Experimental Verification of Numerically Optimized Photonic Crystal Injector, Y-Splitter, and Bend

Melanie Ayre, Tim J. Karle, *Member, IEEE*, Lijun Wu, Tom Davies, and Thomas F. Krauss

Abstract—We present the experimental measurement of a photonic crystal (PhC) device comprising an injector, Y-splitter, and 60° bend. The complete device consists of a 9-μm-long injector tapering down from 5 μm into a triangular-lattice-of-holes single-line defect waveguide with period $a = 430$ nm and 36.2% air filling factor (corresponding to a radius over period (r/a) ratio of 0.30), an optimized Y-junction, 60° bend and output injectors, with a total device footprint of 30 μm. This is etched into a GaAs/AlGaAs heterostructure using chlorine/argon chemically assisted ion beam etching (CAIBE). An erbium-doped fiber amplifier (EDFA)-based source and Fabry–Perot technique are used to characterize the device. The device displays a bandwidth of approximately 110 nm in the 1.55 μm window, and a transmission of 70% relative to the same length of 5-μm-wide waveguide. This is compared with three-dimensional finite-difference time-domain (3-D FDTD) results, which have a bandwidth and transmission of 120 nm and 75%, respectively. The highlight of this paper is the close agreement of the numerically optimized complete microcircuit with its experimental equivalent, and the significant improvement in bandwidth over previous work on Y-junctions.

Index Terms—Chemically assisted ion beam etching (CAIBE), Fabry–Perot resonance, input coupling, numerical optimization, photonic crystals (PhCs).

I. INTRODUCTION

TWO-DIMENSIONAL (2-D) photonic crystals have the potential to provide much-needed miniaturization in integrated optics. By etching holes into a III–V semiconductor waveguide heterostructure, optical confinement is provided via refractive index in the vertical direction and via the photonic bandgap phenomenon in the plane of the device. The single-line-defect (W1) geometry of photonic crystal (PhC) waveguide is particularly favored for many applications, as it can be designed to be monomode in the frequency region of interest. A wide range of measurements on many types of waveguide on this scale have previously been presented—[1]–[3] are a small selection. For effective routing and for building up

more complex circuits, however, other devices than straight waveguides are required. Bends and Y-splitters are of particular importance here.

The first generation of such devices represented a simple approach, namely, deleting holes in the PhC lattice in the desired forms. Unfortunately, such devices demonstrate both poor transmission, which improves with fabrication technology and the design of input couplers, and also very low bandwidth, which is, perhaps, intrinsic to the design [4], [16]. By considering the field patterns in such devices via both the finite-difference time-domain (FDTD) and eigenmode expansion (EME) modeling techniques, a number of methodologies for improving this bandwidth have been proposed. To oversimplify somewhat, bends can be improved via the use of a mirror plane at the point of the bend, and Y-splitters by the addition of holes to reduce the area of the junction cavity. Both simulated and experimental results for a variety of such devices, designed by hand, have been reported [5], [6], [16]. The most recent of these, [16], displays very limited bandwidth, in the region of 30 nm, which clearly demonstrates the need for a more sophisticated numerical approach. In PhC circuits, device transmission depends very strongly on the substrate form, and in particular, it is challenging to reliably determine the excess loss per component in Y-splitter devices. For the heterostructure-type substrate studied in this work, the best total performance of a Y-splitter at the time of writing is a total transmission of 80% with respect to a comparable W1 waveguide [6]. Please note that in this paper, we report a similar value for the transmission of an entire microcircuit, consisting of Y-junction, bends and injectors. This is even more noteworthy as we normalize to a 5-μm-wide ridge waveguide rather than the commonly used comparison against a W1.

Additionally, the matter of access to the submicron scale of PhC devices must be considered. The best method for connecting the access waveguide to the PhC is an active area of research, the main issue being the mismatch at the waveguide-PhC interface and the resulting loss due to reflections and scattering. The standard techniques employ an adiabatic taper which, in order to work effectively, requires a length of the order of one hundred microns or more. This length is in stark contrast to the miniature scale of PhC devices. Hence, there is considerable motivation to reduce the access length to PhC devices [7]–[10].

