Photonic crystal interfaces: A design-driven approach

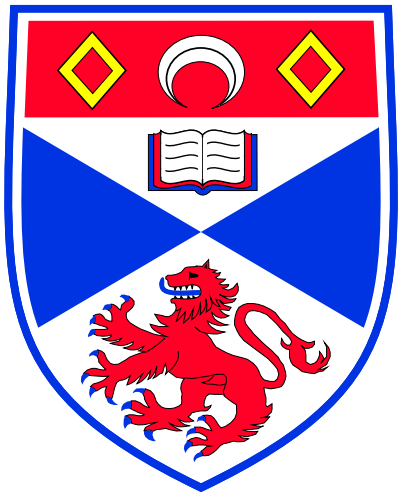

A thesis presented by

Melanie Ayre MEng.

to the

University of St Andrews

in application for the degree of

Doctor of Philosophy

June 2006

Abstract

Photonic Crystal structures have been heralded as a disruptive technology for the miniaturization of opto-electronic devices, offering as they do the possibility of guiding and manipulating light in sub-micron scale waveguides. Applications of photonic crystal guiding - the ability to send light around sharp bends or compactly split signals into two or more channels have attracted a great deal of attention. Other effects of this waveguiding mechanism have become apparent, and attracted much interest - the novel dispersion surfaces of photonic crystal structures allow the possibility of "slow light" in a dielectric medium, which as well as the possibility of compact optical delay lines may allow enhanced light-matter interaction, and hence miniaturisation of active optical devices. I also consider a third, more traditional type of photonic crystal, in the form of a grating for surface coupling.

In this thesis, I address many of the aspects of passive photonic crystals, from the underlying theory through applied device modelling, fabrication concerns and experimental results and analysis. Further, for the devices studied, I consider both the relative merits of the photonic crystal approach and of my work compared to that of others in the field. Thus, the complete spectrum of photonic crystal devices is covered.

With regard to specific results, the highlights of the work contained in this thesis are as follows:

- Realisation of surface grating couplers in a novel material system demonstrating some of the highest reported fibre coupling efficiencies.
- Development of a short "injecting" taper for coupling into photonic crystal devices.

- Optimisation and experimental validation of photonic crystal routing elements (Y-splitter and bend).
- Exploration of interfaces and coupling for “slow light” photonic crystals.

Declaration

I, Melanie Ayre, hereby certify that this thesis, which is approximately 40,000 words in length, has been written by me, that is the record of work carried out by me, and that it has not been submitted in any previous application for a higher degree.

date signature of candidate

I was admitted as a research student in September 2002 and as a candidate for the degree of Ph.D. in September 2003; the higher study for which this is a record was carried out in the University of St.Andrews between 2002 and 2006.

date signature of candidate

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of Ph.D. in the University of St. Andrews and that the candidate is qualified to submit this thesis in application for that degree.

date signature of supervisor

Copyright

Acknowledgments

In August 2002, as an exciting beginning to my doctoral studies, I was given the opportunity to travel to the town of Ascona, on the banks of Lago Maggiore in Switzerland, to attend a summer school on two dimensional photonic crystals. Having little experience of any kind of optics, I remember being alternatively amazed and baffled by the other delegates, the apparent breadth and depth of their knowledge, and their willingness to engage in vigorous debate with their elders and betters. I am particularly indebted to Maria Kotlyar, then a fellow St Andrews student, who helped me to hide my nervousness and inexperience!

Arriving finally in St Andrews, I found myself part of a big research group under a mostly absent professor. Thankfully, with so many people around, I was well placed to start learning my new subject. I'd like to thank Rab Wilson, my post-doc, for helping me to see just how much I'd need to know about, if not actually know, and Mikey Settle and Steve Neale, who started their PhDs with me, just for being themselves. Donald Brown and Simon McGreehin helped me to see the importance of attention to detail, and Will Whelan-Curtin the importance of blind leaps of faith! There are now 35 current and former members of Prof. Krauss's Microphotonics Group, and if they'd all been chosen solely to help me towards gaining my PhD they couldn't have been better.

None of the work in this thesis could have been done alone. In many of the chapters, I have specifically acknowledged the people I have collaborated with on particular topics, and their valuable contributions. But there is one individual without whom none of this would be possible - I give you George Robb, technician extraordinary, fixer of stuff and worker of miracles, and the only person I've ever seen keep a steel rule in his pocket just to stir his tea. Honourable mention for his new assistant Steve Balfour, apprentice magician

with a hard act to follow.

As for the elusive Prof. Krauss, well, I'm not sure I could've chosen a better PhD supervisor than Thomas. Since he's so hard to get hold of, on technical matters I found myself asking the real experts for help first. This has meant that when I've finally managed to catch him, my opinion is well enough informed that I can be like one of those obstreperous delegates myself, trying to find the line between useful debate and pointless argument, and walk that knife edge carefully. For me, this has helped to maximise the benefits of having a well-known and travelled supervisor, always jetting off but returning full of enthusiasm, with new ideas and approaches to try. Although sometimes it's been exhausting. So let me admit here what may not have been clear at the time - Thomas, at least sometimes, I really thought you were right!

On a personal note, I'd like to thank everyone who's put up with me while I've been writing, especially Paul Cruickshank and Dave Bolton, both for constantly reminding me why it's good to do a PhD, and making it easier; Elaine, Matt and John, who can always make time to watch a movie; Ruth and Christina for endless cups of tea; and my parents and grandparents, who've been desperately waiting for me to stop being a student.

Finally, there's a line I've kept hearing about St Andrews University. Apparently its students either meet their future spouse or become alcohol-dependent. Although I'm still thinking about which, I'd like to thank someone else I first got to know in Ascona, and one of the first to tell me that story: Tim Karle.

Contents

Chapter 1

Introduction

In this work, I will present the results of my research into Photonic Crystal (PhC) devices, undertaken in Prof. T.F. Krauss' Microphotonics and Photonic Crystals Group at the University of St Andrews. In this introductory chapter, I will briefly explain the potential role of Photonic Crystal devices, and then outline the rest of this thesis - as I have studied a number of significantly different devices, the literature reviews are presented in their relevant context.

1.1 The role of PhCs

The two major technological drivers for modern industry are computer chips and communications systems, but it is very difficult to effectively integrate the two. Computer chips, or silicon-based integrated circuits (ICs), are hugely complex mass-produced devices, and have millions of features on a scale of around 100nm, and communications between these chips currently have a bit rate of several gigahertz. Optical technologies for long distance communications can go much faster than this - 40GHz systems are deployed and there are roadmaps for 100GHz, but the devices are bulky and expensive,

and designed to operate with fibre optic systems. A single functional device may have a length of several centimetres, which is unthinkable to electronics system designers! There are also high density optical technologies, with a feature size of the order of 1μm, which are more suited to integrate with ICs, and this integration is an active area of research.

For many years it has been recognised that optical communications and electronic data-processing technologies must converge. The ICs are increasingly limited by communications both on and between chips, and the demand for broadband internet is fueling an expansion in communications markets, albeit without the headlines of a few years ago. This convergence has proved rather more difficult than was hoped, with optical systems repeatedly proving more demanding than improved electronics, but the major players in the silicon industry still believe that optics will ultimately be necessary to the continued growth of the IC markets and technologies[1]. There is much still to do to enable effective use of optics technologies in ICs - a lingering issue being that modern high density integrated optics do not exist on silicon. Instead, indium phosphide and gallium arsenide are the choice semiconductors where active functionality is required, and silica-on-silicon, which is convenient for use with optical fibre, is prefered for passive systems. Although the search for silicon-based alternatives continues, the current thinking is to integrate suitable optics materials into the silicon platform where required.

The context of this work is routing light effectively, and coupling between different kinds of optical devices, which will be essential technologies for future electro-optic integrated circuits. The fabrication technologies required are highly compatible with the planar processing of ICs, and PhCs are presented as a means of reducing the size of conventional electro-optic devices, both because they have very small waveguiding geometries, and because the intrinsic

dispersion properties of PhC devices are a strong contender for enhancing light-matter interactions [2, 3], thereby further reducing the device footprint. Particularly, I have studied the principles and techniques that can be used to design a range of photonic crystal devices, taking into account the practical matters of fabricating and characterising physical devices. I present a complete solution, taking light from a fibre into a semiconductor waveguide, then down to the small scale of PhCs, and finally coupling into a PhC device and using it for routing of light.

1.2 Thesis structure

In chapter 2, I discuss the background needed to understand the operation of PhC devices and their underlying theory, with a particular regard to justifying the choices that have been made for the designs presented in later chapters. Thereafter, we concern ourselves with devices and introduce the design techniques used, and both their strengths and weaknesses. As I have focussed on practical devices, it is necessary to actually fabricate them, to gain an understanding of the constraints that this implies. Hence, in chapter 3, I will discuss the fabrication technologies that have been used, concentrating particularly on the techniques that have been developed for effective PhC device fabrication here at St Andrews. In some ways my fabrication work is unusual, because for this thesis I have studied and made devices in the three major optics material systems: InP, SOI and AlGaAs. Each of these has their application in optics technologies, and also presents its own fabrication challenges. Many techniques only work for specific systems, as is explained in chapter 3.

After this introductory section, we finally move into the devices that have

been studied, beginning in chapter 4 with a surface grating coupler on InP that very successfully couples light from an optical fibre on to a chip. The principles behind the design are elaborated, before concentrating on the specific fabrication effort to realise these devices. Finally measurement results are presented. This work has been undertaken in collaboration with Universiteit Gent, in Belgium.

These couplers are relatively large devices, dictated by the 6μm mode diameter of standard single mode fibre. This means that the coupler itself has a footprint of around 12 μm^2. PhC devices at telecommunications wavelengths have a mode diameter of around 500nm. A standard adiabatic taper between these two scales is over 500μm in length. This represents the major part of the device footprint in many cases, and hence is undesirable. In chapter 5, we use multimode interference techniques to try and reduce this length, successfully designing and characterising a taper which is only 9μm long. I also present some work which has been undertaken to try and improve the coupling into a PhC, making use of impedance matching techniques, and suggest a design which appears to work in theory. I also discuss a Mach-Zehnder Inteferometer device design which was carefully studied in terms of coupling, which would require a successful interface design to couple light efficiently and without reflection or loss, and hence show effective use of the anomalous dispersion of PhCs.

The last of the device chapters makes use of PhC principles and numerical design methods, to successfully design, fabricate, and characterise a PhC Y-splitter and bend. As this was a major topic in the PhC community for some years, the relevant literature is reviewed in detail, tracing the evolution of designs for PhC routing devices. We show that the method used here is the best to date, with a design created using optimisation software and very careful

reproduction of the design. This was carried out in collaboration with Photon Design, an electromagnetics software company based in Oxford, UK.

Finally, I offer some conclusions on the success of this project, and the viability of designing PhC devices for real applications, along with some suggestions for further work.

Chapter 2

Theory and modelling of photonic crystals

In this chapter, I will introduce Photonic Crystals (PhCs) in a conceptual way, explaining the origins of the effects that I will refer to continually in chapters 4, 5, and 6. I will then introduce the modelling techniques used to study the devices presented in these chapters. Both time domain and frequency domain techniques have been used, and will be compared. In general, time domain methods are simpler to understand and use, but frequency domain techniques can be more powerful, and give greater insight into problems. The key with both methods is to determine a dispersion diagram which agrees with a transmission model.

2.1 Introduction to photonic crystals

Many introductions to PhCs start with Maxwell's equations, and then progress from there. Although I have developed techniques for extracting information from many different models of PhCs, I have never written a program to solve

Maxwell's equations to gain this data, and so it seems more appropriate to start with the concepts I work with. However, to use such programs effectively it is necessary to understand the mathematics and algorithms behind them, so I will return to Maxwell shortly.

We start with a PhC as a periodically patterned dielectric, for example a Bragg stack in a 1D system, or opals in 3D. The key discovery was made independently by Yablonovitch and John, who both showed that a periodicity in all of the dimensions of a given space gives rise to an omnidirectional bandgap, because no states exist for propagation in any direction, for certain wavelengths. Further, they showed that the periodicity must be on the order of the wavelength for this bandgap to exist [4, 5, 6, 7, 8]. To do this, it was necessary to find periodic solutions to Maxwell's equations, invoking the Floquet-Bloch theorum and many concepts from solid-state theory, particularly those of Brillouin Zones and Bloch Modes.

In saying that there are wavelengths where no states exist in PhC structures, we imply that there are wavelengths where states do exist. As these are the solutions to Maxwell's equations, they necessarily define the propagating waveguide modes of the system - we commonly refer to these as lattice modes. They have been found to have interesting dispersion surfaces, and as a result the use of these lattices has been proposed and demonstrated for such devices as super-prisms and super-collimators [9, 10, 11].

The first demonstrations of an omni-directional bandgap created by a PhC were in the microwave regime, with the PhC on a macroscopic scale. Nevertheless, the idea behind PhCs was to control spontaneous emission, and so many groups moved quickly towards optical devices. As part of this, it was found that fully 3D PhC devices are very challenging to fabricate, particularly on the sub-micron scale necessary for telecommunications wavelengths. Instead, the

concept of PhC slabs was developed by Joannopoulos [4] - using PhC effects in the plane, but waveguiding via total internal reflection (TIR) reflection in the third dimension. Although still challenging, as will be demonstrated in chapter 3, fabricating devices in planar PhCs is much simpler than in 3D crystals, as techniques developed for planar processing of conventional semiconductor devices can be adapted for this purpose. Hereafter, any reference to PhCs will strictly mean PhC slabs.

Planar PhCs inherit many properties from the slab waveguides they inhabit, and so strictly they should be modelled as complete 3D devices. However, because it is possible to separate these effects, the majority of the modelling work both in this thesis and in the PhC community is performed in 2D only, assuming that the third dimension is either infinite and uniform, or of zero size. Any simulations performed in 3D will have this noted. Of course, it is necessary to compare the results of 2D and 3D models, especially when assessing fabricated devices. For many structures the only major difference is a slight shift in frequency of spectral features.

To permit comparison between 2D models and real devices, we must have methods to determine the effects of the third dimension. Most importantly, we work with the concept of the lightline, which is superimposed on the dispersion diagrams that are calculated as $\omega = ck_0$, or $u = k/n_{cladding}$ in the normalised units typically used. As even a PhC mode cannot be expected to propagate faster than light, states above this line on the dispersion diagram are inherently lossy - that is, they are quasi-guided or radiation modes.

There are two common types of PhC slabs, classified by Baets as HOP and LOP devices respectively [12]. These names refer to high and low out-of-plane contrast, that is whether the TIR guiding is like that used in a semiconductor laser (or of course an optical fibre), with a small index difference between core

and cladding (LOP), or like a photonic wire, with high Δn (HOP). LOP devices are essential when one aims to integrate with active functionality. In this case, all operation must be above the lightline, but it is hoped that sufficient gain can be applied to overcome intrinsic losses. HOP devices, which can operate below the lightline, are much more suitable for waveguiding applications.

2.1.1 Photonic Crystal Slabs

For all PhC devices, the first point of reference is the bandstructure. In the next sections, I will describe ways in which these can be calculated, but here I will simply present the results, for 2D bandstructures. With PhC slabs, the first decision is whether to use a lattice of high index rods in a low index medium, or low index rods (i.e. holes) in a high index medium. It is the properties of rod systems that first caused excitement about PhC devices, as these have a wave guided in air; however these structures have no TIR confinement and so very long rods are needed. These are again very challenging to fabricate, and simpler fabrication is one of the motivations for slab PhCs. Hence, here I will exclusively study the properties of lattice-of-holes PhCs.

The next question is, what kind of lattice to use? It has been shown that triangular lattices have broader bandgaps than square lattices, essentially due to increased symmetry[13]. Hence we choose this system for most of the studies here. The exception is for the work on 2D grating couplers discussed in chapter 4 where it is essential that the basis vectors for a unit cell are orthogonal, and hence we use a square lattice. Figure 2.1 shows a typical bandstructure for a triangular lattice of holes[1].

[1]Specifically, air holes through a material of refractive index 2.89, which is the effective index value for the SOI slab used in later chapters, with a normalised hole radius of 0.28. These values are used throughout unless otherwise stated.

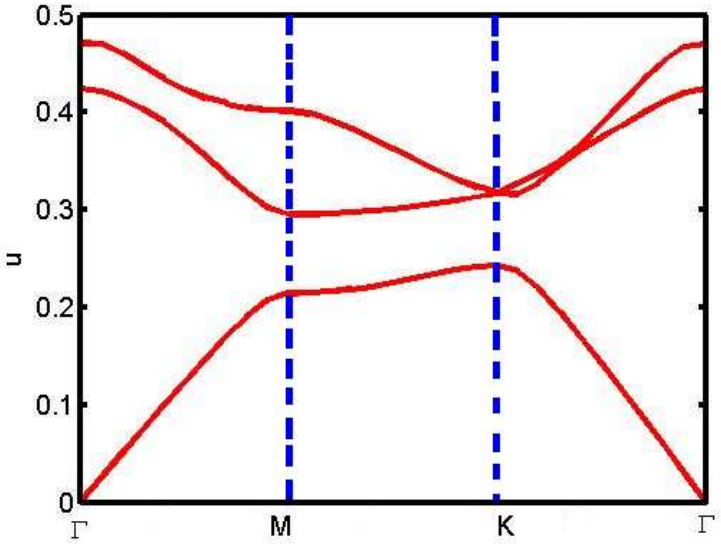

Figure 2.1: *A typical dispersion diagram or bandstructure for a triangular lattice of holes-type photonic crystal. The ordinate traces the edge of the irreducible Brillouin zone for this structure. Γ, K and M are the vertices of the IBZ, following standard symmetry labelling. This is the normalised k vector. The abscissa is normalised frequency, $\omega a/2\pi c$ for a PhC with lattice period a. The bands represent the allowed (w, k) states. The calculation method is described later in this chapter.*

This shows the allowed states for propagation through the lattice. However, for the devices of interest here, we wish to propagate through a defined waveguide. To permit this, we can introduce defects into our PhC lattice. The most common approaches are to remove lines of holes in the ΓK (nearest neighbour) direction. If we remove one line of holes, we create a so-called W1 waveguide; three lines gives a W3 waveguide, and so on. I have focussed on the W1 system. By removing holes, we introduce defect states inside the PhC bandgap, and the W1 defect introduces only a single odd-symmetry mode. As we wish to make compound devices, single mode operation is a strong

advantage.

To calculate the bandstructure for such a waveguide, we use a supercell consisting of many periods of the unit cell, and then delete a hole to form the defect. However, because we wish free propagation in the ΓK direction, the supercell must be only one period wide in this vector. As a result, we must have a rectangular supercell, containing several holes, and hence we now calculate a projected bandstructure. In figure 2.2, we calculate the bandstructure for identical parameters to the above, but now using a rectangular basis. The first figure is much simpler, but the second can be understood in terms of zone folding from the first. The point indicated as K' is the high symmetry K point for the supercell in the triangular basis, and for our defect waveguides, it is at this point that folding effects occur. For the mathematics, see [14], or for a geometric approach [15]. Finally, in figure 2.3, we show a typical bandstucture for a W1 waveguide, with the various regions of interest indicated. The bandstructure is restricted to ΓK, because we are solely interested in propagation in this direction, hence the presence of an omnidirectional bandgap is not important.

In the figure, the air and silica lightlines are shown - the lines $\omega = ck$ and $\omega = ck/n_{SiO_2}$ for suitably normalised ω and k. These are used to express the effects of guiding in the neglected third dimension of the slab PhC waveguide. That is, they show the limits of guiding by total internal reflection - modes calculated in a 2D system but above this line are quasi-guided or leaky, whereas those below are truly guided in three dimensions. For a air-bridge structure, only the air lightline is relevant, whereas for a supported membrane such as SOI the cladding lightline is the determining one. Both of these are HOP type structures. For LOP structures, the refractive index of the cladding is the same as the effective index of the slab. Hence the lightline on the bandstructure

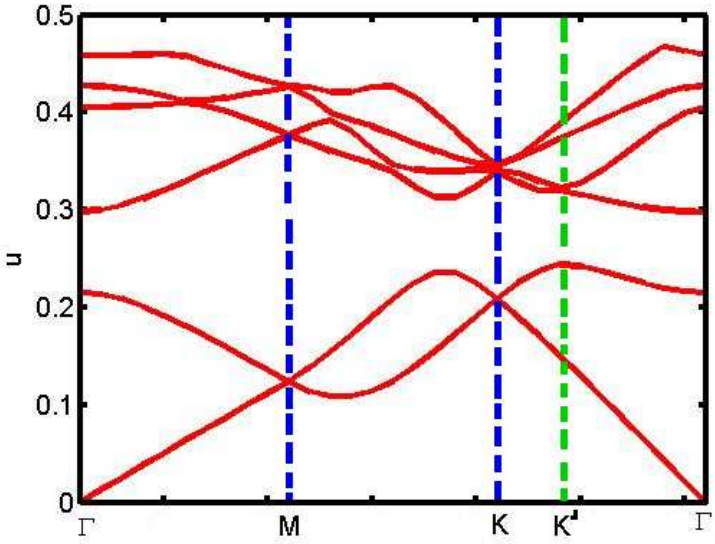

Figure 2.2: *Bandstructure for a triangular lattice of holes, using rectangular basis vectors as opposed to the natural unit vectors in the previous figure. The supercell now contains two holes, so some otherwise degenerate bands are duplicated at other points in the dispersion diagram. Appropriate folding of the Brillouin zones returns this figure to the simple form above. The point marked K' is important, as this is the projected point equivalent to the K point for a unit cell containing a single hole.*

closely follows the fundamental mode on its first pass between Γ and K, before it folds back. As a result the folded-back PhC mode in the band-gap is not TIR guided at all, but instead radiates for all frequencies.

The fundamental mode of the HOP W1 waveguide crosses this lightline, thus in transmission experiments the transmitted power reduces as frequency increases. However, for this mode the transmission also decreases as the frequency is reduced below the lightline. This is a coupling effect which will be discussed more fully in chapter 4. In simple terms, the group index of the band is $\partial\omega/\partial k$, and near the fold the group index tends towards zero. The effect

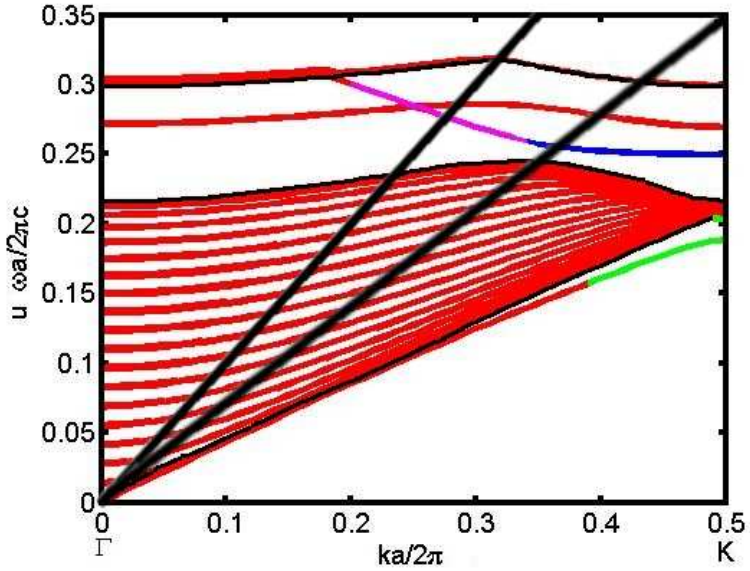

Figure 2.3: *Bandstructure for a W1 waveguide in a triangular lattice of holes, using rectangular basis vectors. We now only consider the Brillouin zone between* Γ *and* K*. The fundamental mode is magenta for normal propagation and blue in the slow light regime; the fundamental mode below the bandgap is green (see chapter 4), and the lattice regions are outlined in black. The two heavy black lines are the air and silica lightlines.*

becomes more pronounced as the mode interacts with the lower band edge of the uniform lattice as well as with itself. The precise points of these two interactions are determined by the fill-factor; but the frequencies are similar. This is discussed in detail at the end of chapter 4. As a result, coupling into the PhC mode from any external mode becomes more challenging, and simultaneously extrinsic losses increase, because the optical length of the waveguide is greater. However, this so-called slow light zone is also one of the major advantages to guiding in a PhC - because non-linear effects are equally enhanced by the increased interaction length. The part of the mode with "normal" group velocity, coloured magenta in figure 2.3, is referred to as

the fast or index-guided mode, whereas the blue slow-light mode is gap-guided. The transition between these two regimes is smooth, and the reasons for them are explained in detail in [16].

2.2 Maxwell's equations

After the conceptual introduction above, we now return to mathematical analysis, starting with Maxwell's equations. These are given here for reference, particularly for the descriptions of bandstructure calculations using planewave and finite-difference-time-domain (FDTD) methods which will be discussed in subsequent sections.

$$div(\mathbf{B}) = 0 \tag{2.1}$$

$$div(\mathbf{D}) = \rho \tag{2.2}$$

$$curl(\mathbf{H}) = \partial_t \mathbf{D} + J \tag{2.3}$$

$$curl(\mathbf{E}) = -\partial_t \mathbf{B} \tag{2.4}$$

Using the standard symbols - $\mathbf{E}$ is the electric field, $\mathbf{H}$ the magnetic field, $\mathbf{D}$ the electric flux density, and $\mathbf{B}$ the magnetic induction. For linear materials $\mathbf{B} = \mu\mathbf{H}$ for magnetic permeability μ, and $\mathbf{D} = \epsilon\mathbf{E}$ with $\epsilon = \epsilon_0\epsilon_r$ representing electric permittivity. ϵ_0 is the free space term; commonly I will work with the refractive index $n = \epsilon_r^{0.5}$ for pure dielectric materials. ρ and J represent free charge and current densities respectively. For the systems studied in this thesis, we assume that there are no free charges, hence:

$$\rho, J = 0 \tag{2.5}$$

that the magnetic permeability $\mu_r = 1$, and that the dielectrics are purely real, therefore:

$$div(\mathbf{H}) = 0 \tag{2.6}$$

$$\epsilon\, div(\mathbf{E}) = 0 \tag{2.7}$$

$$curl(\mathbf{H}) = \epsilon\, \partial_t \mathbf{E} \tag{2.8}$$

$$curl(\mathbf{E}) = -\partial_t \mathbf{H} \tag{2.9}$$

2.3 Practical numerical modelling of photonic crystal slabs

In the introduction to PhCs above, I have stated that there are periodic solutions to Maxwell's equations. As with normal waveguide modes, these periodic solutions are transcendental, and so must be found using numerical methods. There are a variety of techniques for doing this, which have different strengths and weaknesses in terms of application for designing PhC based devices. It is this application which I will discuss here, rather than either the details of these techniques, or even the theoretical considerations leading to these solutions. For the latter, refer to [4].

The methods of solving these problems can be separated into two major categories, namely time domain and frequency domain methods. In general, time domain methods represent a brute force approach, where one chooses some approach to solve Maxwell's equations in discrete steps over a grid, whereas frequency domain methods apply some insight into the problem, and obtain a much faster numerical solution to simplified problems. For our various device designs presented in the other chapters, we always use a combination of

these methods to determine useful structures. Further, one can be confident in one's results only when they can be obtained using a variety of methods.

In this section, I will consider first frequency and then time domain methods, and for each of these discuss how bandstructure and transmission spectra can be determined using a variety of commercial and freely available packages. This only covers techniques I have used significantly, as an overview of other alternatives is provided in a subsequent section. For each of the techniques here, the aim is to explain the criticalities in its use, and the relative merits of the method.

2.3.1 Frequency Domain Techniques

We begin by considering frequency domain methods - in general, to use these, it is necessary to understand the problem and the expected result, to try and calculate an intelligent and (semi-)analytic solution. The benefit of this prior effort is that the result is relatively easy to interpret, as it will be in terms of modal eigenvalues and eigenvectors, and further using some analytic results can both improve accuracy and reduce simulation times. The disadvantages are that the method used must be targeted to the system of interest; there are no completely general frequency domain techniques. Also, results may often be for single frequencies or k-vectors and only give steady state responses, making it difficult to see how a system performs over time[17]. It is often claimed that frequency domain methods are faster, but this is only true for specific systems, and frequently not the case for PhC device modelling.

2.3.1.1 Bandstructure Calculations

In the frequency domain, there are two methods available for calculating bandstructure, namely eigenmode expansion and planewave expansion of the

basis vectors that describe the dielectric structure and the fields. However, there is a strong preference for the planewave technique, as the dependent variable is the k-vector, whereas in eigenmode the dependent variable is the frequency - and therefore it is a resonance-finding method, which we prefer to implement in FDTD (see next subsection). For a complete discussion of the eigenmode method, see reference [18].

For the planewave method, there are two options readily available at St Andrews, namely the MIT photonic bands package (known as mpb) [19, 20], or Photon Design's bandsolver package [21]. Although the detail of the code is of course different, the underlying algorithm is the same, and so the packages are often used interchangeably. To confirm this, figure 2.4 compares the (2D) bandstructures calculated for a typical W1 waveguide[2] in each package. mpb uses a completely periodic supercell technique, although the dielectric function is discretised in real rather than Fourier space. Photon Design, on the other hand, allows the use of an absorbing boundary condition, but only in the out-of-plane direction for the 3D bandstructures of PhC slabs. It is commonly thought otherwise! As a result, the supercell size used for 2D calculations is significant. In mpb we typically use a supercell with 18 periods of PhC transverse to the propagation direction [3], but the graphical nature of Photon Design encourages smaller supercells. In the figure, I have used the different typical parameter values for resolution and supercell size, and as a result the calculations are more different than would be expected based on Photon Design's validation documents[4].

[2]The parameters are the same as in chapter 4

[3]Determined by S. Boscolo of the University of Udine to have no significant coupling between parallel waveguides in the real space representation for this resolution.

[4]Validated against published results by S.G. Johnson, the author of mpb [19].

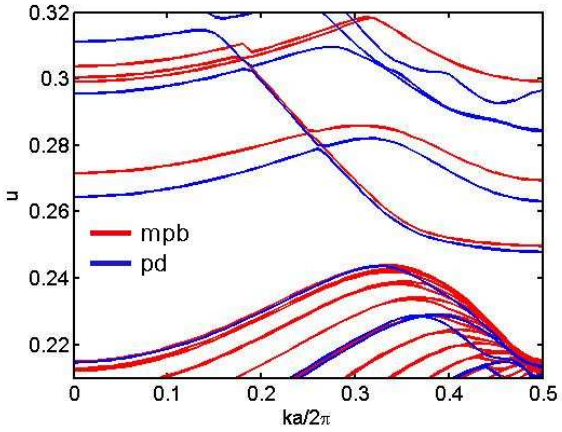

Figure 2.4: *Comparison of 2D bandstructures calculated with the mpb and Photon Design bandsolver packages, for a W1 waveguide. The calculation is performed with a grid resolution of 20 in mpb, which is optimal, and 36 in Photon Design as only powers of 2 are permitted. The nature of the packages encourages different supercell sizes, so for accurate results care must be taken with convergence. The different band positions apparent in this figure are mainly related to the different supercell sizes.*

Planewave algorithm For a detailed mathematical treatment of the algorithm, the reader is referred to [19], and the relevant manuals. Here, I will try to provide a sketch of the process, mostly derived from the manual for mpb.

In a planewave model, the bands of a PhC are computed as eigenmodes of the dielectric system, using Helmholtz's equations:

$$curl\,(\frac{1}{\epsilon}\,curl\,\mathbf{H}) - (\frac{\omega}{c})^{2} = 0 \tag{2.10}$$

$$\mathbf{E} = -(\frac{ic}{\omega\epsilon})\,curl\,\mathbf{H} \tag{2.11}$$

We must also add the constraint that the magnetic field must be transverse, that is $div(\mathbf{H}) = 0$. Using Bloch's theorum, we can write the magnetic field $\mathbf{H}$

as:

$$\mathbf{H} = \mathbf{H_k}(x)e^{i\mathbf{k}.x} \tag{2.12}$$

Where $\mathbf{H_k}$ is a periodic function. This is equivalent to saying that the calculation has periodic boundary conditions. To solve 2.11 numerically, we then expand $\mathbf{H_k}$ in a Fourier basis, as a sum of planewaves. To multiply by the inverse dielectric function $\mathbf{H_k}$ is Fourier transformed into real space, and the desired PhC discretised accordingly. We then return to the frequency domain to solve the eigenvalue equation, using iterative methods as we only desire a small number of the eigenvalues of the matrix. The key to a good planewave algorithm is in the discretisation of the dielectric function, because this impacts strongly on the time to convergence in the iteration. In both mpb and Photon Design, it is difficult to know exactly which dielectric function has been used, because the discretisation is optimised in many ways. The number of planewaves required to achieve convergence is strongly dependent on the disctetisation of the dielectric function, but we do not explore htis further here, as we only wish to use the available tools effectively. Because we calculate eigenstates, the overhead to calculate field profiles and group velocities is minimal. Hence, we can have all of this information from one calculation, which is not possible using time domain techniques.

3D Bandstructures We can also extend this to 3D, but for PhC slabs, there is a problem. As discussed, mpb simulations necessarily give solutions which are periodic in all directions, and so we must also consider an out of plane periodicity which is very unphysical[5]. To do this, we can define our slab

[5]I have been unable to test 3D bandstructure calculations in Photon Design due to computational limitations.

structure in a large supercell, with a large volume of air above and below. In principle, for a low out of plane contrast system, we could instead fill space with a high index material. For the high contrast case (e.g. SOI), we can then easily compute the true modes, those that are fully guided by the slab (i.e. below the light-line). The leaky radiation modes can also be calculated, but this is more complex. Consider: the mode tail for the radiation mode extends out of the slab, and only becomes negligible at a certain distance. If there is any significant field, it may couple into the slab mode in the next repeat of the device, or it may reflect from the interface and come back. Either way, the result is non-physical. To reduce this problem we need to use very large supercells, and the computational burden becomes beyond our capabilities. Instead, it is possible to perform the calculation with relatively low air volumes. In this case, we have to consider modes guided in the substrate, modes guided in the air, and glancing angle modes, all of which are purely a result of the boundary conditions. This can be done, by outputting the mode profile in some way (e.g. the flux density in the guiding layer) and filtering the bandstructure data accordingly, but this has not been attempted. It was found easier to use an FDTD method with absorbing boundary conditions (see later). An alternative approach is to insert an absorbing medium into the out-of-plane dielectric profile, as has been done in Photon Design's Crystalwave, however I have not evaluated this.

2.3.1.2 Transmission and Phase

Also in the frequency domain, we can calculate the transmission properties of a device. The techniques which can be used are quite general, but intelligent algorithms can take advantage of periodicity to reduce overall calculation time. Here, our choices are the freely available program CAMFR [18, 22], and Photon

Design's FimmProp and Omnisim frequency domain engine [21]. Of these, only FimmProp is a fully 3D solver, and it has only been used incidentally in this project.

The CAMFR and FimmProp codes are based on the principles of Eigenmode expansion, detailed in [18]. The structure is discretised into transverse slices, which are then stacked together in the propagation direction to form a device. The solver calculates the modes for each slice, propagates the mode (actually simply rotates the phase) for the desired length, and then performs the overlap integral between each adjacent slice. This is implemented in a scattering matrix algorithm, to take account of forwards and backwards propagating modes in a numerically stable fashion (compare transfer matrix). For simple waveguide devices and grating couplers, this is a very useful approach, because it scales with the number of different slices, rather than the area covered by the device.

