well structures, are not very suitable. They have a material Q that typically ranges between 10 and 40, which is therefore the highest cavity Q that should be used. There is one active material system, however, that combines the narrow linewidth normally associated with atomic emitters and the high gain achievable in semiconductors. It is the system known as self-organised quantum dots, e.g. the nominally InAs quantum dots grown by the Stranski–Krastanov method [94]. These dots, at low temperature, have the ability to trap electron–hole pairs very efficiently and can thereby fully quantise the electron wave function, thus approaching the ideal scenario of quantum boxes. The resulting emission spectrum is an assembly of delta-function-like lines, with very high Q 's typically above 5000. Using such dots, pillar microcavities have shown Purcell-enhancement factors around 5 [91], which clearly demonstrates the feasibility of the principle.

The disadvantage, even for state-of-the art quantum dots, is that the size distribution is still relatively broad, i.e. the self-organised dots nucleate in different sizes, which leads to a broad distribution of emission wavelengths and relatively low spectral density. This means that, when placed into a high Q cavity, there are either no or only a few dots emitting at the right frequency. The second problem is the large carrier overspill that is observed at room temperature, which means that the carriers are not properly contained by the dots, but can escape and bypass the radiative recombination process. Therefore, the active layer appears to most carriers simply as a strongly corrugated quantum well rather than a series of quantum boxes.⁹ These two problems

⁹ A good example for this argument is the state-of-the-art threshold current density reported for quantum dot lasers [95,96]. At moderately cryogenic temperatures, when all the carriers are confined to quantum dots, the current density can be extremely low, around 10 A/cm². At room temperature, the threshold current density for emission from the same quantised state is approximately an order of magnitude higher, indicating that most carriers are not strictly confined to the quantum dots any more.

are not intrinsic, however, and may possibly be solved by the clever use of technology, e.g. seeding the dots or growing them close enough together in the vertical direction.

A problem that does, however, set limits to the practically useful Q is that of overlap of the emission line with the cavity resonance. Let us assume that it was possible to grow InAs dots, or anything else that has a very narrow, well-defined, emission line, in a uniform and controlled manner. The task would then be to make a cavity that resonates at exactly the same wavelength. This cavity would also have to be quite small and therefore be fabricated with Ångstrom-scale accuracy (see also Section 4.4 on WDM applications). This is possible, in principle, but becomes more difficult as the Q of the cavity increases.

4.1.2.2. Small volume cavities. The alternative approach to maximising (5) is to use cavities with very small volume, an approach that has so far only been explored theoretically. By studying a variety of possible defect configurations in a 2-D lattice etched into an approximately $\lambda/2$ thick semiconductor membrane, a volume $V' = 2$ cubic half-wavelengths was established as the tentative minimum [97]. The Q of these cavities varied between 10–100, depending on the type and size of the defect. This low value can be understood by considering the full spatial configuration; although the light is strongly confined to the defect in the plane of the photonic crystal slab, it can radiate freely in the weakly confined, vertical direction (Fig. 15).

Fig. 16 shows an example for the physical realisation of such a small cavity. To give a numerical example, let us consider small volume cavities in the InP/InGaAs system. This material system is known [98] to suffer relatively low surface recombination velocities, which means in practise that the non-radiative recombination at the exposed surfaces of the photonic lattice is tolerable. Using the minimum cavity volume of $V' = 2$ mentioned above [97] and assuming that the low Q of the material (approximately 10) is matched to the cavity, this yields, according to Eq. (6), an overall enhancement factor of 3.




Fig. 15. Illustration of the operation of a photonic crystal microcavity; light is confined to a small volume in the plane defined by the photonic lattice but free to radiate otherwise. The hexagonal section shown here can be understood as the unit cell of a large array of cavities, which is how these devices may eventually be used.




Fig. 16. A photonic crystal microcavity based on a semiconductor membrane. The structure was fabricated on a GaAs/AlAs heterostructure using steam oxidation. The confining layer delaminated after oxidation, resulting in a free-standing membrane. The microcavity, in the centre of the picture, is formed by an enlarged region between three holes [97].

4.1.2.3. Discussion. Both approaches, high Q or small volume, as far as they have been explored to date, yield similar results, i.e. an enhancement factor f_p between 3 and 5. The obvious route towards increasing f_p is therefore to combine the two approaches. The high Q method suffers from the lack of suitable room-temperature emitters and creates the problem of matching the narrow linewidths of the cavity and the material. The low volume approach, in contrast, is unsuitable for most commonly used semiconductors, because of surface recombination. Both constraints are material-related, and can therefore only be addressed by materials engineering, for example via further developing the growth of self-organised InAs dots. By making the dots larger, their high Q can be traded off against increased confinement [99]; by growing successive layers of dots in close proximity, the electron wavefunctions between different layers of dots can be made to overlap and the confinement should increase to the extent that carriers remain trapped even at room temperature, thereby avoiding carrier overspill and eliminating surface recombination. Moderate Q 's (e.g. 100–300) and volumes (e.g. $V' = 5$ –10) would then make a possible device easier to fabricate, yet still yield Purcell factors between 10 and 40, which is very respectable indeed. A more controlled approach to quantum dots, as provided by “seeding”, where the nucleation point of the dot is predetermined by lithography [100], could also help to improve the spatial overlap between the dots and the cavity mode. Of equal importance is the need to improve the orientational overlap between the dots and the cavity mode, which yields the factor 3 in Eq. (5) and is lacking in “standard” dots as shown by Gerard et al. [91].

A further consideration is that of carrier injection. The smaller the device, the more difficult it is to direct the current flow to the region of interest, and *only* there, because

current injection into the non-active parts of the matrix decreases the efficiency of the device. Once the carriers have arrived in the active part of the lattice, the next problem is that of non-radiative recombination at one of the many etched interfaces. III–V semiconductors, particularly GaAs, suffer badly from surface recombination, with InGaAs/InP and quantum dot systems, as already mentioned, being the notable exception [98,101].

An alternative to optimising the V 's and Q 's of the cavity [102] is to use the “bare” lattice as a light emitter. Very recent work [103] has shown that a structure similar to the one shown in Fig. 16 can provide significant enhancement of emission even without the intentional defect. The enhancement without defect/cavity takes advantage of the leaky nature of the conduction-band modes. The light first couples to these modes and then radiates out of the structure, which is yet another way of realising the aim of the exercise, i.e. increasing the extraction of light.

4.2. Lasers

An argument often made in the context of photonic crystals is that they have the ability to lower the threshold of semiconductor lasers. This is a very desirable prospect, because one of the major hurdles for work towards denser photonic integration is that of the removal of excess heat generated by the finite laser threshold. Some applications, e.g. photonic interconnects or large strips of devices for laser printing, envisage the use of thousands or tens-of-thousands of lasers [104]. If each device dissipated a milliwatt, the Watts of heat thus generated, in a small area, would cause considerable heatsinking problems.

Where does the laser threshold come from, and what can photonic crystals do about it? Essentially, this is a mode counting argument, because typically only one of the many available modes in a resonator leads to lasing. If we imagine a laser with a “box”-type cavity of volume V and refractive index n , consisting of material with gain bandwidth $\Delta\nu$ and operating at frequency ν , the number of modes in the cavity N_{cav} can be expressed as [105]

$$N_{\text{cav}} = V \frac{8\pi}{c^3} n^3 \nu^2 \Delta\nu. \quad (7)$$

For a typical semiconductor laser device, this amounts to about 10^4 – 10^5 modes, of which only one is the desired lasing mode, so much light is wasted before stimulated emission can occur. Furthermore, because the 10^4 – 10^5 non-lasing modes are still present after the onset of lasing, they cause noise. Reducing the number of available modes would therefore reduce both the threshold and dramatically modify the noise characteristics of the laser. The factor that quantifies this mode ratio and thereby determines the laser threshold is given by the following expression, a ratio known as the β -factor;

$$\beta = \frac{\Gamma \lambda^4}{8\pi V n^3 \Delta\lambda} \quad (8)$$

with Γ as the confinement factor (ratio of gain volume to cavity volume), using Eq. (7) and converting from frequency to wavelength. It is the β -factor that puts photonic crystals into the scene and motivates the interest in defects, because defects can place modes into an otherwise forbidden spectral regime. If, for example, one mode is artificially created in the forbidden regime by a defect, the β -factor assumes a value of unity, because the single allowed mode of the system is also the lasing mode. Making a laser is then, in principle, very simple, because all we need to do is to design the crystal so that its defect frequency coincides closely enough with the gain peak of the respective material.