Recently, there has been a lot of interest in the PhC community in the benefits of numerical optimization techniques to further improve the transmission and bandwidth of such injectors, bends, and Y-splitters. Here, we present the experimental results of a complete optimized PhC device comprising all these elements.

M. Ayre and T. J. Karle are with the Microphotonics and Photonic Crystals Research Group, School of Physics and Astronomy, University of St. Andrews, St. Andrews, Fife KY16 9SS, U.K. (e-mail: ma37@st-and.ac.uk; tk19@st-and.ac.uk).

L. Wu is with the Institute of Nano Science and Technology, The Hong Kong University of Science and Technology, Hong Kong, China (e-mail: phljwu@ust.hk).

T. Davies is with Photon Design, Oxford OX4 1TW, U.K. (e-mail: tdavies@photond.com).

T. F. Krauss is with the Microphotonics and Photonic Crystals Research Group, School of Physics and Astronomy, University of St. Andrews, St. Andrews, Fife KY16 9SS, U.K. (e-mail: tfk@st-and.ac.uk).

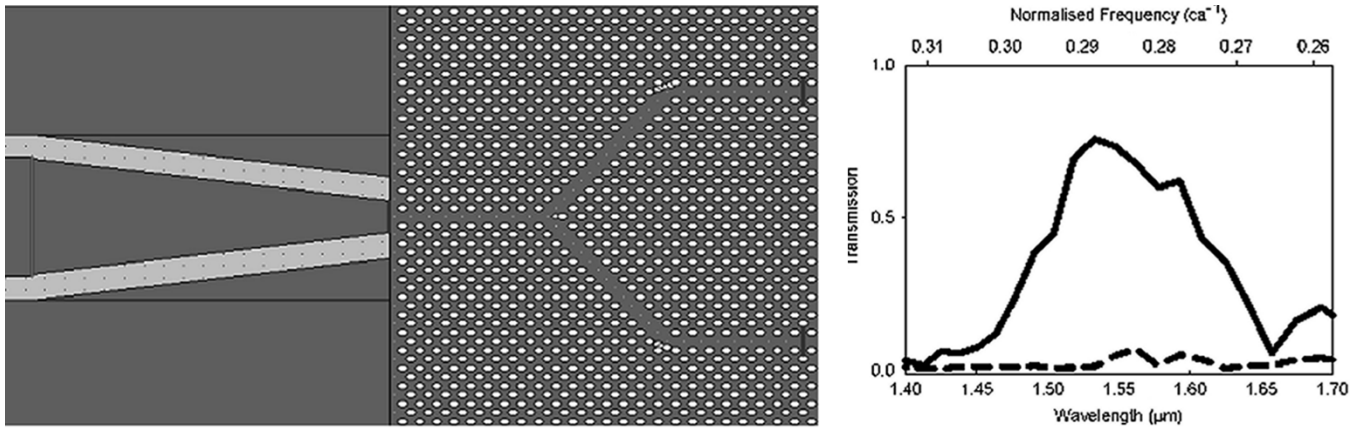Digital Object Identifier 10.1109/JSAC.2005.851169

Fig. 1. Simulated structure and its transmission spectrum (solid line). This is fully 3-D FDTD in Crystal Wave [11]. The peak transmission is 76% over both outputs as compared with the fundamental mode of the 5 μm input waveguide. The 3 dB bandwidth is approximately 120 nm. The dotted line shows the same simulation for an optimized injector but simple PhC Y-splitter. The peak transmission here is ~10%, and the bandwidth ~30 nm, which highlights the striking improvement afforded by the optimization.