For our photonic crystal waveguides, however, there are some significant disadvantages to this method. As we like to make round holes, we must discretise heavily into a staircase approximation to the circles, which reduces the benefits of slice-scaling. More important, however, are the boundary conditions. In practice, our device is surrounded by at least tens of microns of unpatterned dielectric, but simulating this empty space is undesirable. As a result, we bring a boundary close to the edges of the device, which must be positioned and designed so as not to affect the simulation results. For eigenmode techniques, the choices of boundary conditions are electric and magnetic walls, and perfectly matched layer (PML) boundaries. An electric wall behaves as a metal box surrounding the device, numerically forcing the electric field to be perpendicular to it; as a result this is a symmetric boundary. Conversely, a magnetic wall is an anti-symmetric boundary, because the electric

field is forced to zero. Both of these methods are simple to implement, and both can be easily used to divide a simulation in two along the propagation direction, forcing propagation in odd or even modes respectively. However, at the edges of the simulation, these methods can cause reflections back into the simulation space, that in real life would decay away.

PML To solve this problem, we have the perfectly matched layer[23]. Essentially this is an absorbing layer placed between the edges of the device and the wall. This is more difficult to implement in practice, but is a necessary technique. However, the mathematics of absorption cause the problems associated with PML. If we consider a real evanescent mode of the form $E_0 e^{\omega t}$, and multiply by an absorbtion $e^{-\alpha L}$, then the amplitude will reduce; and we can set α and L such that there is no significant field at the boundary. But, for example, assume that we have a propagating photonic crystal mode with a significant intensity in the lattice and hence close to the boundary. We now have a field of the form $E_0 e^{i\omega t}$, which will decay more slowly, and then reflect from the wall and re-enter the simulation. Depending on the method in which the PML is defined, it can even experience gain in some circumstances. Essentially, then, we need to first of all run a simulation for a single transverse slice of the structure, and vary the PML position, width and strength until the result is stable. A photonic crystal already has a high transverse width, and so the already large simulation must become even greater for reliable simulations in this scheme. Ideally, one must have a good idea about the modes and their properties before starting such a simulation.

As this shows, there is a significant amount of prior effort involved in setting up a frequency domain simulation of a photonic crystal device, it can be very worthwhile. Because we calculate the modes, the steady-state field

profile is instantly available, as is the amplitude and phase at any point in the simulation. This data is, of course, also obtainable in the time domain, but suffers more strongly from numerical noise, and must be extracted via arduous post-processing.

Photon Design's Frequency Domain Engine In addition to these, we have Photon Design's 2D frequency domain engine. This is a proprietary finite element technique with an adaptive grid, rather than a mode solver. However, it provides results similar to CAMFR: the mode at a given frequency of the entire system (not slices), and the relative amplitude and phase between any two points. This program is very simple to use, unlike CAMFR where discretisation must be taken care of by the user, and often solves systems very quickly. In particular, this program has been used extensively in my work on tapers (chapter 4). Unfortunately no information is available on the algorithm it uses, so I cannot describe it. However, results for simple systems compare well with those calculated using FDTD in the same package.

It was also hoped to use this package to complement the PhC coupling results in chapter 5, particularly to prove the effectiveness of the tapered crystals discussed there as a solution to the slow-light coupling problem; however, the predicted memory requirements for the necessary simulations approach 10GB RAM, hence it was not possible to complete this work.

2.3.2 Time Domain Methods

In contrast to the frequency domain methods above, time domain simulations are a heavy handed approach, because to use them little physical insight is needed - instead, they are a brute force method to solving Maxwell's equations. However, they do have many advantages, because their operation is simple, so

it is relatively easy to understand what your simulation is doing. Calculations can be set up quickly, and changing systems is easy. However, only the amplitude of a field is readily calculated, saying little about mode properties. To compensate for this, for PhC devices it is important to start and end time domain calculations in a simple ridge waveguide, so that at least you can calculate an overlap integral with the waveguide mode. Many groups consider FDTD results as a numerical experiment.

2.3.2.1 Bandstructure Calculations

For the time domain methods, it is more usual to start by discussing transmission techniques, as these are the usual application of this method. However, for the sake of consistency I will start with bandstructures.

As with plane wave techniques in the frequency domain, to calculate a dispersion relation using FDTD we start with either a unit cell or supercell, as desired. However, this is where the similarity ends. Instead of converting the dielectric profile into its k-space representation, which then becomes infinitely periodic, we consider it in real space. In this case, the types of boundary we need to consider are very different. We must consider a supercell which is conformal with the grid, hence a square or rectangle for normal codes (FDTD algorithms on other lattices are possible, but not common). We then choose to make each parallel pair of boundaries either periodic or absorbing (symmetric and antisymmetric boundaries are also possible, in which case we have a virtual parallel boundary outside of the simulation domain). An absorbing boundary is a PML, as in the frequency domain case; the implementation is different but the effect is the same. This time, though, both propagating and evanescent waves are readily absorbed. The periodic boundary is more interesting, and helps to illuminate the physical meaning of the dispersion diagram.

Firstly, though, we must understand how the FDTD algorithm works. We have previously shown Helmholtz's equations 2.11, which is derived from Maxwell's equations by assuming solutions which are sinusoidal in time, and said that it can be solved numerically. In FDTD, this is done in the following way [24]:

1. discretise the dielectric profile onto a regular grid
2. implement the Yee cell[6]
3. impinge a magnetic (electric) field
4. calculate the electric (magnetic) field via the curl equations
5. increment the time step
6. calculate the magnetic (electric) field via the curl equations
7. repeat from 4

To illuminate this, refer to Hagness [24] or Min Qiu [25], from either of whom FDTD codes and algorithms are available[7]. I have worked almost exclusively with commercial codes, and so will offer only the limited perspective needed as a user, rather than developer, of FDTD methods.

Beginning with 1, we have to consider our device not as a slab with some holes etched into it, but rather as a fine mesh, in which different cells have a different dielectric constant. In our device, we have $\epsilon_{dielectric}$ and $\epsilon_{background}$, but in a 2D calculation we use the effective index of the slab ϵ_{slab}, and ϵ_0 for an air background. Each cell in the mesh overlaid on the device will contain either air, or slab, or a proportion of both, and we model the "both" with

[6] explained subsequently

[7] also MEEP from MIT, recently released

an intermediate dielectric constant. The size of the mesh must be adjusted to give a sufficiently good representation of the structure, via convergence testing, however *period*/20 is a good rule of thumb for photonic crystal devices. Commercial codes also implement sub-gridding to increase the resolution of small features.

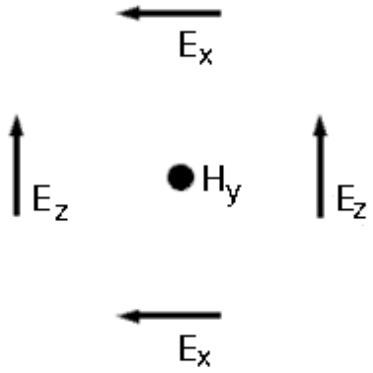

Figure 2.5: *One period of a 2D Yee cell. The electric and magnetic fields are calculated on separate interspersed grids.*

Then, we consider the Yee cell [24]. This is the key to the FDTD algorithm. A typical Yee cell schematic is shown in figure 2.5. Note that the electric and magnetic fields are calculated at different points, such that every position on the electric field grid is surrounded my magnetic field points. This makes calculation of curl equations very simple! Interpolation can be used to calculate the field values on the other grid if necessary. An incident field (3) is introduced by adding a value at a chosen (range of) points, and again this must be done slowly. To have a cw experiment, the incident field must be introduced by slowly increasing the value at a point over several time steps, proportional to the desired wavelength. According to Fourier theory, the longer the ramp time the greater the spectral purity, but even so the system shows a transient

response and so must be run for a long time to a steady state.

The remaining steps (4 onwards) calculate the evolution of the field as it propagates away from the launch. This is numerically simple in principle, especially on the Yee cell. The curl is the circulation of the field, so recalling:

$$curl(E) = \nabla \times (\mathbf{E}) = \begin{vmatrix} \hat{i} & \hat{j} & \hat{k} \\ \partial_x & \partial_y & \partial_z \\ E_x & E_y & E_z \end{vmatrix} \tag{2.13}$$

which expands as:

$$curl(\mathbf{E})) = \hat{i}(\partial_y E_z - \partial_z E_y) - \hat{j}(\partial_x E_z - \partial_z E_x) + \hat{k}(\partial_x E_y - \partial_y E_x) \tag{2.14}$$

If we consider a one dimensional system for the sake of simplicity, then ∂_x and ∂_y are zero, and we obtain:

$$curl(\mathbf{E})) = -\hat{i}\partial_z E_y + \hat{j}\partial_z E_x \tag{2.15}$$

assuming propagation in z.

Now, recalling Maxwell's curl equations:

$$curl(\mathbf{E}) = -\frac{1}{c}\,\partial_t(\mathbf{H}) \tag{2.16}$$

$$curl(\mathbf{H}) = \frac{\epsilon}{c}\,\partial_t(\mathbf{D}) \tag{2.17}$$

The magnetic field at $t + \Delta t$ can be written as:

$$\mathbf{H} = c\,(\hat{i}\partial_z E_y - \hat{j}\partial_z E_x)\,\Delta t \tag{2.18}$$

hence:

$$H_x = \frac{E_y^{z+\Delta z/2} - E_y^{z-\Delta z/2}}{\Delta z} \times c\Delta t \tag{2.19}$$

$$H_y = -\frac{E_x^{z+\Delta z/2} - E_x^{z-\Delta z/2}}{\Delta z} \times c\Delta t \tag{2.20}$$

where Δz is the spatial step, or grid size. The $\pm\Delta z/2$ refers to the interleaved grid of the Yee cell. Given **H**, we can write a similar equation for **E**. These can be extended to a full three dimensional form on the same principles. We can then continue to compute **E** and **H** alternately for as long as is desired.

In this equation, Δt and Δz cannot be arbitrary - instead, they are linked by the Courant stability condition. Essentially Δz cannot be greater than the distance that light can propagate in a time Δt, in the appropriate medium.

So, the result of the FDTD algorithm is to reduce the process of solving Maxwell's equations to simple arithmetic. Of course interesting effects arise, making FDTD an active area of research - see *www.fdtd.com* to see the current scope of this! There are several sources of complexity, including boundary conditions, efficient computations, and reducing the sheer volume of computer memory required to perform these calculations over a significant device area, amongst others.

FDTD Bandstructures After this whirlwind tour of the principles of FDTD, we return to the question of bandstructures and periodic boundary conditions. As we clearly introduce a frequency into the system with the incident field, the boundary must allow a method to find the k-vectors for any particular frequency. In chapter 4, we will consider a k-vector as an angle, and it can also be considered as a phasor. Dealing with the simple case of propagation along a waveguide, we know that $\mathbf{k} = \sqrt{k_x^2 + k_y^2 + k_z^2}$ is a particular angle, and that k_z is the component of that angle in the direction of propagation. Then, since $k = 2\pi n/\lambda$ this describes a rotation of the phase of the mode. Hence, the role of the periodic boundary condition is to map the

field at the leading edge of the supercell to the field at the trailing edge, plus a phase rotation according to the desired k-vector. The simplest case is at the Γ point, where the rotation is zero.

Previously, we have said that a field must be introduced slowly, but at a boundary we introduce on the left hand edge exactly the field calculated on the right hand edge at the previous time step. We do this by adding this set of values to the existing set before calculating the next evolution of the field. We keep doing this, and advance the time step until we reach a steady state condition. For frequencies where the both edges of the simulation are in phase, the system resonates, but all others will die away rapidly. We can determine these frequencies by using an input field with a broad spectral range, and monitoring the variation of the field with time. The Fourier transform, or better yet harmonic inversion [20] as the signal is decaying, yields the frequencies. This is then repeated for other k-vectors.

There are several advantages and disadvantages to this method. At first sight, the disadvantages seem overwhelming. Firstly, it lacks the mathematical elegance of the planewave method as described in the previous section. To find the field profile the simulation has to be repeated at the appropriate frequency. The results are noisy, even using harmonic inversion, so repeated simulations with narrower launch bandwidths have to be used to home in on precise values. The group velocity isn't intrinsically calculated. FDTD is a brute force method, and hence strictly limited by the RAM available, and the post-processing required to extract the frequencies is similarly computationally expensive. Finally, we have to work in real space with real frequencies.

However, the advantages are compelling. Firstly, as outlined above, the method of calculation is both simple and intuitive. Secondly, the real space and real frequency approach means we can easily focus our efforts on the parts

of the bandstructure we're interested in. More important, though, are the possible boundary conditions for bandstructure calculations. By using a PML in the transverse direction, we can have a supercell significantly smaller than those required for the plane wave method. Further, we can extend to 3D trivially (excepting memory requirements). By also using PML in the out of plane direction, we do not have the box mode problems that exist with plane wave, so we can solve for guided resonances above the lightline as easily as guided modes below.

2.3.2.2 Transmission and Phase

As with the bandstructure discussed above, in the time domain we typically calculate transmission using the FDTD approach. We do this simply, using the same algorithm as above, but instead of a supercell and periodic boundary conditions we lay out the dielectric profile of the entire device.

The two commercial codes we typically use, Photon Design's Omnisim [21] and RSoft's FullWAVE [26] [8], have slightly different implementations of the basic FDTD algorithm. With Fullwave, we can only determine the net field at any given point, whereas Omnisim gives access to the forwards and backwards going flux as well. However, with either approach, it is challenging to obtain accurate phase information, yet this is necessary to study the effects of the interfaces of our devices. This is discussed more fully in chapter 5. To obtain this information, in principle one could simply look at the phase from the Fourier transform of a pulsed simulation. Unfortunately, this phase is not well defined, because only the pulse centre-frequency is correctly launched in the appropriate well-defined waveguide mode, and we need a very high degree of accuracy in the phase, because the group velocity effects we are interested in

[8] A good name, because FDTD is a full-wave, as opposed to quasi-static, approach.

are calculated using the second derivative of this phase [15]. The only method which has been found to be acceptable is to run cw calculations to a steady state, and then fit a sine wave to the last few hundred periods of the response. Then we must scan the required spectral range in sub-nanometer steps. This is hugely time-consuming, and has been a major limitation on the theoretical work in chapter 5.

2.4 Conclusion

In this chapter, I have explained the terms used to describe the PhC design work later in this thesis, such as lattice regions and slow light, and showed the origin of these effects. Thereafter, I have discussed in detail the calculation methods that will be used throughout the rest of this thesis to design devices and study their behaviour. These methods are derived very briefly; with emphasis on the practical application of the methods used, their individual strengths and limitations. I stress again, no single method is completely to be trusted, instead one seeks agreement from a variety of calculations.

Of the computational techniques and packages discussed in this chapter, the tools most used in the device chapters are FDTD and PWE, using FullWAVE and mpb respectively. The time domain calculated transmission spectra from FDTD accord most closely to the form of experimental results, but the PWE bandstructure is the basic design tool for PhCs, as it is used to choose the operating point of the device, and shows the mode interactions which are necessary to understand the form of the transmission spectra. Frequency domain transmission spectrum methods have also been described, as these have been used as a complement to the FDTD calculations. Similarly, the process for calculating a bandstructure using time domain techniques has been

explained, as this provides useful physical insight into PhC device behaviour.

Although this chapter is short on pictures, there are many examples throughout the device chapters of the results of these techniques, and the success of some of the device designs show that a suitable combination of these methods can be very valuable.

Chapter 3

PhC fabrication techniques

3.1 Introduction

In this section, the microfabrication techniques used to fabricate the various devices described in the later chapters will be discussed. This comprises a wide variety of both materials and processes. However, the same basic techniques are used for all of these, and the same issues must be considered. The chapter is laid out in the order of typical process flow: substrate choice/source (external, as there is no growth in-house) hard mask and resist deposition, pattern definition and transfer, and (dry) etching. Many of these processes are somewhat tedious, and many more are very standard procedures, which will be rather glossed over. Furthermore many procedures have been developed by colleagues, and form part of our in-house knowledge, but are highly relevant here. Finally, of course, I have been heavily involved in studying some areas. To try and make my own involvement clear, I will use "we" for the former, and "I" for the latter types of procedures.

3.2 Substrate choice

As there is no wafer growth at St Andrews, this section is included merely for completeness. For the devices detailed in other chapters, I have used three particular wafers - SOI with 220nm silicon over 1000nm (or more) silicon oxide on a silicon substrate; 200nm Q1.22 InP on an InP heterostructure designed to be removable after wafer bonding; and a complex AlGaAs heterostructure. These have been sourced from various suppliers, and the wafer designs form part of other projects. Clearly, the choice of wafer is dictated by the application, and governs the fabrication processes.

In each of these, the critical part is that the guiding layer (top for InP and SOI, buried for AlGaAs) is single moded at telecoms wavelengths, which has been verified in each case by the wafer designers, F. van Laere at Universiteit Gent, T.J. Karle and T.F. Krauss at St Andrews. The layer structures are shown for SOI and AlGaAs in figure 3.1; for InP refer to chapter 4.

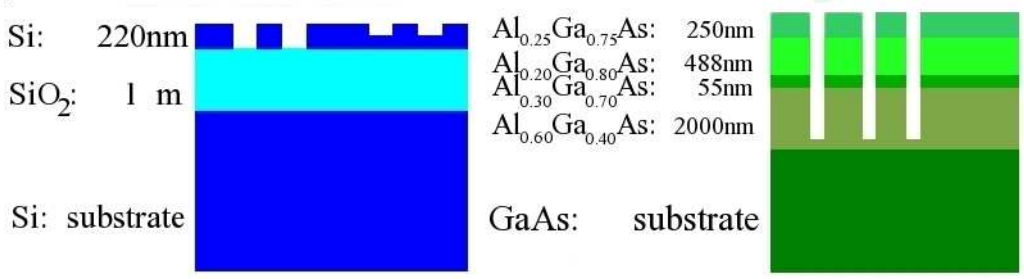

Figure 3.1: *Layer structures for the SOI (left) and AlGaAs (right) used for the devices in later chapters (not to scale). The layer compositions are indicated as well as typical etch depths.*

With the fabrication tools available at St Andrews, we have two basic types of process, silicon-based and non-silicon based. These have different capabilities and methods, but in general we can achieve similar results on

either system, albeit by different methods.

3.3 Masks

A mask is an intermediate step between the device you design and the device made. For Si processes we typically use the lithographic resist layer as the etch mask. I have tried using an oxide based hard mask as part of the process development. However, in our RIE the mask etches away faster than the silicon for PhC features of 100s of nm size, and as such is useless.

For non-Silicon (heterostructured AlGaAs and InP) fabrication, etching the pattern requires a hard mask to be deposited on the substrate. This is because the ion-beam etching is very selective to dielectric masks and the etch depth required may be great ($> 2\mu$m), and perhaps most importantly, the etch chemistry necessitates a high temperature process. In the case of shallow InP etching, I did attempt to use only a soft mask. Although it actually survived the very short etch time, it seems to have softened such that it flowed back together, deleting the pattern!

3.3.1 Hard masks

Typically, we use one of two hard masks - PECVD silica; or sol-gel hydrogen silesquioxane (HSQ, but commonly referred to by the brand name, FOX). The former is deposited for us by collaborating groups, and is a very typical process in PhC fabrication. We specify the mask thickness, for a desired etch depth. For a typical 250nm diameter hole 300nm mask gives an etch depth of up to 2μm on AlGaAs, which is elegantly sufficient; but a 200nm mask barely permits a 1μm etch. This non-linearity seems to be related to the etching process - in some way, the thicker mask seems to compact and become more

resistant, but the thinner mask does not. However, this is conjecture.

The sol-gel process has been developed for use in house, and involves a spin-coating technique to deposit the film. The chemistry of the sol-gel is not relevant here, as we use a commercial product *FOX-14*. When suitably spun and baked (at very high temperatures, up to 500^{o}C), it provides a silica-like hard mask. Each layer is typically 140nm thick, and in principle multiple layers can be deposited, to give the required mask thickness. However, I have not found this material to show the non-linear selectivity of the PECVD material, so a thicker mask is required. The HSQ is preferable in cases where electrical contacts are required at a later stage in the processing [27].

3.3.2 Lithographic resists

A resist is (typically) a polymer film deposited via spin-coating and baking at low temperatures, typically 100 to 200^{o}C. The film is altered in some way by the interaction with photons and/or electrons, by becoming either more (positive) or less (negative) soluble in its developer solution. As such, resists are used for pattern definition, and as etch masks for some processes. As with hard masks, the critical parameter is thickness, although in this case I found that multiple layers cannot always be used to increase the film thickness. Instead, solute concentration or spin speed must be altered appropriately. This would seem to be solvent-related. Specifically for the ebeam[1] resist polymethylmethacrylate(PMMA), our standard solution is a mixture of commercial PMMA and a thinner. With hindsight, it seems clear that the PMMA will readily re-dissolve in the thinner!

The typical soft masks are this PMMA and the proprietary ZEP-520A (commonly, ZEP) which are ebeam resists. No significant optical lithography

[1]short for "electron beam", and used throughout to mean "electron beam lithography"

has been used here, so optical masks will be neglected. With both of these resists, it is necessary to consider the required sensitivity to the ebeam exposure. Although this decreases only slightly with mask thickness, it is easier to accurately define a PhC with a thinner mask. Of course, possible etch depths increase with thickness. For the narrow lines and large holes needed for my InP work, it is necessary to have a very thin mask, to achieve sufficient control of feature size. On the other hand, for my SOI etching, it never became possible to make a thick enough film of PMMA to reliably etch the silicon layer well, and as a result it was necessary to switch to ZEP. This will be discussed further in the etching section.

3.4 Lithography

The desired pattern is defined in the chosen resist in the lithography step. As I typically need sub-100nm resolution, I use ebeam. The major pattern definition tool used in this project is a Raith Gemini electron beam system. Although installed at St Andrews several years prior to the beginning of this project, a huge amount of detailed knowledge and experience was required to make the devices for this project, and has been developed over the course of (although not as part of) this project. Almost all of the credit goes to D.H Brown and L. O'Faolain; but again I will emphasise my contributions.

Some additional electron beam lithography was also performed at Glasgow University on their Leica EBPG-5 semi-commercial tool.

I have also been involved in creating patterns for DUV lithography performed at IMEC, Belgium, as part of the EU projects PICCO, ePixNET, and Funfox. However, the limitations of the DUV process are very restricting for the device designs I have created, and no useful measurements have

resulted, so I will not discuss this further here. For details of this process, see [28].

3.4.1 Electron beam lithography

Ebeam, more formally known as electron beam lithography, is perhaps the most important technology for PhC fabrication. By using high energy electrons rather than low energy photons to expose the resist, the degree of control over pattern definition considerably improves, and as a result the minimum attainable feature size decreases.

Although a detailed review of the theory of ebeam lithography is beyond the scope of this thesis, a discussion of the process from a user's perspective is illuminating - since this is the level at which we have mastered our system. For details of the theory, see [29], and the system manuals. Also worth a read is the user manual which has been developed in house. I then present a discussion of the techniques that I have helped to develop, which have been critical to fabrication success in this project.

3.4.1.1 Process flow

This is provided for reference purposes, as it is very specific to our current system. However, it is useful to provide context for terms which will be used later.

1. Load sample
2. Wait at least one hour, to allow stage drift to stabilise
3. Measure beam current
4. Set step size and dwell time

5. Correct focusing, stigmation and aperture alignment at very high magnification (greater than 100,000 times)

6. Burn contamination spot [2]

7. Align writefields, ideally using a contamination spot

8. Expose pattern

Doing all of this correctly can take quite some time! When writing small numbers of large features, this process can be very simple. But to successfully write a set of photonic crystal devices, the process breaks down into a large number of smaller steps, each with their own complications. Subsequently, I will discuss the interesting ones.

3.4.1.2 Beam currents, beam speeds, and step sizes

In the process flow step 3, "measure beam current" sounds very simple. In fact, as a process, it is! However, the next part, setting step sizes and dwell time is rather less so, and depends strongly on the measured current. These values affect both resolution and stitching issues. The former is important for my work on elliptical holes, the latter for making access waveguides to bring the PhC-matched waveguide mode up to an easily measurable scale.

Consider firstly resolution: we operate with a fixed beam current, and write on a resist with a fixed clearing dose, where dose is equal to current $\times$ time /area, typically in μAs/cm^2. Clearly, there is a physical limit to

[2] This is very specific to Raith systems - it translates to exposing a single point and watching the behaviour of the beam current on a picoammeter. As the point is exposed - burnt - its conductivity changes and the current falls. The rate of current decrease and the size of the resulting contamination spot show how well the focussing etc. have been done, and helps to fine tune these.

how quickly the beam can move, or be switched off. As a result, there is a minimum area that can be exposed - this is known as the step size for the raster scan. Initially, our ebeam lithography was done with the beam area and hence current constricted by a thirty micron aperture. In figure 3.2, the pattern was defined using the minimum step size possible with this aperture. The pattern is an elliptical hole with a constant air filling factor, replicated over a square lattice. The bottom left ellipse is at a 45 degree angle to the lattice vectors, and is rotated across the lattice such that the ellipse is vertical at the top left and bottom right. The important point is this: in the SEM image, the air filling factor is not constant. The reason for this is the discretisation involved with the raster scan at the minimum step size; which is sufficient to vastly distort the filling factor.

To solve this problem, we clearly must further reduce the step size. To do this, the only choice is to reduce the current. We do this by reducing the aperture, and hence the beam diameter. Figure 3.3 shows a pattern for a real device, written with 30μm, 20μm, and 10μm micron apertures. The 10μm pattern has the highest definition, but takes longest to write.

However, as is common with a complex technology, solving one problem tends to enhance another. In this case, the downside is very obvious - by writing with a smaller current, we increase the dwell time (the time the beam stays on one pixel) to keep the exposure dose constant. As a result, we increase the total time it takes to write the device! If there was only a small photonic crystal to consider, this would be an increase from a tiny duration to a short time, but there are also waveguides to consider. With this small aperture, the time to write a complete set of devices becomes quite unfeasible. To solve this problem, we change the aperture during the write, a relatively complex process that has only recently been developed to a useful level. Again there is

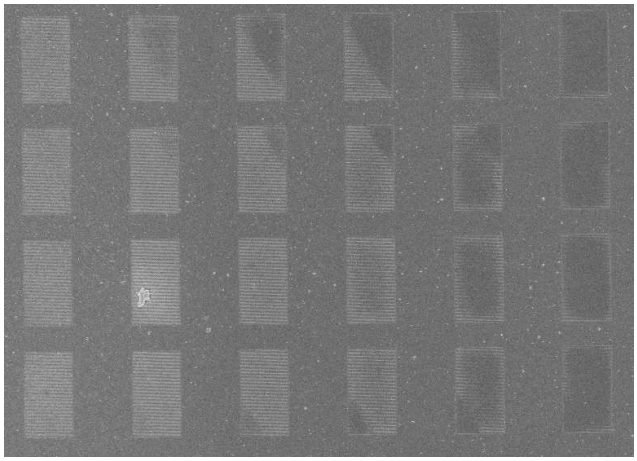

Figure 3.2: *SEM image showing the effect of discretisation. The dose is increased from left to right. In each block, we have a lattice of identical elliptical holes, which are increasingly rotated in each row and column of the block. In each of the four rows in the image, the rotation angle of the first hole in the block is increased. As a result, the over-exposed area moves for each row. This image conclusively demonstrated that the minimum step size achievable with a large aperture cannot be used to write holes with a controllable shape.*

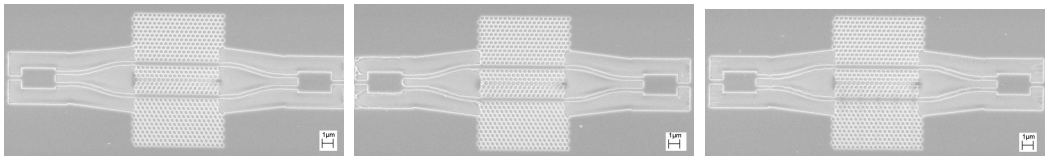

Figure 3.3: *A real device, which needs definition of hole shapes and hence very accurate lithography. The top and bottom halves contain PhCs with elliptical holes, with eccentricity 1.5 and rotation of 10^o and 50^o respectively. The ebeam aperture is increased from left to right. The definition of the hole shapes decreases markedly with the increasing aperture.*

a disadvantage; firstly that the system needs to be focused individually for the two apertures, and secondly that they need to be brought into alignment with each other. The former effectively doubles the set up time, and the latter can introduce errors as perfect alignment is not trivial. However, these are much less troublesome than the step size and writing time problems, so changing apertures is now commonly accepted as a viable compromise.

Stitching Moving on to the issue of stitching, which is part of the "complexity" mentioned above. Briefly, to write a pattern, you can define a small area called a writefield, which can be exposed by raster-scanning the beam. The area over which the beam can be rastered is mechanically limited, and also computationally limited. The computational part is trivial - a digital to analogue converter is used to convert between the digital file and the analogue motion of the beam, and the DAC only has a certain number of bits and hence quantisation levels. Hence, regardless of the current/ aperture issues, the minimum step size is related to the size of the writefield. The optimum for our system, considering many other variables, is 100μm - to write larger patterns, the stage must be moved. The stage position and motion are controlled via feedback from laser interferometers, with a mean theoretical error of 40nm. To achieve this, there are alignment routines that must be followed. However, this 40nm error is not typically seen in ridge waveguides that are written, because there are other effects that dominate.

It is necessary to pay attention to the speed of the beam's raster scan. It can happen, that for a given small step size, that the beam has to move faster than it physically can, and as a result the edges of write fields tend to get missed out. This means that the DAC changes faster than the beam can be deflected, particularly at the edges of the write fields. Instead, care

is taken to keep the beam speed below 5mm/s, the lower the better. For a small aperture, doing this requires increasing the step size above the minimum possible - so again, we use an aperture changing technique, using a moderate aperture and large step size (hence low beam speed) to write long waveguides, and a small aperture and small stepsize to write fine details. Long photonic wires, which need both fine control and good stitching are best avoided, unless an additional technique I have discovered is used: break all large areas up into many contiguous small areas! The software then adds a beam settling time to the exposure for each polygon, reducing the effective beam speed over the waveguide, and hence improved stitching, regardless of step size.

There is an additional problem that has been found, that for many years was thought to be a stitching error. Again, this is a software limitation. In my process flow above, one step is "leave for one hour, to allow thermal drift to stabilise". But, in practice there are two effects, the drift of the stage as it settles to the mean temperature of the system, which affects focussing and stigmation, and also the beam drift, which is continuous and periodic, with a period of several minutes. An example of the error is found in figure 3.4.

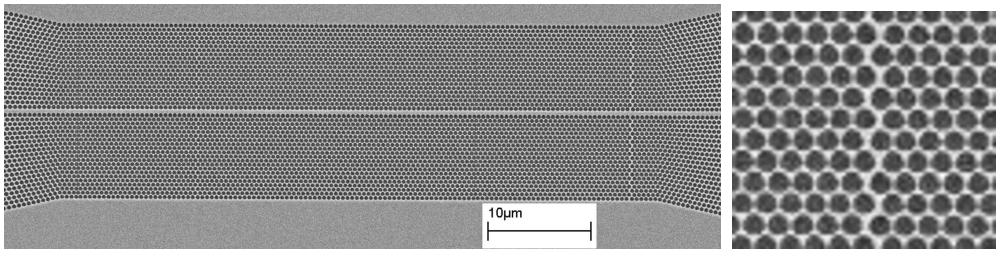

Figure 3.4: *Pseudo-stitching errors caused by out of order writing. The left hand image is a cutaway from the right, taken at the same scale but reproduced here at a higher magnification, to make the defect more obvious.*
The device pictured is a PhC with increasing crystal period at each end of a W1 waveguide, discussed in chapter 5.

This occurs when the PhC pattern is not written in a canonical order - so two adjacent rows of holes are written at different points on the bean drift cycle instead of continuously. This is a significant error, on the order of 200nm in the worst case - because a 5nm change in the position of a row of PhCs is used to create a so-called heterostructure cavity [30]. The error is actually caused by the memory limitations on the proximity correction process - as the whole pattern cannot be held in memory at once, the pattern order gets jumbled! Figure 3.5 shows the same device written correctly. Proximity correction is discussed in the next section.

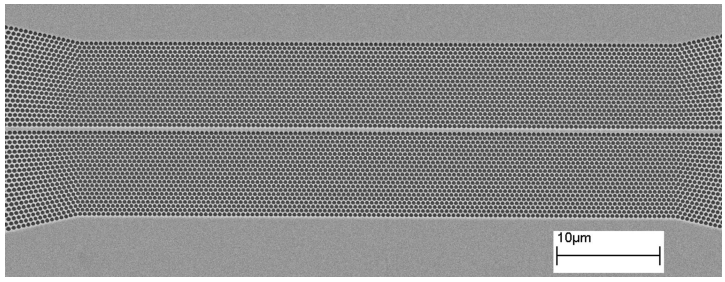

Figure 3.5: *The same device, with the holes written in canonical order. The beam drift no longer induces stitching errors. I designed these devices, but thanks to G. Pagnotta for writing the code which fixes the problem.*

An alternate method to remove all stitching errors is a technique known as shot-shifting which is also used to remove statistical disorder in the PhC lattice, but has not been not used here [31, 32].

3.4.1.3 Spots

For lithography with this system, the limiting factor on the techniques that can be used is the spot visibility. As explained in the process flow, to test the alignment of the system one exposes a single point, and then examines the result, to perfect the focus and stigmation. This spot then becomes

the smallest useful feature for write-field alignment, and as such the greatest accuracy in alignment is gained when using this spot. To use aperture changing techniques, this accuracy is necessary.

In figure 3.6, there are two images of these spots, one on SOI and one on InP. The much greater visibility of the SOI spot lends itself to complex writing techniques, but with the poor visibility InP spot, we are restricted to 20μm or greater apertures, and so cannot use aperture changing techniques. This is significant for the work in chapter 4.

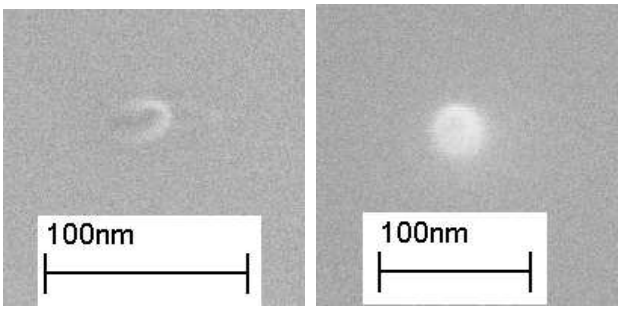

Figure 3.6: *Spots - on InP on the left, and on SOI on the right. The much greater visibility of the SOI spot makes writing on SOI easier, and also allows the use of multiple-aperture techniques. For the InP, we are restricted to apertures greater than 20μm, because with smaller apertures contamination spots are not visible.*

3.4.1.4 Dose to target

In this part, we briefly discuss the technique of dose to target, also known as adjusting the design file to get what you really want. For a fixed dose one reduces the physical aperture - the size of the feature on the mask, to get the target feature size. On first thought, this technique should not be necessary with ebeam lithography - after all, the dose can be directly controlled. However, it does prove to be very useful in many cases.