The problem is that, as yet, there are no suitable crystals that can both reject all modes in a particular frequency range and exhibit gain at the same time, at least in the visible and near-infrared parts of the spectrum. In the microwave regime, related experiments have already been performed [76], where an antenna, when placed near the defect of a 3-D crystal, showed much improved radiation characteristics over the free-space case; equally, such antennas have been used to detect radiation where the effect of the crystal was quantified in terms of detection efficiencies [76]. A power enhancement factor of up to 3450 has thereby been demonstrated, which highlights the dramatic improvement possible with the PBG-approach.

These developments are all very encouraging, but we should also ask the question of how low a threshold is really needed in a laser. There have been predictions of sub- μA thresholds: are values as low as this really necessary? Many low-threshold laser applications are aimed at some form of communications applications, so we shall use a communications example to suggest an answer.

A back-of-the-envelope calculation shows that 1 W of light at 1 μm wavelength consists of around 5×10^{18} photons/s, so there are 5×10^{15} photons/s in 1 mW. If we assume a typical system loss of 30 dB between transmitter and receiver and a data rate of 10 Gbit/s, a speed that is already routinely achieved in research laboratories, at 50% duty cycle, we are left with 250 photons per bit. Considering that tolerable bit error rates are in the 10^{-8} – 10^{-10} range, these 250 photons are just about sufficient to ensure error-free data transmission, so 1 mW is the minimum power required for such a system. Lasers that operate at 1 mW would typically have threshold currents of the order of a few 100 μA , so lower thresholds than this would not be practical.

Clearly, these numbers can be somewhat adjusted to suit a particular argument, but the essence is that data transmission requires a certain number of photons to be reliable,¹⁰ and extremely low thresholds are therefore not sensible. Furthermore, thresholds around 50 μA are already being achieved [107] with vertical cavity lasers (VCSELs), due to their small volume, so one wonders whether the photonic bandgap effort really is necessary.

The first answer is that VCSELs can only be made in a select few laser materials, so photonic crystals, in principle, allow low threshold operation to be obtained for a wider variety of materials, e.g. the commercially important 1.55 μm wavelength

¹⁰ Unless specialist techniques such as “quantum cryptography” are used [106], which is possible, but unlikely to become the mainstream method of communication.

window, where VCSELs are elusive because of the low refractive index contrast available in the InP-material system. Another argument concerns noise; the above discussion assumes that many photons are needed for reliable data transfer. By reducing the number of available modes in the laser cavity, photonic crystals can reduce the fluctuation in power distribution across the available modes and thereby the emission noise. Finally, there are many applications that require high density arrays of light emitters and that impose different constraints on the individual devices. These array applications would clearly benefit from lower threshold currents, e.g. in laser printing or in data transfer over a large number of channels.

4.2.1. Examples of microcavity and photonic crystal based lasers

The type of laser with the highest beta-factor demonstrated so far has, in fact, nothing directly to do with periodic structures; it is the microdisk laser, nicknamed “thumbtack” laser after its peculiar shape. Values for β between 0.3 and 0.4 have been estimated for these [108] and for the “photonic wire” lasers [109] that use a similar circular geometry. Microcavity lasers based on defects in photonic crystals or small active areas surrounded by periodic structures should appear soon and compare favourably with the existing circular geometries. Configurations suitable for lasing have already been reported [52,53,79,110]. Other types of laser that employ periodic structures for feedback are the “horizontal” VCSELs that have now also been demonstrated, i.e. edge-emitting lasers with short high-reflectivity mirror stacks [111,112]. The mirrors rely on the same principles as waveguide-based photonic microstructures, i.e. they are deeply etched and use a waveguide for confinement in the dimension normal to the propagation direction. The best measured mirror reflectivities are around 95% [112] and the shortest lasers of this type yet realised are 42 μm long, amongst the shortest edge-emitters demonstrated to date. Projected threshold currents are below 100 μA and therefore as low as that of the best VCSELs [111,113].

Most of these initial demonstrations are very encouraging and anticipation is high for more to come. A serious problem with operating very small lasers, however, is the issue of current injection, even more so than with LEDs, because the current densities in lasers are typically an order of magnitude higher. Optically pumping these tiny devices, for proof-of-principle demonstrations; is therefore one thing, but efficiently directing the current flow is another, particularly when the active area is only a few μm^2 and is, ideally (from an optical point of view), surrounded by insulating dielectric material and/or air.

4.3. Other materials

The choice of active materials discussed so far was based on the philosophy of taking a good light emitter and making it better, hence the almost exclusive reference to III–V semiconductors. How about taking a bad light emitter, that may have other advantages, and making it competitive with III–V semiconductors? There are a variety of more unusual material systems in which electroluminescence has already been demonstrated, even in the visible wavelength range, which makes them applicable to display and possibly laser applications. All of these materials have some

distinct advantage over III–V semiconductors, be it cost, compatibility with silicon microelectronics or ease of fabrication. As examples, we shall discuss three different cases which highlight the issues involved:

4.3.1. *Porous silicon*

In this context, it is the “microporous” variety of silicon that we are interested in, i.e. the sponge-like skeleton of silicon veins on a nanometre scale [114]. Microporous silicon has the (for silicon) remarkable ability to emit light in the visible wavelength range. External quantum efficiencies (in electroluminescence) of around 1% have already been realised [115] and, being based on silicon, the host material is cheap, readily available and may eventually be used in conjunction with microelectronic circuits. We do not see much scope from the photonic band gap point-of-view, however, simply because the structure mainly consists of air. If regular, periodic structures were etched into this material, the effective refractive index change between the “spongy” host and air would be too low to cause the strong effects required for photonic band gaps to open up.

4.3.2. *Erbium-doped silicon*

The advantages of the silicon host quoted above also apply here. Furthermore, erbium is the key to providing gain in silica-fibres, and amplifiers based on erbium-doped silica are now widely used in the fibre optic telecommunications industry. So, combining silicon and erbium seems to be a winning concept. There is a fundamental catch, however: The bottleneck of the Er:Si system is the slow energy transfer between the relevant energy levels in the silicon and the erbium, which leads to long radiative lifetimes, so τ_{spont} is measured in microseconds [116] rather than the nanoseconds we are used to in III–Vs. This is a rather fundamental limitation. Is there anything photonic crystals can do to improve matters?

The answer is an unconvincing “may be”. The emission of Er:Si at 1.54 μm , being based on a transition from the excited, inner 4f shell of the Er^{3+} configuration, is very sharp and therefore fulfils the criterion for a high Q emitter discussed in Section 4.1.2.1. Purcell-type enhancement of the emission is therefore possible, which, in principle, would speed up the radiative transition and make a brighter, more efficient light emitter. It is highly questionable, however, whether the Purcell-effect can be exploited sufficiently to make up for the several orders of magnitude difference in lifetime. We therefore, at present, do not think that the combination of PBG + Er:Si is particularly promising.