II. DEVICE DESIGN AND OPTIMIZATION

The optimization routine takes full account of the material heterostructure used in our experiments. A single-mode AlGaAs heterostructure (250 nm $Al_{0.25}Ga_{0.75}As$/488 nm $Al_{0.20}Ga_{0.80}As$ / 55 nm $Al_{0.30}Ga_{0.70}As$ / 2000 nm $Al_{0.60}Ga_{0.40}As$) that has previously shown good experimental performance was used in our experiments. A global algorithm in combination with 2-D FDTD and 2-D EME is used to optimize an injector from a 5-μm access guide to the PhC W1 waveguide, using a maximum length of 10 μm as the major design constraint. Similarly, a 60° bend is optimized by altering a mirror plane to maximize transmission; and the Y-splitter via adjusting the size and position of a tuning hole. The details of the design process have previously been published [11], [15], which describes a Y-splitter on a membrane of refractive index 2.5. The design process was then repeated for our wafer and a 1.55 μm operating point. The optimized design has a transmission of 97%. 3-D FDTD was used to verify the suitability of the designs, to ensure that 3-D effects such as out-of-plane diffraction do not dominate the response. In each case, the 3-D FDTD gives a peak transmission of 76% for the device. The 3 dB bandwidth is approximately 120 nm. This simulation lacks an output taper due to memory limitations, so a slightly higher transmission in the simulation than the experiment can be expected—no interface-induced losses will be induced at the output. As shown by Felici in [15], the taper has a bandwidth of the order of 300 nm, so the output taper is not expected to disturb this. The simulated structure and spectrum are shown in Fig. 1. To demonstrate the benefit of such optimization, the result of a 3-D FDTD simulation for a Y-splitter with optimized injector but "simple" PhC is also shown for comparison.

III. FABRICATION

The fully optimized device has been fabricated into the aforementioned AlGaAs/GaAs heterostructure, grown by metal- organic chemical vapor-phase deposition (MOCVD) at the National Centre for III-V Semiconductors, Sheffield University. A 300 nm silicon oxide is then deposited over the wafer using plasma enhanced chemical vapor-phase deposition
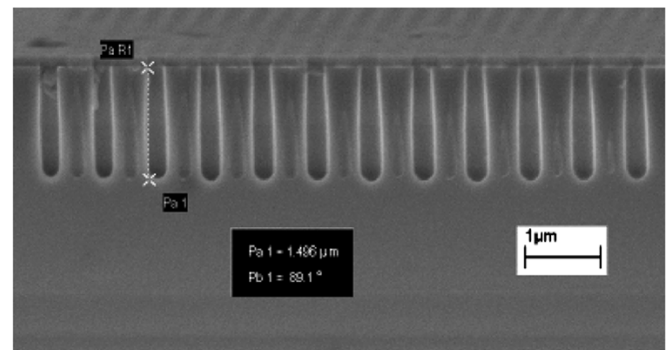

Fig. 2. Optimized etching of a test sample. The depth, sidewall verticality, and smoothness of hole bottoms is particularly important.

(PECVD), followed by a 200-nm-thick layer of PMMA as electron-beam resist. The resist was patterned using a Leica EBPG-5 Beamwriter at the Nanoelectronics Research Centre, University of Glasgow. The pattern was then transferred into the oxide hard mask via reactive ion etching (RIE) using a fluorine process.

The deep etching of the PhC holes is achieved using chemically assisted ion beam etching (CAIBE) in a chlorine/argon process. Sufficiently deep etching is achieved using a balanced low current—high voltage regime, which has been shown to readily produce good sidewall verticality and low roughness [12]. This is carefully optimized to give the required quality of etching. This is necessary to maximize agreement with the simulated design. After deep etching, the sample is returned to the RIE to remove any oxide mask that may remain, as this can provide an additional source of scattering loss. The etch profile achieved in a test sample on identical material to the device is shown in Fig. 2.