It all comes down to stray electrons. With a PhC lattice, backscattered electrons and other secondary exposure mechanisms give rise to a low level exposure away from the feature currently being written, no matter how carefully the dose is selected. This low level exposure is also developed, so a feature is always larger on a sample than on the mask. How much larger is related to the developing process, but as examining the results here is challenging, it is simpler to have a consistent developing process and control the hole size lithographically. To manage the backscattering effect, it is quantified as a background level for the nearest neighbour holes. In a triangular lattice, a central hole receives a background dose contribution from each of its six nearest neighbours, whereas one on the edge of a pattern may only have two contributions, and an isolated hole none. As a result, a proximity correction algorithm assigns a low dose to the central hole, a larger one to the edge hole, and larger yet to the isolated case. This is easy to understand, but to calculate the effects requires a very computationally intensive algorithm [33].

As a result of this computational expense, limits are placed on the fractioning of the pattern. A PhC hole is treated as a single element. In general, this is unimportant - every part of the hole has the same background dose. However, this does mean that holes of differing shapes are treated equally, and the resolution of hole size is limited. We can reduce the exposure dose and hence the hole size, or we can use a dose to target approach and reduce the exposed area for each hole. Although both are useful, experimentally it was found that the dose to target approach gave better control and repeatability. A further solution is to vary the mask thickness, using a thinner and hence more sensitive resist, which improves the resolution of the hole size. However, this impacts the etching as well as the writing process, and is best avoided where possible. After several iterations, it becomes possible to reliably write

the desired pattern.

This dose to target effect is also useful for isolated lines and one dimensional gratings, simply as over-exposure tends to give the smoothest possible lines. Photonic wire bends are similar. The effect of raster scanning on a bend is to make it rather uneven because of the square grid at the step size; so the width varies around the bend. To make a smooth bend, so the width varies around the bend. To make a smooth bend, over-exposure is necessary, and to make a smooth bend of the desired width, dose to target techniques must be used. Some examples of photonic wire bends are shown in figure 3.7.

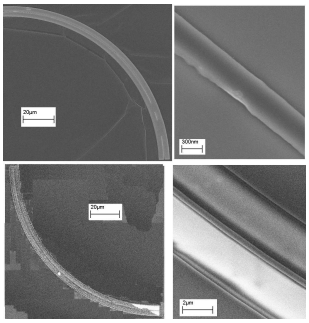

Figure 3.7: *A variety of bends. The first is written using a large step size, and critically exposed. In some areas, the bend closes up completely. The zoomed image shows the roughness of the bend. The second bend is written using an overexposure technique, and is both continuous and smooth.*

3.5 Pattern transfer and etching

We now have a design, and have defined the pattern in the chosen resist. The next step is to transfer the pattern into the hard mask (if used), and etch into the sample. We can again divide this section into heterostructure and silicon based processes, as these are significantly different. As with previous sections, the intent here is not the detailed discussion of chemical processes, but instead

to explain the parts where significant development work has been done, and where my processes provide substantial benefits.

3.5.1 Heterostructure devices

For this project, two significant heterostructure devices are involved, namely a PhC y-splitter et al. (chapter 6) on AlGaAs, and a surface grating coupler on InP (chapter 4). The techniques used in etching these two devices are fairly similar. In both cases, we use PMMA as the resist, and transfer into a hard mask using RIE (Reactive Ion Etching) in a fluorine chemistry, and then etch the pattern into the heterostructure using CAIBE (Chemically assisted ion beam etching) in chlorine chemistry. For the AlGaAs patterns our concern is to maximise the etch depth, whereas achieving exactly the target etch depth is important in the InP case.

The pattern transfer step is the most standardised of all the fabrication processes used here; and poor results are usually due to external effects. There are two limits to this, firstly that the RIE machine is used with other etch gases which can cause contamination that affects the etch recipes, and secondly mask thickness. Although the standard recipe is consistent, the selectivity is 1:1 at best. Using a thicker hard mask to get a greater CAIBE etch depth requires a thicker resist, and this lowers resolution in the pattern. The etching is done in CHF3 gas, taking about 15 minutes for 300nm of silica hard mask.

CAIBE is a more variable process, and requires continuous optimisation to achieve the desired results. This is particularly true for AlGaAs heterostructure, as the layers with high aluminium concentration are sometimes found to act as etch stops for recipes optimised for pure GaAs. Usually, one tries to achieve the maximum possible etch depth with vertically edged holes, and our machine is very successful at this [34, 35, 36]. To do this, one controls the

sample temperature, the velocity and number of argon ions for the mechanical etching component (via beam voltage and current)and the relative number of argon and chlorine ions available for etching. There are several regimes that can be optimised to give good results, some examples of which are shown in figures 3.8 and 3.9. For a detailed discussion of this, refer to M.V. Kotlyar's thesis [37], and L. O'Faolain's thesis [38].

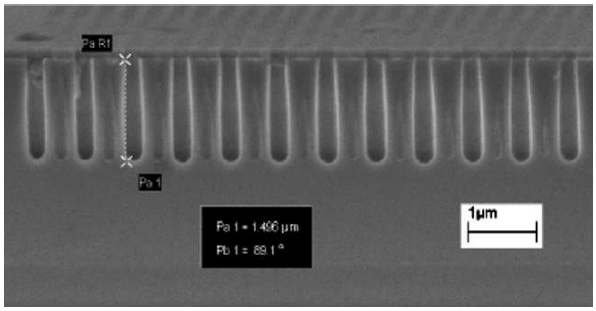

Figure 3.8: *Good quality AlGaAs etching in CAIBE for the Y-splitter devices in chapter 6. The etch depth and sidewall verticality are the critical parameters to optimise.*

3.5.2 Silicon etching

The silicon etching here is primarily SOI rather than silicon wafers, and the typical SOI is a 220nm Si guiding layer on 1μm buried oxide (BOX), over a silicon substrate. For PhC devices, Bogaerts et al. [28] have shown that simply etching holes into the top silicon layer, and neglecting the BOX, gives the best results consistent with the additional difficulty in transferring the pattern into the oxide. One can also make a silicon membrane, by using a wet HF process to undercut the PhC device. Although there are arguments for and against this, it has not been done in this project.

Silicon is etched in a fluorine chemistry RIE process, using a mix of CHF_3

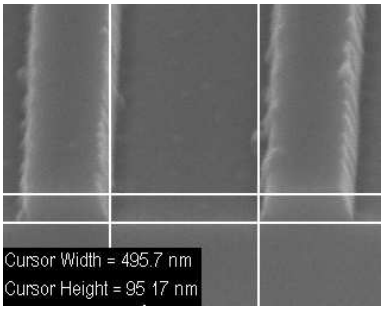

Figure 3.9: *Shallow InP etching in CAIBE, for an attempt to realise the 1D gratings in chapter 4. The etch depth in the figure is close to 95nm. In this case, we have to accurately hit a target etch depth while maintaining good etch-floor roughness and sidewall verticality. This is the best result I could achieve, which manages two of the three requirements. However, CAIBE is not well suited to shallow etching, and so we have changed to an ICP system for these devices. This is discussed further in chapter 4.*

and SF_6 gases. The SF_6 gas etches the silicon, but it does so in an isotropic fashion which is not very conducive to PhC holes with vertical sidewalls. CHF_3 is added to provide directionality; it reacts with the product gases to form a polymerisation layer, and so prevents etching outwards.

The challenge with silicon etching in this kind of process is in the masking. In the pattern transfer process, we can use 200nm of PMMA and etch for 15 minutes, which indicates that PMMA resists CHF_3 gas well [39]. Using a similar etch pressure but replacing half of the CHF_3 flow with SF_6, we find that the 200nm mask lasts less than a minute, and the silicon is only etched to around 100nm. All values given here are very approximate, as it depends strongly on hole size, and the accuracy of depth measurement is limited - the SiO_2 BOX layer is an insulator, so charging effects in the SEM reduce precision. By using over 600nm of PMMA, an etch depth of 200nm is achievable for holes of diameter greater than 250nm. Using this much resist makes it difficult to

make small holes, and difficult to accurately control hole size. Clearly, an alternative solution is desired.

Although several attempts were made to alter the etch chemistry, none were successful. Instead, the resist type was changed, from PMMA to ZEP, following the practice of several other groups fabricating PhCs in Si. This improves the situation somewhat, although the result still is not perfect. We can reliably etch holes of diameter greater than 200nm all the way through the top Si layer. For a typical target r/a of 0.27, this gives a minimum period of 370nm, which is acceptable for devices designed to operate at 1550nm, but only just. Further increases in resist thickness are not desirable, so work is in progress on further improving the process. An example of etching is shown in figure 3.10.

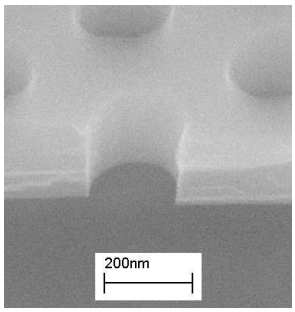

Figure 3.10: *RIE etching of PhCs in SOI, using a mixture of SF_6 and CHF_3 gasses. The aim is to etch through the 220nm top Si layer only, while maintaining sidewall verticality.*

With SOI fabrication, there is one other etch process of critical importance, namely substrate removal. This is necessary to achieve good quality cleaved facets, which are often the limiting factor in characterising SOI devices. This can be done either mechanically using an alumina paste, or chemically in a solution of KOH; both require a huge degree of care and attention in relation to the perceived simplicity of the process! However, the yield of the process

has improved with the experience of the group, by learning which details must be taken care of[3], and now only a few samples are lost.

3.6 Additional processes

As well as the processes described here, there are many others that are routinely undertaken, and that I have helped to develop, such as those for sample cleaning. These fall under the heading of tedious in the introduction to this chapter. However, there are additional interesting processes, particularly the wafer bonding used in making the InP grating couplers, which are very specific to particular devices, and as such are discussed in the relevant chapters.

3.7 Conclusion

In this chapter, I have discussed fabrication techniques used in making the devices in the other chapters, specifically showing the problems I have encountered, and solutions discovered by myself and colleagues. The most significant technologies, or at least those that I could most influence, are the ebeam lithography used to define the device patterns, and the dry etching used to transfer these patterns into the various material systems used. I have done lithography and etching on InP, GaAs, and SOI platforms, each of which presents their own different challenges. Both my own work and that of my colleagues towards overcoming these has been presented. With due care and attention, we have the tools to fabricate almost any PhC structure designed to operate at telecommunications wavelengths, that is for lattice pitches $> 350nm$. However, I have also shown that this is not trivial, and

[3]for example, in KOH the sample must be level in the solution!

that the realities of fabrication must be considered at an early stage in the design process. There are many further SEM images of devices that have been successfully fabricated throughout this thesis.

Chapter 4

Coupling into microphotonic devices

4.1 The coupling problem

In previous chapters, we have seen that PhCs have the potential to form useful devices, and have examined techniques for modeling and fabricating PhC structures. Now we will consider the use of these devices in combination with current technologies, particularly optical fibre based systems. Although the small footprint of PhC devices is attractive in terms of integration, especially with CMOS technologies, from an optics perspective isolated PhC devices present a significant challenge. This is illustrated in figure 4.1.

In this figure, the respective profiles for the fundamental modes of standard single mode fibre and a typical SOI photonic wire (500nm $\times$ 220nm) are shown, to scale. The wire mode displays wings, because it is not fully confined within the high index material, and it is highly elliptical compared with the perfectly circular fibre mode. Hence we have two distinct coupling problems, one of size and one of shape.

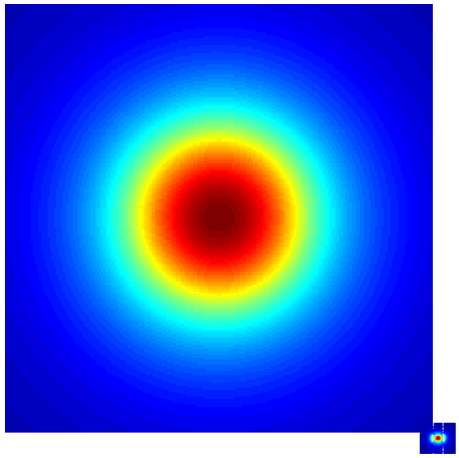

Figure 4.1: *Scale comparison of a standard optical fibre mode (left) and an SOI photonic wire mode(right). The difference in mode size is clearly apparent, but notice also the shape.*

PhC modes, particularly when gap-guided, have different profiles again, both transverse and longitudinal. This will be discussed later. We approach coupling as three distinct problems, mode size, mode shape, and effects of periodicity. This aids understanding of the problem while optimising the use of limited computational resources.

In this chapter, I will present the various approaches in the literature used to solve the mode-size coupling problem, before focusing on the approach we use as part of the EU Funfox collaboration, namely surface grating couplers (GCs). Clearly this size problem is not specific to PhC devices, and has been widely studied. I will briefly review the theory behind grating coupler design, discussing relevant modeling work; then device fabrication and characterisation. Within this collaboration, design and characterisation are the responsibility of F. van Laere, our collaborator at Universiteit Gent, Belgium, whereas my role is in lithography and the limits it imposes on

design [1]. However, because these grating couplers fit so well both in terms of understanding PhCs and my focus on coupling, I have studied the design, and so much of the work presented here is my own, albeit perhaps duplicated. For a detailed discussion of the purpose and design principles behind surface grating couplers, the reader is referred to the PhD thesis of D. Taillaert [40].

The reader should also be aware that the majority of work on this kind of GC is on an SOI platform. However, within Funfox, we use Indium Phosphide material. The particular challenges of working with this material are also discussed in this chapter. InP is desirable for the active devices considered elsewhere in the project, and we aim to integrate the utility of GCs with these.

4.2 Fibre couplers

As discussed above, figure 4.1 shows the disparity in size between a fibre mode and a photonic wire mode, and we wish to couple a significant fraction of the light between these two modes. There are two conventional approaches - firstly, apply an anti-reflection coating to the facets of the photonic wire die, and secondly use an intermediate coupling step. AR coating is a well known technology, and substantially increases input coupling but does not compensate for the mode size. Possible intermediate steps include the use of a high power lens, drawing the fibre into a taper, or creating a taper on the die. Of these, the first two can be highly effective, but are not suitable for mass production, due to the high skill level required and low repeatability. Fabricating a taper on the die is a better solution here, and a variety of methods have been published by various groups.

[1] Etching and wafer bonding are largely performed at Universiteit Gent.

The ideal solution is to taper the guiding layer transverse to the propagation direction, in two dimensions, up to a suitable size[2]. However, this would require specialist wafers as well as greyscale lithography/etching. This has been studied in detain in reference [41], which proposes a modified and relatively feasible solution, although not compatible with current standard wafers. Instead, it is relatively simple to taper a waveguide in plane, using either an adiabatic (hundreds of microns) taper or an optimised short taper (chapter 5 and [34]). For the low out-of-plane contrast substrates initially popular for PhC devices, this is a reasonable solution as the mode can also expand vertically to some extent; but for membrane materials such as SOI the waveguide mode becomes highly elliptical. Further, in both cases broad (several micron) waveguides are significantly multimoded when etched to the typical extent.

A more effective, albeit more complex solution has been demonstrated by both NTT[42] and IBM[43]. This is shown in figure 4.2, taken from reference [43]. Instead of increasing the width of the SOI wire, the width is reduced gradually to the minimum that can be fabricated, typically < 100nm. Then, a second lithographic step is used to define a polymer layer over the end of the taper. The facet is cleaved through the polymer. This structure is known as a "tips taper". Although the input coupling is significantly improved, this method is not commonly used even where the process has been perfected, simply because of the inconvenience of the second lithography.

[2]This has been patented by Bookham Technologies

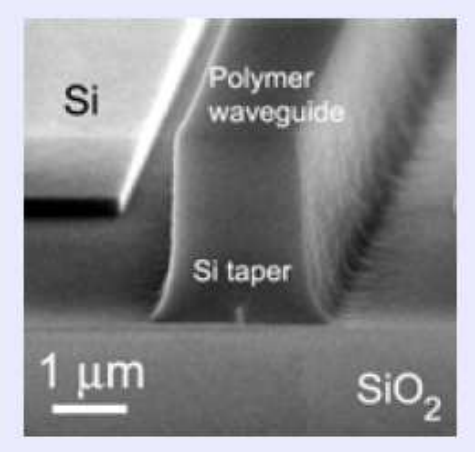

Figure 4.2: *Tips taper with polymer overlay, from [43]. To improve coupling between standard single mode fibre and micro-photonic waveguides, an intermediate step is used, as shown in this SEM. The waveguide is narrowed to the sharpest possible point, then a polymer layer is defined over the point. The polymer approximately matches the fibre mode index and size, then the mode is slowly coupled down into the silicon waveguide.*

4.3 Surface grating couplers

Instead of these edge coupler approaches, in this chapter, we will consider the use of a photonic crystal as a surface, as opposed to edge, coupler. Aside from the benefits of greater coupling efficiency which will be shown, GCs have two major advantages over edge couplers. As a cleaved facet is no longer necessary, wafer-scale testing of photonic devices becomes possible, which is significant for industrial processes. The other advantage is more subtle. Most existing optical networks are transparent to polarisation, i.e. signals have arbitrary, and potentially varying polarisation states. However, the asymmetry of typical photonic wires means they have different responses to TE and TM polarisations. PhCs are even more polarisation-dependent - the lattices of holes considered here do not even exhibit a bandgap for TM light! On the

other hand, a 1D grating will only couple light that is transverse to the grating into a waveguide, and it will do so into a TE-like mode. Using a square lattice PhC essentially gives us a superposition of two orthogonal 1D gratings. Each of these will couple as before, so we can capture TE and TM light from the fibre and couple it into the desired TE polarisation state, in two separate waveguides. This is illustrated schematically in figure 4.3. This concept has been developed by D. Taillaert [40].

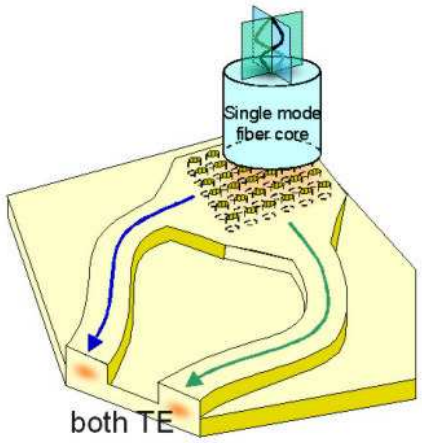

Figure 4.3: *Schematic of a "'polarisation diversity coupler"', taken from [40]. The orthogonal polarisations in the fibre couple into orthogonal waveguides, but with the same polarisation state in each.*

4.3.1 Grating coupler theory and device design

In the limits of spatial extent and index contrast, a grating is a PhC and a PhC is a grating. A compact grating coupler in a high index contrast system, however, is an admixture of the two. We will approach the design of this device using conventional grating coupler theory, but we will apply it using the numerical tools favoured for photonic crystals discussed in chapter 2. The other devices in this thesis are approached from an almost entirely numerical

basis, so the relevant grating theory is briefly introduced here. For the fine details, I again refer the reader to [40].

We begin by considering light incident on a plane dielectric boundary, and recall Snell's law, depicted in figure 4.4.

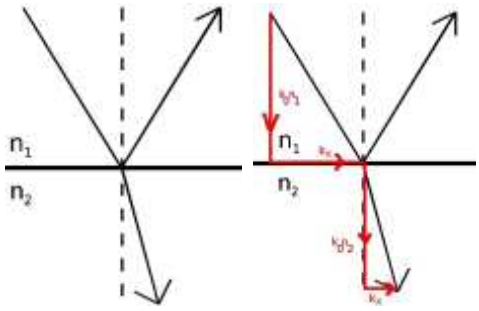

Figure 4.4: *Snell's law in simple form on the left, and with k-vectors superimposed on the right.* $n_2 > n_1$

Traditionally we write this as $n_1 sin\theta_1 = n_2 sin\theta_2$; but in k-space we can write:

$$\mathbf{k}_0 n_1 = \mathbf{k}_0 n_2 \tag{4.1}$$

Hence the component of the k-vector parallel to the dielectric interface is conserved, as shown in figure 4.5. Because we now sketch in k-space, this is more obvious than in the previous figure. The vector fields are also shown.

Then, we can introduce a periodicity, and hence the Bragg condition again:

$$\mathbf{k}_0 n_1 = \mathbf{k}_0 n_2 + m\mathbf{G} \tag{4.2}$$

where m is an integer, $|\mathbf{G}|=2\pi/\Lambda$ and Λ is the grating period. This reduces back to Snell's Law for $\mathbf{K}=0$. As a result, we get propagation at various angles in the second medium as shown in figure 4.6.

Cast into waveguide terms, we get

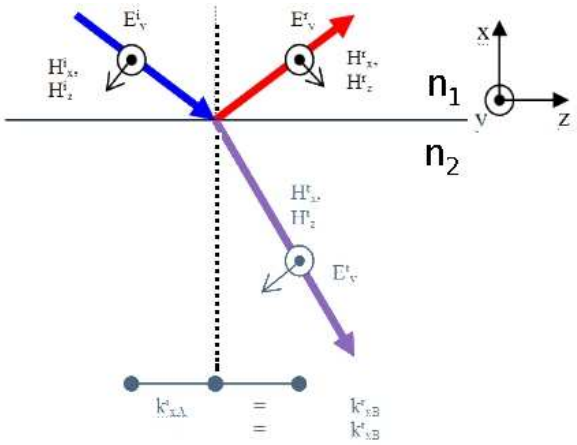

Figure 4.5: *Snell's law again, this time in k-space and indicating k-vectors. The incident medium n_1 is lower index, and n_2 higher index, which implies that the k-vector in n_2 must be longer to conserve momentum.*

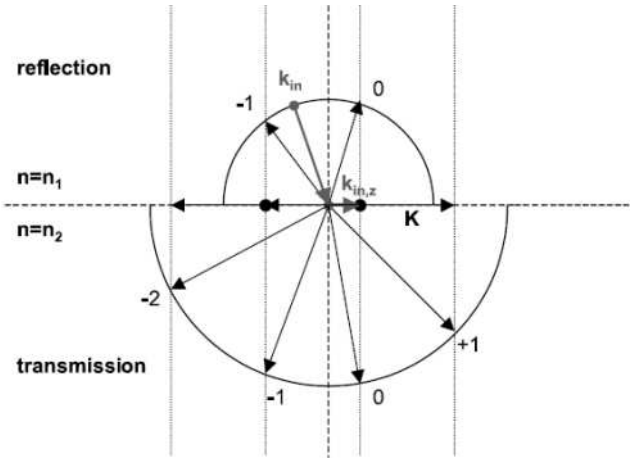

Figure 4.6: *Imposing a periodicity on the dielectric interface allows propagation at various angles, both reflected into n_1 and transmitted in n_2. The particular angles depend on the grating period and the incident wavelength. From [40].*

$$\mathbf{k}_0 n_1 = \mathbf{k}_0 n_2 + m\mathbf{G} \tag{4.3}$$

If we now introduce another (non-periodic) boundary beneath the grating, some of these angled-propagation rays are reflected back up, and therefore undergo total internal reflection and become quasi-guided modes. Figure 4.7 shows such a mode. For a single wavelength, the possible propagation angles are discrete. We have air above the grating, and an angle of incidence perpendicular to the phase fronts shown in the top of the figure. The grating couples this into the propagating mode along the centre. From the grating, the light can then again radiate, either back up into the air, or down into the substrate. Here, k_x and $k_0.n_2$ are conserved, so for a higher index substrate (silica with n=1.45 in the figure) the propagation angle increases. Once beyond the end of the grating, the quasi-guided mode couples readily into a fully guided waveguide mode. The provenance of this figure will be explained subsequently.

As explained in the introduction to this chapter, we use an Indium Phosphide membrane material system - InP(m). The layer structure and refractive indices are shown in figure 4.8. The grating is defined between the InP and BCB (Benzo-cyclo-butene, a transparent resin) layers, rather than on the top surface, for processing reasons, which are discussed in the relevant section. The fabricated devices are tested using telecoms range laser sources; for which a grating period on the scale of half a micron is suitable. The equations above show that we have three independent variables once the wavelength is fixed: the grating period Λ, and the duty cycle and etch depth, which control the index contrast. With these three variables, there are many possible solutions to achieve any desired grating operation, but the important part is to find a stable solution with respect to fabrication tolerances. To this end, we will fix the duty cycle at 50:50, and the etch depth at 120nm to find

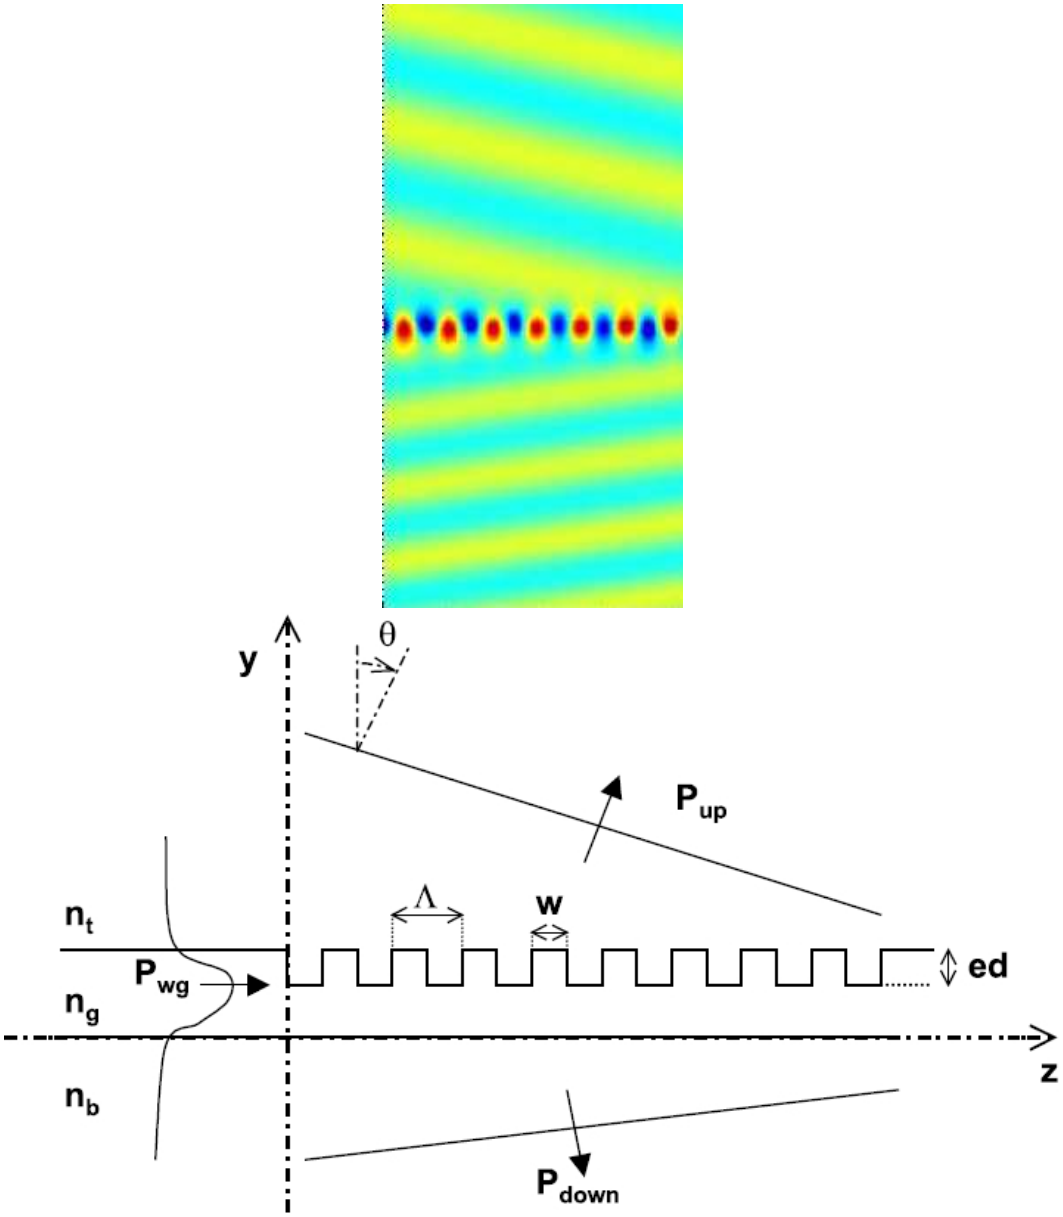

Figure 4.7: *Top: A grating coupler in operation. This was calculated using FDTD in a supercell technique, and then multiple periods created by rotating the phase appropriately - so there is no loss across the grating. However, frequency domain tools are also appropriate here. Bottom: Schematic of GC operation, with the input signal coupling up and down from the grating.*

a design[3].

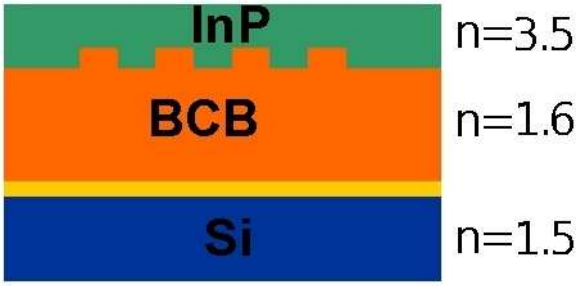

Figure 4.8: *InP(m) layer structure. The thickness of the top InP layer is fixed at 300nm, but the other layer thicknesses are variable. The yellow layer is a high reflectivity film which is used in some device designs. The substrate is typically silicon, but GaAs is used where cleaving of the bonded structure is required.*

We now input a unit cell of this structure into a bandstructure solving program. The high reflectivity layer and silicon substrate are neglected. Any of those discussed in chapter 2 would be suitable, but I have chosen FDTD in RSoft's FullWAVE in this case, with periodic boundary conditions. For a problem of this nature, we must consider the real thickness of the InP layer, so performing the calculation in normalised units is not helpful. Instead, to maximise accuracy, we must start with a suitable period, or else iterate the calculation. We choose a grating period of 660nm[4], and calculate the bandstructure shown in figure 4.9. The red zigzag shows the grating mode, the blue line is the air (upper cladding) lightline and the green line the BCB substrate lightline. Because the contrast is low, the bandgap is not apparent

[3]For the actual design by F. Van Laere, numerical optimisation methods and CAMFR[44] have been used to find an optimum combination of the three parameters, but here I wish to present a simple approach, to highlight the design choices. The optima are slightly different to the parameters stated.

[4]Obviously a fix, as I have the final solution from my collaborator.

in this simulation - but theory demands its presence - remember that gratings are used as very narrow line width filters, too! However, this is not important; as we are not trying to create a filter, it is better to stay away from bandgaps.

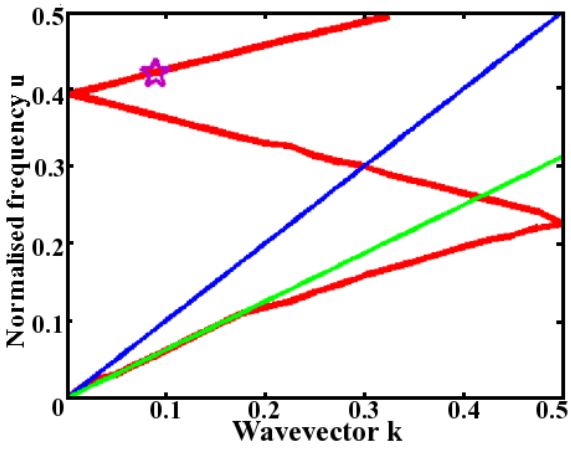

Figure 4.9: *A dispersion diagram for a grating with 50:50 duty cycle, 660nm pitch and 120nm etch depth in our InP(m) system. The red line is the guided mode, which folds back on itself; although theory demands that a bandgap open up at the folds, the gap is too small for the resolution of the simulation. The blue and green lines are the air and substrate (BCB) lightlines respectively, and the star indicates our chosen operating point of 10º from vertical (Γ point) emission.*

Clearly, we need to operate above the substrate lightline, because surface coupling requires the grating to radiate out of plane. If we work at the Γ point $k = 0$, then the out of plane radiation is vertical. In this case, light incident on the grating can couple equally into forwards and backwards propagating modes of the slab waveguide. Instead, we choose a 10^o angle of incidence, and tilt the fibre appropriately. This is shown in the schematic 4.10. Note that this will have consequences when we move onto 2D, polarisation-diversity GCs.

Using $ka/2\pi$ for the design period and wavelength, $10^o \equiv k = 0.9$. The first fold of the bandstructure has already been ruled out as non-radiating. The

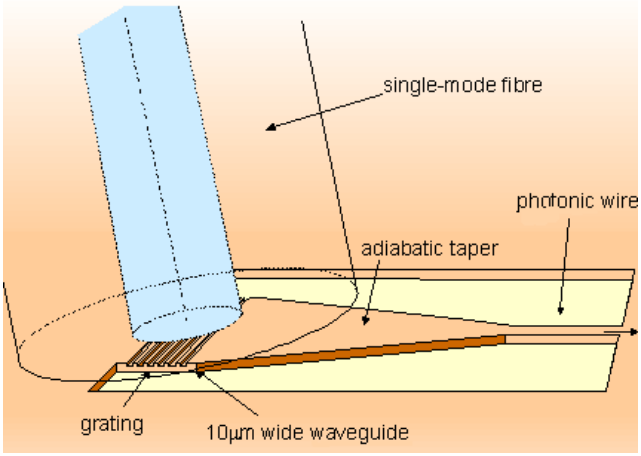

Figure 4.10: *A schematic of a GC in operation, showing the fibre being tilted to ensure coupling in the forwards direction only.*

second fold radiates backwards with respect to the propagation direction of the system. Intrinsically this is an acceptable solution, but inconvenient for testing practical devices - a simple device, comprising 2 GCs and a waveguide between them, needs two fibres for characterisation, and these would intersect in this geometry. Instead, we choose the third band, and the operating point marked by the star in figure 4.9, which permits the characterisation scheme in figure 4.10. The required grating pitch can then be read as 660nm from the abscissa of the dispersion diagram. If the required value was more than a few percent different to the value used in the bandstructure calculation, the simulation should be repeated with the new pitch, as discussed previously[5]. To calculate the field profile previously shown in figure 4.7, we change our FDTD simulation to contain several repeats of the supercell (or cycle the field data in another program), launch a sinusoidal input at the target frequency, and then capture the field once the simulation has run to a steady state. Again, there are many other appropriate methods.

[5]however, the fix is in.

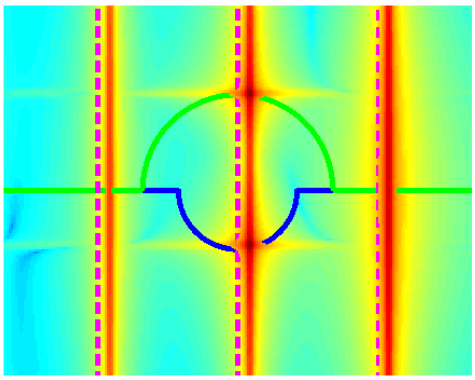

Figure 4.11: *Spatial Fourier transform of the field profile in figure 4.7, which is a single frequency calculation. Superimposed are the semi-circles which show the allowed wavevectors of the grating - compare figure 4.6. The differing radii of these circles are due to the different media above and below the grating. The dotted magenta lines are also superimposed, and represent the Brillouin Zone boundaries of the grating. Everything else in the image results from the spatial Fourier transform - the bright red lines are the quasi-guided mode, repeated in other Brillouin Zones due to the periodicity of the grating. The dark red/black spots are the plane waves above and below the grating, and their intersection with the quasi-guided mode demonstrates the coupling.*

It is instructive to take the spatial Fourier transform of this field profile, as shown in figure 4.11. The magenta dotted lines and green and blue semi-circles are super-imposed. The dotted lines indicate the boundaries of the Brillouin zones, which is equivalent to saying the half-widths of the supercells in k-space. The green and blue semi-circles are the allowed wavevectors for this grating for varying frequencies (compare the schematic figure 4.6). Of course, the wider semi-circle belongs to the air modes, and the narrower radiation into the substrate. Everything else in the image is contained in the SFT, and the horizontal breadth of the features is a consequence of the Fourier

transform of a finite number of repeats of an infinite grating. The red vertical lines are the waveguide-like mode. As discussed in chapter 2, a translation-invariant waveguide does not have Brillouin zones, but the periodicity of the grating folds the mode out. The horizontal yellow lines are the plane-waves propagating above and below the grating, obviously lying on their appropriate semi-circles. But, these are not repeated outside of the irreducible Brillouin zone. The significant point of this diagram, however, is in the dark red/black spots, which show the point at which the plane waves couple to the waveguide mode.