4.3.3. *Organic materials, PPV in particular*

The third example describes a rather different material system, i.e. the group of electroluminescent organic materials based around para-phenylene-vinylene (PPV in short). The luminescence spectrum of these materials is relatively broad ($\Delta\lambda = 100\text{--}200\text{ nm}$, Fig. 17) and, depending on the chemical composition, the spectrum can be tuned over most of the visible wavelength range [63]. The attraction with these materials lies in the fact that they can be spin-coated onto any suitable carrier (e.g. glass), so they are cheap and very versatile. The main area for applications




Fig. 17. Photoluminescence spectra from a film of the light-emitting polymer MEH-PPV. Dotted line: on a planar substrate. Plain line: on a corrugated substrate, showing additional peaks due to Bragg-scattering of waveguide modes. Courtesy of I.D.W. Samuel, University of Durham.

so far has been in broad area LEDs; using some form of microstructuring to enhance the efficiency (see Section 4.1.1) therefore appears as an obvious recipe for improvement. The question that then arises concerns the fabrication, i.e. the microstructuring of the polymer, because PPV cannot be easily patterned using standard techniques. Two of the fabrication methods demonstrated so far are to spin PPV onto a prepatterned substrate [117,118] or to use high-pressure embossing [65], the latter being particularly promising for achieving the 100 nm-type resolution required for photonic crystals in the visible.

What can we hope to achieve? (a) Purcell-type enhancement is not an option because of the broad emission bandwidth, whereas (b) enhanced extraction seems a promising mechanism to employ. As discussed in Section 4.1, the external efficiency of an LED scales as $1/n^2$. With a refractive index of around 2, PPV has some scope for photonic crystal based light extraction enhancement, although not as much as the high-index crystalline semiconductors. Placing PPV into a microcavity or onto a corrugated surface (Fig. 17) is a possible realisation. Due to Bragg scattering, the emission is spectrally redistributed, i.e. some wavelength range within the gain spectrum is enhanced whilst others are suppressed [118]. This would be one way of obtaining the spectral purity required for true-colour displays. (c) Another option would be to embed PPV into a 3-D photonic crystal. The self-organised materials, such as opal or titania, (Section 3.1.6) offer themselves, but it is not yet clear whether the low refractive index contrast between these materials and PPV can be engineered sufficiently¹¹ to obtain strong enough effects that justify the required effort.

¹¹ The refractive index of the silica spheres that form opal is 1.45, that of titania is around 2.5 and PPV varies in the range of approximately 1.7–2.2, depending on its chemical nature. An index contrast of about 2 : 1 is required for a photonic band gap to open up [21], although significant benefits may already be felt for lower index contrasts.

4.4. Photonic VLSI

Since the inception of integrated optics, it has been argued that photonics would follow in the footsteps of the silicon-based electronics revolution and create circuits with a high level of integration. As data rates keep growing, miniaturisation in photonics is becoming an ever more pressing issue, because more channels have to be fitted into a limited space, and the delay times between widely separated active elements are becoming increasingly intolerable. To date, there is little progress on this front, and most photonic devices are still made as discrete elements. At best, integration involves a small number of devices on the same chip, performing a limited range of functions. Why is this so? One reason is the disparate size of the different functional elements and their connections. Typical integrated optical circuits, even ones that are only passive and incorporate waveguides, bends, junctions and couplers extend over millimetres rather than microns in size. Waveguided light does not readily negotiate sharp bends, so transitions need to be gradual and, couplers typically have long interaction lengths with respect to the wavelength of light.

Here, photonic crystals offer a possible solution. Driven by the inventive work of the group at MIT [9,11,119], there is a variety of concepts for low-loss waveguide bends, junctions and couplers. These concepts are based on the recognition that light confined to a waveguide within a photonic crystal can only follow the waveguide, since there are no radiation modes available, i.e. there is nowhere else for it to go to. Sharp bends and large-angle Y-junctions have been computationally “demonstrated” with virtually no excess losses [120,121]. These are impressive results that certainly point in the right direction.

In practise, however, life is not so simple. The MIT group’s work tends to use waveguiding in air as an example, where the single-mode waveguide width is on the order of the lattice constant. Measurements on “air waveguides” in the microwave regime have indeed shown impressively low waveguide and bending losses [122] and thereby confirmed the general principle. These measurements can be misleading, however, because metal plates are generally used for confinement in the third dimension, an option not available in optics. If the metal plates are not present, the transmission drops significantly, i.e. to 35% in one particular measurement [123]. Hence, optical waveguiding in air can only be realised with photonic crystals that are either three-dimensional or that extend infinitely in the third dimension. Neither of these assumptions is easy to fulfil, so the nearest immediate realisation is a “conventional” waveguide structure surrounded by a 2-D lattice. Another reason against “air waveguides” is that they are only desirable for some specialist applications that require high power transmission or “odd” wavelengths where low-loss dielectrics are difficult to come by. In general, waveguides are made out of high-index material because they transport light to and from active photonic devices (e.g. lasers, modulators, detectors) that require light-matter interaction.

The above argument therefore makes a case for dielectric waveguides, so how about dielectric waveguides with photonic crystal boundaries, i.e. photonic crystal waveguides, PCWs? The width of such a waveguide has a single mode cut-off of a few 10s of nm [124] i.e. it becomes undesirably multi-mode if wider. If the waveguide has to be

that narrow, it usually incurs considerable scattering losses, since the holes or pillars that constitute the lattice can never be perfectly smooth or perfectly aligned in size and/or position. Measurements on such waveguides in the optical regime have not yet been performed, to our knowledge, but related data [29,42] suggests that a few unit cells of photonic crystal waveguide already imply a noticeable loss. There is no doubt that this loss can be reduced, but it is not at all obvious that it can be brought down to a level where cascading a series of functional waveguide elements, requiring 100s of unit cells, becomes a practical option.

This argument must also take account of the fact that “classical” waveguides¹² can show curvature losses reduced to “negligible levels” for 3 μm radius curves [51], and exhibit very high scattering losses when narrow [51], due to the high index contrast. This result begs two controversial questions: (a) why do we need photonic crystal waveguides if we can realise tight bends without them, and (b) considering that the scattering loss for the high-index classical waveguides is large, why should photonic crystal waveguides that are an order of magnitude narrower, not suffer from the same loss? Admittedly, PCWs have, by design, no modes to scatter into, but unless the photonic crystal is perfect and three-dimensional, this argument is not absolute. We will have to wait for experimental results before we can assess this point more authoritatively.

Another issue for photonic VLSI is that of filtering, an application often mentioned in the context of photonic crystals. Very sharp filtering characteristics have indeed been observed with waveguide-based high Q cavities [30,32] despite associated losses. Similarly, very compact photonic lattice-based add/drop multiplexers have been proposed [125]. The hope is that this type of device will be able to perform operations such as wavelength division multiplexing on a few square microns of surface area, instead of the square millimetres or centimetres currently required [126].

The question here is whether PBG-based filters will be able to meet their target, since there is not much use for a filter with a good performance if the wavelength does not match. In the simplest case, PBG-type filters operate on the principle of a defect that is intentionally being introduced into the lattice, i.e. an area of the lattice that is typically bigger or smaller by a quarter wavelength, $\lambda/4n$, n representing the material refractive index.¹³ So the size of the defect is $d = \lambda/4n$, and its tolerance is $\Delta d = \Delta\lambda/4n$. For a worst-case scenario, let us consider what this translates to for the densest channel spacing in a wavelength division multiplexing (WDM) system, which is 100 GHz or 0.67 nm at 1.55 μm . The refractive index of the relevant InP material is 3.17 at that wavelength, so a size accuracy of $\Delta d = 0.67 \text{ nm}/(4 \times 3.17) \approx 0.5 \text{ \AA}$, would be required, clearly a formidable task. Considering these tight tolerances, filters based on single

¹² The waveguides referred to here are classical in the sense that guiding relies on total internal reflection, but they are also comparable to photonic crystal waveguides in the sense that they are based on silicon on silica waveguides, one of the high-contrast type of material systems envisaged for photonic crystal waveguides.

¹³ Integer multiples of a half-wavelength, $\lambda/2n$, may be added, but this addition does not influence the argument, since the main impact on the wavelength response come from the first $\lambda/4n$.

defects will not be fabricated in the foreseeable future. The experimental demonstrations of such filters that we have seen already [30,32,42] underline this point, because the target wavelength was missed by 10 nm or more [30,32] and it was shown what a large tuning range can be achieved with a small change in cavity size [42].