As shown in Fig. 3, devices have been fabricated, with a normalized hole radius of ~27%. Fabrication close to design (30.3%) is particularly important for such an optimized device, because the usual lithographic tuning to access the desired frequency range cannot be straightforwardly extended to the injectors, as they do not simply scale with the lattice constant. However, the injector is expected to have a very large
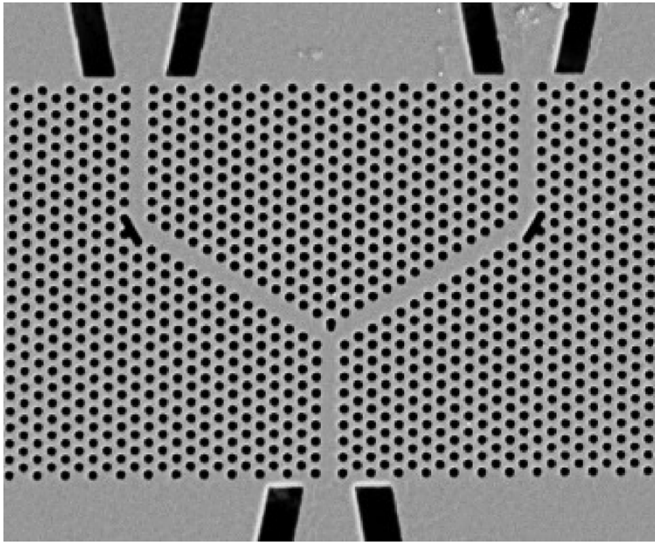
Fig. 3.   Final device as fabricated. Only the PhC end of the injector is shown.

bandwidth, so this slight deviation in hole size is acceptable. Complete devices including injectors scaled by plus and minus five percent have also been fabricated, to allow some broad-spectrum characterization.

IV. CHARACTERIZATION

To characterize the Y-splitter device, an EDFA-type source with a bandwidth of 50 nm around a center wavelength of 1550 nm was used. The response was recorded using an optical spectrum analyzer (OSA) at a resolution of 0.01 nm over the 50 nm bandwidth. This is sufficient to resolve the shape of each fringe of the Fabry–Perot response of the device, between cleaved facets. This Fabry–Perot response comprises multiple cavities [13]. However, the complete die cavity, i.e., the cavity set up by the two cleaved facets, can be expected to dominate—as the device design naturally includes the minimization of reflection at the internal interfaces—that is, maximizing the transmission minimizes spurious reflections. This can be confirmed via the period of the Fourier transform of the Fabry–Perot fringes, see next.

The actual transmission of the device can only be obtained by appropriate normalization. The spectral response of the source, coupled via single-mode fiber to the OSA is recorded, and all subsequent measurement spectra are normalized against this. Thereafter, there are a number of choices for level normalization—for example, a bend is commonly compared with a W1 of the length of the "stretched out" bend. This is appropriate for simple, single element devices. However, because a complete device is of interest here, including the in-coupling from the injector, we decided to normalize against a 5-μm-wide ridge waveguide, referred to as the reference device. This device is equivalent to the PhC device access waveguides, and is fabricated on the same die, hence having the same dominant cavity length. Moreover, the very short PhC device can be assumed to have negligible propagation loss due to the PhC, and we can therefore assume that interface losses will dominate.

To establish the bandwidth of the device, we only consider the peaks of the Fabry–Perot response as they represent the

transmission of the device in the absence of interference effects. These points are then normalized against the source spectrum.

V. RESULTS AND DISCUSSION

Fig. 4 shows the normalized transmission spectrum of the left and right ports of the Y-splitter. The flat tops of the central set of measurements show that the 3 dB bandwidth of the device exceeds that of the measurement window. We now turn to the lithographically tuned devices, and the normalized-frequency response. Although the 5% tuning chosen was too great to allow the spectra to stitch together, this shows that the frequency response does fall off eventually. The response is too complex to rigorously determine the 3 dB bandwidth, particularly in relation to the scaling issues mentioned above, but we estimate it to be approximately 110 nm. As mentioned above, the injectors do not scale with lattice constant, and so this spectrum is somewhat distorted with respect to the 3-D FDTD simulation. However, the estimated bandwidth is less than 10% below the simulated result, which is astonishingly good agreement.