Finally, we consider the design of a 2D grating for use in the polarisation diversity approach. For this, we need to have a 3D supercell for the calculation of the bandstructure, and either a 3D model or a clever combination of 2D models for numerical optimisation. The 3D model is possible in FDTD, but the large volume that would need to be simulated is prohibitive even before considering the many iterations necessary for numerical optimisation. CAMFR [44], the program used for the real designs, is limited to 2D in its public release. Instead we use a mixed strategy. We choose to keep the same lattice period and air-filling factor as in the 1D designs, so a 50% duty cyle becomes a normalised hole radius r/a of 0.4. Optimisation in 2D from this gives a design with a lattice period of 640nm, 50% duty cycle and an etch depth of 110nm.

4.3.2 Improving the efficiency of the designed coupler

In the previous section, I have explained the design of an appropriate GC. However, the coupling efficiency has not been discussed. There is a good reason for this - published GC designs often choose to calculate the efficiency of the grating in a form to suit their results. To add rigour, R. Baets' group at Universiteit Gent have adopted a standard method of calculating coupling

efficiency, as the overlap integral of the output of the coupler device with the mode profile of single-mode fibre. Many groups cite their grating efficiency as all of the light which is coupled out; but this is not appropriate for fibre-coupler designs.

The substrate was neglected in the bandstructure calculation, as it is not relevant to the simple operation of the GC. However, for the real device efficiency, the substrate, and particularly the thickness of the BCB layer are essential to the calculation. I will now discuss the parameters critical to improving the efficiency of the device, and state efficiencies for some optimum choices. In the absence of a substrate, the efficiency is of the order of 30%.

Figure 4.7 clearly shows a major cause of this poor efficiency, which is technologically demanding to correct in real devices. The propagation into the substrate constitutes a loss mechanism for both input and output GCs. In principle, this is a simple problem - place a reflector of some kind at a distance of $\lambda/2$ below the waveguide. In theoretical calculations for a simple SOI grating, Taillaert [40] shows an increase from 55% with an optimised insulator layer thickness, which gives Fresnel reflection from a single interface, to 82% with a two-pair Bragg stack, which shows just how important this is! One of the challenges with our polymer membrane is controlling the BCB thickness to this optimum, as will be discussed under device fabrication. Further, we add a gold mirror between the BCB and the substrate, to create the high-reflectivity layer shown in figure 4.8. This possibility is a major benefit of the fabrication methods employed.

Although an output angle of 10^o has been chosen, in some cases it is more convenient to use vertical fibres, particularly where packaged devices are envisioned. An alternative solution to prevent coupling to backwards propagating waves is to add a rear reflector to the grating, ideally a facet

etched all the way through the guiding layer. Again, though, this is demanding, because the spacing between the end of the grating and the reflector is phase-critical. Alignment tolerances demand that this be done in one etch step, and we cannot control two different etch depths from one run[6]. Instead we can use a shallow-etched grating, with a different period such that it reflects. We then need a longer reflector. With either of these techniques, however, the bandwidth of the coupler is reduced, so tilting the fibre is the preferred method.

Finally, we consider the longitudinal profile of our radiation. Ideally, this should be identical to the emission profile of the fibre, to maximise coupling. To do this, we study the propagation loss (to the radiation modes) across the grating, and then apodise the grating (by changing period, duty cycle or etch depth across the grating) to optimise the overlap with the fibre mode.

To summarize, there are three simple causes of intrinsic inefficiency in simple gratings, namely radiation into the substrate, back reflection, and poor overlap with the fibre mode. Although the methods of solution have been proposed above, the detailed design requires numerical calculations and optimisation algorithms. Note that the periodicity of the device remains a significant, and so minimisation of the numerical problem can be enhanced by using software that can take advantage of this. CAMFR [44] has been designed to solve grating problems such as these, but in general eigenmode expansion techniques are preferable over the FDTD I have used for the design so far. Using CAMFR for SOI and implementing all of the improvements discussed here,i.e. bottom mirrors, rear reflectors, and apodised gratings, Taillaert's best GC design has a theoretical coupling efficiency $> 95\%$.

[6]Etch depth is always proportional to feature size, so making two different etch depths is simple. It is the precise control that is not possible.

4.4 Fabrication considerations

Using the techniques and tools discussed previously, we can make a theoretical design for an almost perfect GC. A reflective layer in the substrate is the single most important addition, and the positioning of this is important, as we need to achieve the Bragg condition in the vertical direction. Then, we must apodise the grating to maximise the overlap with the fibre mode, and add a critically spaced rear reflector to prevent backwards coupling.

In theory, all of this is relatively simple. In this section, I will discuss the steps that have been taken to find the current working surface coupler design used in the project, outlining the fabrication steps that have been determined. I will also discuss the planned improvements for the next generation of devices.

4.4.1 Lithography

Firstly, I consider the GC layout, the top view. I have simply tried to fabricate the simple grating that has been selected as the basic design, with a grating period of 660nm, and a fill factor of 50%, to get a handle on the techniques to use and the limits of the fabrication processes. The best first result is shown in figure 4.12.

The result is very different to the design, particularly with regard to the line width /duty cycle. As discussed in chapter 3, this pattern has been proximity corrected using software optimised for periodic structures, hence apodised gratings would be significantly more challenging - hence we do not try to make them at this stage. By reducing the resist thickness and the exposed linewidth, we can improve the uniform grating significantly. The next figure 4.13 shows the optimisation pattern used to find the parameters for 1D and 2D GCs.

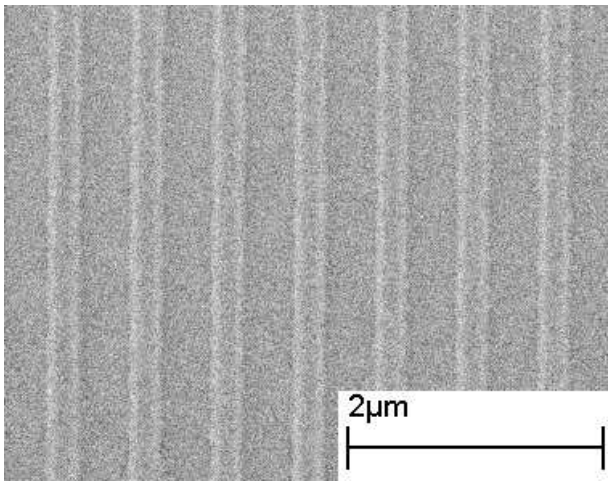

Figure 4.12: *A uniform grating, with design period 660nm and duty cycle of 50%, for first guess fabrication parameters. The pattern is only etched into the hard mask, hence the poor contrast and image quality.*

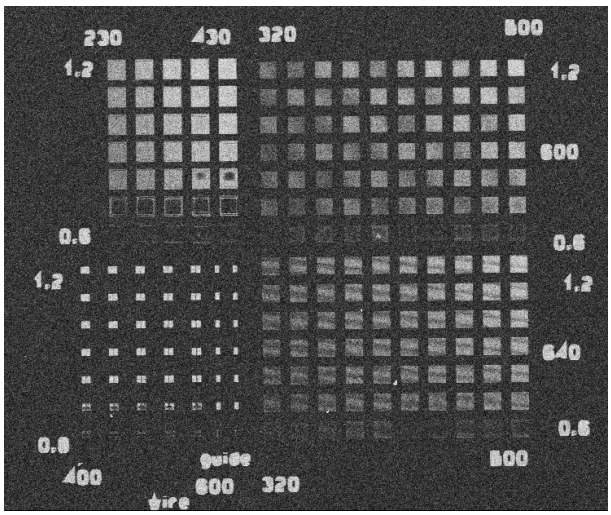

Figure 4.13: *Optimisation pattern, for 1D grating (top left), isolated line (bottom left), and high and low period 2D gratings (right). In each block, the writing dose is increased from left to right, and the gdsii-file feature size increased from bottom to top. By examining these carefully, we can determine which sets of parameters lead to the desired features.*

We repeat this three times on separate samples, as we also need to consider the repeatability of results. These are shown in figure 4.14. Although only one pattern from the optimisation layout is shown, it was necessary to characterise all of them.

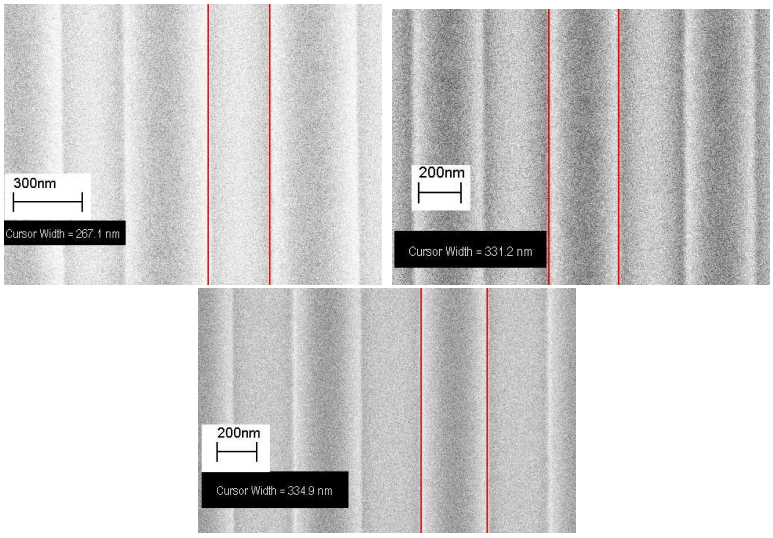

Figure 4.14: *Three repeats of the target grating, fabricated using identical parameters but on separate samples. This illustrates the effects of unavoidable, uncontrollable process variations.*

As before, the contrast makes it difficult to accurately determine the edges of the line, and hence its width; these are estimated to be ±20nm. However, the variation between the three repeats shown is greater than this error. Based on this result, together with all of the other devices in the set, we choose a dose factor of 0.8 on the grating, and a line width of 280nm to achieve the target 330nm. The terms and mechanisms are explained in chapter 3.

Perhaps more importantly, these results demonstrate the difficulty of making a uniform grating reliably. For the device set that is needed to completely characterise the operation, the total writing time is over 24 hours, so we want a robust design. Therefore, we choose to neglect apodising the

grating entirely - this can be simply added once a reliable optimized process for uniform GCs has been found.

As implied above, we use the same approach to determine writing conditions for 2D GCs. One effect is seen that does not occur for 1D gratings, shown in figure 4.15. The reason for this effect is intrinsically related to the ebeam lithography. For our chosen grating with 640nm period and 0.4 r/a, the separation between holes in the lattice direction is 128nm. For the triangular lattice used for the other devices in this thesis, typical parameters of period 400nm and r/a =0.3 give a minimum hole separation of 160nm, and this 30nm difference is crucial. To achieve good definition of the holes in the triangular lattice, I typically use an aperture of 10μm and a 6nm step size. However, this is for SOI, a good material for lithography. For the InP used here, we have to use the larger 20μm aperture, which passes a larger current for the same step size, and hence the feature resolution and controllability is significantly reduced. This is explained further in chapter 3 and [38]. Furthermore, lattices with high filling factors are intrinsically difficult in ebeam lithography. This has been studied by R. Wuesst at ETH Zurich, using a similar ebeam system to ours, and figure 4.16 shows his results. Essentially for large holes, the beam width and local proximity effects between one hole and its neighbours[7] make the process window in which a high fillfactor hole can be accurately fabricated very narrow[33]. Considering the repeatability shown for the 1D gratings, we simply have to accept this, and try to err on the side of holes smaller than the target size. Holes which become too large connect, and then our lattice of holes becomes a lattice of rods. As all of our modeling is for holes, this is undesirable. As with the 1D gratings, we will try to improve the yield of this

[7]The nanopecs proximity correction software used at St Andrews does not take the size or shape of holes into account!

process before considering apodising the gratings.

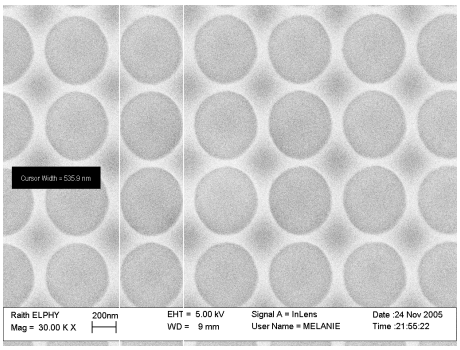

Figure 4.15: *Lattices with high air filling factor (hence r/a) are very sensitive to minor variations in the fabrication process and especially to astigmatism in the beam. Stigmation effects enhance deviation from the round hole, leading to breakthroughs between adjacent holes.*

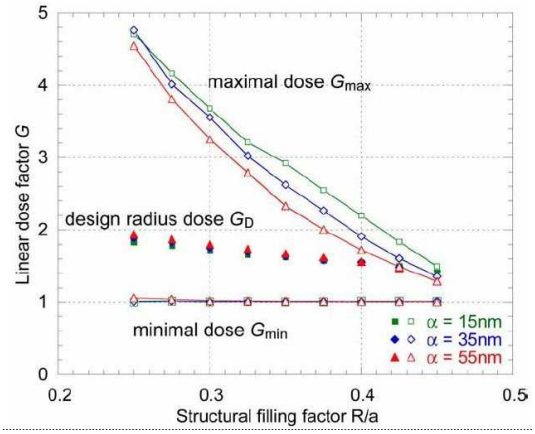

Figure 4.16: *Courtesy of R. Wuesst, ETHZ. The separation between the minimum and maximum dose lines is the process window in which the desired hole size can be successfully fabricated. α is a parameter used to model the ebeam beam, broadly controlled by the aperture for our system. The important result is, that fabricating holes with a high r/a is challenging and hard to repeat.*

4.4.2 Etching

The next task in fabricating our designs is to consider how to realise the required etch depth. This has been stated above as 120nm for 1D gratings and 80nm for 2D gratings, but the essential part here is control. As discussed in chapter 3, our customary etching processes for III-V materials have maximising depth as the guiding principle, so this is a new challenge. Etching InP to a high quality finish is well known to be demanding [37]. For the requirements of this project, it was not possible to find a good combination of process parameters in our CAIBE, perhaps unsurprisingly as CAIBE is a method for enhancing etch depth and aspect ratio over conventional IBE. Instead, a process has been optimised on Universiteit Gent's ICP system, which is much better suited to shallow etching, with results shown in figure 4.17.

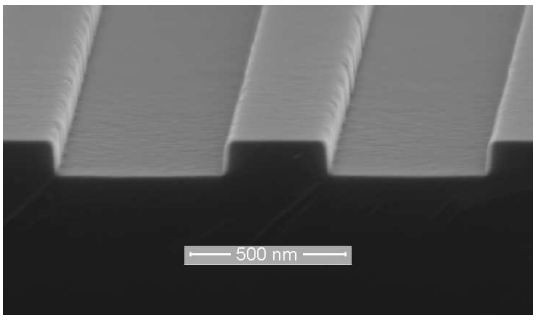

Figure 4.17: *Shallow ICP etching of InP. It is necessary to control the etch depth for the GCs very tightly, as this alters the operating frequency, while retaining vertical sidewalls and smooth floors.*

4.4.3 Wafer bonding and high reflectivity layers

In figure 4.8, I have shown the material system used in this project, and mentioned that this choice of system has implications for the design, particularly with respect to bottom mirrors. However, this membrane of InP

is not a conventional substrate for integrated optics. In this section, I will discuss both the choice of this layer structure, and how it is fabricated.

As discussed in the introduction, InP is desirable for active devices at telecoms wavelengths, and we wish to integrate active and passive functionality. However, most previous work on compact grating couplers has been undertaken on SOI, a membrane material. The high contrast between the guiding layer and substrate is necessary to the formation of a guided mode under the grating, but conventional InP devices are of the buried heterostructure type, with low out-of-plane contrast defining the guiding layer. To resolve this discrepancy, we would ultimately design a wafer with two guiding layers, one on the surface, with a sacrificial layer beneath, and then the buried heterostructure lower down. We then use wafer bonding as discussed in the next paragraph and selectively etch away the wafer so that we are left with a membrane where required for fibre coupling, some intermediate coupling region, and the entire buried heterostructure for active devices. For GCs, then, we need only concern ourselves with the top guiding layer.

SOI wafers are made using the "smartcut" wafer bonding technique, which is a direct wafer-to-wafer method: Two silicon wafers are oxidized such that the total oxide thickness is equal to the desired buried-oxide layer. One of the wafers is implanted with hydrogen ions, at a distance below the Si/SiO2 interface corresponding to the top silicon thickness. The two oxidized surfaces are cleaned and smoothed to the atomic level, and then brought into proximity. The oxide layers chemically bond to each other. Finally, the conjoined wafers are heated, which both anneals the oxide bonds and causes the implanted wafer to cleave along the plane of the implantation. This is a demanding technology! To insert an in-plane DBR stack, as proposed by Taillaert [40] one repeats this process several times, for the required Si and SiO2 thicknesses.

Figure 4.18 shows the alternate wafer bonding technique, which we use to create our InP(m) [45]. Although direct wafer bonding can be successfully used with InP as one of the wafers, we use the less demanding adhesive bonding method. One begins by preparing the two wafers to a flatness better than the desired adhesive thickness (which can be several microns for some applications - and hence permits the use of a patterned substrate). Then spin-coat the substrate with adhesive, stick down the top surface of the guiding wafer, then cure the adhesive [8]. The unwanted substrate of the guiding wafer can then be removed chemically or mechanically. Modifications to this process to include a bottom mirror are simple - for example, we can deposit gold on the top of the substrate - however the cost is that this process is much less well controlled than the hi-tech smartcut. We use BCB as the adhesive for its planarising properties, but other choices are available. Additional benefits of BCB include low refractive index, low curing temperature, and good resistance to solvents. This process has been developed at Universiteit Gent [46], and recently adapted to include bonding in vacuum, which increases the control of BCB thickness.

4.4.3.1 Bonding process flow

The flexibility allowed by this adhesive bonding process gives rise to a few possible variations on the process flow for our GCs. The main question is whether to define our pattern on the host wafer - at step *a* in the schematic flow - and then bond the etched devices, or instead bond the wafers and then do pattern definition and etching - i.e. at step *e*?

Considering first the bond-first procedure, the bonding process is easier with unpatterned substrates, as the gratings can house air-pockets, which are detrimental to device performance and may cause the bonding to fail. The

[8] Plus a selection of pre-curing, soft-baking, and temperature-ramping steps

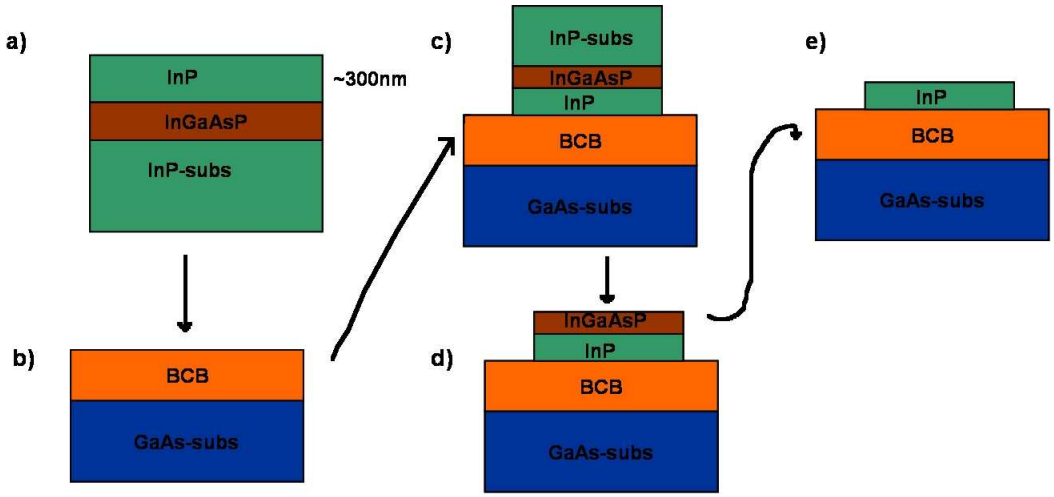

Figure 4.18: *The process flow for creating our InP(m). The InGaAsP layer is sacrificial, as it chemically etches in different acids to InP. We begin with two wafers, one as a "handler wafer" substrate, and the other the epitaxial InP with the top guiding layer. The adhesive BCB is spun onto the substrate, then the guiding layer stuck down. The substrate of the epitaxial material is removed mechanically, and the sacrificial layer chemically, to create a smooth surface. The green layers are stochiometric InP, brown InGaAsP, orange BCB and blue the GaAs substrate.*

lithography becomes more challenging however: figure 4.19 shows a bonded substrate which has been set up to begin ebeam lithography. The poor image quality makes good lithography more difficult than on an InP wafer, and as discussed above we are already limited by poor lithography. Figure 4.20 shows another issue - the InP membrane layer is very thin and fragile, so even loading the sample for ebeam can shatter it. Additionally, it is necessary to deposit a hard mask to etch the InP, and this process requires heating. The spin-on-glass typically used as a hard mask (see chapter 3) must be heated to 500^{o}C; but the BCB degrades above 350^{o}C. This last can be solved by using a PECVD mask, as the PECVD process has a lower maximum temperature.

With the etch-first method, processing becomes identical to standard planar processing for InP, which is convenient. The lithography discussed above has been performed in this etch-first process. However, to finally choose this process, refinements to the bonding technique discussed above were necessary - bonding in vacuum removes the possibility of trapping air in the gratings. Also, this decision must be revisited for any possible integration of this device, as burying the patterned surface is not always desirable.

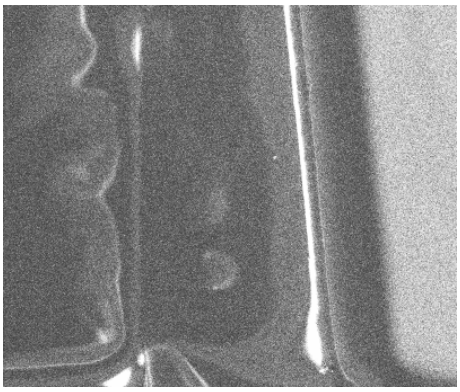

Figure 4.19: *Bonded substrate being prepared for ebeam lithography. On the left is BCB, on the right, bonded InP. The BCB is an insulator, hence we see charging. That is, the image quality is poor and cannot be readily improved, and the image quality obtainable is a limiting factor on the quality of the lithography.*

4.4.4 Bottom reflectors

I have explained that adding a bottom reflector to the GC can improve performance significantly, as it minimises radiation into the substrate. The reflection needs to have the correct phase relation with the grating, and hence the relative position of the mirror is important. We use the simplest available method for making a mirror, depositing gold on to the appropriate surface.

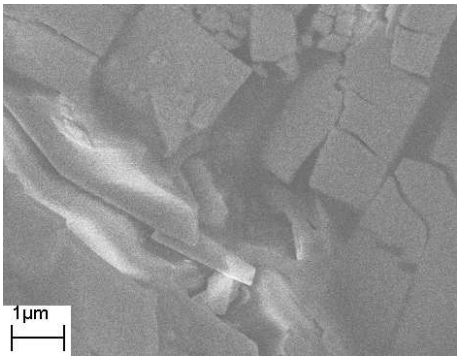

Figure 4.20: *The bonded InP layer is very thin and fragile, and shatters even under the light mechanical stress applied while loading a sample, the results of which are shown in the image. The stress is from the positioning of the grounding pins, which are necessary on samples containing an insulating layer.*

Then we must control the BCB deposition to a thickness of $(2N + 1)\lambda/4$. Typically we choose N=1, as this gives a convenient thickness. Note that this value is not exact for the BCB, as the InP thickness must also be taken into account. This process does work, but not reliably. The difficulty is in the actual bonding step, because we must apply pressure to ensure a good contact. As a result, the thickness can be changed. Another benefit of the new vacuum bonding process is that this pressure is controlled. Further, in BCB the required thickness is approximately 300nm [9], whereas a good bond from the simple process requires several microns of BCB, particularly with a patterned surface. A more controllable, but more complex procedure is as follows: deposit a low index dielectric (e.g. silica) on the etched surface to the critical thickness. Put the gold on top of this, then spin a thick (uncontrolled) layer of BCB, and bond as usual. It is important that the low index dielectric material has good planarising properties and can also be deposited by an accurately controllable method such as PECVD. Both of these processes are still being developed.

[9] approximate as the refractive index is sensitive to the curing process

A device prepared using the first process is shown in figure 4.21. This is a bonded-SOI rather than InP device [47], but the principle is the same.

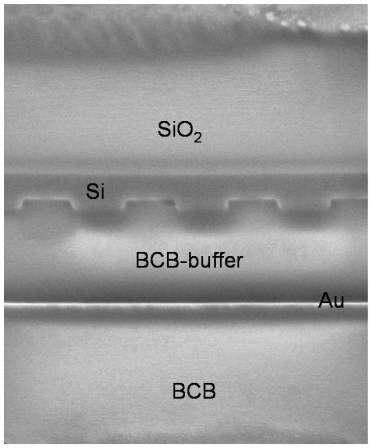

Figure 4.21: *A GC with bottom reflector made using adhesive wafer bonding, with the BCB thickness controlled to give the correct phase relation between the transmitted and reflected waves. The grating is etched into the top surface of an SOI wafer, then inverted into the BCB; then the substrate and BOX layer of the SOI removed. The fabrication principles are the same as for the InP gratings discussed here. Courtesy of F. van Laere.*

4.5 Characterisation and testing

I have now discussed how to design and fabricate GCs, and the parameters that must be considered to do so. Next, I will discuss the testing procedures for these devices. Firstly, we must decide what parameters we wish to be able to extract from the devices, and then lay out an appropriates set of structures. In particular, we would like to know our grating efficiency and centre frequency, the two parameters that can show how closely the real devices compare to the designs.

The measurement setup used to characterise these devices consists of a tunable laser, polarisation rotator, a polarisation maintaining fibre, the sample, output fibre, and a photodetector. Therefore, characterisation is fibre-to-fibre; which gives the simplest and most realistic expression of the coupling behaviour [10]. For 1D devices the PMF is not strictly necessary. The basic device consists of two gratings connected by a waveguide. Mechanical considerations make the minimum waveguide length approximately 2mm. The total efficiency can be expressed as $I_{output} = \eta_{in}\eta_{out}e^{-\alpha L}I_{input}$; where $\eta_{in} = \eta_{out}$ is the efficiency of each individual grating. Assuming negligible fibre losses, I_{input} is known as a function of wavelength. The device loss, for a waveguide, is modelled as $e^{-\alpha L}$, which can be determined by making waveguides of several different lengths, and using the cutback method. This is important because we expect the waveguides to be quite lossy - the fabrication process is optimised for the gratings, not the waveguides.

Ultimately, we would want to use a GC to couple light from a fibre into a photonic wire or photonic crystal, typically in the fundamental mode. To assess the efficiency of coupling into this mode, we must include a photonic wire as a mode filter, plus appropriate tapering structures, and then we must also account for the loss of the wire. If we use four waveguide lengths to determine losses via cutback, we now need twice this many devices to calculate the coupling efficiency of the designed grating.

We also need to determine the centre frequency. Ideally, this would be the same for all the gratings already needed, but to improve accuracy, we also fabricate gratings with a slight lithographic tuning, and another fillfactor as insurance against fabrication tolerances.

For the polarisation diversity couplers, we need all of this in two dimensions,

[10] Characterisation by my collaborator F. van Laere.

hence both waveguide and wire bends, hence more calibration structures. A typical layout to characterise a 2D single-mode polarisation diversity splitter is shown in figure 4.22. The lithography time tends towards days rather than hours. Although this is dominated by the time to write the waveguides rather than the actual GCs, it has largely been found impractical to use optical lithography to reduce the writing time, as photonic wires, and photonic wire bends in particular, are beyond the reach of the available tools.

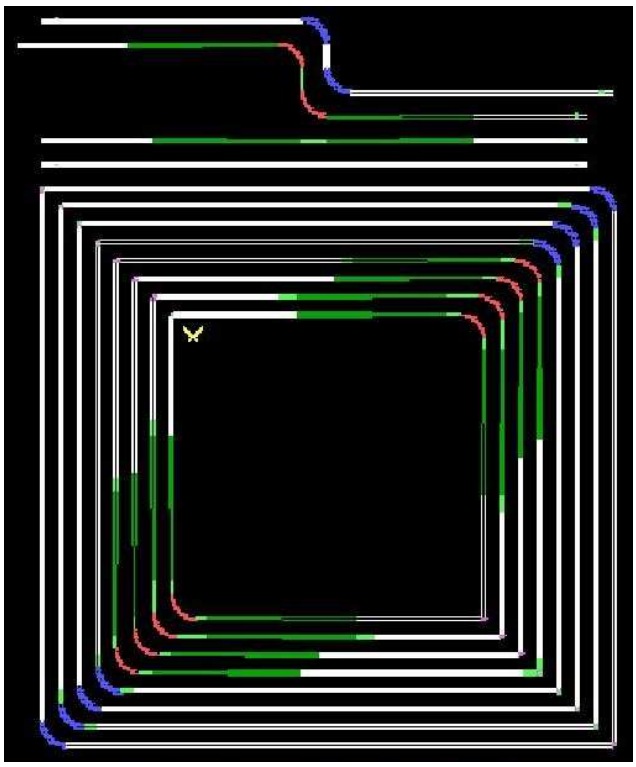

Figure 4.22: *Mask layout showing all of the devices necessary to fully characterise the efficiency and centre frequency of a single polarisation diversity GC design. The different colours represent different types of structure, such as photonic wires and ridge waveguides, and their bends. The pattern is approximately 2mm wide, and the actual GCs are not visible at this scale.*

Although the process is now fixed, we do not yet have measurement

results on complete 2D polarisation diversity couplers. However, we do have results for 1D gratings, which have one of the highest coupling efficiencies ever recorded [48]. The best devices make use of the gold mirror and wafer bonding techniques developed in this project, with almost 70% coupling efficiency [47]. However, these are for SOI[11]. Accordingly, I will show the InP results here. The characterisation has been performed by F. van Laere. An SEM image of the structure is shown in figure 4.23, and the normalised spectra in figure 4.24.

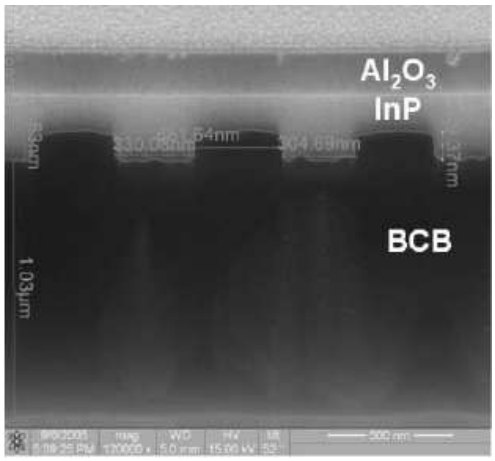

Figure 4.23: *SEM image of a FIB cross-section of a measured GC. The lithography is mine, performed at St Andrews. In particular, this image demonstrates the success of the wafer bonding, as the grating is fully filled with BCB. The grating linewidth and etch depth are also found to be suitable.*

The AR coating in these results is a simple step towards increasing the upwards coupling, not otherwise discussed here as the thickness is not reliable and the refractive index poorly characterised. The gold mirror technique is much preferable, and we expect results comparable to those from the SOI structures soon. Figure 4.25 shows the coupling efficiency with a gold mirror

[11] and use lithography and etching developed as part of the PICCO project in the DUV fab at IMEC, Belgium [28].

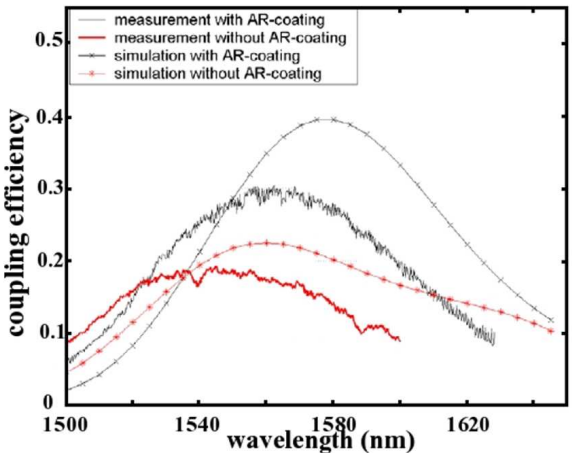

Figure 4.24: *Coupling efficiencies for the measured InP GCs. The peak coupling is around 30%. These are first generation devices, with 1D gratings, made using the etch-first procedure discussed above, with no bottom reflector. These results agree very well with modeling results.*

added (the AR coating is no longer useful), where an attempt has been made to control the thickness of the BCB layer. The quoted efficiency here is 56%, although simulations suggest a 78% efficiency for the best BCB thickness [49].

Further work on these devices is in progress, both to improve the fabrication and hence the coupling efficiencies further, and to integrate the GCs with other devices. In particular, we are evaluating methods to make the lithography more repeatable, such as using other resists; trying to deposit the gold for the grating mirrors solely under the grating as a means of reducing waveguide loss and increasing bonding yield, and optimising the etching for 2D as opposed to 1D structures.

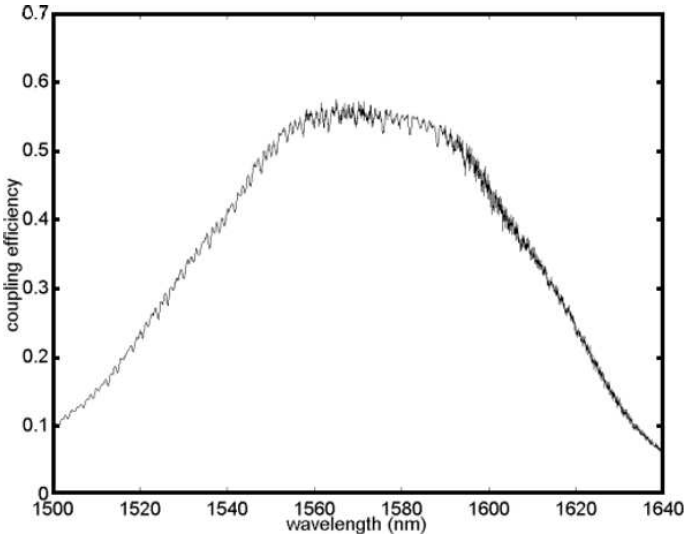

Figure 4.25: *Coupling efficiencies for 1D InP devices with a gold mirror added, bonded in vacuum using the simple method - i.e. attempting to control the BCB thickness. The peak efficiency is around 56%.*

4.6 Conclusions

In this chapter, I have discussed the difficulties associated with coupling to nanophotonic devices from optical fibres, and showed our preferred solution. This is to use a surface grating coupler, which has the benefits of a small footprint and wafer scale testing, as well as the major advantage of presenting a polarisation diversity approach. The theory underlining this has been explained, from a simple perspective, and this has been tied to the PhC concepts developed in chapter 2, as well as those which will be used for the devices discussed in subsequent chapters. The major development is the addition of a bottom mirror to the design, which reduces radiation to the substrate, one of the major sources of loss. However, the main body of my work here dwells on the practical considerations which impinge on final device design - and particularly the lithographic challenges. The criticalities

of the fabrication have been discussed, and the most demanding parts found to be accurate lithographic reproduction of the high fill-factor gratings, and controlling the thickness of the BCB used for the wafer bonding such that it gives the correct phase relation for reflections from the bottom mirror. Finally, I have shown measurement results. The best published GC devices use techniques developed by my collaborators as part of this project, and I have helped to fabricate devices with coupling efficiencies of $> 50\%$ in our InP system, and developed methods that are expected to give 2D polarisatoin diversity couplers with fibre to single-mode waveguide efficiencies of $> 80\%$, the target for the project.