This was a worst-case study, however. Many systems will use a wider channel spacing than 100 GHz and thereby relax the fabrication tolerances. Multiple defects can be used to tailor the wavelength response of the filter, which, in a two-dimensionally periodic system, can also be used to create wavelength and angular-selective channels. Moreover, ingenious designs that use cascaded groups of gratings are appearing [127] and have already shown suitability for WDM applications. Also, the wavelength dependence does not necessarily have to arise from a defect, i.e. it is well known that there is strong dispersion at the photonic band edge, leading to very wavelength-sensitive diffraction (“superprism”) effects [50,61]. So, what might seem to be an insurmountable hurdle at first sight can be tackled with intelligent design.

4.5. *Surface plasmons*

Many device designers (e.g. in the area of semiconductor lasers) spend effort on avoiding interaction between the optical field in their devices and metal films, e.g. contacts. This is easily understood, because, at optical frequencies, the absorption losses in a metal can be large. In fact, the reality that there are no perfect metallic reflectors for optical waves is one of the motivations for making photonic crystals, because in many cases they can simply be thought of as providing high-reflectivity mirrors.

On the other hand, metals offer exciting phenomena that can be usefully applied to light emitting devices. Surface plasmons, for example, the quanta of electron oscillations at metal–dielectric interfaces, have been used to explain the “supertransmission” effect observed recently [128], where a transmission of around 4% was measured through a thin metal film perforated with holes of 2% area fill-factor and size well below cut-off (150 nm diameter holes at 1.55 μm wavelength). The best explanation put forward so far is that the incoming light excites the plasmon on one surface, which then couples to the plasmon on the other side of the metal film. This second plasmon subsequently couples to radiation modes, which makes the whole process appear as if the incoming light had been transmitted directly through the metal film [129].

These surface plasmons are a form of surface wave that occurs at dielectric interfaces, and are closely related to waveguide modes. When the metal film is corrugated, the surface plasmon experiences band gaps [130] in the same way as a mode does in a corrugated waveguide. Surface plasmons are usually associated with thin films, but they can also be observed in small metal spheres [131,132]. Recent work at the University of Rochester, USA, where resonances in small metal particles have been studied, has shown that the detection efficiency of thin-film photodectors can be increased by over an order of magnitude (in a limited wavelength range) if the surface is covered with metal islands [133,134]. The explanation that the authors propose is that the incoming radiation couples to the metallic resonance and then to




Fig. 18. Illustration of the interaction between incident plane wave, metal nanoparticle and waveguide mode [133,134].

the waveguide mode of the thin absorbing film. This mechanism increases the interaction length between the incident light and the film, and thereby greatly enhances the detection efficiency [Fig. 18].

In the two examples just discussed, the metallic resonances have been used to transfer light between different media. The question therefore arises whether one could not use these resonances to transfer light from an active semiconductor to the outside world, and thereby address the problem of extraction quoted many times throughout this paper. We believe that there is scope in such an approach, which could be realised by matching the frequency of the plasma resonance to that of the emission, i.e. the frequency corresponding to the band gap of the semiconductor. Also, by positioning the metal spheres in some deterministic manner, e.g. in a configuration resembling a photonic lattice, a certain degree of control over the emission/coupling properties of this structure should be possible.

Using metals in structures to mediate light transfer between different media is only the tip of the iceberg, however; the resonances of metallic nanoparticles can be exploited further via their mutual interaction, an area of research that has hardly been explored and opens up an entirely new chapter in the way periodic structures can be used [134]. Another avenue is that of self-organised assembly, where the interaction between plasmon resonances on the surface of metallic nanoparticles can be used to aid the formation of periodic structures [135].

4.6. Quasiperiodic structures

Just when we think we know everything there is to know about periodic photonic structures, we encounter yet another type of arrangement that is as rich, if not richer, in its physics and in novel concepts as the square and triangular lattices we have been studying so far. These structures cover everything between “diatomic” and completely disordered, i.e. “amorphous” lattices and are discussed in increasing degree of disorder.

4.6.1. Diatomic lattices

Diatomic lattices, i.e. lattices that consist of two types of feature (e.g. differently sized holes etched into the material) have already been evaluated theoretically [136] and show larger band gaps than their “monoatomic” counterparts.

4.6.2. Chirped gratings

The term “chirp”, according to the *Encyclopaedia Britannica*, describes the “characteristic short sharp sound especially of a small bird or insect” and implies that the frequency of the sound varies during the call. In the context of periodic structures, a chirped grating is one whose period varies with length. Chirped gratings have found a variety of uses in integrated optics, e.g. as dispersion-compensating elements in fibre optic communication systems [137]. In the context of photonic crystals, chirping has recently been (re)discovered as a method of extending the width of the bandgap, e.g. in order to give one-dimensional Bragg stacks a more omnidirectional response [138]. The idea is that different locations in the periodic structure respond to different frequencies and so, by stacking the different periodicities, a wider frequency range can be addressed overall. This is an elegant method that has great potential to widen the parameter-range over which photonic bandgaps can be achieved, but it also trades off the frequency response vs. the penetration depth. The increased penetration depth results in reduced spatial confinement, which is critical for some applications such as microcavity-enhanced LEDs (Section 4.1.2) and must therefore be considered.

4.6.3. Quasicrystals

Quasicrystals are interesting hybrids between amorphous, disordered materials and the perfectly regular crystals discussed so far. They are disordered in the sense that they are not based on a single repeating unit cell, and regular in the sense that the sequence of features that makes up the structure follows particular geometric rules. Fractals are one example for this type of arrangement, and structures based on the Fibonacci and Cantor sequence have recently been compared to one-dimensionally periodic layer stacks [139]. The fractal structures displayed similar features to their periodic counterparts, such as bandgaps and confined states, but featured the most confined states near the band-edge rather than the middle of the bandgap, a property that adds variety to the toolbox of, e.g., a filter designer.

Two-dimensional quasicrystals are also being explored in our group, in particular structures based on Penrose tiles. Penrose lattices are based on two diamond-shaped primitives that, following certain rules, perfectly tile the plane [Fig. 19]. The interest in this particular configuration stems from its ten-fold rotational symmetry, a degree of symmetry that could not be achieved with a perfectly periodic structure.¹⁴ Why is symmetry important? The reason for the difficulty of achieving full photonic bandgaps lies in the fact that the light, as it travels through the crystal along different directions, experiences different effective periods. In a square lattice, for example, the

¹⁴The highest degree of rotational symmetry achieved with a periodic structure is six-fold, i.e. in a triangular or hexagonal close-packed lattice.




Fig. 19. A Penrose-tiled lattice and its practical realisation, whereby a hole is etched at every intersection point. The straight lines superimposed onto the structure illustrate the planes that lead to the regular diffraction picture shown in Fig. 20.

period is $\sqrt{2}$ longer for travel in the direction of the diagonal than for travel along the sides. So, by increasing the rotational symmetry from four to six to ten, one could hope that full bandgaps are also more easily achieved, and should be possible over a wider range of parameters or in materials with lower refractive index. Theoretical work on related structures [140] indeed reports the existence of density-of-states minima, pointing to possible bandgaps, and higher-order effects whereby successive DOS minima are spaced by the intrinsically beautiful ratio of the golden mean (approximately 5.8). Experimental work, in the optical regime, is only in its infancy, but we have already demonstrated that the diffraction pattern from a grating based on Penrose tiles (Fig. 20) displays clear peaks, as if it were crystalline, and yet exhibits the ten-fold symmetry that is not possible with conventional regular crystal structures. Another point to mention is the communality between these quasicrystals and the chirped gratings mentioned in Section 4.6.2. Both types of structure are not strictly periodic and achieve their improved frequency or angular response by increasing the penetration depth of light into the structure. Using a different terminology, one can also say that there are more Fourier components in the spectrum, but each component is weaker than it would be in a perfect lattice [141].