In Fig. 4, it is obvious that the transmitted intensity varies significantly between identical devices. This is due to the sensitivity of this intensity to the input coupling. Hence, simply considering the intensity is a poor way to characterize the device transmission. Therefore, to determine an actual loss figure for the device, we consider the Fourier transform of the response. As shown by Talneau [14], each individual cavity in the device creates a signature peak in the Fourier transform. The high resolution of the OSA allows the detection of multiple round trips through these cavities. The first and highest peak translates to the first pass of the light through the device—from the source into the input facet, through the device, and out at the output facet. Part of this signal is reflected at the facet and, thus, passes through the device again. This repeats at the input facet, and then continues at each facet in turn until the signal level passes below the noise floor. These multiple passes correspond to a time delay and, hence, give rise to the subsequent peaks in the Fourier response. However, each subsequent peak has also experienced the device loss, whether scattering or out-of-plane diffraction, twice more. If we assume loss of the form $\exp(-\alpha L)$, we can obtain α per round trip by fitting to the Fourier response. The transmission can then be determined independent of insertion loss.

For devices with a single-input and single-output port, it is clear that the measured intensity corresponds to the response of the complete device. Aside from the effects due to coupling at the input facet, the transmission derived from the Fabry–Perot response should agree with this. Perhaps not so obvious, is that this also holds true for a multiple-port device such as a Y-splitter. For the intensity measurement, the total device throughput is the sum of the two arms, but the Fabry–Perot derived transmission for the whole device can be derived from either arm. This is because the peak positions correspond to a path length, and heights to the loss over that path, which is approximately equal regardless of the actual path for each subsequent round trip. For an asymmetric device, where the relative path difference is significant, it would be possible to
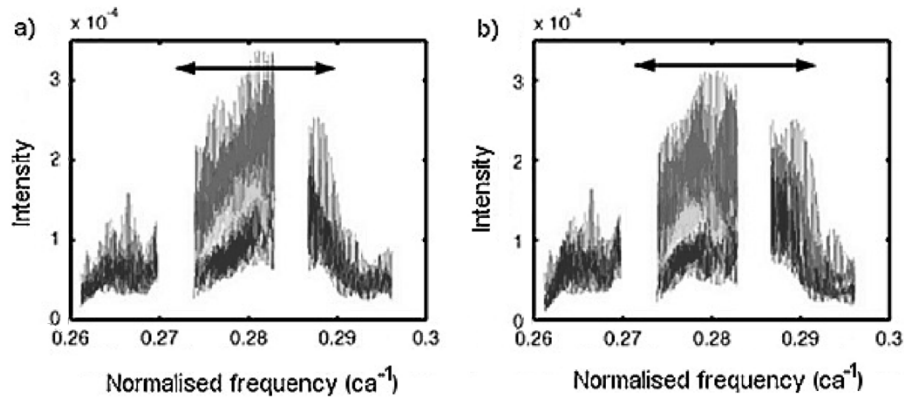
Fig. 4.   Transmission spectrum versus normalized frequency for (a) left port and (b) right port of the Y-splitter. The different colored lines represent measurements on different devices. For the designed lattice period, the range is equivalent to that shown in the simulated results in Fig. 1. The left-hand data set in each image represents the devices lithographically tuned by -5%, the central set scaled as designed, and the right-hand set the devices lithographically tuned by $+5\%$. The arrows indicate the approximate 3 dB bandwidth. Considered in relation to the designed period, this gives a bandwidth ~ 110 nm. Note also the variation in measured intensity between multiple copies of the same device.
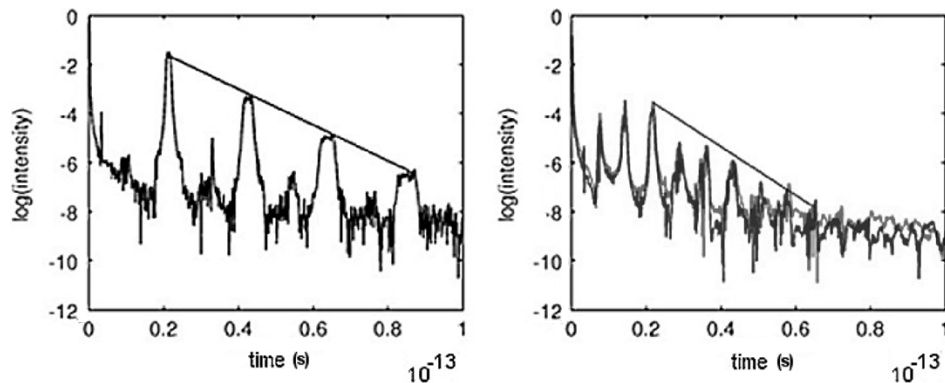