Chapter 5

Photonic crystals and coupling

In the preceding chapter, I have discussed the design, fabrication and testing of a PhC device that can be used to significantly increase coupling between an optical fibre and micron scale waveguides, and in the next chapter I will discuss the successful design and characterisation of some sub-micron period PhC devices. In this chapter, I discuss the intermediate step, coupling from several micron wide waveguides into PhC devices with a period of approximately 0.5μm. I will show that the modes propagating in periodic and translation-invarient waveguides can be very different, and thus the coupling between these also needs consideration. There are three sections, each representing very different work, but all hinging on these problems - how to effectively couple into PhC devices. In the first section, I discuss the numerical optimisation of a short taper, which is an order of magnitude smaller than the conventional taper for its material system, showing the steps to develop the design. Although this taper was used to couple into PhC devices, specifically those in the next chapter, the periodic nature of PhCs has been neglected in its design. Accordingly, in the following section I show the difficulties with coupling into periodic structures, focussing on W1-type PhC waveguides, particularly in

the valuable highly dispersive slow-light regime. I show some techniques for modelling this coupling, which it is hoped will help to take a step back from intensive numerical modelling for PhC interface design. Some structures which potentially improve this coupling are also shown. Finally, in the third section, I present some work on tuning PhC dispersion for device design, the performance of which is severely limited by coupling effects.

5.1 Optimisation of short tapers

In the previous chapter, I have discussed the possible methods for coupling between a fibre and a die, and the limitations of the various methods. Nevertheless, in the lab we often do use the low efficiency technique of using a lens to focus light from a fibre on to the cleaved facet of a ridge waveguide. This gives us at least some light propagating along such a waveguide. Considerations of both the numerical aperture of the waveguide and the mechanical limitations of alignment mean we typically fabricate such waveguides with a width of 3 to 5 μm[1]. Waveguides for GCs are 10-12μm wide. However, the photonic wires and PhCs that we wish to study have a feature size of approximately 500nm, and so we must couple between the two. The most common structure for this purpose is known as a spot size converter. Conventionally, these devices are straight tapers, with their change in width versus length approaching the adiabatic limit, to maximise the conversion between the fundamental mode of the broad ridge waveguide and the microphotonic scale. A taper angle of $<1^o$ is common, which gives a 200μm taper for the SOI work presented later in this chapter. These long tapers present problems - they are difficult to model

[1]Waveguides <1μm are possible, but demand a highly skilled practitioner, high transmission devices, and a lot of time.

accurately, due to staircasing effects and the sheer volume of simulation that is required, and in fabrication terms the large footprint is an issue.

As part of the numerical optimisation work mainly discussed in chapter 6 it was decided to make a design for a short taper-based spot size converter (also called an injector). This is to be used for coupling into a W1 PhC operating in an index guided mode, that is one that matches well to a photonic wire of comparable waveguide size. The terms here are explained in chapter 2, and the difficulty of coupling between PhCs and photonic wires is discussed in the next section of this chapter. Here, we show how the optimisation was performed, but then present further analysis to show the principles of the device operation.

This device design is for AlGaAs heterostructure more fully described in chapter 6, but the majority of the modelling work in this chapter is performed in 2D using an effective index approximation, so the precise layer structure is not important.

5.1.1 Optimisation

For the optimisation of this short taper, the large end has been chosen as a 5μm ridge waveguide, and the small end is the microphotonic circuit discussed in the next chapter. In the final device, this output would be a mode propagating out of the PhC, but we need to simplify the problem, both to understand the various effects we see, and to reduce the computational load. Thus, we begin by choosing a photonic wire to serve as an artificial input to the PhC, which we will then use in place of the PhC as the output of the taper. This is simple problem for devices operating at this point in the bandstructure, because the PhC mode profile is similar to that of the photonic wire, so good coupling can be readily achieved. We have chosen a PhC of lattice period 440nm, operating at a centre wavelength of 1.55μm, and an artificial input wire of

width 640nm. In figure 5.1 we show the field profile for a mode launched using the fundamental of this artificial input. On the right of the figure we again have this wire, and calculate the field across the wire[2]. Using Photon Design's FimmProp program, the transmission is 99.5%.

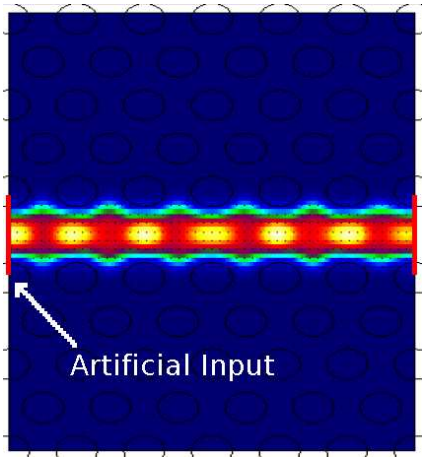

Figure 5.1: *Transmission through a PhC using the photonic wire chosen as artificial input. The sensor on the right (yellow line) measures the transmitted field, recording a transmission of 99.5%. The purpose of this simulation is to determine the effectiveness of the particular photonic wire as an input to the PhC, so that we can perform subsequent calculations to assess the coupling to this wire rather than the PhC, which significantly reduces the size of these calculations.*

We now begin to optimise the taper, for the parameters shown in figure 5.2 - the taper length, width, and offset from the ridge waveguide. The artificial input is the yellow line on the left of the figure, and it remains constant throughout. The optimisation is done in the frequency domain at a single frequency, making use of a symmetric boundary along the symmetry

[2]The optimisation has been performed to maximise the total field rather than the power in the desired mode. This is unavoidable due to software capabilities at the time, but is not ideal. The consequences will be discussed subsequently.

plane of the device. This restricts the available mode set to those with odd symmetry, which is valid because we would not expect significant excitation of even modes experimentally. Figure 5.3 shows the transmission versus taper length for the optimised width and offset, with the chosen length indicated. The optimisation was performed using a global algorithm, hence considering all possible combinations of these parameters.

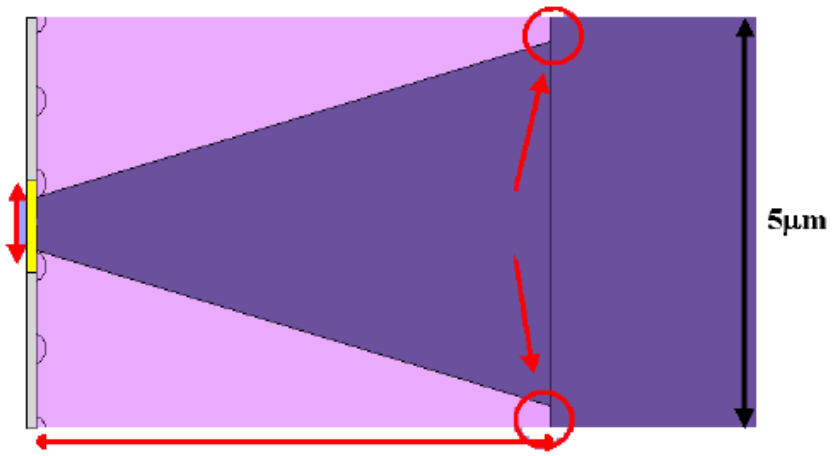

Figure 5.2: *Layout and optimisation parameters for the short taper. The artificial input is the yellow line on the left, which is not changed, and the launch is at the arrow on the right. Note that propagation in this figure is from right to left.*

5.1.2 Analysis

Given the design above, it seems worthwhile to study it further, to gain some physical insight into the operation of this taper. We focus on coupling into the "artificial input" wire, as this is simpler than the coupling into the PhC. This analysis has largely been done using the FD engine of CrystalWave. We begin by calculating the transmission of the taper in this program. Figure 5.4 shows the field profile across the taper. From this, it is apparent that the device

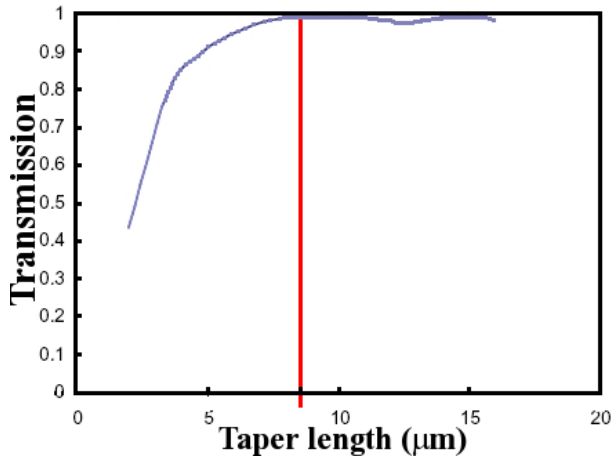

Figure 5.3: *Field in the artificial input vs. taper length, for the optimised taper width and offset. The chosen length is indicated, with a transmission of 99% in 2D. This simulation was run in the reverse direction to the optimisation, with the fundamental mode of the ridge waveguide as the input.*

cannot be considered as an adiabatic taper, slowly converting the fundamental mode - instead, the taper makes use of higher order modes at some intermediate widths, and as such must be considered as an MMI (Multi-Mode Interference device).

5.1.2.1 Transmission spectra

We now calculate the transmission spectrum of the device, recalling that the optimisation has been done for a single wavelength. The transmission is expressed as the percentage coupling between the fundamental mode of the 5μm waveguide and the wire. Figure 5.5 shows the radiation fields at the end of the taper. This shows that the location of the mode monitor must be placed carefully, away from surface states associated with the end of the taper. The transmission spectrum is shown in figure 5.6.

This is an alarming result - the peak transmission is not >99% as

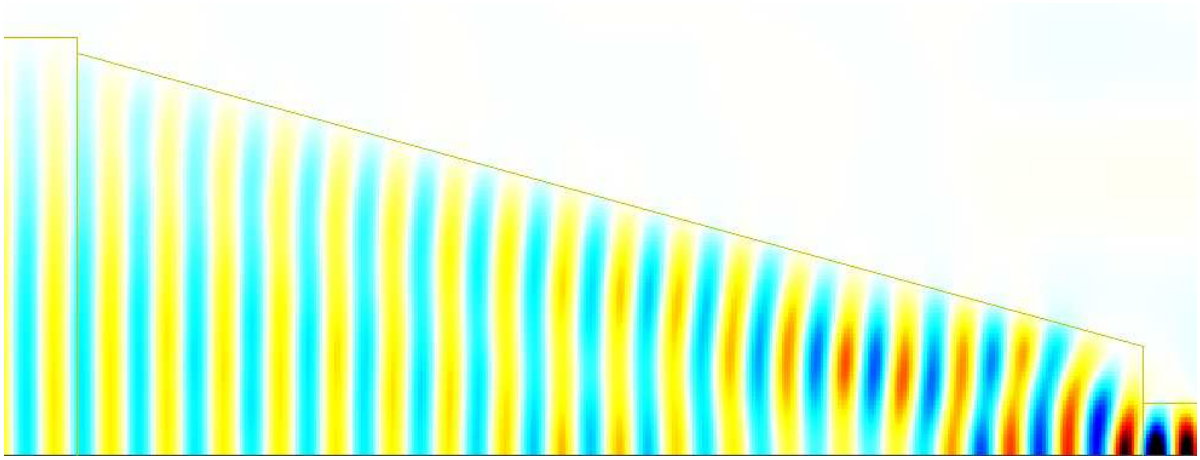

Figure 5.4: *H_y field profile across the taper which has been optimised according to the scheme presented. An adiabatic taper would show a field with steadily increasing intensity along the centre line as the taper narrows, as the energy in the field is confined in a reducing volume; however this field shows the nodes associated with higher order modes. Hence we conclude that this optimised device operates on the principles of a multimode interference coupler.*

the optimisation results suggest, and the peak is not at the frequency of optimisation. We have calculated the transmission into the fundamental mode of a wire, and this is the wire that we have designed as an artificial input into our crystal, so this is the obvious thing to monitor, and to optimise.

We repeat the calculation, this time monitoring the positive flux, i.e. the forwards-going intensity, through the line of the mode monitor. This result is shown in figure 5.7. This is normalised against the input flux. Comparing these two figures, we see that there is significant power in either higher-order or radiation modes, particularly at the design frequency. However, it is still not clear what has been optimised, as we still do not see the expected result.

There is one final option we can explore, namely the field at the end of the taper. This must be monitored at a single point inside the waveguide,

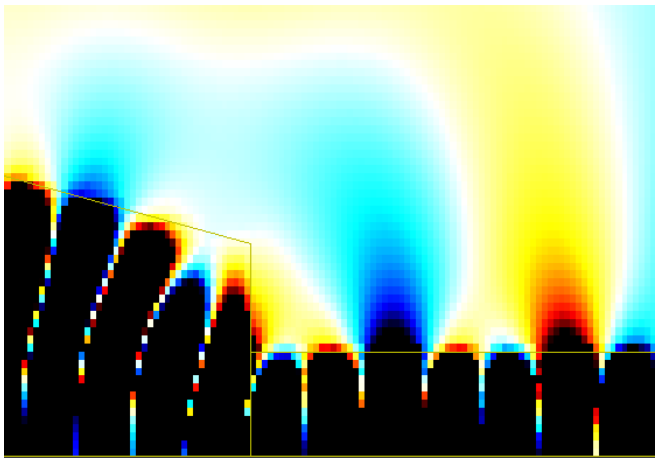

Figure 5.5: *H_y field profile at the end of the taper in the previous figure. The colour thresholds have been changed compared to the previous figure to show the radiated field. From this, we can see that there are significant and complex surface states surrounding the end of the taper. Ergo, the position of the output monitor must be considered carefully to avoid considering contributions from these standing waves in the device transmission.*

and as the device operates as an MMI that point is in the centre of the waveguide exactly at the end of the taper. The result is shown in figure 5.8. The scale is arbitrary relative to the input power of the simulation, but this is constant for all wavelengths. The different mode volumes at the input and output of the simulation make normalisation to the field at a point at the input inappropriate. However, we do now have the peak transmission at the desired wavelength, and conclude that this must be related to the optimised variable. As discussed, the overlap with the fundamental mode is the most appropriate variable to optimise, but this may not have been possible with the earlier software versions used in the device design. We conclude that the taper design as presented represents a good short taper, but perhaps not the optimum one for the parameter set given.

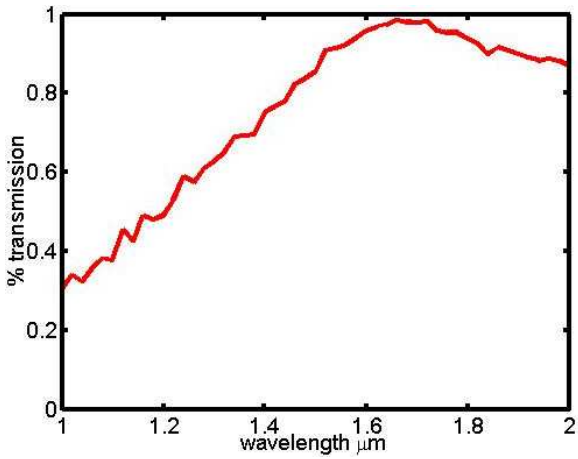

Figure 5.6: *Transmission in the fundamental mode of the wire, as normalised against transmission in the fundamental mode of the ridge waveguide. This gives the effective coupling efficiency of the optimised taper as a function of wavelength. The peak transmission is not >99% as predicted by the single frequency optimisation.*

5.1.2.2 Peak positioning

In the previous section I have discussed which variable has been optimised in the taper design, and shown that the taper has an MMI-like response. For a simple rectangular MMI, it is well known that [50]:

$$L_\pi = (4n_0W^2)/(3\lambda_0) \tag{5.1}$$

Where L_π is the characteristic length of the MMI, n_0 is the effective index of the fundamental mode, W is the width of the MMI and λ_0 the free space wavelength. Thus, for our trapezoidal MMI, we expect that changing the length (and hence the taper angle) will change the peak coupling wavelength. This is demonstrated in figure 5.9. Essentially, this shows that the device operates on diffraction principles, and hence that an analytical approach may be more beneficial than numerical optimisation. Further, the 3dB bandwidth

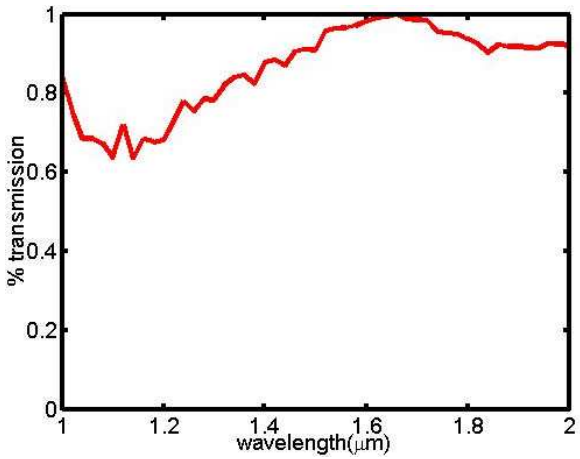

Figure 5.7: *Transmission as above, but in terms of the forward-propagating (positive) flux through the monitor plane, as normalised against the same quantity in the ridge waveguide, as opposed to the fundamental modes in the previous figure. The results in this figure and the previous one are quite different, which implies significant transmission outside of the fundamental mode.*

of these devices is around 500nm, which gives a lot of margin for error in the design and fabrication.

5.1.2.3 Complete device response

For our optimised taper, we now study the complete response of the device, in terms of its coupling with the photonic wire used in the optimisation and with the target photonic crystal. Using the FD engine of crystalwave, we simulate these two structures, shown in figure 5.10. The overlap with the fundamental mode is calculated for the ridge waveguides at the left and right, and the transmission stated is the ratio of these. By calculating over a wide range, we can see if the peak device transmission occurs at the optimised position or at the peak of coupling into the fundamental mode. The results are shown in

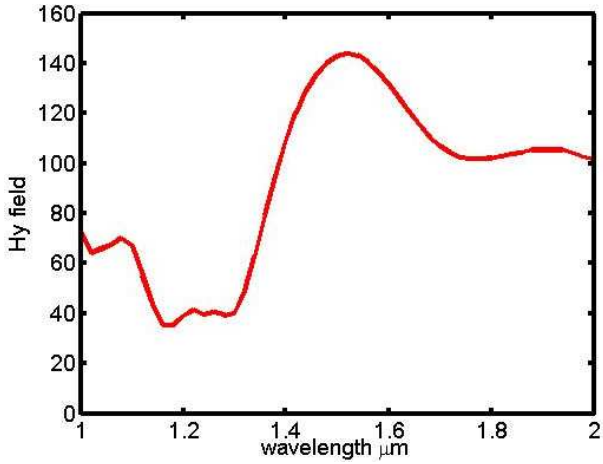

Figure 5.8: *Transmission as above, in terms of the H_y field at the centre of the monitor across the wire. This is not normalised, but instead is relative to the (constant) intensity of the same field component at the centre of the ridge waveguide. The results now agree with those of the optimisation, but this is the least appropriate method of determining the coupling.*

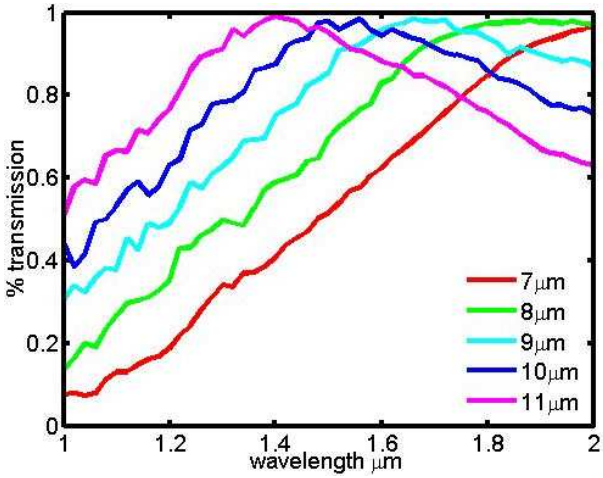

Figure 5.9: *Variation of taper transmission spectrum (for the fundamental mode as in figure 5.6 with taper length. The peak wavelength decreases as the taper length increases. Essentially, this figure demonstrates that numerical optimisation may not be the best method to use to design such a taper.*

figure 5.11.

For the structure using the wire that was optimised to be the artificial input, it seems that we have only found a local maxima on the edge of the broad MMI response. However, once we insert the crystal, the response is significantly better. This is borne out by the experimental results discussed in the next chapter for photonic crystal devices including additional optimised elements, where in many cases the transmission of the photonic wire was found to be significantly worse than even these more complex devices.

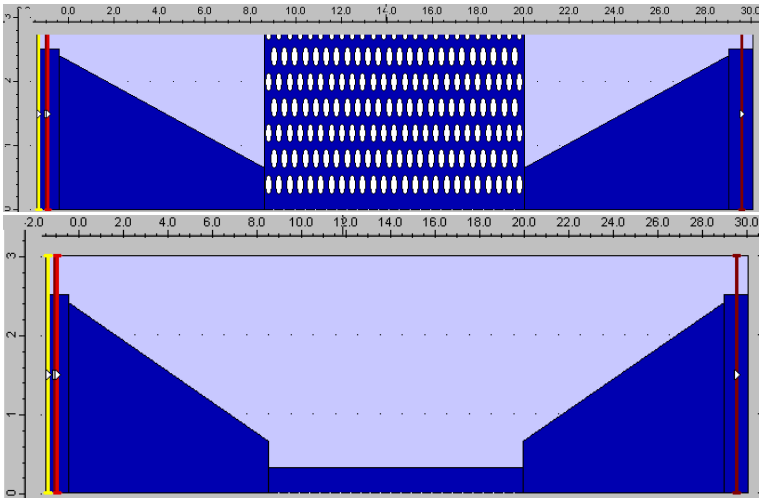

Figure 5.10: *Final taper devices that were simulated to confirm the performance of the complete device, that is without using the artificial input structure to minimise the simulation. To be certain we are considering the relevant response, we use identical 5µm ridge waveguides for the input and output of the simulation, and calculate the mode overlap integral.*

From this result, we can draw a few obvious conclusions. Firstly, numerical optimisation is no substitute for intelligent device design! One must be very aware of the principles of operation of both the device and modelling software being used. Secondly, coupling into high group velocity PhC modes is easy, and that 200μm tapers are not necessarily the best solution, especially considering

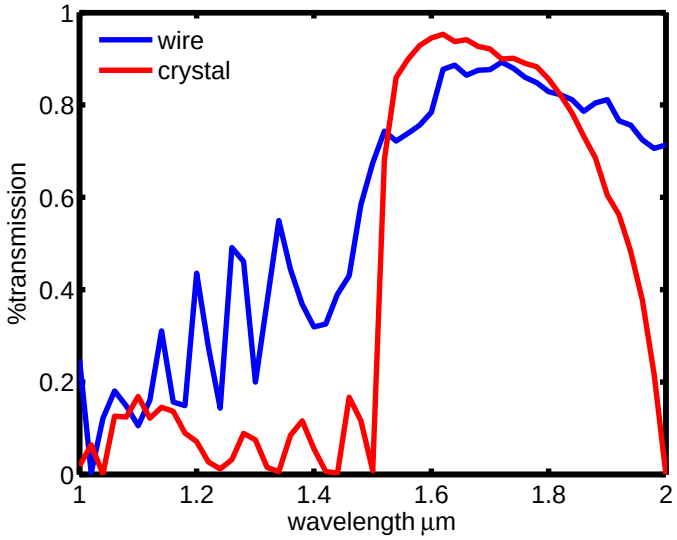

Figure 5.11: *Transmission spectra for the devices in figure 5.10. Unexpectedly, the transmission through the PhC is higher than that through the photonic wire. This is probably due to the coupling of the taper into higher order modes shown in previous figures - these are cut off in the wire but couple into propagating states into the crystal, where the mode profile is wider. This agrees with the experimental results presented in the next chapter.*

their own challenges to realisation. From the transmission data, it is not possible to tell if we are effectively coupling into PhC slow light, but then that was not the intention for this device design. This issue is discussed further in subsequent sections.

5.2 Photonic crystal interfaces

After coupling into a slab waveguide using a PhC, and showing a compact method of spot-size conversion for coupling into the index-guided mode of a PhC, we now address coupling into "slow light"; gap-guided PhC modes which

are noted for their low group velocity. This will be done in a number of ways, and using a number of techniques of increasing complexity and accuracy. For this section, and the rest of the chapter, all results are for the SOI system.

We begin by showing the problem. In figure 5.12, we show a transmission spectrum overlaid on its bandstructure, for an arbitrary W1 waveguide. The transmission is calculated in 2D FDTD using FullWave, using the fundamental mode of the waveguide shown for the input and output, to calculate the amplitude and phase of the transmitted field. The methodology for this involved calculation is discussed in chapter 2. The blue and green lines are the transmitted and reflected field respectively. The orange dots describe the fundamental mode of the PhC, and the pink dots the even mode. These are calculated by PWE using mpb. The red line indicates the dispersion of the fundamental mode of the photonic wire ridge waveguide. Note that the transmission begins to decrease at the point where the W1 and wire modes begin to diverge.

As we have also extracted the phase from the FDTD calculation, we can use the following formula to determine the group index of the device as a function of wavelength[15]:

$$N_g = -\frac{cT}{L} - \frac{\Delta\phi}{2\pi}\frac{\lambda^2}{\Delta\lambda} \tag{5.2}$$

Where cT/D is a retarded potential: the time in microns that the wave must propagate from the source to the sensor divided by the distance between them. This is necessary because we can only determine a relative phase modulo 2π rather than the absolute phase using this method. Using a Hellman-Feynmann algorithm, the group velocity can also be determined from the PWE bandstructure. Figure 5.13 shows the results.

There are two interesting results from these figures: firstly, the group

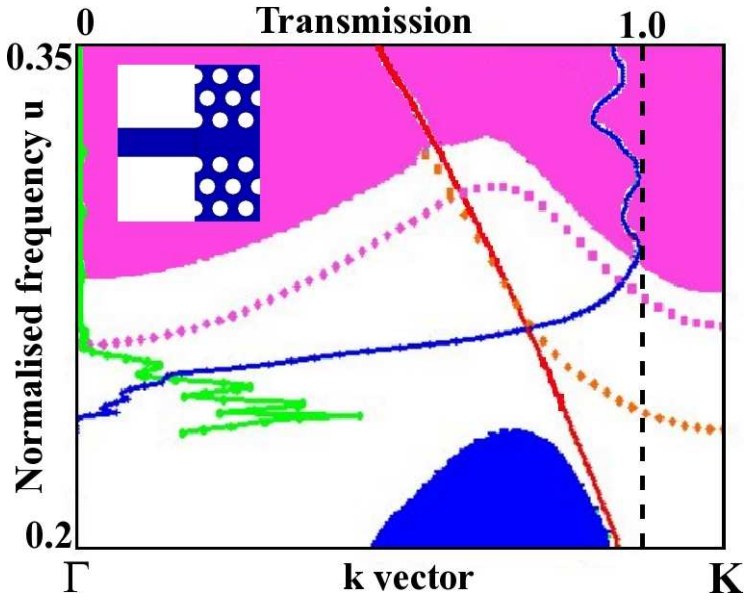

Figure 5.12: *Overlayed 2D FDTD transmission spectrum and bandstructure for a SOI W1 terminated as shown in the inset (top left). The green and blue solid lines are the reflection and transmission, the orange and pink the first two modes of the PhC, and the red line the mode of the input ridge waveguide shown. Note that the transmission decreases instantly the waveguide mode and PhC mode diverge.*

velocity for an un-terminated supercell, and a supercell terminated in a waveguide, are very different. This is significant for any attempts to use the "slow-light regime" of these PhCs. Secondly, for the finite length device, as the band flattens off for increasing wavelength, we see significant ripple in the group velocity, but the transmission appears to fall off smoothly. Surely we should see some repercussions to this ripple effect? Throughout the rest of this section, we discuss this matching problem, and some possible solutions to it.

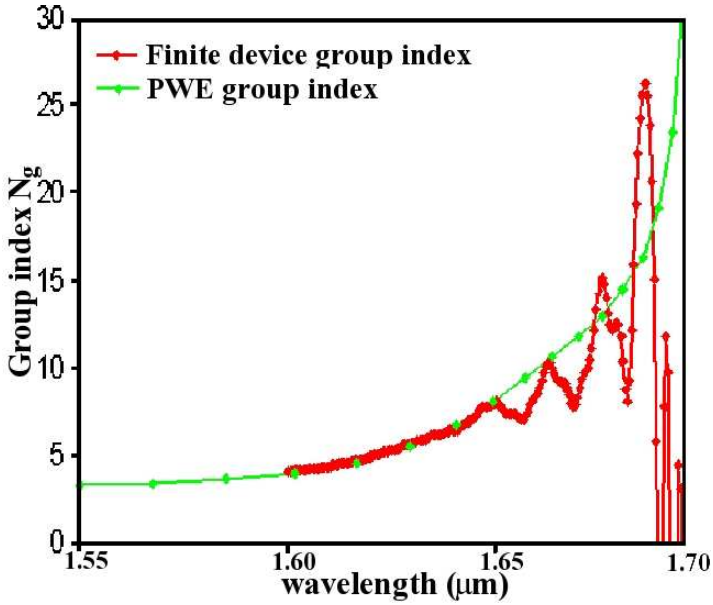

Figure 5.13: *Group velocity of the W1 fundamental mode in figure 5.12, calculated using FDTD for a finite length device (red), and using PWE for the infinitely periodic supercell, using the Hellman-Feynman algorithm which is part of the mpb package (green). The response for the finite device oscillates about that for the infinite structure; with oscillation frequency increasing as the group velocity decreases.*

5.2.1 Impedance effects

As discussed in chapter 2, there are two common ways in which PhC devices are studied, namely planewave expansion (PWE) bandstructures and finite-difference-time-domain (FDTD) transmission spectra. But, as shown above, PWE misses essential termination effects, and FDTD is too demanding in computational resources to make very accurate calculations of fine features. Instead, we pursue a mixed strategy to study the effects of device termination, and add frequency-domain transmission techniques to get good phase resolution. In particular, I will make use of the impedance strategy

pioneered for use with PhCs by S. Boscolo [51].

For any electromagnetic wave propagating in a translation-invariant medium, we can define a characteristic impedance Z from the ratio of the transverse electric and magnetic fields for any mode:

$$Z = \frac{\mathbf{E}}{\mathbf{H}} \tag{5.3}$$

$$Z = \frac{\sum_{i,j} \mathbf{E}_{x,y}}{\sum_{i,j} \mathbf{H}_{x,y}} \tag{5.4}$$

Which reduces to

$$Z_{\mathrm{TE}} = \frac{\sum_{x} E_x}{\sum_{x} H_y} \tag{5.5}$$

for the TE modes in a 2D slab. For the fundamental mode of a ridge waveguide of similar size to our PhC defect waveguide, i.e. a photonic wire, we will use $Z_{\mathrm{TE}} = 110\Omega$. This is equivalent to the impedance of free space divided by the effective index of a SOI photonic wire mode, which varies negligibly with both frequency and wire width over the ranges of interest.

For a PhC equation 5.5 needs to be re-formulated, because the waveguide is not translation invariant. There are two methods in the literature, from Boscolo [51] and Biswas [52]. Both make use of the Poynting vector to weight the fields, but Biswas' (more recent) method is better justified and simpler to calculate. The impedance is expressed as the ratio of the energy density to the power flux:

$$Z = \frac{\sum_{i,j} \mathbf{E}_{x,y}{}^2}{\sum_{i,j} \mathbf{E}_{x,y} \times \mathbf{H}_{x,y}} \tag{5.6}$$

which becomes:

$$Z = \frac{\sum_i |E_x|^2}{\sum_i E_x^* H_y} \tag{5.7}$$

Using mpb, we can easily determine the fields for a PhC supercell for any band and frequency. The calculated impedance for each slice of a supercell discretised with 20 slices transverse to the propagation direction is shown in figure 5.14. Note that the discretisation controls the number of planewaves in the bandstructure calculation, where the actual number of planewaves used here is 20× the area of the supercell.

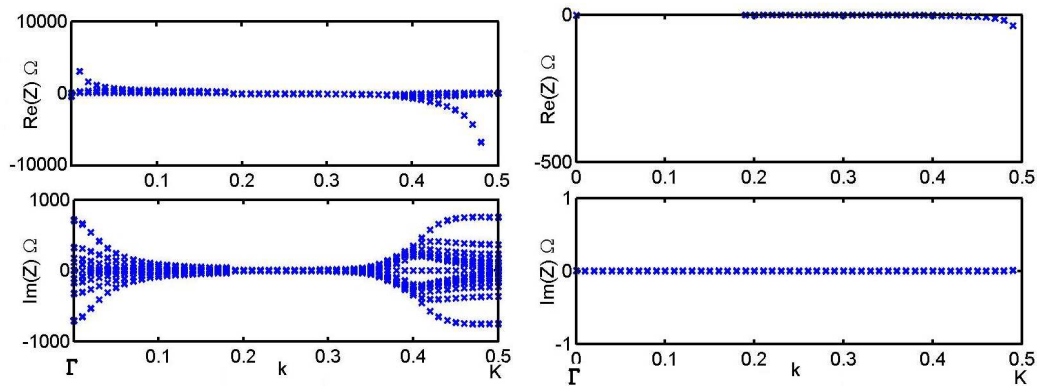

Figure 5.14: *Impedance in Ω versus k for the fundamental mode of a typical W1 waveguide. For low k, the mode is above the bandgap, and has been neglected. For high k, the mode folds and so it moves into the slow light regime, where matching becomes difficult as the impedance becomes high. In the left hand figure the impedance is plotted for each of the 20 transverse slices of the supercell. In the right hand figure the impedance is plotted for the symmetry plane only. Note the different scales. We find that the impedance is purely real only at the symmetry plane, so this will be used henceforth.*

In each transverse slice, the dielectric profile is different, so it is unsurprising that the impedance changes with position. However, the impedance only

becomes purely real at the high symmetry planes, that is through the lattice sites, and so we choose this value to represent the impedance of the infinite PhC supercell. Note that the plotted values are largely negative. This is a consequence of folding the PhC bandstructure back into the first Brillouin zone; and as such the negative sign can be discarded. As expected, the impedance values are very similar to that chosen for the ridge waveguide where the PhC mode is propagating quickly, but the impedance tends to infinity as the group velocity tends to zero.