4.6.4. Photonic “amorphs”

We cannot complete our journey through the non-crystalline photonic world without mentioning that light localisation has even been observed in totally irregular structures, i.e. a photonic “amorph” rather than a photonic crystal [142]. A powder




Fig. 20. Diffraction from the type of Penrose lattice shown in Fig. 19. The same structure as the one shown in Fig. 19 was etched into glass and examined in normal incidence with a He-Ne laser. The ten diffraction peaks in the outer ring illustrate the ten-fold rotational symmetry of the lattice. This clearly indicates that an irregular lattice with no repeating unit cell can have a higher degree of symmetry than a perfect triangular lattice with its three-fold rotational symmetry. Courtesy of P. Millar, University of Glasgow.

consisting of irregular particles of GaAs was able to localise light in a similar way as a defect confines it in a perfectly periodic structure, which again shows that light confinement does not necessarily require periodicity, but only sufficient index contrast.

All examples in this section show that there is much more to photonic crystals than perfectly ordered structures, and that the response of the structure can be improved by an apparently counter-intuitive reduction in symmetry or regularity. The penalty for the increased wavelength-range or angular response of most structures, however, is the loss of confinement, i.e. reduction in the Fourier components, which means that the penetration depth of light into the structure is larger.

4.7. *Quantum cascade-based light emission*

One of the most exciting new concepts in optoelectronics is that of quantum cascade (QC) based light emission. QC lasers, pioneered by Capasso's group at Lucent Technologies [143], have shown impressive progress and have already demonstrated (pulsed) room temperature operation. QC lasers rely on intraband transitions, i.e. transitions between different states in the same or neighbouring quantum wells and typically a multitude of these transitions is cascaded in order to provide sufficient gain. In most cases, InGaAs/InAlAs multilayers are used, but there has also been rapid progress in the GaAs/AlGaAs system [144,145]. The wavelength range demonstrated so far is between 3 μm and 17 μm .

Why do we mention QC emission in the context of photonic crystals? Firstly, the emission is entirely TM polarised, a consequence of the selection rules of the




Fig. 21. “Quasi-3D” bandgap for the TM oscillation in QC lasers. The radiating dipole is represented by the arrow in the centre of the picture. Its radiation cone is indicated by the circle in the left-half of the picture, illustrating the fact that the emission is mainly into the waveguide plane. The right-half of the picture represents the angular extent of the photonic bandgap, i.e. the range of angles for which light emitted by the dipole experiences a band gap. Most of the radiation experiences the band gap, so the microstructure acts on most of the modes, almost as if it were 3-D.

intraband transition. TM polarised light can be understood as originating from vertically oscillating dipoles, so there is no emission normal to the quantum well plane, and most of the light is emitted into the plane and can therefore be captured by a waveguide. This relaxes the requirements for achieving a photonic bandgap, because even a 2-D structure that does not exhibit a bandgap for the full angle of 4π steradian will control most of the light present in the system. Furthermore, it is easier to achieve a large bandgap for one polarisation than for two, so “quasi 3-D” control over the photon modes can be expected with a microstructure that is only two-dimensionally periodic. This high degree of control over all of the existing modes can, for examples, be used to increase the probability of light being emitted into the lasing mode [Fig. 21].

From a fabrication-point of view, the fact that the wavelength is in the mid-to-far infrared reduces the pressure on the pattern generation, simply because suitable microstructures are bigger and can therefore easily be made by photolithography. Electrically, the unipolar nature of the quantum cascade operation eliminates the possibility of surface recombination, a phenomenon that plagues the GaAs/AlGaAs system in particular. Surface recombination is caused by the diffusion of minority carriers to exposed surfaces, so with only one type of carrier involved in the transition, surface recombination is not an issue.

5. Conclusions

We have reviewed the current range of activity in the field, emphasising semiconductor-based photonic crystals because we feel that they are particularly promising due to their high refractive index and light emitting properties. We have also highlighted some of the material problems that occur with semiconductor-based photonic crystals, such as broad emission bandwidth and surface recombination, which hamper cavity enhancement and compromise luminescence efficiency, respectively. Some of these problems may be resolved with quantum dot active layers, due to their atomic-like emission characteristics and carrier capture capabilities.

On the other hand, the large variety of material systems currently under investigation indicates that epitaxial semiconductors are not the only candidate for successful implementation of photonic bandgap concepts. The diversity of the field is still increasing, and work will only start to focus on a few, more promising approaches, once these have been clearly established. Self-organised materials such as opal, once the issues of homogeneity, back-filling with optically active material and refractive index contrast have been resolved, could well become an important approach. The family of electro-optic polymers (PPV, for example) could also become very important indeed, particularly since the ability to microstructure these materials has now been demonstrated.

As far as light emitting diodes are concerned, i.e. the devices used in the introduction to highlight the promise of photonic crystals, we do not think that the benefit of photonic bandgap concepts will be felt as readily as was initially anticipated. The promise of cavity enhancement has now been substantiated both theoretically and experimentally [91,97], but the effect is not (yet) as big as one would have hoped for. Other solutions, however, are appearing, for example the recently observed emission enhancement from a bare lattice [103]. Many other avenues can be explored, and a clever combination of approaches may finally yield the +50% external (electrical) efficiencies that we would hope to see. This combination may entail selective oxidation, microstructuring or metal nanoparticles, all designed to enhance the desired modes or to suppress, extract, or recycle the modes that would otherwise appear as losses on the balance sheet. The challenge then lies in working out the detail of how to combine the relevant effects for device applications.

On the laser front, photonic crystals have a more difficult task to prove their worth because the field is more mature and devices have already gone a long way. If we reverse this statement, however, we can argue that semiconductor lasers can possibly benefit significantly for the very reason that photonic band gap concepts have not really been applied yet, and all the benefits of “total mode control” that photonic crystals offer have yet to bear fruit. The possible benefits apply to the somewhat controversial argument of lowering the threshold current (see Section 4.2), since not every application involves high bit-rates and long-haul transmission, and the important issue of lowering the noise. It also applies to the more exotic types of laser such as quantum cascade-based devices that are not endowed with the high gain values that we are used to from the III–V direct band gap materials. If every photon counts, because the radiative transition has to fight hard against the competition from non-radiative processes, we should try to ensure that every photon goes the right way. The QC system, because of its unique properties, may well turn out to be a prime system for the demonstration of PBG effects in lasers.

As far as the photonic crystal geometry is concerned, both lasers and LEDs point in the same direction: we need better mode control, i.e. we need to be able to control the propagation of light in all directions, not only in-plane. Omnidirectional mode control would allow us to suppress the undesired leaky and waveguide modes in LEDs and consequently increase the extraction efficiency. In lasers, 3-D control would help to increase the β -factor and hence reduce the threshold current and noise.

The realisation of photonic VLSI, i.e. the combination of a large number of functional elements on a small space, does not appear as appealing now as it might have looked at first sight and there are several issues that need to be addressed, such as tolerances for accurate filter designs and losses for elongated waveguide sections. Dispersive elements that exploit a particularly steep section of the band structure (“superprism”) seem also very promising, but there are issues of beam quality and accurate routing that need to be resolved. As a whole, the area of photonic VLSI will become important, however, because there are real industrial needs related to the ever increasing number of channels in communications systems and the associated problems of chip size and delay times. This drive alone may well be sufficient to attract the amount of brainpower needed to tackle the problems involved.

Overall, the photonic band-gap revolution, which is what was hoped for a few years ago, has matured into a photonic band-gap evolution. New applications emerge constantly, but there seems to be no particular aspect of photonic crystals that yields quick and easily applicable results. There is, however, much to be gained from careful design and detailed insight into the particular problem at hand.

Acknowledgements

The authors would like to thank C.J.M. Smith, C.N. Ironside and S.G. Romanov at Glasgow University, D. Labilloy, H. Benisty and C. Weisbuch at Ecole Polytechnique, Palaiseau, as well as M. Boroditsky and E. Yablonovitch at UCLA for fruitful discussions. We would also like to express our gratitude to the Royal Society and the EPSRC of the UK, and the ESPRIT-programme of the European Union for financial support.