Fig. 5.   Fourier spectra derived from measured Fabry–Perot fringes. On the left, the spectrum for a 5-μm-waveguide reference device is shown, on the right for on the left (lighter line) and right (darker line) ports of the Y-splitter ports of the Y-splitter. The slope between the primary peaks is shown in each case, which indicates the loss per path through the device. For the reference device, this is ~ 9.2 cm^{-1}. The peaks for the two Y-splitter ports superimpose, indicating the similarity of path length.

assign a peak to each possible path and compute the transmission for each part of the device. The peak corresponding to a single pass will only be seen in the response of the appropriate arm. For our symmetric structure, the two paths through the left and right ports of the Y-junction are of approximately equal length and the interference peaks, therefore, superimpose—and hence, we can consider the first peak.

Fig. 5 shows the Fourier transform of the transmission spectra of an as-designed device, on the right, and that of a reference waveguide on the left. A very large number of peaks are apparent on the device transform with respect to the reference, with a much closer spacing. The spacing between the peaks in the reference spectrum translates to the device length between the cleaved facets; but the Y-splitter shows additional peaks between these. We denote the device peaks that correspond to the device-length spacing as primary peaks and use these to calculate the loss as described above. This yields an average total transmission of 70% over the two arms of the device. This is similar to the 40% transmission per arm with respect to a W1, previously reported by our group [6]. The key improvement here, however, is that the performance figure is now normalized to a low-loss ridge waveguide rather than to a W1, and the bandwidth is remarkably improved to greater than 50 nm. It

also includes a 9 μm injector as opposed to a $\sim 100 - \mu$m adiabatic taper.

The additional peaks must represent an amalgamation of the internal cavities in the device—by making the input and output guides of differing lengths, these two cavities are spatially distinct in the Fourier spectrum. This indicates that there is a significant reflection from some part of the device that requires further investigation. The presence of these additional peaks in the spectrum shows the limitations of the optimization that was performed—although this may be either because the optimization was not fully 3-D, or because the 5-μm access waveguides do not launch the correct single mode into the injector, which is difficult to assess at this point.

VI. CONCLUSION

We report the fabrication and measurement of a highly optimized design of a complete device, comprising a short injector, and W1-type Y-splitter and 60° bend. Characterization with a broadband source at a center frequency of 1.55 μm shows that the bandwidth of the device is greater than that of the source. This 50 nm EDFA-type source being commonly

used in communications applications, demonstrates that such a photonic crystal device can be used for transparent splitting. The flat-band response of the device is even more remarkable considering the fact that the structure is a composite consisting of a compact injector, W1 waveguide, Y-junction, and bends. A multiple-pass Fourier technique has been used to extract the transmission of the device, which, as normalized to a 5 μm waveguide, is 70% in total. This shows that numerical optimization can have substantial benefits for device performance. Additionally, the advantages of 3-D FDTD simulation are shown—the experimental (70%) and simulated (75%) transmission agrees very well, as does the bandwidth (110 nm experimental, 120 nm simulated). This shows that once out of plane scattering is taken into account for photonic crystals, reality can be simulated, which should considerably reduce the cycle time for future designs.

REFERENCES

[1] Y. Sugimoto, Y. Tanaka, N. Ikeda, Y. Nakamura, K. Asakawa, and K. Inoue, "Low propagation loss of 0.76 dB/mm in GaAs-based single-line-defect two-dimensional photonic crystal slab waveguides up to 1 cm in length," *Opt. Express*, vol. 12, no. 6, pp. 1090–1096, 2004.

[2] Y. A. Vlasov and S. J. Mcnab, "Losses in single-mode silicon-on-insulator strip waveguides and bends," *Opt. Express*, vol. 12, no. 8, pp. 1622–1631, 2004.