We can now calculate the impedance mismatch between PhC devices and ridge waveguides, for cases where the interface is along the high-symmetry plane. To aid comparison with transmission spectra calculated using standard methods, we use the bandstructure to convert the k axis into normalised frequency. We assess the mismatch in terms of the reflection co-efficient ρ:

$$\rho = \frac{Z_L - Z_0}{Z_L - Z_0} = \frac{|Z_{PhC}| - Z_{RWG}}{|Z_{PhC}| + Z_{RWG}} \tag{5.8}$$

where the subscripts L and 0 denote the conventional load and terminal impedances, and PhC and RWG the photonic crystal and photonic wire ridge waveguide impedances respectively. For the W1 above, the reflection co-efficient and the group velocity are plotted in figure 5.15. The group velocity is included to try and quantify the term slow light in this case. Although the reflection co-efficient calculation does not include group velocity, it is clear from this figure that the two are related.

Using this reflection co-efficient data and the Fabry-Perot transmission formula:

$$I(\lambda) = 1 - \left(\frac{1-R}{\sqrt{R}}\right)^2 \times \frac{R}{1 + R^2 - 2Rcos(\phi)} \tag{5.9}$$

and assuming:

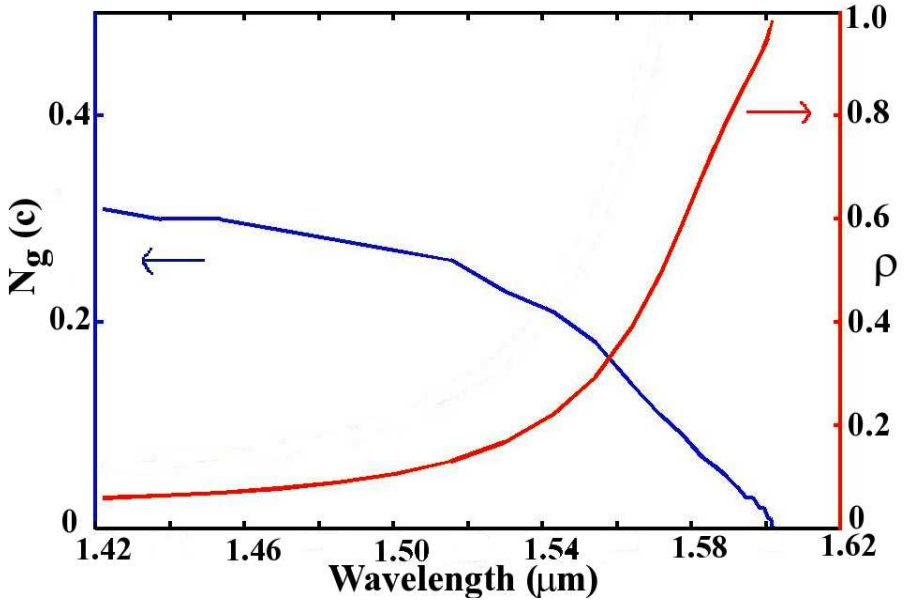

Figure 5.15: *Group velocity as a fraction of c (left axis), for an infinite W1 waveguide and the reflection co-efficient between this waveguide and a photonic wire ridge (right axis), calculated using data from mpb, as a function of wavelength. The relationship between these two data is in the field profiles, which are calculated using the eigenvectors which accompany the eigenvalues which define the bandstructure which is then differentiated to give the group velocity, and the fields are used to calculate the impedance.*

$$\phi = 2nd\frac{2\pi}{\lambda_0} \tag{5.10}$$

We can calculate the transmission of our PhC device for this matching condition using only bandstructure information - by definition the effective index n is given by k/u when these values are normalised. The result is shown in figure 5.16. This does not look like any of the W1 transmission spectra shown elsewhere in this thesis, either simulated or measured. As the bandstructure and the calculated impedance are smoothly varying functions, they have been interpolated, giving a resolution in this figure of approximately 0.1nm, much

greater than the resolution we can achieve with normal numerical calculations. As an example, the first spike in the figure has a 3dB bandwidth of less than 0.3nm and a Q of 678, so would not be seen in FDTD unless the simulation was designed to find it!

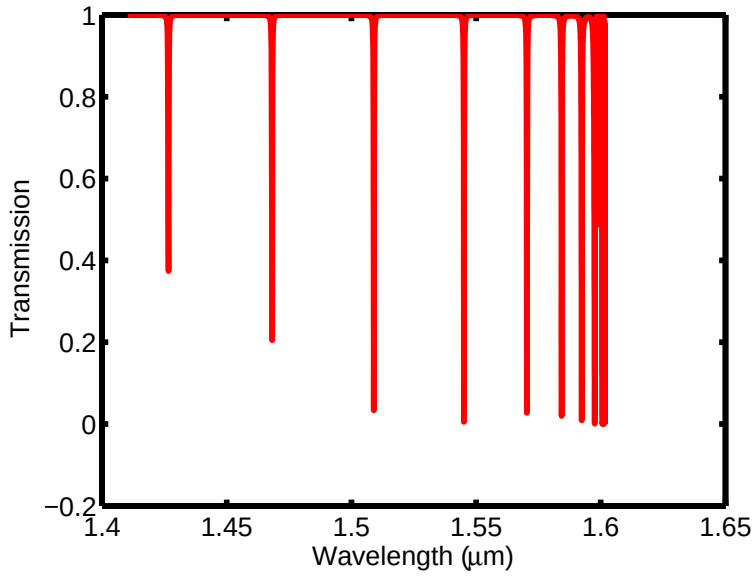

Figure 5.16: *Transmission spectrum calculated from the bandstructure, for a 20 period PhC, assuming the lattice period is 400nm, that the termination is on a symmetry plane (i.e. through the centre of a row of holes), and that the input guide has a constant impedance of 110Ω. This looks very different to the FDTD transmission spectra presented previously - the reason for this is that we only permit one mode in each of the PhC and the photonic wire, whereas in reality we have coupling to many modes, particularly where there is refelction from the PhC interface. The high Q spikes in this figure do however correspond in frequency to the slight dips seen in FDTD transmission spectra, which indicates that impedance mismatch may be the cause of these dips.*

It can be clearly seen that this spectrum is significantly different to that calculated by the FDTD simulations. The reason for this, is that the

impedance analysis assumes that there is only one mode in each of the wire and PhC regions. This is technically true, but not an appropriate assumption, because a PhC is a distributed device - a single row of a PhC has a very different response to a long block. As a result, significant energy is coupled into radiation modes in both transmission and reflection, which is not taken into account in this impedance analysis. Nevertheless, this approach is very useful for understanding device effects.

5.2.2 A solution to the mismatch

One solution which has been proposed to this impedance problem, which limits the coupling into the useful slow light regime, is to use an adiabatic taper inside the PhC. In practice, we cannot approach the adiabatic limit for two reasons, firstly that this would make the device very long and secondly gridding resolution means we cannot position holes with such a high degree of accuracy. On the other hand, we have shown in the first part of this chapter taht good results can be achieved with a taper transition which is far from adiabatic. Accordingly, we have developed the structure shown in figure 3.5, which chirps the period in the propagation and transverse direction in 2.5nm steps[3]. Referring to the impedance analysis above, this chirp simply moves the position on the dispersion diagram, so that we can in principle begin by coupling into a fast Bloch mode of the PhC, and then approach the slow light region stepwise.

To validate this device, a number of methods have been attempted - impedance analysis combining each of the different steps, transmission FDTD for phase recovery, and frequency domain simulations using Photon Design's

[3]The limit to the gridding resolution for the fabrication method employed in this case; typical grid resolution values are between 1 and 10nm.

FD engine. The first has not yet been successful, as again we have to deal with the distributed PhC response. In principle FDTD phase recovery works, as demonstrated previously in figure 5.12. However, to be clear about the effect of the tapering regions, we need the PhC to be significantly longer than them, to give the light enough time in the slow light regime to have a marked effect on the group velocity. Of course we simultaneously have to increase the running time. The simulation becomes too big and time consuming for the available resources. Similarly with the FD engine, the simulation does not fit into the available RAM.

The only remaining approach would be to actually measure the device! This has been attempted, but fabrication has not yet been accurate enough to tune the band-edge into the region we can measure. In figure 5.17, I show the transmission spectra for a device with this chirped interface, and a W1 identical except for this interface (that is, we use the same relative wire termination. It is hoped that tihs can soon be reproduced experimentally.

5.3 Other studies - PhCs with elliptical holes for AWG design.

During the course of this project, it was intended to study the effects of PhC dispersion at an effective index rather than group index level, and use a PhC to form a miniaturised arrayed waveguide grating (AWG) device, as has been proposed by [53]. At first sight, this seems a reasonable proposal, given that commercial AWG devices are of the order of tens of centimetres footprint. The limit to the size is essentially the index contrast, and hence the length of arm needed to give the desired phase change, and with a PhC we can in principle have a very high effective index and hence very short arms, leading to a very

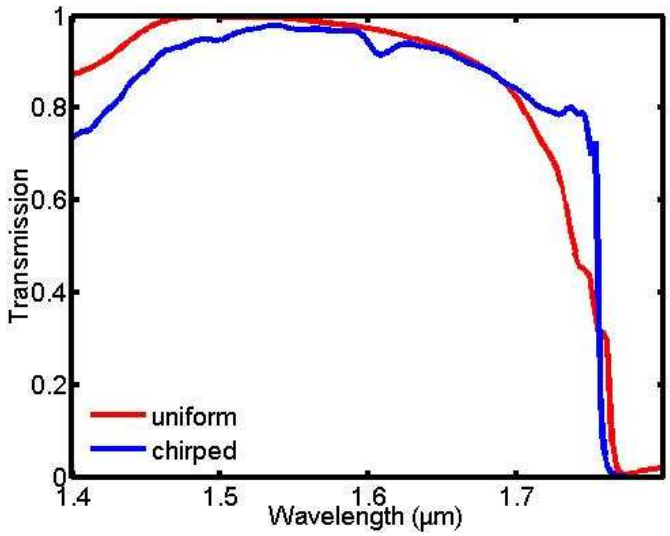

Figure 5.17: *Comparison of FDTD transmission spectra for a W1 with the wire interface shown previously (red), and the same length of W1 waveguide with a chirped PhC region added between the W1 and the photonic wire (blue). The transmission at the region of the spectrum known to display slow light characteristics is improved, but it has not yet been proved that this is increased transmission into slow light modes.*

small device footprint. This work was not successful for a variety of reasons, particularly that too many parameters must be considered and hence it is rather over-ambitious target. However, significant modelling and fabrication effort has gone into this project, and so the work is presented here, as finally coupling effects are the limiting factor.

5.3.1 Phase control for PhC based arrayed waveguide gratings

An AWG is a device that uses interference effects to spatially separate incoming light into many different output channels, depending on the wavelength of

the light. Their operating principle was first proposed by M. Smit at TU Eindhoven. A narrow input waveguide is tapered outwards into a star coupler, and diffracts to illuminate the input to the arms of the device. Conventionally, these arms are identical apart from their length, which increases by a fixed amount between arms, and hence there is a phase difference at the ends of these arms. These are followed by an output coupling region where the differing $\Delta\phi$ between the arms for differing waveguides tilts the output phase fronts and so steers each wavelength into a different output channel.

In our PhC device, we assume that the design of the input and output coupling regions can be borrowed from conventional devices, and study ways in which we can replace the phase-difference generating arms[4]. In [53], this is done simply by increasing the number of PhC periods between each arm, which is simple but not convenient. Instead, we consider ways in which we can keep the length of the arms the same, but change the dispersion relation of the PhC to give a different effective index and hence a different phase length[5]. Obvious methods are to increase the lattice constant of the PhC, so that we operate on different parts of the dispersion curve, and to vary the hole diameter of the crystal. The former is neglected on two grounds, firstly that the length of each channel will be different, and secondly because any input/output structure would need to change, which is not thought to be ideal.

In a purely theoretical design, varying the hole diameter is a good solution. However, we wish to consider practical effects. As discussed in chapter 3, we have an error of $\pm$20nm in any SEM measurement of hole sizes, and therefore

[4]This is not a good assumption, as star couplers are diffraction-based devices, and hence the size is significantly larger than any PhC device we could envisage.

[5]Again, not a good idea - although we can have variable dispersion, this is very non-linear, both with frequency in a single arm and with the physical parameters between different arms, so no simple method could be used to design the AWG.

the lithographic control of the hole size will be less than this. For a typical 400nm PhC period, allowing only this level of discretisation gives a step size of 0.025r/a. The change this gives in the bandstructure is shown in figure 5.18. In practice we can achieve much finer control over this, by carefully calibrating the exposure, but this is not desirable when a wide range of r/a values are required.

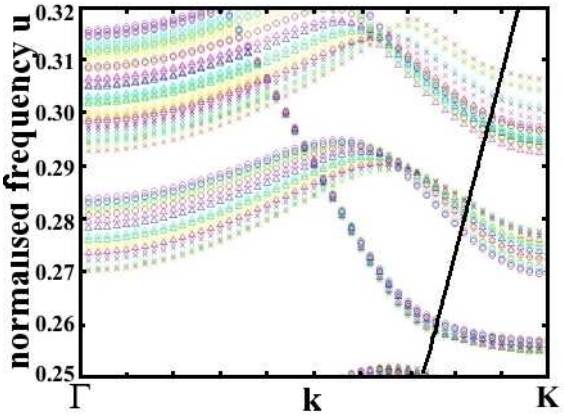

Figure 5.18: *Bandstructures for W1 waveguides with holes of radii between 0.26 and 0.35 r/a. A silica lightline has been superimposed. Although on first glance the change in normalised frequency seems small, it is significant in terms of measurement wavelength - for a 400nm period, the width of the fundamental mode distribution here is over 100nm. Furter, fine control of hole size is challenging.*

5.3.1.1 Hole-shape tuning

Instead, I have studied the effects of changing the shape of the holes, while retaining a constant filling factor. To do this in practice has required the development of small aperture/step size techniques, as discussed in chapter 3. Instead of circular holes we make ellipses, and rotate the axis of the ellipse with respect to the propagation direction. This breaks the symmetry of the crystal

- for example, the symmetry plane with purely real impedance used in the analysis above no longer exists, except of course in the cases where the ellipse axis is either parallel or perpendicular to the propagation direction. Hence we do not attempt any impedance analysis in this design. As was shown in chapter 2, it is the interaction with the projection of the fundamental lattice's K point into the bandstructure that determines the dispersion of the fundamental PhC mode; and by reducing the symmetry to purely translational, we can alter this projection. However, I have not pursued this theoretical basis.

Instead, we choose the fillfactor studied in the impedance work above, as it has significant bandwidth below the lightline, and then calculate the bandstructures for various ellipses. As reducing the symmetry also has the effect of narrowing the bandgap [13], we choose an eccentricity of 1.5. The bandstructures are shown in figure 5.19. From this, we can see that much finer control of the dispersion diagram is possible.

Next, we need to consider these devices in practice: we know that PhC waveguide transmission begins to fall off as soon as the group velocity differs from that of a ridge waveguide due to coupling effects, and yet to minimise the footprint of any potential device we need to operate in parts of the bandstructure with low group velocity. We choose two particular ellipses, with eccentricity of 1.5 and angles of 10^o and 50^o respectively, and calculate their transmission spectrum using 2D FDTD as shown in figure 5.20. The same input waveguide is used for each device.

This figure starts to suggest that designing a complete AWG device is unrealistic - because if the transmission is not equal in each arm, the interference pattern becomes more complex, and conventional methods of coupler design cannot be used. For conventional AWGs, pattern errors leading to unbalanced transmission are one of the major sources of imperfection in

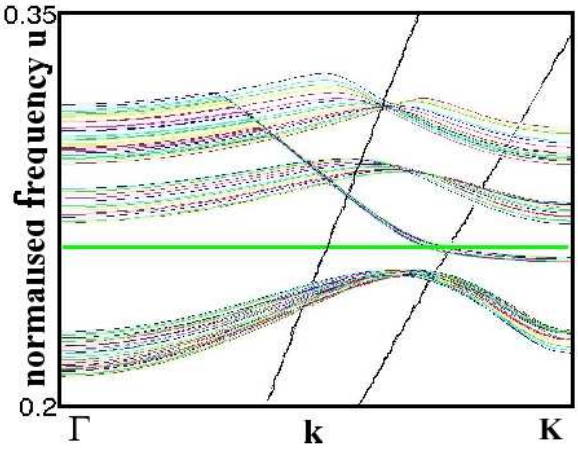

Figure 5.19: *Bandstructures for rotated ellipses, for 10º steps in rotation angle between 0 and 909º. The green line is the chosen operating point, selected such $u/a = 1.55\mu m$. The black lines are the air and silica lightlines respectively. We now have a much narrower range of u values forthe fundamental mode than for the hole size tuning in the previous figure, and we maintain a constant filling factor.*

operation, particularly in terms of crosstalk. However, we wish to explore the effects of these elliptical holes further, and to show the effects of the coupling limitations. Accordingly, we simplify the target device, and instead consider a passive Mach-Zehnder Interferometer (MZI), which uses these two ellipse angles to create a differential phase between the two arms, creating a notch-filter response.

5.3.2 Crosstalk

The next step is to determine the necessary separation between these two arms. To do this, we create two W1 waveguides through a block of crystal, and launch a pulse from a wire waveguide into one arm. Then, we monitor the overlap with the wire mode for transmission through the other arm. The

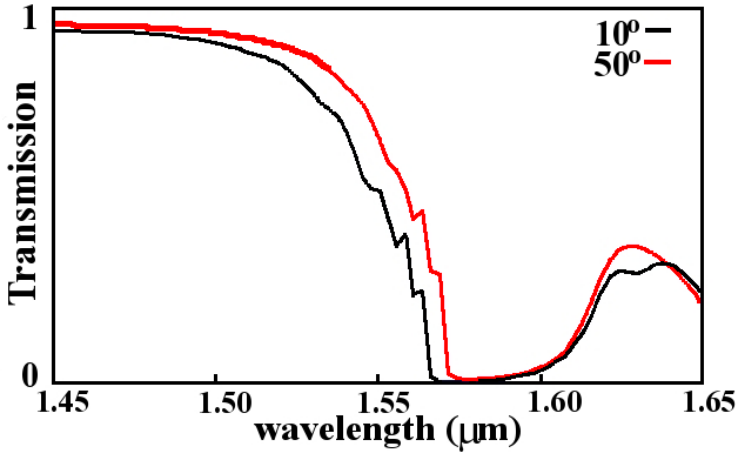

Figure 5.20: *Transmission spectra calculated using 2D FDTD, for 10° (black) and 50° (red) ellipses. Although both have similar peak transmissions, the 3dB point for the transmission on the band edges is separated by around 100nm. But to use the differential group velocities of these ellipses, we need to keep similar transmission levels. This figure demonstrates one of the major difficulties with PhC device designs.*

results are shown in figure 5.21. We select a separation of 7 rows on this basis. The four-port response for this separation is shown in figure 5.22.

5.3.3 MMIs as 1:2 splitters

Although using techniques described earlier, we can recover the phase response from this device either directly from the bandstructure or from transmission simulations, to discover the filter response we must also consider how the two outputs will be interfered. The 7 row PhC spacing is too great for diffraction effects to be useful, otherwise we could characterise the device response in terms of beam steering. Instead, we choose to use an MMI for both the input and output of the device. A simple rectangular MMI can be designed using equation 5.1 above, where the width W can be chosen independently. However,

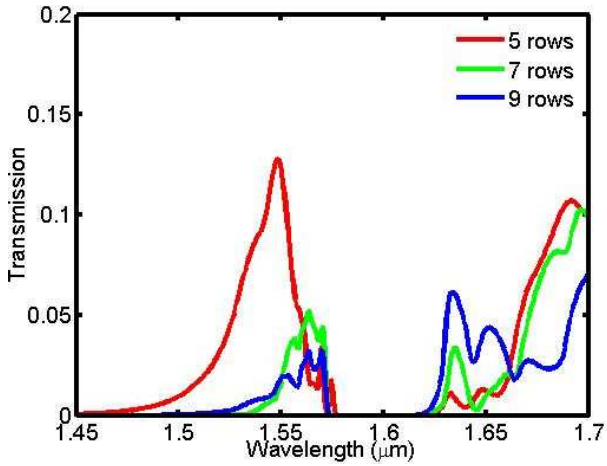

Figure 5.21: *Crosstalk for varying numbers of rows separating two W1 waveguides, calculated using 2D FDTD. Only the forwards transmission in the coupled arm is shown. Note the y axis scale is expanded. From this graph, we decide that 7 rows is an adequate separation.*

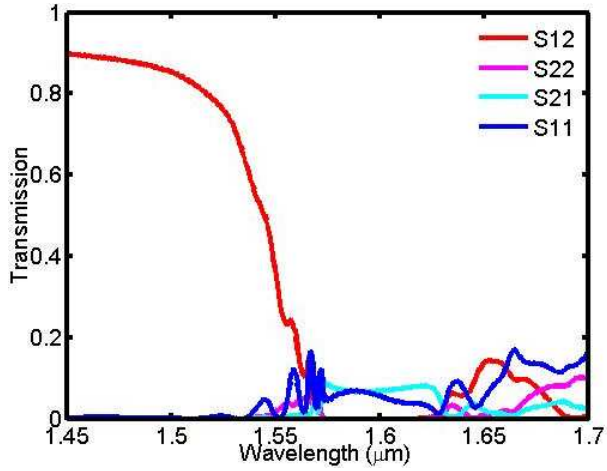

Figure 5.22: *Crosstalk for 7 rows separation, showing transmission and reflection in both arms. Note that the total transmission does not sum to unity, as we consider mode overlaps rather than total power in each direction.*

because this equation has been derived for low index contrast systems [54], some numerical tweaking is necessary. Here, I chose a width of 1μm, which works well in simulations - the transmission spectrum for the MMI is shown in figure 5.23. The resulting length is just over 3μm[6].

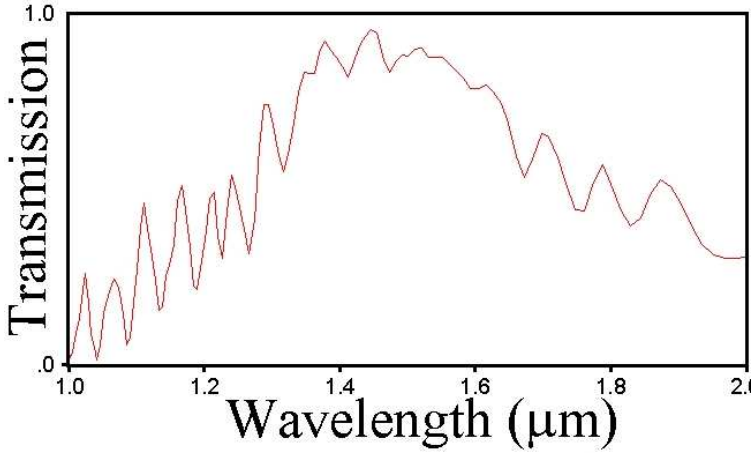

Figure 5.23: *2D FDTD transmission spectrum for the selected MMI. The transmission plotted is that measured in one output arm, and then multiplied by 2 to give the total. The 3dB bandwidth is over 300nm, which was initially thought to indicate a roubustness to fabrication tolerances.*

Although the bandwidth of this device is well over 300nm, significantly greater than our PhC, it only works in simulations. In fabricated devices, there is no error tolerance - a 20nm size error shifts the centre frequency by 50nm, and achieving this accuracy takes much process characterisation. When we also have to make the ellipses correctly, this becomes too time-consuming. A larger MMI is suggested, but due to the geometric scaling, this rapidly becomes larger than the PhC devices.

Remaining with simulations, we now drive this MMI backwards, changing the relative amplitude and phase presented to the two output arms, and

[6]Which would make it one of the world's smallest MMIs, if it worked in practice.

recording the transmission at the input. Figure 5.24 shows the result. Again, we see that it is necessary to have very similar output transmission, which the two PhCs do not provide.

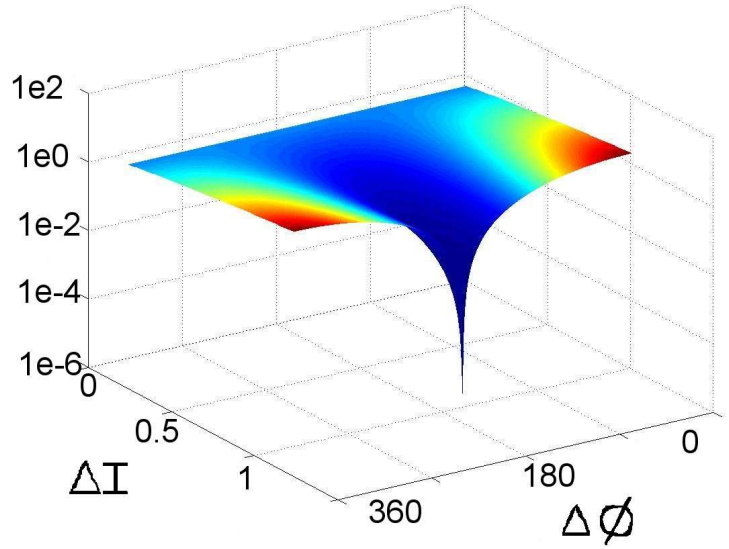

Figure 5.24: *2D FDTD transmission spectrum for the selected MMI, where the device is driven in reverse, with varying amplitude and phase between the arms. This shows the limits of differential transmission and phase to obtain a given extinction level.*

5.3.4 Simulated performance of complete MZI device

Finally, we wish to study the response of the entire device, which is shown in figure 5.25. Using the effective index determined from the bandstructure, we can predict that a device with a small number of periods will show the desired response. However, varying the device length increases the chances of seeing the desired notch-filter response within the available measurement bandwidth; but also increases the necessary simulation volume. To counter

this, we calculate the amplitude and phase for each of the individual elements, and determine the final response by multiplying these together as appropriate. The final simulated response is shown in figure 5.26. From this we can be confident that the elliptical arms do indeed have a different amplitude and phase relation. However, the slow-light regime, where the transmission falls off, shows neither good rejection nor high finesse peaks, which does not give confidence for using PhCs with slow light operation, for passive devices of this kind.

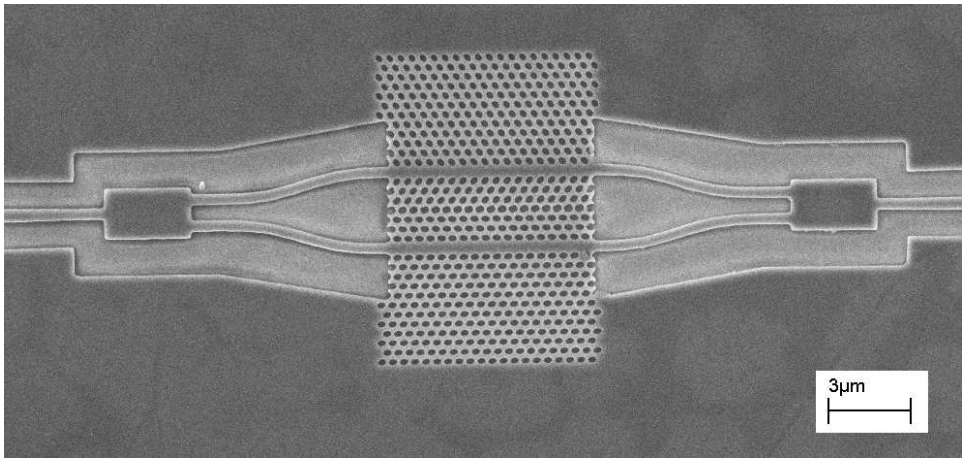

Figure 5.25: *An SEM image of the complete device being studied, consisting of input and output MMIs and PhC arms containing ellipses of different angles.*

5.3.5 Experimental results

As has been intimated, the device discussed in this section has not been successfully fabricated, essentially because it has too many features which are very sensitive to fabrication tolerances. As an example, the device shown in figure 5.25 has too large a fillfactor, which moves the operating point of the device both out of the MMI passband, and the available measurement

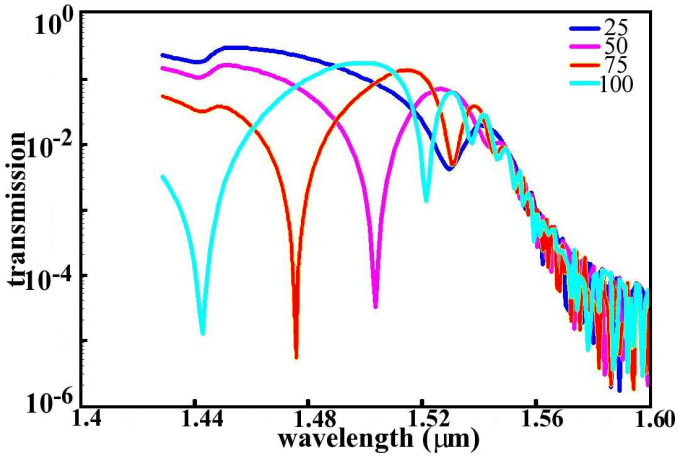

Figure 5.26: *Simulated complete device response, calculated by convoluting the response of the various elements. Note that this figure has a logarithmic scale, to show the strength of the rejection when the two arms are out of phase.*

bandwidth. However, work on this device has been very important to developing some of the fabrication techniques in chapter 3.

Nevertheless, it has been possible to make some preliminary transmission measurements on W1 waveguides with elliptical holes, which show that the transmission spectrum varies slightly with the ellipse angle, as has been predicted in the simulations above. The first example is shown in figure 5.27. Note the axes here, where we plot a log transmission against normalised frequency. The vertical scale makes the cut-off much more obvious, but the normalised frequency axis gives the clue to the results - this cannot be the designed W1 band. Instead, these measurements are for the high-fillfactor devices pictured, and show the PhC operating below the lattice bands. The experimental technique is fully described in chapter 6. With very careful calculation of bandstructures using 3D FDTD, we can duplicate this movement in the mode position, which gives a strong indication that the ellipse W1s would work as predicted, if they were fabricated correctly.

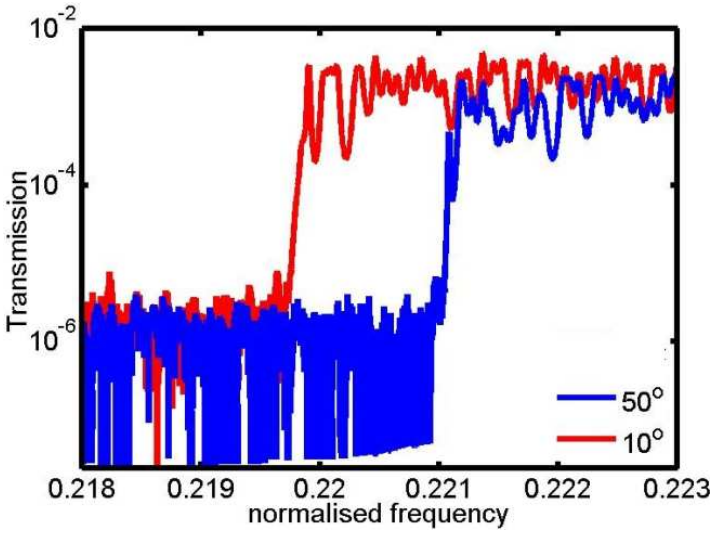

Figure 5.27: *Experimental transmission spectra for two W1 waveguides, with different angles of elliptical holes. The ellipses have different transmission responses, which is good, but the response is not at the normalised frequency expected. Analysis suggests that this transmission is actually below the crystal band - fabrication errors have shifted the device spectrum a long way from the target.*

Chapter 6

Realisation of optimised photonic crystal Y-splitters and bends

As has been mentioned previously, one of the strengths of PhC devices is their ability to confine a signal to a small volume, and route, bend, and split it. In this chapter, I discuss why this functionality is necessary for high density integrated optics, and the benefits a PhC approach provides. The literature of PhC bends and splitters is thoroughly reviewed, before presenting the methodology of our own device design, its fabrication details, and experimental results. Although the actual design work was carried out by T. Davies of Photon Design, and the initial fabrication and characterisation by L. Wu and T.J. Karle of this group, my contribution is significant - particularly depth of study, and analysis and interpretation of results. The majority of the simulation results shown in this chapter are my work, and I have also done supplementary fabrication and characterisation.

The work discussed in this chapter has been published as [34].

6.1 Integrated optics and routing

Conventionally, routing in optics means fibres - whether a signal must be transmitted across the lab or across the Atlantic. For long distance communications, the fibres need many support systems, such as multiplexers and amplifiers. Some of these are actually implemented in modified fibres - the Erbium-doped Fibre Amplifier (EDFA), for example. These can be readily spliced into the cable. For those functions which are not, it is necessary to route the signal, usually on chip, to the appropriate device.

At its logical extreme, it is envisaged that this kind of integrated optical routing could be used for inter-device communications in electronic integrated circuits, where data communication over copper is increasingly becoming a limiting factor in processor design [1]. Particularly for this kind of application, optical routing needs to be compact, with sharp bends. For any function utilising phase effects - modulation, for example - we also need to be able to split the signal into two (or more) parts. In the following sections, I will discuss the current state of the art in integrated optics routing, and the limitations of routing in standard waveguides, and then contrast PhC approaches.

6.1.1 Conventional on-chip routing

To bend a waveguide, one can begin simply by selecting the desired radius of curvature. For each specific waveguide geometry, this translates to a given degree of excess loss associated with the bend. In the simplest visualisation, to "keep up", the light propagating around the outside of the bend has to travel faster than that on the inside, as it must go further in the same time. For this to be possible, the guided wave must become evanescent, and so some energy radiates away as the light propagates around the bend.

In the conventional, mature silicon on silica technology, the very low index contrast requires bending radii on the order of a few centimetres [55][1], to have acceptably low losses. This cannot be improved in this material system. To have smaller bends it is necessary to change the platform, as the limiting factor is the index contrast. In SOI technology, where we have silicon $n = 3.5$ clad by silica $n = 1.45$ and air, and our typical single mode waveguide is 450 $\times$ 220nm, 90^o bends with of 1μm radius and bending loss <0.1dB have been demonstrated [57]. In this case, however, the waveguide mode is very sensitive to the edge roughness of the waveguide, which induces excess loss over and above the bending-induced losses.

As well as bending waveguides, it is often useful to split off part of a signal, whether for modulation applications, where a relative phase shift is induced between different branches of the same signal, or for monitoring, where a small part is detected (and hence absorbed). Here we will only discuss 50:50 splitting, that is dividing the signal into two equal arms, but other ratios can be viewed as an extension of this.

There are three conventional methods that can be used. The simplest is to branch off the waveguide into two parts. To minimise reflections, the branching angle must be low, and to usefully separate the arms, we must introduce waveguide bends. In low contrast (large) waveguides, this splitting can be done relatively simply, but at photonic wire scales it is lithographically challenging to fabricate the sharply pointed narrow trench that results from the low branching angle.

Secondly (in order of increasing complexity), there is the multimode interference coupler (MMI), discussed further in chapter 5. In low contrast systems simple MMIs work very well, and even in high contrast systems

[1]State of the art, or any optics textbook for the theory, e.g. [56]

high performance robust devices can be designed. However, there are still fabrication tolerance issues, unless the dimensions are of the order of 100s of microns.

Finally, one can use a directional coupler to separate off any desired fraction of signal. Again, these are wonderful devices on a macro scale, but become challenging to fabricate when small. For all of these devices, however, the state of the art is much advanced over that for PhCs, but nevertheless we will pursue PhC based solutions to signal splitting and bending.