References

- [1] *Harnessing Light: Optical Sciences and Engineering for the 21st Century*, National Academy Press, Washington DC, 1998.
- [2] V.V. Bezotosnyi, P.V. Bulaev, V.A. Gorbylev, I.D. Zalevskii, N.V. Markova, Y.M. Popov, A.A. Padalitsa, *Quantum Electronics* 28 (1998) 292–293.
- [3] R. Jager, M. Grabherr, C. Jung, R. Michalzik, G. Reiner, B. Weigl, K.J. Ebeling, *Electronics Letters* 33 (1997) 330–331.
- [4] D. Botez, L.J. Mawst, A. Bhattacharya, J. Lopez, J. Li, T.F. Kuech, V.P. Iakovlev, G.I. Suruceanu, A. Caliman, A.V. Syrbu, *Electronics Letters* 32 (1996) 2012–2013.
- [5] E.M. Purcell, *Physical Review* 69 (1946) 681.
- [6] S. John, *Physical Review Letters* 58 (1987) 2486–2489.
- [7] E. Yablonovitch, *Physical Review Letters* 58 (1987) 2059–2062.
- [8] E. Yablonovitch, *Journal of the Optical Society of America B-Optical Physics* 10 (1993) 283–295.
- [9] P.R. Villeneuve, M. Piche, *Progress in Quantum Electronics* 18 (1994) 153–200.
- [10] T. Baba, *IEEE Journal of Selected Topics in Quantum Electronics* 3 (1997) 808–830.
- [11] J.D. Joannopoulos, R.D. Meade, J.N. Winn, *Photonic Crystals-Molding the Flow of Light*, Princeton University Press, Princeton, NJ, 1995.
- [12] J. Maddox, *Nature* 348 (1990) 481.
- [13] E. Yablonovitch, K.M. Leung, *Nature* 351 (1991) 278.

- [14] E. Yablonovitch, T.J. Gmitter, K.M. Leung, *Physical Review Letters* 67 (1991) 2295–2298.
- [15] K.M. Ho, C.T. Chan, C.M. Soukoulis, *Physical Review Letters* 65 (1990) 3152–3155.
- [16] E. Yablonovitch, T.J. Gmitter, R.D. Meade, A.M. Rappe, K.D. Brommer, J.D. Joannopoulos, *Physical Review Letters* 67 (1991) 3380–3383.
- [17] E.R. Brown, C.D. Parker, E. Yablonovitch, *Journal of the Optical Society of America B-Optical Physics* 10 (1993) 404–407.
- [18] J.G. Fleming, S.Y. Lin, *Optics Letters* 24 (1) (1999) 49–51.
- [19] C.C. Cheng, A. Scherer, *Journal of Vacuum Science & Technology B* 13 (1995) 2696–2700.
- [20] C.C. Cheng, A. Scherer, R.C. Tyan, Y. Fainman, G. Witzgall, E. Yablonovitch, *Journal of Vacuum Science & Technology B* 15 (1997) 2764–2767.
- [21] R.D. Meade, K.D. Brommer, A.M. Rappe, J.D. Joannopoulos, *Applied Physics Letters* 61 (1992) 495–497.
- [22] D. Cassagne, C. Jouanin, D. Bertho, *Physical Review B-Condensed Matter* 53 (1996) 7134–7142.
- [23] U. Grüning, V. Lehmann, C.M. Engelhardt, *Applied Physics Letters* 66 (1995) 3254–3256.
- [24] H.B. Lin, R.J. Tonucci, A.J. Campillo, *Applied Physics Letters* 68 (1996) 2927–2929.
- [25] S. Rowson, A. Chelnokov, C. Cuisin, J.-M. Lourtioz, *Proceedings IEE (Pt. J., Optoelectronics)* 145 (1998) 403–408.
- [26] U. Grüning, V. Lehmann, S. Ottow, K. Busch, *Applied Physics Letters* 68 (1996) 747–749.
- [27] H.W. Lau, G.J. Parker, R. Greef, M. Hölling, *Applied Physics Letters* 67 (1995) 1877–1879.
- [28] T.F. Krauss, R.M. DeLaRue, S. Brand, *Nature* 383 (1996) 699–702.
- [29] D. Labilloy, H. Benisty, C. Weisbuch, T.F. Krauss, R.M. DeLaRue, V. Bardinal, R. Houdre, U. Oesterle, D. Cassagne, C. Jouanin, *Physical Review Letters* 79 (1997) 4147–4150.
- [30] T.F. Krauss, B. Voge, C.R. Stanley, R.M. DeLaRue, *IEEE Photonics Technology Letters* 9 (1997) 176–178.
- [31] D.N. Atkin, P.S.J. Russell, T.A. Birks, P.J. Roberts, *Journal of Modern Optics* 43 (1996) 1035–1053.
- [32] J.S. Foresi, P.R. Villeneuve, J. Ferrera, E.R. Thoen, G. Steinmeyer, S. Fan, J.D. Joannopoulos, L.C. Kimerling, H.I. Smith, E.P. Ippen, *Nature* 390 (1997) 143–145.
- [33] M.D.B. Charlton, S.W. Roberts, G.J. Parker, *Materials Science and Engineering B-Solid State Materials For Advanced Technology* 49 (1997) 155–165.
- [34] J.M. Gerard, A. Izrael, J.Y. Marzin, R. Padjen, F.R. Ladan, *Solid-State Electronics* 37 (1994) 1341–1344.
- [35] T. Krauss, Y.P. Song, S. Thoms, C.D.W. Wilkinson, R.M. Delarue, *Electronics Letters* 30 (1994) 1444–1446.
- [36] V. Berger, O. GauthierLafaye, E. Costard, *Electronics Letters* 33 (1997) 425–426.
- [37] A. Yi-Yan, C.D.W. Wilkinson, P.J.R. Laybourn, *IEEE Journal of Quantum Electronics* QE-16 (1980) 1089–1092.
- [38] R. Zengerle, *Journal of Modern Optics* 34 (1987) 1589–1617.
- [39] K.D. Choquette, K.M. Geib, C.L.H. Ashby, R.D. Twisten, O. Blum, H.Q. Hou, D.M. Follstaedt, B.E. Hammons, D. Mathes, R. Hull, *IEEE Journal of Selected Topics in Quantum Electronics* 3 (1997) 916–926.
- [40] J.M. Dallesasse, N. Holonyak, A.R. Sugg, T.A. Richard, N. Elzein, *Applied Physics Letters* 57 (1990) 2844–2846.
- [41] J. O'Brien, O. Painter, R. Lee, C.C. Cheng, A. Yariv, A. Scherer, *Electronics Letters* 32 (1996) 2243–2244.
- [42] C.J.M. Smith, T.F. Krauss, R.M. De La Rue, D. Labilloy, H. Benisty, C. Weisbuch, U. Oesterle, R. Houdre, *Proceedings IEE (Pt. J, Optoelectronics)* 145 (1998) 373–378.
- [43] T. Baba, N. Kamizawa, M. Ikeda, *Physica B* 227 (1996) 415–418.
- [44] C. Youtsey, R. Grundbacher, R. Panepucci, I. Adesida, C. Caneau, *Journal of Vacuum Science & Technology B* 12 (1994) 3317–3321.
- [45] V. Lehmann, H. Foll, *Journal of the Electrochemical Society* 137 (1990) 653–659.
- [46] H. Masuda, H. Yamada, M. Satoh, H. Asoh, M. Nakao, T. Tamamura, *Applied Physics Letters* 71 (1997) 2770–2772.
- [47] K.D. Choquette, H.Q. Hou, *Proceedings of the IEEE* 85 (1997) 1730–1739.