[3] Y. Sugimoto, N. Ikeda, N. Carlsson, K. Asakawa, N. Kawai, and K. Inoue, "AlGaAs-based two-dimensional photonic crystal slab with defect waveguides for planar lightwave circuit applications," *IEEE J. Quantum Electron.*, vol. 38, no. 7, pp. 760–769, 2002.

[4] Y. Sugimoto, N. Ikeda, N. Carlsson, K. Asakawa, N. Kawai, and K. Inoue, "Light-propagation characteristics of Y-branch defect waveguides in AlGaAs-based air-bridge-type two-dimensional photonic crystal slabs," *Opt. Lett.*, vol. 27, no. 6, pp. 388–390, 2002.

[5] H. Benisty, S. Olivier, C. Weisbuch, M. Agio, M. Kafesaki, C. M. Soukoulis, M. Qiu, M. Swillo, A. Karlsson, B. Jaskorzynska, A. Talneau, J. Moosburger, M. Kamp, A. Forchel, R. Ferrini, R. Houdre, and U. Oesterle, "Models and measurements for the transmission of submicron-width waveguide bends defined in two-dimensional photonic crystals," *IEEE J. Quantum Electron.*, vol. 38, no. 7, pp. 770–785, 2002.

[6] R. Wilson, T. J. Karle, I. Moerman, and T. F. Krauss, "Efficient photonic crystal Y-junctions," *J. Opt. A—Pure Appl. Opt.*, vol. 5, no. 4, pp. S76–S80, 2003.

[7] D. W. Prather, J. Murakowski, S. Y. Shi, S. Venkataraman, A. Sharkawy, C. H. Chen, and D. Pustai, "High-efficiency coupling structure for a single-line-defect photonic-crystal waveguide," *Opt. Lett.*, vol. 27, no. 18, pp. 1601–1603, 2002.

[8] T. D. Happ, M. Kamp, and A. Forchel, "Photonic crystal tapers for ultracompact mode conversion," *Opt. Lett.*, vol. 26, no. 14, pp. 1102–1104, 2001.

[9] A. Talneau, P. Lalanne, M. Agio, and C. M. Soukoulis, "Low-reflection photonic-crystal taper for efficient coupling between guide sections of arbitrary widths," *Opt. Lett.*, vol. 27, no. 17, pp. 1522–1524, 2002.

[10] P. Pottier, I. Ntakis, and R. M. De La Rue, "Photonic crystal continuous taper for low-loss direct coupling into 2-D photonic crystal channel waveguides and further device functionality," *Opt. Commun.*, vol. 223, no. 4–6, pp. 339–347, 2003.

[11] [Online]. Available: www.photond.com

[12] M. V. Kotlyar, L. O'Faolain, R. Wilson, and T. F. Krauss, "High-aspect-ratio chemically assisted ion-beam etching for photonic crystals using a high beam voltage-current ratio," *J. Vac. Sci. Technol. B*, vol. 22, no. 4, pp. 1788–1791, 2004.

[13] D. Hofstetter and R. L. Thornton, "Loss measurements on semiconductor lasers by Fourier analysis of the emission spectra," *Appl. Phys. Lett.*, vol. 72, no. 4, pp. 404–406, 1998.

[14] A. Talneau, M. Mulot, S. Anand, and P. Lalanne, "Compound cavity measurement of transmission and reflection of a tapered single-line photonic-crystal waveguide," *Appl. Phys. Lett.*, vol. 82, no. 16, pp. 2577–2579, 2003.

[15] T. Felici, A. Lavrinenko, D. Gallagher, and T. Davies, "New design rules for planar photonic crystal devices obtained using automatic optimization, leading to record efficiencies," in *Proc. ECOC-IOOC*, vol. 5, 2003, pp. 66–67.