6.1.2 Routing in photonic crystals

In contrast to large radius waveguide bends and small splitting angles of conventional integrated optics routing, we can consider the PhC equivalents. In principle, using a 3D PhC with omnidirectional bandgap, an intrinsically lossless bend can be made, with a bend radius of effectively zero [58, 59] at lattice angles. The crystal's bandgap prevents propagation in any direction other than along the defect, and around the bend, so light can only be transmitted or reflected. However, effective fabrication of such 3D structures is challenging, and this effect has not yet been demonstrated experimentally. Note, however, that the actual bend site represents an area larger than that of the defect waveguide, acting as a cavity. Therefore, although the system is lossless, transmission is resonant, and so may be narrow-band. PhC Y-splitters operate in the exactly the same way - instead of introducing a bend defect into the crystal, one introduces one with three paths which intersect at the same point.

If we have a truly 2D photonic crystal, with semi-infinite extent along the axis of the crystal, we can achieve the same effect - intrinsically lossless bending along the lattice directions of the crystal. This works with guiding due to the

photonic band gap effect only - that is, for systems of pillars in air [60]. Figure 6.1 shows a field plot for this effect.

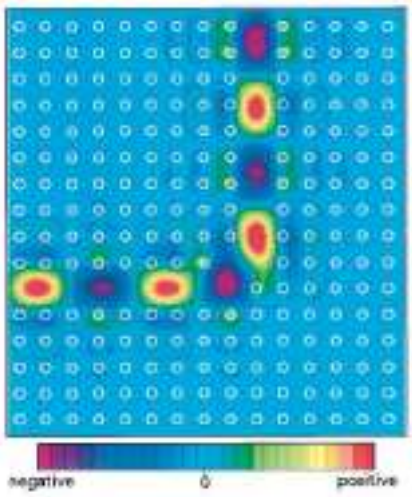

Figure 6.1: *A 90º Photonic crystal pillar bend, taken from [60], showing the field profile. Theoretically, such bends provide significant advantages over standard waveguide bends. The wave is solely band-gap guided in a purely 2D system.*

For the fabrication technologies available here, it is much preferred to use the opposite system of PhCs, namely etching holes into a dielectric structure. In this system, we do not have the simple elegance of the pillars in air bends, and hence their intrinsic good performance. However, it was felt useful to investigate PhC bends in this more realisable system. When this was first begun, the low etch roughness and hence low loss, small bending radius photonic wire bends mentioned in the previous section were not yet possible, so it was thought that PhC bends may be superior[2]. Similarly Y-splitters - improvements in lithography permit further miniaturisation of photonic wire Y-splitters, whereas for PhCs they can be made with the same technology as for straight PhC waveguides.

[2]Current thinking favours the much simpler photonic wire bends.

For our standard triangular lattice of holes structure, we can make an intrinsically lossless PhC waveguide, a W1 operating below the lightline. However, making a bend or Y junction violates the symmetry that allows this waveguide to be lossless, because we change the periodicity and hence the k-vector, and as a result operate intrinsically above the lightline. As a result, we do not have total-internal-reflection guiding, and so must have out of plane radiation - the device is lossy.

On the other hand, as discussed in the previous section, all waveguide bends are intrinsically lossy, so this in itself is not a reason to give up. Further, PhC bends have zero effective bending radius, so sharp corners are possible, unlike for photonic wires. In the rest of this chapter, we make use of simulation tools and optimisation routines to design a photonic crystal bend (and Y-splitter) with acceptable transmission and bandwidth.

6.2 Design of integrated PhC coupler, Y-splitter and bend

In the previous section, I have briefly described why a photonic crystal can be used to route light, and how this can be done. In this section, I again separate routing in photonic crystals into bends, Y-splitters and couplers into the photonic crystals (which are of course necessary adjuncts to using PhCs for routing). For the bends and splitters, I will discuss the literature - also known as "other people's designs", and then our design and how we arrived at it. Tapers are discussed in chapter 4, but have been used as part of the final device presented in this chapter. In subsequent sections, on fabrication and characterisation, we will consider a total device - a short taper as a coupler, Y-splitter, pair of bends to re-orient each branch of the transmission back to

a straight-through, and output couplers. An SEM of the final design for this is shown in figure 6.2.

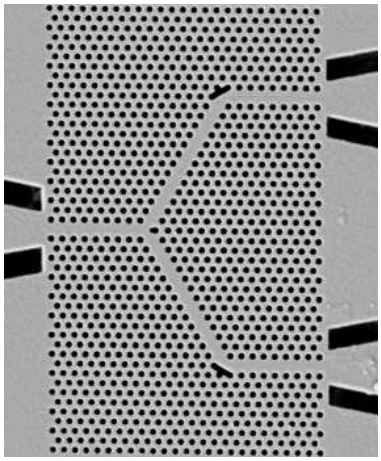

Figure 6.2: *SEM image of the PhC device, comprising short tapers, a Y junction and 60º bends. Its design, fabrication and characterisation are discussed in this chapter.*

Details of simulation and fabrication techniques have been discussed in the relevant chapters; although some additional theories are discussed here.

6.2.1 Evolution of PhC bends

Work has been published on the subject of Photonic Crystal bends by a variety of groups, with both theoretical discussions and measurement results fairly well represented. However, a large part of the published theoretical discussion is based on explaining experimental results, so I will proceed from simple to complex bends, citing theory and practice as appropriate. Note that many of these papers also discuss work on Y-splitters, which will be presented in the next section.

If we consider the basic photonic crystal waveguide to be a line defect created by deleting a row of nearest neighbour holes, then it follows that the

simplest bends occur by deleting holes along other lattice directions. Thus, for a square lattice, one naturally makes 90^o bends, and for a triangular lattice, 60^o or 120^o bends. Bends from the ΓK direction to ΓM, e.g. 45^o in a square lattice, are also possible, but intuitively are not expected to have high transmission or bandwidth. The most likely bend configurations are shown in figure 6.3.

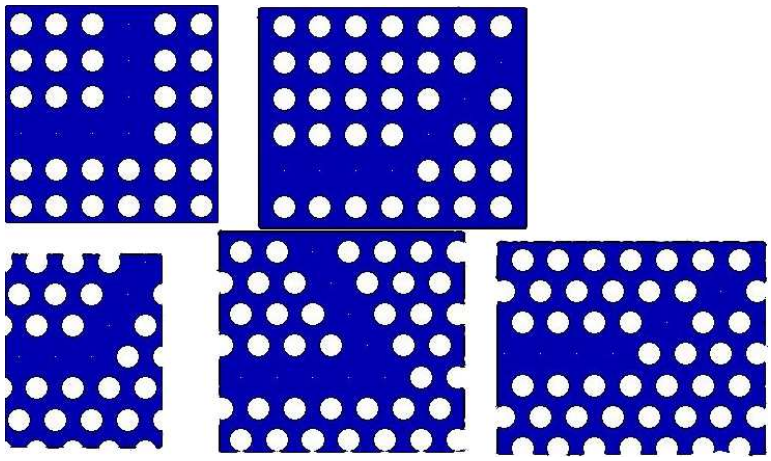

Figure 6.3: *Schematic of bends in square and triangular lattices.*

The first detailed experimental work on such bends was carried out in the highly-multimode W3 system, in a low-out-of-plane-contrast (or heterostructure) substrate by Olivier et al. [61]. The first demonstration of light transmission around a bend in a PhC was, by Baba et al. [62], and the first simulations by Mekis [60] albeit for rods in air. In Olivier's paper, the simple bend is called the "nonresonant bend", which is somewhat at odds with the terminology I will use. The 2D simulation and experimental results show a complex behaviour, that I would describe as resonant - but, with different peaks representing resonances for different modes in the bend. To improve the response - make it less complex - the bend cavity was subsequently altered by reducing the bend volume, and instead making a cavity with a simpler set of modes. Olivier et al. label this their resonant bend, which clearly uses a

defined resonance and coupling between the waveguides and the cavity. The transmitted intensity is not available due to the characterisation method used, and the bandwidth, as befits a simple resonant response, is narrow. For all work in photonic crystals, actually getting any light through the devices is a challenge, both due to the coupling issues discussed in the previous chapter and the high scattering-induced losses in early devices, before the fabrication technology was well developed. Although with hindsight it seems obvious that single mode waveguides will lead to simpler solutions, multimode systems such as W3s (which have wider waveguides, and less interaction with roughness in the holes, hence easier coupling and lower losses) drastically improved the chances of getting a measurable response!

For the next improvement over simple PhC bends, we remain with Olivier, in [63]. The waveguide type remains the same. In this paper, the techniques of conventional optics are used to improve bend performance: firstly, in a straightforward bend the field profile centre tends to shift outwards, such that a mismatch between the straight and curved sections can be observed. This can be compensated by enlarging the bend waveguide. Secondly, a low-Q resonator may be added at the inner corner of the bend. This becomes very complex for a multimode PhC waveguide, and as such, is adapted by Olivier et al. in some sense as slowing down the bend - so, the 60^o bend is broken down, intuitively, into some 30^o sections. For the best bend in this paper, a transmission around the bend in the fundamental mode of around 70%, and a bandwidth of 45nm are achieved. This represents a significant improvement - both a quoteable transmission and a reasonable bandwidth! Moreover, this is proof that conventional techniques are applicable to PhC devices. Talneau et al. [64] show similar results, on an InP rather than AlGaAs heterostructure at about the same time - this is based on moving holes around

in the bend - effectively the same approach, but without the justification in terms of conventional methods. The complex index approximation, which can be used in 2D models of PhCs to approximate 3D effects, is extended to bends in this paper.

Both of these approaches are extended in [65], which considers PhC bends on AlGaAs and InP platforms, on W1, W3, and W5 waveguides. Again, the conclusion is that the bends are better in larger waveguides, due to lower loss (mode overlap with the holes is smaller, reducing scattering losses) and easier insertion, however regardless of this the best peak bend transmission is less than 90%, additional to the loss of an equivalent straight PhC. This is for a W3 in an AlGaAs heterostructure.

The rest of the literature on PhC bends concerns work on membrane structures, as opposed to heterostructures. Recalling the lightline arguments presented in 2, these structures can be designed to operate free of intrinsic diffraction losses, and as such can be expected to have lower losses. The first experimental report of a PhC bend in membrane is by Chow et al.[66] in GaAs slabs, and shows near perfect transmission, over a very narrow bandwidth, and explains that this narrow bandwidth is due to the bandstructure. On an air-bridge GaAs membrane, Sugimoto et al. [67] have made simple W1 PhC bends, with very high quality fabrication. These are probably the best simple PhC bends that can be made - and they have a bandwidth of less than 30nm! Actual transmission values are not recorded. Sugimoto et al. have also published further work on integrated PhC devices including bends [68], but these remain simple bends on air-bridge structures (their interest is in integration and active properties), with very similar results.

Next, we have optimised bends from Miao et al. [69], again on GaAs membrane. Three bends are studied in this work, a simple one for comparison,

a bend with a mirror reflector at the outer corner, and a bend with a mirror reflector and deleted inside-corner hole. These are both fabricated and simulated in 3DFDTD. The last of these exhibits a 140nm bandwidth with less than 1dB loss per bend. The simple bend has a peak transmission at a similar level - but over the same narrow bandwidth as in [67]. The average loss over the 140nm bandwidth for the simple bend is approximately 10dB.

Finally, we consider Frandsen et al. [70]. Based on [63, 61, 65, 69], it is clear that non-simple bends have much higher bandwidth and transmission than simple ones, regardless of material system or waveguide structure. Frandsen follows this to its logical conclusion, and uses topology optimisation schemes to maximise the response of his W1 bends on SOI. In practice, this means that the sizes and positions of the holes close to the bend were carefully chosen. These show less than 1dB additional loss about a 30nm bandwidth, and also resonant transmission in another part of the spectrum. This has been continued by Tetu [71], with further optimisation to find a single bend design with high transmission and wide bandwidth for both TE and TM polarisations.

Chutinan et al. [72, 73] present theoretical work on a different method of improving bend transmission. They change the waveguide by putting small holes on lattice sites at the bend and for several periods on either side. This has the same effect as the methods above - reduce the waveguide volume to reduce the number of modes, and has the same result - broader bandwidth and higher transmission. Again, though, numerical design gives the best results, and so this is an alternative, as opposed to better, method.

6.2.2 Design process for our bend

In the following, we discuss the principles of design for our PhC bend. In this case, we choose to operate, not in a membrane, but instead in a low

out-of-plane contrast system, for which the waveguides are not intrinsically lossless. Although the above shows clearly that the best bends are in air-bridge structures, heterostructures are advantageous if active functionality is desired from the integrated optics routing, and so this approach is also very valid.

We start with an AlGaAs heterostructure waveguide, which supports a single mode at 1550nm. We pick[3] the convenient pitch of 440nm (easy to fabricate and with a bandgap centred around 1550nm), and consider the hole size. We choose r/a =0.33 to start with, and determine the effective index of the substrate. Further, we choose to work in the W1 regime - given that we will also be designing an efficient coupler to these waveguides, the reduced number of modes in the bend presents a stronger advantage than simply greater transmission from a wider waveguide. From this, it is possible to determine the bandstructure, as shown in figure 6.4.

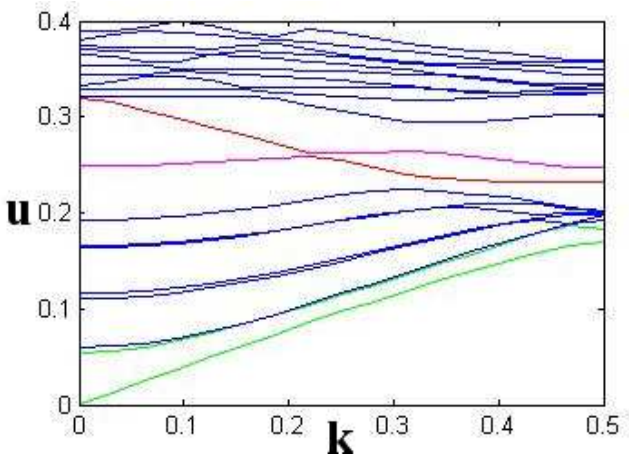

Figure 6.4: *Calculated 2D bandstructure for a W1 in an AlGaAs heterostructure, with r/a = 0.33. Fundamental mode is red in the band-gap and green elsewhere, lattice states are blue, odd mode is magenta.*

[3]after several iterations...

However, at the actual time of device design, no bandstructure tool was available - and so a mode-solving tool was used - with the result in the subsequent 6.5.

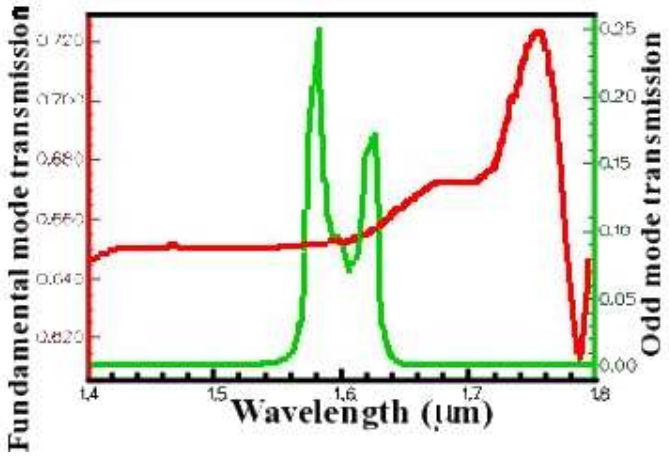

Figure 6.5: *At the time these devices were designed, no bandstructure tool was available to us. Instead, it was necessary to calculate a mode spectrum for a real period and wavelength. The fundamental mode is red and has the intensity on the left hand axis, the odd mode is green and is plotted against the right hand axis. From this figure it was decided that the target 1.55μm operating point was too close to the odd mode, so the lattice must be changed.*

Whichever of these we use, we can see that for this combination of lattice pitch and hole size our chosen wavelength coincides with the odd mode of the crystal. Clearly, single mode operation is desirable, and so we choose to shift the hole size slightly, to $\tilde{0}.31\ r/a$, to give us a single mode regime (The decision to shift the hole size rather than the period is a consequence of the available tools - in a bandstructure tool, period is normalised out, so no further calculations are necessary but in a transmission model, recalculation is required either way, and in this CAD package, hole size is the easy option). Note, now, we have found a fabrication constraint - if the hole is 5% too large we will

again operate in the multi-mode regime, which will have a significant impact on performance - so we need to err on the low side! The changed bandstructure is not plotted, as the result of the changed design is not obvious over the scale in figure 6.4. Figure 6.6 shows the change in the mode calculation.

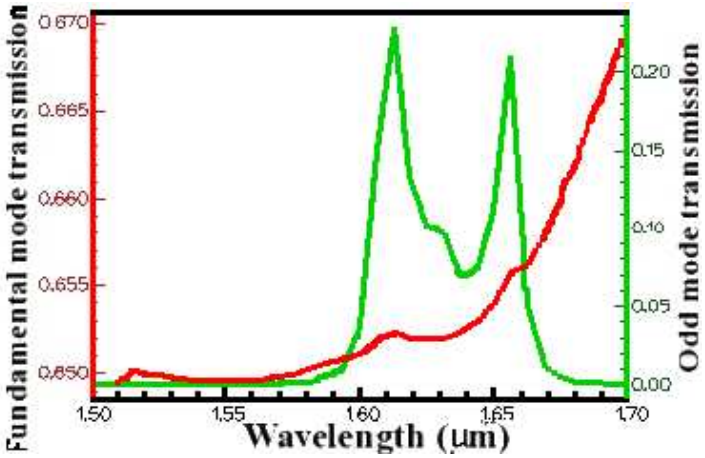

Figure 6.6: *Mode spectrum as in the previous figure, for the same period but slightly reduced hole radius - without a bandstructure tool, it is as easy to tune the device by changing hole size as period. The fundamental mode is red and on the left hand axis, the odd mode is green and on the right hand axis; this time the operating point is sufficiently far from the odd mode.*

Figure 6.7 shows the 2D FDTD transmission spectrum of 20 periods of this crystal, calculated by launching a planewave pulse inside the crystal and taking the positive flux also inside the crystal 20 periods away from the launch. This is not an accurate transmission spectrum, but merely indicates the operating point on a familiar curve.

Now, we put a bend in our waveguide, at 60^o to the propagation direction. The transmission is simulated using 2D FDTD as before, giving the results in figure 6.8. This is shown at the same wavelength scale as for the W1, to permit comparison. The subsequent figure 6.9 shows the transmission over the

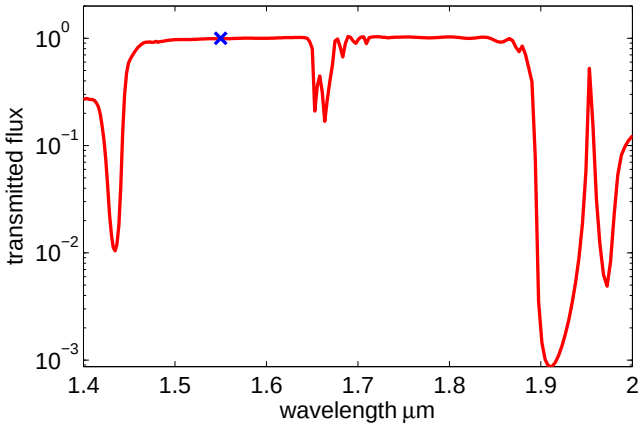

Figure 6.7: *The FDTD transmission of the crystal, calculated using CrystalWave, for the structure determined in the previous figures. The cross indicates the centre frequency of the input pulse for the simulation, at our 1.55µm operating point. The dip between 1.6 and 1.7µm is the signature of coupling into the odd mode, and our operating point is well away from this.*

wavelength range of interest.

Note the form of the transmission - although very high at the design wavelength, it falls off very quickly, with a 3dB bandwidth of only 10nm - showing that the transmission is resonant in nature. Recall that the 2D model assumes an infinite vertical dimension, and hence the PhC is intrinsically lossless. In the previous section, we discussed a simple bend fabricated by Sugimoto et al. [67, 68] on a GaAs membrane. Our simulation result agrees well with both the simulation and measurements in these papers - so this bend is showing the expected result.

To improve the bandwidth with high transmission, i.e. the 1dB or 3dB bandwidth, we could consider the bend as a cavity, and treat the problem in terms of impedance matching (see chapter 5). The impedance mismatch between the waveguide and the cavity can be thought of as giving rise to the

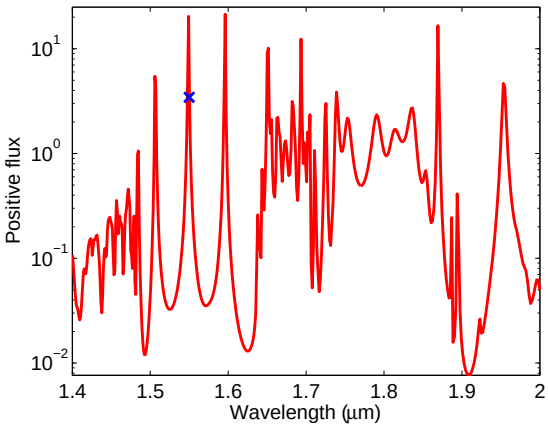

Figure 6.8: *The FDTD transmission of a simple bend, calculated using CrystalWave. The cross indicates the centre frequency of the input pulse for the simulation, at our 1.55μm operating point. Not the many complex resonances, as compared to the straight W1 waveguide in the previous figure.*

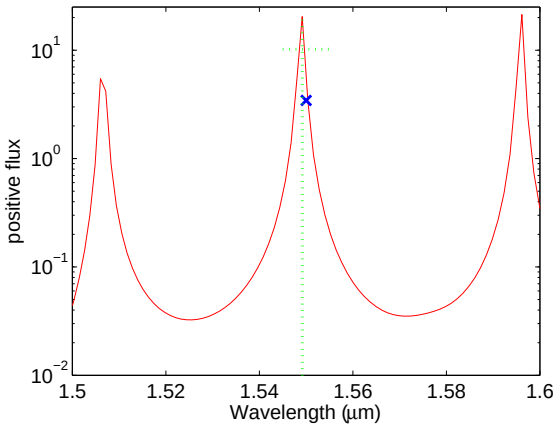

Figure 6.9: *Again, the FDTD transmission of a simple bend, calculated using CrystalWave. The axes are rescaled to allow us to determine the bandwidth of the resonance. The cross indicates the centre frequency of the input pulse for the simulation, at our 1.55μm operating point. The dotted lines show the centre of the resonance and the 3dB level.*

resonance effect. Instead, we need a detuned device, so that we have good transmission - equivalent to matched impedances - over a wide bandwidth.

Without considering the mathematics, let us pause and look at the geometry of the system. As previously mentioned, in simple terms the problem is that the bend has a higher volume than the waveguide, and is not single moded. So, to improve the response, we need to reduce the waveguide volume at the bend. As shown experimentally in [69] one way to do this is to deform the crystal at the outside of the bend, to form a continuous reflector. Alternately, the hole at the inside of the bend can be moved [69, 65].

We could now follow the method of Boscolo et al. [51], and consider the tuning of these two kinds of defects at least semi-analytically. Instead, we choose to pursue a so-called inverse method, and use optimisation software in combination with numerical simulation. Using CrystalWave's frequency domain engine[4], we simulate the structure shown in figure 6.10, with the constraints as shown in the figure - we move the inner hole, and the outer mirror. Note that this is not the ideal termination of the crystal on the, as was later discovered. The angle between the new propagation direction and the surface normal of the "facet" causes diffraction effects and beam steering, and reflection back into the crystal different to that when the bend design is integrated into a complete device. The optimisation is performed at a single frequency, using a global algorithm - one which searches the whole parameter space of the design, rather than trying to home in on a local maxima. In this case, the global algorithm rather than an optimising one such as a genetic algorithm is very useful, as we can be confident that we will find several peaks, some of which will correspond to narrow bandwidth resonant states - such as that found by [70]. We therefore use a pulsed 2D FDTD simulation for the

[4]discussed in chapter 2

parameters relating to the highest peaks, and use this to select the final design. This is verified using 3D FDTD as part of the complete design - computational constraints do not permit 3D FDTD for optimisation! The response is shown in the figure 6.11. Experimental testing will be discussed subsequently.

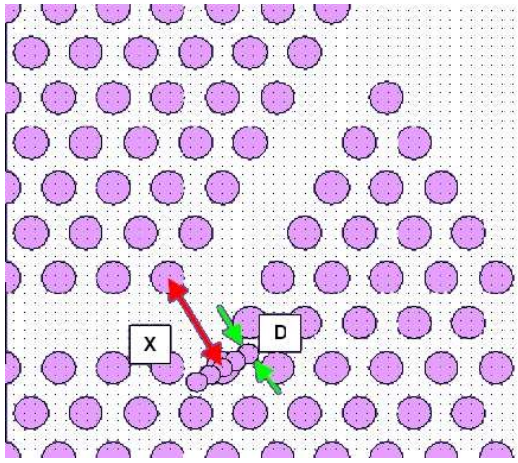

Figure 6.10: *Constraints for bend optimisation. The global algorithm used will try all possible combinations of distances X to the line of small mirror holes, and diameters D of these holes.*

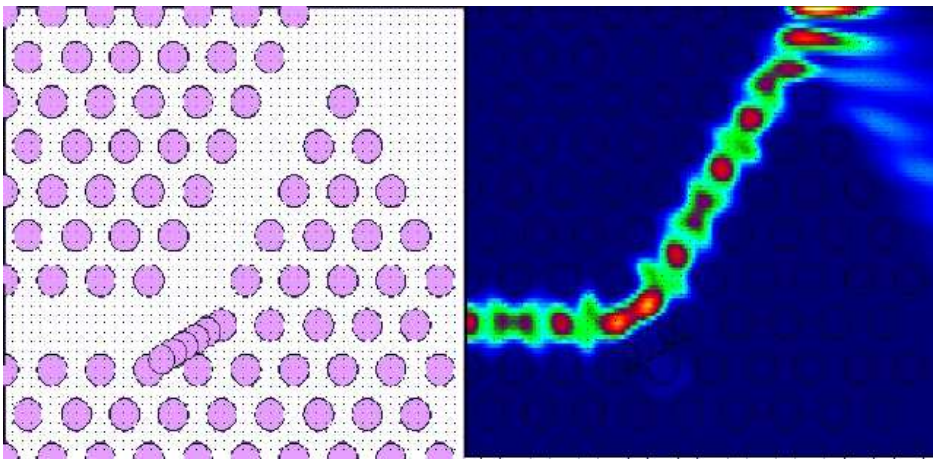

Figure 6.11: *Layout and field plot of the optimised bend, simulated using 3D FDTD.*

6.2.3 Y-splitter designs

We now trace the design of a PhC Y-splitter in the same way as for the bend above. The first theoretical report that a Y junction is possible in a PhC, for a rods in air structure, is in [74]. They note that careful design is necessary, because the Y cavity is different to the waveguide, and so reflections are possible. A mathematical approach to mode matching for resonant transmission is given in [75]. Here, tuning rods are introduced with arbitrary position and size, and adjusted to maximise the simulated transmission. As with the bend, we can start with a simple W1, or W3, or W5... and delete some holes, to make a simple Y. There are also alternate approaches - PhC based directional couplers for power splitting are considered in [76], which could be integrated with the bends discussed above to make an effective Y-splitter - although this would need to be larger than a straightforward Y. Also, a W3-like device, comprising larger holes rather than no holes for the defect waveguide, has been experimentally demonstrated [77]. Finally, coupled-cavity PhC Y-splitters are discussed in [78], in the microwave regime.

Regarding the more realistic air-holes systems, Boscolo et al. present a semi-analytic method based on impedance matching [51] and apply this to Y-splitters [79]. However, the results are necessarily for a single wavelength, and thus can lead to a very narrow-band device.

As with bends, the first experimental characterisation of a W1 PhC Y, in the simple deleted holes configuration, is by Sugimoto et al. [80]. The device is an AlGaAs air-bridge-type membrane. The bandwidth is less than 20nm, and the transmission very low. I would expect this device to be displaying resonant transmission with an almost perfect transmission at the Y, as propagation in the fundamental mode at the Y has been identified. The low measured transmission is likely to be related to input coupling into the photonic crystal

- the paper does not discuss the excess loss due to the bend. This is born out by his subsequent work, still using simple Y-splitters [67, 68]. In the latter, he also shows a small improvement in transmission by retaining the hole at the centre of the Y.

Experimentally, an efficient and broadband Y-splitter in an air-holes heterostructure has been demonstrated by Wilson et al. [81], intuitively using a hole in the centre of the Y cavity to reduce the possible mode volume, and a small tuning hole towards the input side, to remove the threefold symmetry which has been shown to reduce performance. This follows the approach in the theoretical rods-in-air papers above. A 40nm bandwidth with high transmission and 84% relative to a W1 of the same length are claimed.

Finally, Frandsen et al. [70] apply the same topology optimisation scheme to Y-splitters that they have to waveguide bends, giving devices without additional loss over the straight PhC waveguide, over a 25nm bandwidth. The 3dB bandwidth is not noted in the paper.

6.2.4 Design process for our Y-splitter

To allow for subsequent integration, we attempt to design a broadband photonic crystal Y-splitter, using the same basic PhC lattice as for the bend above. As with the bend, the simple Y suffers resonant transmission and narrow bandwidth, as shown in the figure below 6.12, and in agreement with the literature.

For the simple Y-splitter, we can again consider the narrow bandwidth to be a consequence of the impedance mismatch between the large junction cavity (now even larger than the bend) and the chosen W1 waveguide. We therefore pursue a similar strategy to improve the performance - we insert a small tuning hole into the cavity, to reduce the volume, and numerically alter its size and

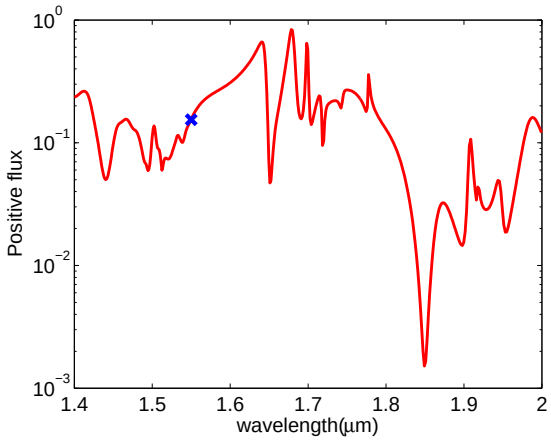

Figure 6.12: *Transmission spectrum of a simple Y-splitter, calculated using Crystalwave FDTD. The limitations of the software gridding mean this is representative, rather than accurate. Several narrow peaks are seen, which indicate resonant transmission utilising the various modes in the cavity of the Y.*

position to optimise the response, following the same verification strategy as above. The optimisation variables and layout are shown in 6.13. This strategy was also pursued by Wilson et al. [81], without the numerical optimisation. As discussed above, they show significant improvements by this method, but with the aid of numerical optimisation ours are better yet.

We follow the same optimisation strategy as for the bend - a single-frequency frequency-domain calculation, combined with a global search algorithm charting the transmission versus the position of the additional tuning hole. We then use a pulsed FDTD calculation to determine the bandwidth for several of the optimum points found, and choose one that has both high transmission and broad bandwidth. The field profile is shown in 6.14.

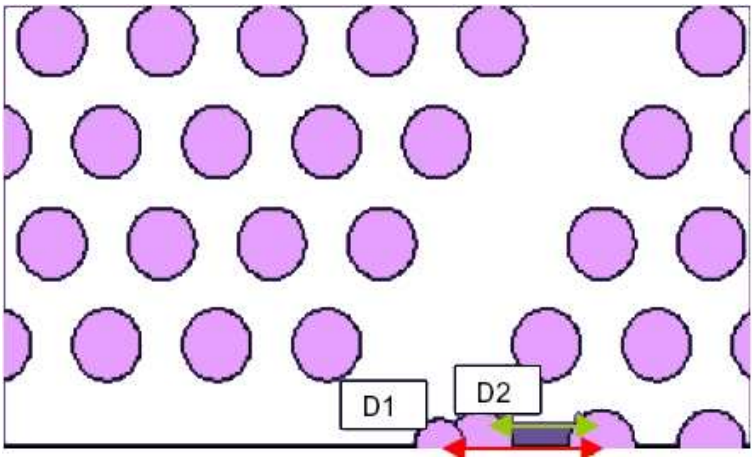

Figure 6.13: *Constraints for Y-splitter optimisation. Note the use of a symmetric boundary at the lower edge, to minimise the simulation size and remove odd modes. The small hole D1 and large hole D2 can be moved independently along the axis.*

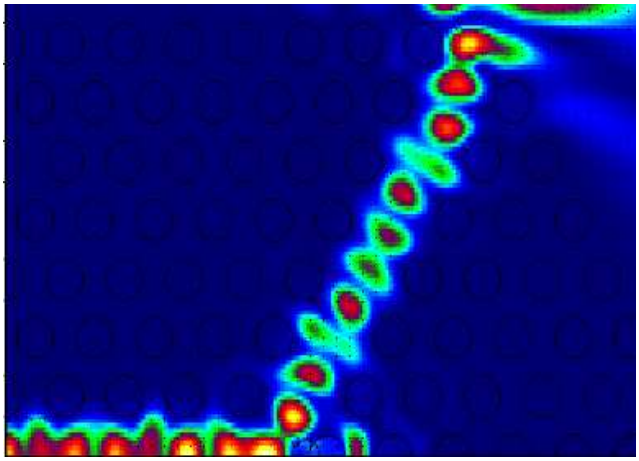

Figure 6.14: *The field profile for the optimised Y-splitter at the 1.55μm operating frequency.*

6.2.5 Tapering and coupling into photonic crystals

To make the Y-splitter and bend discussed above into a characterisable device, we need to find some way to connect them to the outside world, to a scale larger than that of our PhC. The many issues surrounding this are discussed in chapter 4. However, as a part of the design for this device, we have an opportunity to apply the optimisation tools used with the PhC to the coupling problem.

We choose to use 5μm wide waveguides as an intermediate step between optical fibre and our photonic crystal. These are deeply etched in the heterostructure, and support a suitable fundamental (as well as many higher order modes). With the aid of a 60$\times$ lens, we can couple into these fairly readily (the detailed measurement set up is discussed in a subsequent section). Next we must determine a suitable method to couple between this 5μm waveguide and the PhC device. The details of the design are discussed in chapter 5, but a brief overview is included here.

The simplest coupling method is to use an adiabatic taper. Such a taper has to be at least 200μm long, i.e. twenty times longer than our projected PhC device. Instead, we choose to use a trapezoidal multi-mode interference coupler (MMI). Such devices are also referred to as "spot-size converters".

To reduce simulation time, we consider a photonic wire physically matched to our photonic crystal, and then optimise the taper to maximise coupling between the fundamental of the 5μm waveguide and this wire. The taper layout and the optimisation parameters are shown in figure 6.15. As with the PhC, a global algorithm is used for optimisation in the frequency domain. However, memory limitations in crystalwave become significant for this large taper, and methods used to get around them affect the result. Figure 6.16 shows the 2D transmission spectrum calculated using the FD engine, which

gives the general form of the response. Figure 6.17 shows the field profile at the design frequency. The different field patterns in the input and output taper clearly indicate that there is reflection from the crystal.