- [48] M.H. Macdougall, S.G. Hummel, P.D. Dapkus, H.M. Zhao, Y. Cheng, *IEEE Photonics Technology Letters* 7 (1995) 385–387.
- [49] E.L. Hu, C.H. Chen, *Microelectronic Engineering* 35 (1997) 23–28.
- [50] P.S.J. Russel, *Physical Review A* 33 (1986) 3232–3242.
- [51] B.E. Little, J.S. Foresi, G. Steinmeyer, E.R. Thoen, S.T. Chu, H.A. Haus, E.P. Ippen, L.C. Kimerling, W. Greene, *IEEE Photonics Technology Letters* 10 (1998) 549–551.
- [52] A. Scherer, O. Painter, B. D'Urso, R. Lee, A. Yariv, *Journal of Vacuum Science Technology B* 16 (1998) 3906–3910.
- [53] O. Painter, J. Vuckovic, A. Scherer, *Journal of the Optical Society of America B* 16 (1999) 275–285.
- [54] M. Kanskar, P. Paddon, V. Pacradouni, R. Morin, A. Busch, J.F. Young, S.R. Johnson, J. MacKenzie, T. Tiedje, *Applied Physics Letters* 70 (1997) 1438–1440.
- [55] E. Yablonovitch, T. Gmitter, J.P. Harbison, R. Bhat, *Applied Physics Letters* 51 (1987) 2222–2224.
- [56] A. vanGeelen, P.R. Hageman, G.J. Bauhuis, P.C. vanRijsingen, P. Schmidt, L.J. Giling, *Materials Science and Engineering B-Solid State Materials For Advanced Technology* 45 (1997) 162–171.
- [57] A. Ersen, I. Schnitzer, E. Yablonovitch, T. Gmitter, *Solid-State Electronics* 36 (1993) 1731–1739.
- [58] I. Schnitzer, E. Yablonovitch, C. Caneau, T.J. Gmitter, A. Scherer, *Applied Physics Letters* 63 (1993) 2174–2176.
- [59] G. Witzgall, R. Vrijen, E. Yablonovitch, V. Doan, B.J. Schwartz, *Optics Letters* 23 (1998) 1745–1747.
- [60] S. Kawakami, *Electronics Letters* 33 (1997) 1260–1261.
- [61] H. Kosaka, T. Kawashima, A. Tomita, M. Notomi, T. Tamamura, T. Sato, S. Kawakami, *Physical Review B* 58 (1998) 10096–10099.
- [62] S. Kawakami, T. Kawashima, T. Sato, *Applied Physics Letters* 74 (1999) 463–465.
- [63] D. Bradley, R. Friend, A. Holmes, *Philosophical Transactions of the Royal Society of London Series A-Mathematical Physical and Engineering Sciences* 355 (1997) 693.
- [64] S. Brittain, K. Paul, X.-M. Zhao, G. Whitesides, *Physics World* 11 (1998) 31–36.
- [65] J.A. Rogers, Z.N. Bao, L. Dhar, *Applied Physics Letters* 73 (1998) 294–296.
- [66] R.J. Tonucci, B.L. Justus, A.J. Campillo, C.E. Ford, *Science* 258 (1992) 783–785.
- [67] J.C. Knight, J. Broeng, T.A. Birks, P.S.J. Russel, *Science* 282 (1998) 1476–1478.
- [68] N.M. Lawandy, R.M. Balachandran, A.S.L. Gomes, E. Sauvain, *Nature* 368 (1994) 436–438.
- [69] B.T. Holland, C.F. Blanford, A. Stein, *Science* 281 (1998) 538–540.
- [70] J. Wijnhoven, W.L. Vos, *Science* 281 (1998) 802–804.
- [71] V.N. Bogomolov, S.V. Gaponenko, I.N. Germanenko, A.M. Kapitonov, E.P. Petrov, N.V. Gaponenko, A.V. Prokoviev, A.M. Ponyavina, N.I. Silvanovich, S.M. Samoilovich, *Physical Review E* 55 (1997) 7619–7625.
- [72] W. Stober, A. Fink, E. Bohn, *Journal of Colloid and Interface Science* 26 (1968) 62ff.
- [73] G. Kolbe, Jena, 1956.
- [74] S.G. Romanov, N.P. Johnson, A.V. Fokin, V.Y. Butko, H.M. Yates, M.E. Pemble, C.M.S. Torres, *Applied Physics Letters* 70 (1997) 2091–2093.
- [75] S.G. Romanov, A.V. Fokin, V.I. Alperovich, N.P. Johnson, R.M. DeLaRue, *Physica Status Solidi A-Applied Research* 164 (1997) 169–173.
- [76] B. Temelkuran, E. Ozbay, J.P. Kavanaugh, G. Tuttle, K.M. Ho, *Applied Physics Letters* 72 (1998) 2376–2378.
- [77] D. Labilloy, H. Benisty, C. Weisbuch, T.F. Krauss, R. Houdre, U. Oesterle, *Applied Physics Letters* 71 (1997) 738–740.
- [78] D. Labilloy, H. Benisty, C. Weisbuch, C.J.M. Smith, T.F. Krauss, R. Houdre, U. Oesterle, *Physical Review B* 59 (1999) 1649–1652.
- [79] C.J.M. Smith, H. Benisty, D. Labilloy, U. Oesterle, R. Houdre, T.F. Krauss, R.M. DeLaRue, C. Weisbuch, *Electronics Letters* 35 (1999) 228–230.
- [80] D. Labilloy, H. Benisty, C. Weisbuch, C.J.M. Smith, T.F. Krauss, R. Houdre, U. Oesterle, *Physical Review B*(1999).
- [81] R. Ulrich, R.T. Torge, *Applied Optics* 12 (1973) 2901–2908.
- [82] V.N. Astratov, M. Stevenson, M. Skolnick, D.M. Whittaker, S. Brand, I. Culshaw, T.F. Krauss, R.M. De La Rue, O.Z. Karimov, *Proceedings IEE (Pt.J, Optoelectronics)* 145 (1998) 398–402.