[16] Y. Sugimoto, Y. Tanaka, N. Ikeda, K. Kanamoto, Y. Nakamura, S. Ohkouchi, H. Nakamura, K. Inoue, H. Sasaki, Y. Wanabe, K. Ishida, H. Ishikawa, and K. Asakawa, "Two dimensional semiconductor-based photonic crystal slab waveguides for ultrafast optical signal processing devices," *IEICE Trans. Electron.*, vol. E87-C, pp. 316–327, 2004.

Melanie Ayre received the M.Eng. degree (First Class Honors) from the University of Glasgow, U.K., in 2002. She is currently working towards the Ph.D. degree in the Microphotonics and Photonic Crystals Group, University of St. Andrews, Fife, U.K., focussing on design, simulation, fabrication, and characterization of photonic crystal-based devices.

Tim J. Karle (M'02) received the M.Phys. degree (Honors) from the University of Manchester, Manchester, U.K., in 1999, the M.Sc. degree (with distinction) in optoelectronic and laser devices jointly from Heriot-Watt University, Selkirkshire, U.K., and the University of St. Andrews, Fife, U.K., in 2000, and the Ph.D. degree from the University of St. Andrews, in 2005, having investigated the dispersion engineering of planar photonic crystal structures.

He was a Design Engineer with Renishaw plc, conditioning diode laser light for Raman spectroscopy. He is currently with the Microphotonics and Photonic Crystals Research Group, University of St. Andrews, researching the interactions between ultrafast optical pulses and photonic crystals.

Dr. Karle was the recipient of the Neil Forbes/Scottish Enterprise Prize for Excellence in Optoelectronics.

Lijun Wu received the Ph.D. degree from the Institute of Physics, Chinese Academy of Sciences, Beijing, China, in 2000.

In October of 2000, she joined Microphotonics and Photonic Crystal Research Group, University of St. Andrews, Fife, U.K., as a Research Fellow, where she was focusing on photonic crystal "conductors," which are based on the anomalous dispersive properties of the photonic band edge. In November of 2003, she then joined in Hong Kong University of Science and Technology. Her current interest is to fabricate 1-D, 2-D, and 3-D polymeric photonic crystals by laser holography lithography technique and explore their applications.

Tom Davies was born on December 26, 1977, Chatham, Kent. He received the B.Sc. degree (First Class Honors) in mathematics and physics from the University of St. Andrews, Fife, U.K., in 2000.

Since January 2002, he has been with Photon Design, Oxford, U.K., as their Technical Advisor.

Thomas F. Krauss received the Dipl.-Ing. degree in photographic engineering with a diploma thesis entitled "excimer laser etching of polyimide" from FH Koeln, Germany, in 1989 and the Ph.D. degree in semiconductor ring lasers from the University of Glasgow, Glasgow, U.K., in 1992. His Ph.D. work at the University of Glasgow was on semiconductor ring lasers, where he demonstrated low threshold current levels and greatly improved external device efficiency.

He has been involved in integrated optics research since 1987 when he spent a year at IBM T. J. Watson Research Center, Yorktown Heights, NY, to study $LiNbO_3$ waveguides. In 1993, he started work on photonic crystals when he won an EPSRC Research Fellowship in the area of photonic bandgaps, thereby initiating a new field at Glasgow University, followed by a Royal Society Research Fellowship in 1995. He is one of the pioneers of semiconductor-based photonic crystals and made the first demonstration of 2-D photonic bandgap effects at optical wavelengths (1996). In 1997, he spent a year at the California Institute of Technology (Caltech), Pasadena, CA, working on photonic crystal-based lasers and light emitters. He became Personal Chair of Optoelectronics, School of Physics and Astronomy, University of St. Andrews, St. Andrews, U.K., in 2000, where he set up the Photonic Crystal Research Group (now called the Microphotonic Group) and a semiconductor microfabrication facility. He is grant holder of several national and industrially sponsored research projects. He was Coordinator of the European Research Project "PICCO" during 2000–2003. He was Organizer and Chairman of the trend-setting Workshop "PECS3" in St. Andrews in June 2001, and is on the committee of several other workshops and summer schools, including PECS and ECOC.

Dr. Krauss was elected a Fellow of the Royal Society of Edinburgh and Fellow of the Institute of Physics in 2002.