- [83] I. Schnitzer, E. Yablonovitch, C. Caneau, T.J. Gmitter, *Applied Physics Letters* 62 (1993) 131–133.
- [84] M.W. Hodapp, in: G.B. Stringfellow, M.G. Craford (Eds.), *High Brightness Light Emitting Diodes*, vol. 48, Academic Press, San Diego, 1997.
- [85] H. Benisty, H. DeNeve, C. Weisbuch, *IEEE Journal of Quantum Electronics* 34 (1998) 1612–1631.
- [86] H. DeNeve, J. Blondelle, P. VanDaele, P. Demeester, R. Baets, G. Borghs, *Applied Physics Letters* 70 (1997) 799–801.
- [87] D. Delbeke, R. Baets, B. Dhoedt, R. Bockstaele, I. Moerman, P. Van Daele, T.F. Krauss, in *IEEE/LEOS Summer Topical Meeting*, San Diego, CA, 1999.
- [88] T. Baba, R. Watanabe, K. Asano, F. Koyama, K. Iga, *Japanese Journal of Applied Physics* (Pt. 1) 35 (1996) 97–100.
- [89] H. Yokoyama, in: *Spontaneous Emission and Laser Oscillation in Microcavities*, CRC Press, Boca Raton, FL, 1995, pp. 275–310.
- [90] R. Loudon, *The Quantum Theory of Light*, 2nd ed., Oxford, New York, 1983.
- [91] J.M. Gerard, B. Sermage, B. Gayral, B. Legrand, E. Costard, V. ThierryMieg, *Physical Review Letters* 81 (1998) 1110–1113.
- [92] J.M. Gerard, D. Barrier, J.Y. Marzin, R. Kuszelewicz, L. Manin, E. Costard, V. ThierryMieg, T. Rivera, *Applied Physics Letters* 69 (1996) 449–451.
- [93] H. Yokoyama, K. Nishi, T. Anan, Y. Nambu, S.D. Brorson, E.P. Ippen, M. Suzuki, *Optical and Quantum Electronics* 24 (1992) S245–S272.
- [94] L. Goldstein, F. Glas, J.Y. Marzin, M.N. Charasse, G. Leroux, *Applied Physics Letters* 47 (1985) 1099–1101.
- [95] D.L. Huffaker, G. Park, Z. Zou, O.B. Shchekin, D.G. Deppe, *Applied Physics Letters* 73 (1998) 2564–2566.
- [96] D. Bimberg, N.N. Ledentsov, M. Grundmann, F. Heinrichsdorff, V.M. Ustinov, P.S. Kopev, M.V. Maximov, Z.I. Alferov, J.A. Lott, *Physica E* 3 (1998) 129–136.
- [97] R. Coccioli, M. Boroditsky, K.W. Kim, Y. Rahmat-Samii, E. Yablonovitch, *Proceedings IEE* (Pt. J, Optoelectronics) 145 (1998) 391–397.
- [98] E. Yablonovitch, H.M. Cox, T.J. Gmitter, *Applied Physics Letters* 52 (1988) 1002–1004.
- [99] J. Temmyo, R. Notzel, T. Tamamura, *Applied Surface Science* 121 (1997) 63–72.
- [100] E. Kuramochi, J. Temmyo, H. Kamada, T. Tamamura, *Japanese Journal of Applied Physics* (Pt. 1) 37 (1998) 1559–1564.
- [101] J.M. Gerard, O. Cabrol, B. Sermage, *Applied Physics Letters* 68 (1996) 3123–3125.
- [102] P. Rigby, T.F. Krauss, *Nature* 390 (1997) 125.
- [103] M. Boroditsky, R. Vrijen, T.F. Krauss, R. Coccioli, R. Bhat, E. Yablonovitch, *Proceedings SPIE* 3283 (1999) 184–190.
- [104] C.L. Chua, R.L. Thornton, D.W. Treat, M. Kneissl, C. Dunnrowicz, *Applied Physics Letters* 72 (1998) 1001–1003.
- [105] L.A. Coldren, S.W. Corzine, *Diode Lasers and Photonic Integrated Circuits*, Wiley, New York, 1995.
- [106] J. Rarity, *Physics World* 9 (1996) 19–20.
- [107] D.L. Huffaker, D.G. Deppe, *Applied Physics Letters* 70 (1997) 1781–1783.
- [108] T. Baba, M. Fujita, A. Sakai, M. Kihara, R. Watanabe, *IEEE Photonics Technology Letters* 9 (1997) 878–880.
- [109] J.P. Zhang, D.Y. Chu, S.L. Wu, S.T. Ho, W.G. Bi, C.W. Tu, R.C. Tiberio, *Physical Review Letters* 75 (1995) 2678–2681.
- [110] D. Labilloy, H. Benisty, C. Weisbuch, T.F. Krauss, C.J.M. Smith, R. Houdre, U. Oesterle, *Applied Physics Letters* 73 (1998) 1314–1316.
- [111] T. Baba, M. Hamasaki, N. Watanabe, P. Kaewplung, A. Matsutani, T. Mukaiyama, F. Koyama, K. Iga, *Japanese Journal of Applied Physics* (Pt. 1) 35 (1996) 1390–1394.
- [112] T.F. Krauss, O. Painter, A. Scherer, J.S. Roberts, R.M. DeLaRue, *Optical Engineering* 37 (1998) 1143–1148.
- [113] T.F. Krauss, A. Scherer, J.S. Roberts, R.M. De La Rue, in *Integrated Photonics Research*, Victoria, B.C., Canada, 1998, pp. Paper IWB3.
- [114] F. Buda, J. Kohanoff, *Progress in Quantum Electronics* 18 (1994) 201–226.

- [115] A.G. Cullis, L.T. Canham, P.D.J. Calcott, *Journal of Applied Physics* 82 (1997) 909–965.
- [116] L.C. Kimerling, K.D. Kolenbrander, J. Michel, J. Palm, *Solid State Physics-Advances in Research and Applications* 50 (1997) 333–381.
- [117] A. Schulzgen, C. Spiegelberg, M.M. Morrell, S.B. Mendes, P.M. Allemand, Y. Kawabe, M. KuwataGonokami, S. Honkanen, M. Fallahi, B. Kippelen, N. Peyghambarian, *Optical Engineering* 37 (1998) 1149–1156.
- [118] I.D.W. Samuel, B.J. Matterson, M.G. Salt, W.L. Barnes, in *CLEO/Europe - EQEC*, Glasgow, Scotland, 1998, pp. Paper QWF6.
- [119] J.D. Joannopoulos, P.R. Villeneuve, S.H. Fan, *Nature* 386 (1997) 143–149.
- [120] A. Mekis, S.H. Fan, J.D. Joannopoulos, *Physical Review B-Condensed Matter* 58 (1998) 4809–4817.
- [121] A. Mekis, J.C. Chen, I. Kurland, S.H. Fan, P.R. Villeneuve, J.N.-M.D.P.C.M. Joannopoulos, *Physical Review Letters* 77 (1996) 3787–3790.
- [122] S.Y. Lin, E. Chow, V. Hietala, P.R. Villeneuve, J.D. Joannopoulos, *Science* 282 (1998) 274–276.
- [123] B. Temelkuran, E. Özbay, *Applied Physics Letters* 74 (1999).
- [124] H. Benisty, *Journal of Applied Physics* 79 (1996) 7483–7492.
- [125] S.H. Fan, P.R. Villeneuve, J.D. Joannopoulos, H.A. Haus, *Optics Express* 3 (1998) 4–11.
- [126] M.K. Smit, C. van Dam, *IEEE Journal of Selected Topics in Quantum Electronics* 2 (1996) 236–250.
- [127] C.F. Lam, R.B. Vrijen, P.P.L. ChangChien, D.F. Sievenpiper, E. Yablonovitch, *Journal of Lightwave Technology* 16 (1998) 1657–1662.
- [128] T.W. Ebbesen, H.J. Lezec, H.F. Ghaemi, T. Thio, P.A. Wolff, *Nature* 391 (1998) 667–669.
- [129] P.R. Villeneuve, *Physics World* 11 (1998) 28–29.
- [130] S.C. Kitson, W.L. Barnes, J.R. Sambles, *Physical Review Letters* 77 (1996) 2670–2673.
- [131] M. Born, E. Wolf, *Principles of Optics*, 6th ed., Cambridge University Press, Cambridge, 1980.
- [132] T. Klar, M. Perner, S. Grosse, G. von Plessen, W. Spirkel, J. Feldmann, *Physical Review Letters* 80 (1998) 4249–4252.
- [133] H.R. Stuart, D.G. Hall, *Applied Physics Letters* 69 (1996) 2327–2329.
- [134] H.R. Stuart, D.G. Hall, *Physical Review Letters* 80 (1998) 5663–5666.
- [135] M.I. Antonoyiannakis, J.B. Pendry, *Europhysics Letters* 40 (1997) 613–618.
- [136] C.M. Anderson, K.P. Giapis, *Physical Review Letters* 77 (1996) 2949–2952.
- [137] Y. Liu, L. Zhang, I. Bennion, *Measurement Science & Technology* 10 (1999) L1–L3.
- [138] E. Yablonovitch, *Optics Letters* 23 (1998) 1648–1649.
- [139] C. Sibilía, I.S. Nefedov, M. Scalora, M. Bertolotti, *Journal of the Optical Society of America B-Optical Physics* 15 (1998) 1947–1952.
- [140] Y.S. Chan, C.T. Chan, Z.Y. Liu, *Physical Review Letters* 80 (1998) 956–959.
- [141] M. Kaliteevski, S. Brand, R.A. Abram, T.F. Krauss, R. De La Rue, P. Millar, (1999) in press.
- [142] D.S. Wiersma, P. Bartolini, A. Lagendijk, R. Righini, *Nature* 390 (1997) 671–673.
- [143] C. Sirtori, J. Faist, F. Capasso, A.Y. Cho, *Pure and Applied Optics* 7 (1998) 373–381.
- [144] Y.B. Li, J.W. Cockburn, M.S. Skolnick, J.P. Duck, M.J. Birkett, I.A. Larkin, R. Grey, G. Hill, M. Hopkinson, *Applied Physics Letters* 72 (1998) 2141–2143.
- [145] C. Sirtori, P. Kruck, S. Barbieri, P. Collot, J. Nagle, M. Beck, J. Faist, *Applied Physics Letters* 73 (1998) 3486–3488.