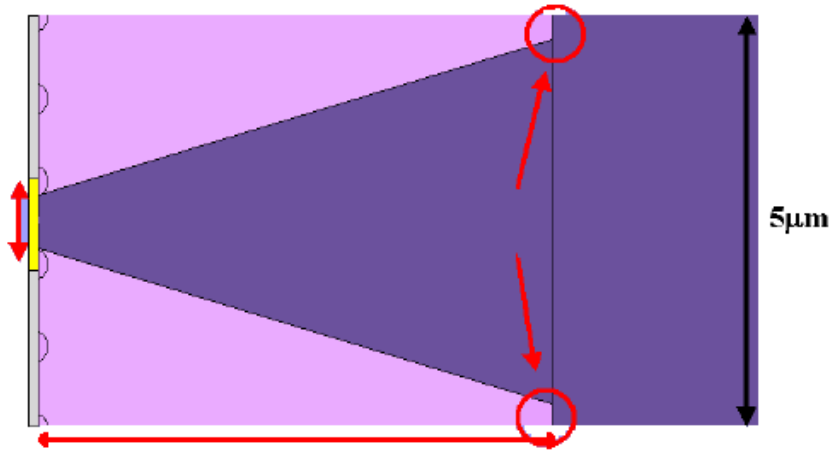

Figure 6.15: *Layout and optimisation parameters for the short taper.*

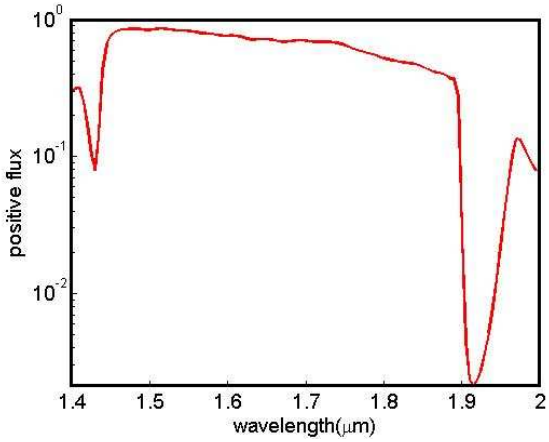

Figure 6.16: *Transmission spectrum for the optimised short taper. This is repeated from and discussed further in chapter 5. Note the transmission bandwidth is large, which indicates a robustness to fabrication tolerances.*

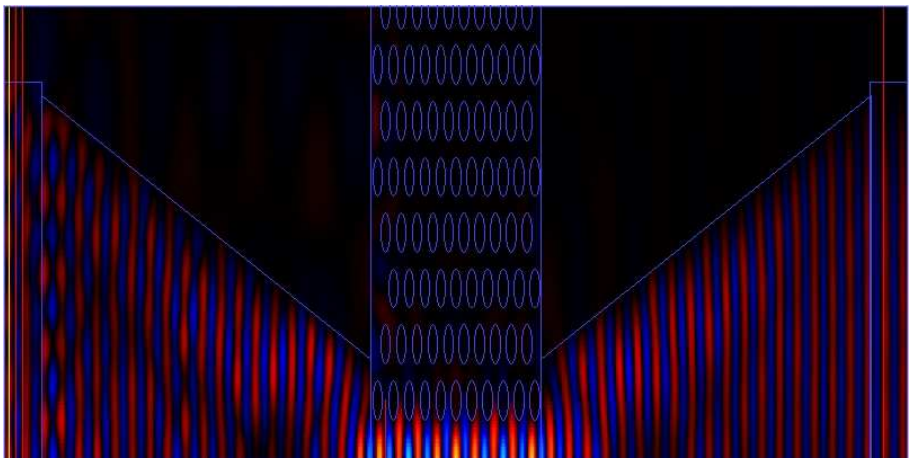

Figure 6.17: *Field profile for the optimised taper coupled to a short W1 waveguide, at the 1.55μm operating frequency. This shows that the taper couples effectively to the fundamental mode of the PhC.*

6.2.6 Verification of design

We now have a complete design for an integrated photonic crystal device - a large-scale waveguide, a spot-size converter, photonic crystal input coupler, and a suitable photonic crystal, complete with defect waveguides, Y-splitter and bends, plus suitable outputs back to fibre. Since fabrication and testing of photonic crystal devices can be somewhat involved, it seems useful to firstly verify the complete design - to check that interaction between components does not cause problems. Moreover, we have used a variety of models, tools, and software packages in the design - so, checking for consistent results seems wise.

We approach verification with a blunt instrument - clever tricks are useful, but we need to show that approximations made are valid. Correspondingly, we go for 3D FDTD. We choose Crystalwave as the verification simulation results shown are from Tom Davies at Photon Design. This approach has two major drawbacks - firstly discretisation in 3D needs many more points than in 2D, all

of which must be held in the computer's memory during calculation. Secondly, since so many points need to be calculated at each time step, the running time is long - and we need to run for a large number of timesteps to ensure a valid result. A more detailed discussion of the pros and cons of this approach can be found in chapter 2, and the notes on this taper in section 5.

In figure 6.18, we show the best compromise structure. We include the photonic crystal input taper, Y-splitter, and a bend in each arm. The output crystal is connected to the matched wire used in the taper design - there is not enough computer memory available to perform the calculation including output tapers.

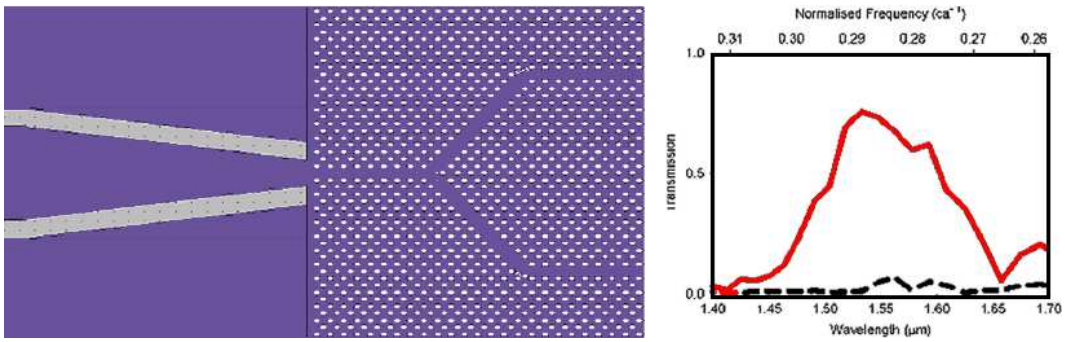

Figure 6.18: *Layout and 3D FDTD result for verification of the optimised taper, Y-splitter and bend, courtesy of Tom Davies (Photon Design) [34]. Memory constraints preclude simulating the output tapers. The peak of the optimised device transmission is over 70%, shown in red, and the 3dB bandwidth is ~150nm. Transmission with the optimised taper, but simple Y-splitter and bend is shown in black.*

At the time of calculation, the software used was limited in that it did not possess a suitable mirror boundary, which could in principle be used to half the size of this simulation. Although this could have provided valuable information, since we could have also considered output coupling, it does

mean that odd modes are included in the simulation - which might matter, in the Y. The simulation result is very angular/discretised in frequency, which implies insufficient resolution in the Fourier transform that gives the spectrum, and hence insufficient running time. This is unavoidable due to the memory management methods used in this package - it was not possible to run the calculation for longer.

Usually, we analyse the results of simulations by considering the overlap integral of the input and output waveguides - and take care that these waveguides are identical. Obviously, though, in this case they are not. Instead, we use the energy that is transmitted, with careful normalisation. Happily, this is implemented by default in the software used. The result is shown next to the simulation layout.

6.3 Device fabrication

We now need to realise our design. The general fabrication techniques that are used are described in chapter 3, here we only discuss the detail.

The device as designed above is fabricated on an AlGaAs/GaAs wafer grown with the desired layers by the National Centre for III-V Semiconductors, Sheffield University, via MOCVD. Further, we make lithographically tuned variations on the pattern and reference structures for normalisation. The lithographic tuning is chosen as $\pm 5\%$, which is too large a step, as is evident in the results - there are gaps in the transmission versus normalised frequency plots due to this. However, ebeam time is a constraint here.

For the results presented here, the pattern was defined in a 200nm PMMA mask using a Leica EBPG-5 Beamwriter at the Nanoelectronics Research Centre, University of Glasgow, and then transferred into a 300nm PECVD

SiO2 hard mask using RIE in a fluorine chemistry. The structure was then etched into the heterostructure using CAIBE in chlorine chemistry. Particular attention is paid to the etch depth to improve comparison with the 3D modelling. In general, loss for PhC heterostructures is proportional to etch depth up to the limit of the vertical mode extent, and so it is necessary to go past this limit - we do not want the measured device response to be buried by excess loss. Results from etch test structures prepared with these devices, show that the PhC was etched to a depth of approximately 2μm - an aspect ratio of around 7:1 for the holes. This is shown in figure 6.19. Finally, the hard mask is removed by returning the sample to the RIE. Again, this is to minimise excess losses, this time due to scattering from the surface of the device. We cleave the fully fabricated sample to have high quality facets for measurement.

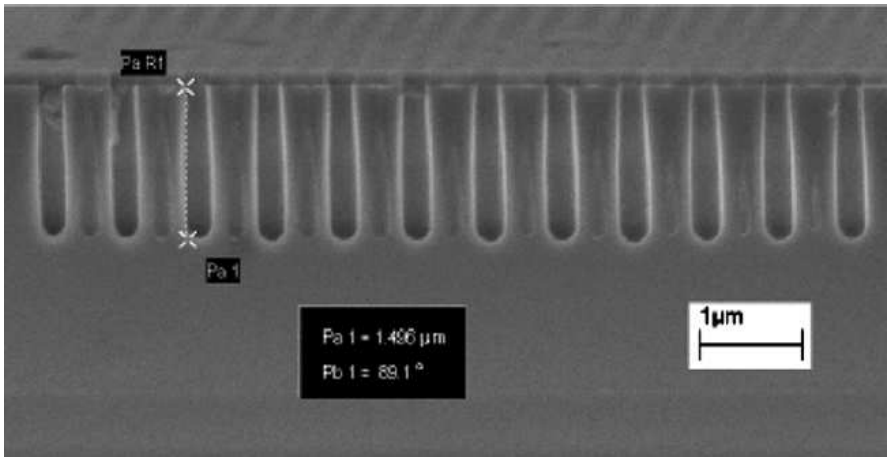

Figure 6.19: *Profile of CAIBE etching of AlGaAs heterostructure, for the optimised devices. The etch depth is approximately 2μm, much greater than the 1.5μm required, and the side wall verticality is good.*

We will discuss in particular the response of two measured devices - a double 60^o bend, and a Y-splitter as shown in the integrated device response above. Top-view SEM images of these are shown in figure 6.20. Three reference

structures are also fabricated - a straight W1, a pair of tapers surrounding a wire, and a 5μm waveguide with no device.

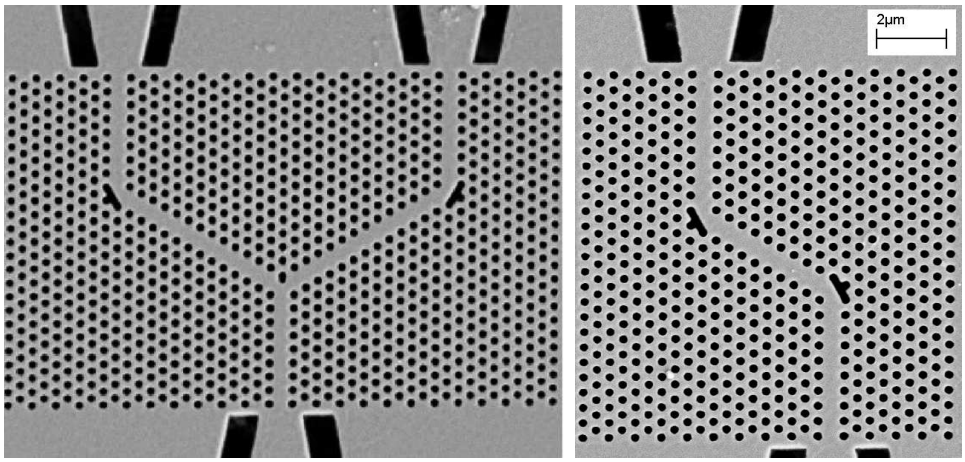

Figure 6.20: *SEM images of the optimised Y-splitter and double bend devices.*

6.4 Measurement setup

To characterize the Y-splitter device,bends and tapers, an erbium doped fibre (EDF)-type source with a bandwidth of 50nm around a centre wavelength of 1550nm was used. The response was recorded using an optical spectrum analyzer (OSA) at a resolution of 0.01nm over the 50nm bandwidth. A Glan-Thomson polariser is used on the output of the line to control polarisation (Hole-type photonic crystals gap-guide only in TE. However, for this structure, an adequate, albeit lossy TIR guide is formed for TM polarisations, so we must take care to filter this from the experimental results.). To couple light into and out of the die, we use 60× aspheric lenses. From Fresnel, we expect an input coupling efficiency on the order of 10%, and the source is unpolarised - hence we are detecting very small signals. At the end of the free-space line, we couple back into single mode fibre using a similar lens, and then use a fibre

splitter between a photodetector and the OSA. Optimisation of coupling is performed by hand, via micrometers and the photodetector response. Extra care is needed for the Y-splitter - optimising transmission in one arm at the expense of the other is not helpful, as it will not reveal the balanced device response.

6.5 Principles of measurement analysis

The purpose of devices such as Y-splitters and bends is to transmit light, over a broad bandwidth and with little loss, in a path other than straight through the device. One can, therefore, simply measure the transmission spectrum and claim that the device is characterised. Bearing in mind the lithographic tuning and 50nm source width mentioned previously, we stitch together the various spectra as shown in figure 6.21. The spectrum here is normalised against that of the source. The high facet reflectivity naturally gives a strong Fabry-Perot response. Since we could claim the fringes have nothing to do with the device behaviour, and are solely related to the cleaved facets, we could simply stitch together the peaks of this spectrum, and claim a bandwidth over 100nm for each device type, and relatively poor transmission.

However, this is not the only way to consider these measurement results. If we again consider our very strong Fabry-Perot fringes, we recollect that this is due to high facet reflectivity. Given the coupling set up, though, we can recall another effect of this reflectivity - high external reflection, too! That is, not a great deal of the available signal can make it into our device, regardless of how well we optimise transmission - and this fact makes the final result incredibly sensitive to input coupling. Thankfully, we can use the Fabry-Perot fringes in such a way, as to establish coupling-independent losses for our devices.

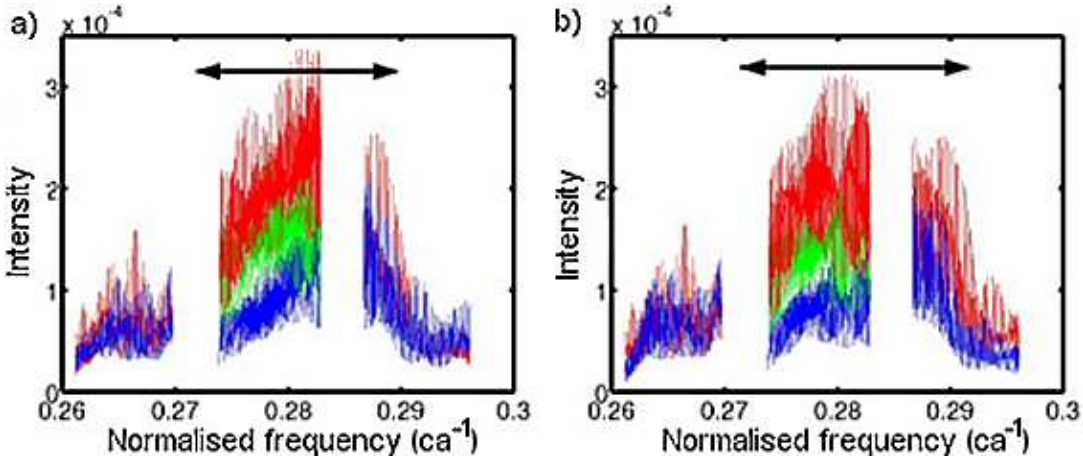

Figure 6.21: *Response of the Y-splitter versus normalised frequency, using the lithographic tuned devices. Note that the spectra do not stitch together, but we can still estimate the bandwidth from this plot.*

As an aside, we mention here the "conventional" method for establishing the loss of photonic crystal devices[5]. This method is commonly known as cutback. One creates PhC waveguides of varying lengths, but with identical input coupling sections (tapers and waveguides, for example). Then plot transmission as a function of device length. With moderately careful design, the loss in the non-PhC sections is almost negligible, and hence effects due to different waveguide lengths can be neglected. The loss due to taper sections, input coupling, and coupling to the PhC is in principle the same for each device, and so can be lumped together in the intercept of the graph, and the remaining line represents the PhC loss as a function of length.

This method does, unfortunately, have a series of drawbacks. For our purposes in this experiment, two in particular matter - firstly, the need for very long devices, which is very costly in terms of ebeam fabrication, and secondly that our device is intrinsically lossy, as propagation is achieved using radiation

[5]Perhaps the accepted method, much better established than the FP method used here!

modes. Although less relevant for us, the "in principle" part regarding the coupling is not actually very good - undetected facet damage, for example, may upset the results considerably. Similarly, for short devices, one can envisage a tunneling type of transmission - not actually through the desired waveguide. Given optical lithography and a statistically significant sample set, this method does give very reliable results - but both of these prerequisites are out of our grasp!

Instead, we use the fringe visibility of the Fabry-Perot peaks to analyse the response of our devices. This method is well known in the characterisation of lasers, and detailed in [82]. To explain this method, we begin by considering a simple etalon. We can obtain the transmission as:

$$I(\lambda) = (1 + \frac{4R}{1-R^2}\ sin^2\ \frac{\phi}{2})^{-1} \tag{6.1}$$

Although this sine formulation is simplest, we obtain the cosine form using the trigonometric identity $cos(\phi) = 1 - 2sin^2\frac{\phi}{2}$, which is preferred in the literature [83], as:

$$I(\lambda) = (\frac{1-R}{\sqrt{R}})^2 \times \frac{R}{1+R^2-2Rcos(\phi)} \tag{6.2}$$

where T is the transmission, R the facet/mirror reflectivity, and ϕ the roundtrip phase change between the mirrors. If we assume normal incidence and negligible phase change on reflection at the mirrors, we have:

$$\phi = 2nd\frac{2\pi}{\lambda_0} \tag{6.3}$$

where λ_0 is the free space wavelength, n the material index, and d the separation between the mirrors.

Next, we apply this formalism to our devices, beginning with the simple ridge waveguide. The effective index of the slab is calculated to be close to

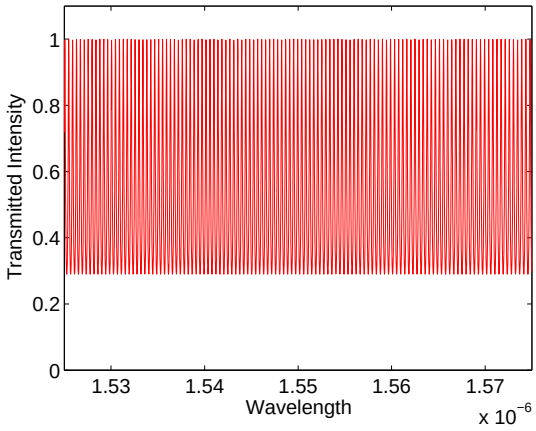

Figure 6.22: *Simulated frequency domain response for the 5μm blank ridge waveguide, using equation 6.2 with an effective index of 3.29, 30% facet reflectivity and 1mm die length. The wavelength range corresponds with the measurement data, to illustrate the analysis below. The simulation does not consider losses.*

3.29, for the relevant wavelengths. We estimate the (field) facet reflection[6] to be 30%. Finally, we choose a die size of 1mm. This leads to the response in figure 6.22.

We can now take this function and Fourier transform it. As we are ultimately interested in loss, which is typically of the form of an exponential decay, all of the Fourier spectra shown here have a logarithmic scale. After Hofstetter and Thornton [82], we make the wavelength step linear in k, hence giving the Fourier spectrum in spatial rather than time units, shown in figure 6.23. This has the benefit of giving sharp Lorentzian peaks, which are very suitable for further theoretical analysis [84], for example. However, in this case, our actual measurements have a wavelength step linear in wavelength, and as a result we Fourier transform into time, giving Gaussian peaks and less

[6]We could also calculate this, but the error is not significant.

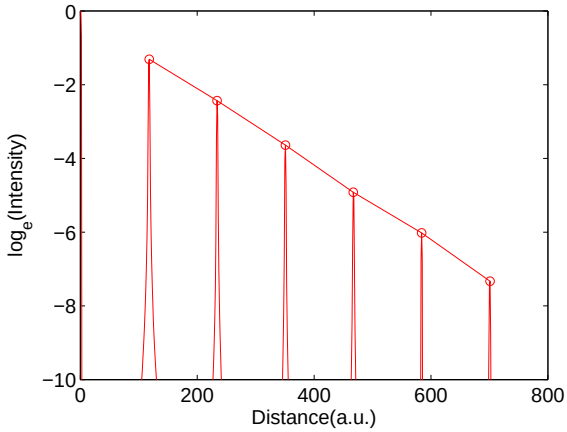

Figure 6.23: *The Fourier transform of the Airy function in figure 6.22, on a logarithmic scale. The Fourier transform is taken using data discretised in k, which gives rise to the sharp Lorentzian peaks shown. The ordinate is in space rather than time. These Lorentzians are well suited to the theoretical analysis of loss.*

mathematical elegance - compare figure 6.24.

In each of these figures, we show a line joining the peaks of the Fourier response. Obviously, the slope of the line is the same in each figure. This gradient is known as the α value, and $1 - e^{-\alpha}$ represents the roundtrip loss of the device. This value is intrinsic, from the Fresnel reflectivity of the facets.

We now refine our model by considering excess loss - which in practice is due to diffraction from the PhCs, imperfect facets, and roughness-induced scattering. This is done using a lumped absorption A at the device facets, which would occur for a real etalon with metallized mirrors. For the lossless case, we have $T + R = 1$, which is modified to $T + R + A = 1$, and hence we can modify equation 6.2 by writing the R terms outside of the prefactor as $Re^{-\alpha\prime}$:

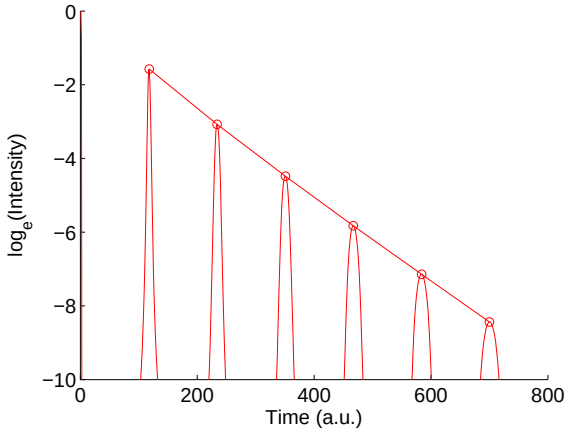

Figure 6.24: *The Fourier transform of figure 6.22, as above. In this figure the Fourier transform is taken from data discretised in λ, which gives a time axis and Gaussian peaks. This agrees with the form of the experimental results.*

$$I(\lambda) = \left(\frac{1-R}{\sqrt{R}}\right)^2 \times \frac{Re^{-\alpha\prime}}{1+\left(Re^{-\alpha\prime}\right)^2 - 2Re^{-\alpha\prime}cos(\phi)} \tag{6.4}$$

where $\alpha\prime$ is the additional loss factor - that is, our fitted loss term α is equal to $\alpha\prime$ plus the intrinsic Fresnel loss, α_f. This model describes the loss analysis that is applied to the experimental devices.

6.5.1 Compound cavities

In the above, we have considered the analysis of our device response in the Fourier transform, almost as if the PhC were only a loss-adding element. However, as discussed in chapter 5 we also have an impedance issue. Therefore, it is wrong to consider our entire device as a single cavity between the input and output facets - instead, as discussed by Talneau et al. [85], we have a compound cavity response. Although the PhC response is necessarily dispersed as opposed to occurring at a single point in the device, we can think of many reflective planes in the device - and so, many cavities. The Y, the bend,

and the W1 sections between them all create cavities. Then there are the tapers, and the super-cavities; for example the input W1 and the Y; or the output W1 and the bend. Each of these cavities makes a contribution to the Fabry-Perot response of the device, this should be identifiable in the Fourier transform. These cavities are physically shorter than the inter-facet cavity, so the relevant peaks ought to be found in-between these. However, in the device design we have maximized the transmission. Hence we should have a relatively low impedance mismatch defining these cavities, although this has not been considered explicitly - so, the facet response should dominate. By comparing the differing complexities of device fabricated, from the simple broad waveguide, through the W1 and to the Y-splitter, we should be able to identify the cavities and locate the impedance mismatches. Note that the very high resolution of the Fabry-Perot fringes will be required to adequately resolve these lower peaks, particularly in the presence of intrinsic radiation losses.

6.6 Results and analysis

In figure 6.21, we showed the experimental response of the Y-splitter device in the frequency domain. Here, we will briefly build up to this result, then start to consider its Fourier spectra and analyse it as explained in the previous section. Ideally, this would then be repeated for the double bend device and the W1s, so that we could extract the loss per element of the complete splitting device. This would be ill-advised for statistical reasons - not enough copies of the device were successfully fabricated and characterised for any rigour in this analysis. Although some other results will be presented qualitatively and discussed, only the Y-splitters will be treated quantitatively.

6.6.1 Reference waveguides

We begin with the simplest fabricated devices, 5μm wide straight ridge waveguides, referred to as "blanks". These will be used for normalisation of the Fabry-Perot loss, as described above - to determine, α_f in equation 6.4. We can assume that the transmission loss is minimal, by comparing to the simulated result discussed above. To determine this accurately, we would need the same degree of accuracy in the mode index and facet reflectivity in the simulation. However, since it is also necessary to consider facet imperfections and fibre-coupling losses, this does not add anything. The reference waveguide with the lowest loss has the best facets and lowest roughness - however, as shown in figure 6.25, there is a lot of variation. We take the average of the best three as being representative, and use this value for normalisation.

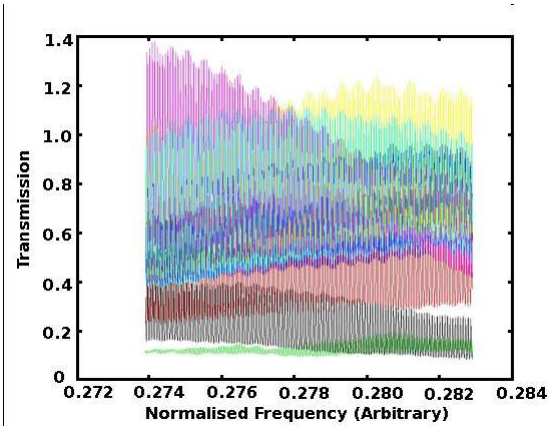

Figure 6.25: *Experimental responses of several copies of the reference blank waveguides in the frequency domain. Each measured response is a different colour. The normalised frequency axis is arbitrary, but equivalent to that used for the PhC devices. Although these blanks are nominally identical, uncontrollable fabrication differences - for example facet quality - give rise to a wide range of transmission responses. Optimisation of coupling is also a factor.*

The other important use for these reference devices is to accurately determine the device length in the time domain - that is, the period of the complete round trip, or the spacing of the facet-cavity peaks. These will subsequently be referred to as "primary" peaks. We show a typical Fourier response of a blank device in figure 6.26.

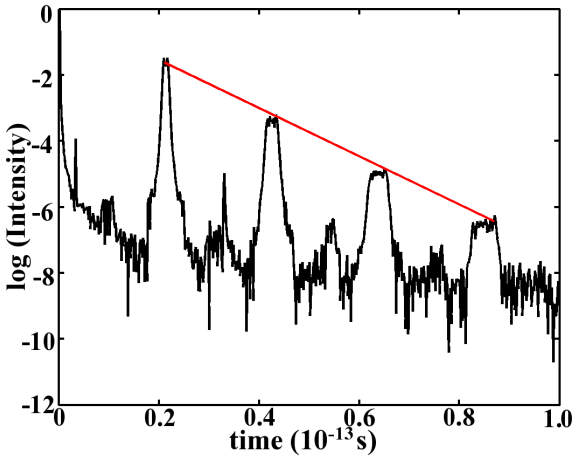

Figure 6.26: *Fourier transform of the response in the previous figure, for the highest transmitting ridge waveguide. The red line is the best fit to the peak positions, and its gradient is the α_f value of the lumped facet and waveguide loss.*

6.6.2 Taper plus wire

The next simplest device fabricated is a pair of tapers connected by a short (approximately 12μm) wire. This was considered as a device in chapter 4, as it was used to optimise the taper design. The Fourier response of these is shown in figure 6.28. The tapers have been geometrically scaled to agree with the lithographic tuning of the crystal, so the measured wavelengths have been converted into normalised frequency to reflect the appropriate PhC periods.

As we only have one wire length, and that a short one, it is not possible to say anything about the taper-induced losses. There are two points to note - firstly, that the primary peaks have the same spacing as for the reference waveguides, showing that the principle of peak selection is good, and secondly that there are two smaller peaks between these (indicated on the figure). These indicate that the impedance mismatch between the tapers and the wire causes a reflection - but that it is relatively weak compared to the facet reflection. This agrees with the simulation results in chapter 4. The results for the different taper sizes indicate that the tapers used have no significant frequency response over the measurement range, so that we can be sure that the frequency response for the complete device is due to the PhC.

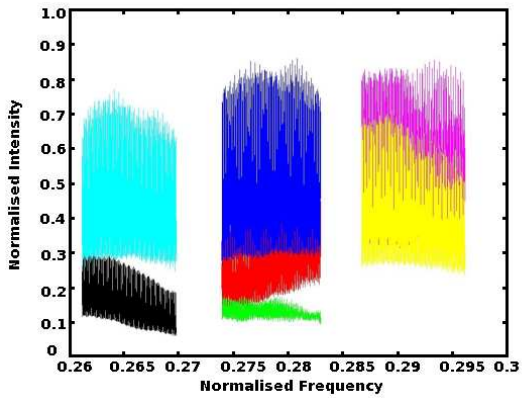

Figure 6.27: *Frequency domain response of many tapers and wires. The conversion from wavelength to normalised frequency arbitrary again as there are no PhCs in the device. It reflects the scaling of the tapers to the different PhC periods in the next section. The intensity is normalised against the reference waveguide selected above. The uniformity of the frequency response indicates that the centre frequency is well away from the MMI cut-off, and hence any spectral signature in the Y-splitter devices is related to the PhC.*

We can also consider the lithographically tuned devices in the frequency

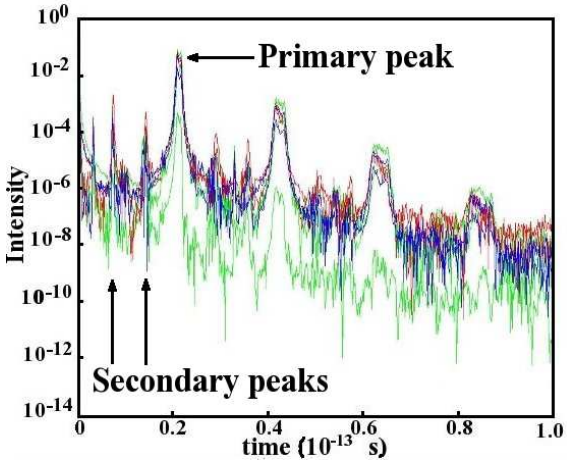

Figure 6.28: *Time domain response - i.e. Fourier transform of figure 6.27 of taper plus wire devices. The indicated primary peak, and its round trip harmonics, occur at the same position as for the blank ridge waveguide, indicating that these are due to the external facet reflection. The smaller intermediate peaks, indicated as secondary peaks have a shorter time interval and hence pathlength, and can be attributed to reflections from either end of the taper.*

domain - 6.27. Although this result is very complex, as it shows measurements on several different devices, it is clear that we have very broad bandwidth tapers, as is expected from their MMI nature and the simulation results. Hence, any narrowing of the spectral response is due to the properties of the photonic crystal, not the taper.

6.6.3 W1 waveguide

Next we consider our W1 waveguide, sandwiched between two tapers. The W1 is the same length as the wire, and as with the wire we can say nothing about its transmission - with only one length, it is not possible to decouple the tapering and coupling losses from the W1 radiation loss. Further, as shown in chapter 5, the coupling into the wire and the PhC are substantially different,

so we cannot assume the difference between the two devices relates to the W1 loss.

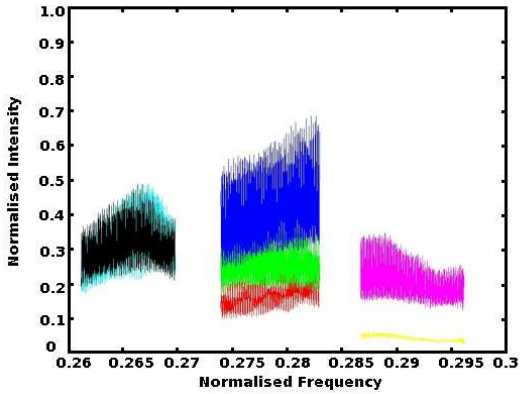

Figure 6.29: *Response of many W1s with tapers as measured in the frequency domain, plotted against their lithographically tuned normalised frequencies. The spectral shape is changed as compared to the wire devices, indicating the cut-off of the W1 band at the low frequency end. At high frequency, comparison to the bandstructure suggests that the additional loss is associated with coupling into the odd mode.*

Again, we see secondary peaks, and rather stronger than those for the wires. This tends to indicate that the impedance matching strategy used works - the taper is matched well to the wire. However the matching to the PhC could be improved. Also, considering the frequency domain response, we see that the PhC is having its effect - since the spectrum can no longer be considered uniform. By comparison to the bandstructure, we can suggest that the lower transmission in the higher period devices may be due to leakage into the odd mode. This also validates the design approach, shifting the crystal so that the centre frequency is away from this odd mode, as it adds loss as expected. The crystal itself imposes bandwidth limitations, even before considering the routing elements.

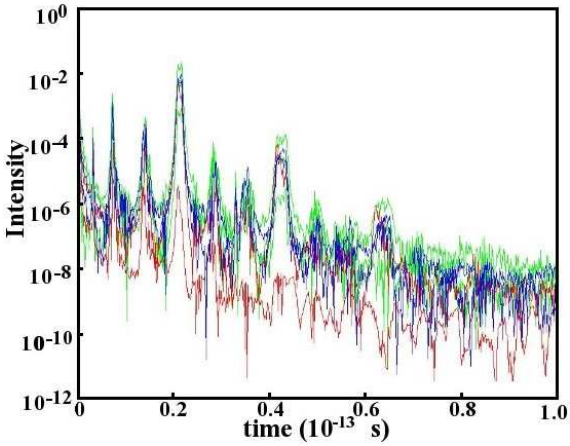

Figure 6.30: *Response of many W1s in the time (Fourier) domain. Note the enhancement of the secondary peaks over the wire device, which indicates greater reflection/impedance mismatch.*

6.6.4 Double bends

Finally, after all these reference structures, we come to a routing element - the W1 waveguide bend. The spectral and Fourier responses are shown in the following figures 6.32, 6.31.

Considering these figures, it is clear that the device loss has increased slightly, but more significantly that the secondary peaks in the Fourier response have become much stronger, with a height comparable to that of the primary peaks. From this, I infer that the bend cavities are reinforcing the taper related cavities - so, with sufficient resolution (which we do not have) peak splitting would be seen. As the W1 lengths are short, it would seem that despite the 0.01nm resolution of our OSA, we still cannot resolve the bend cavities, because the taper mismatch is so high. Of course, making devices with longer (and varied) W1 lengths would also help here, but because the device is intrinsically lossy due to its above-lightline operation, this may come at a cost.

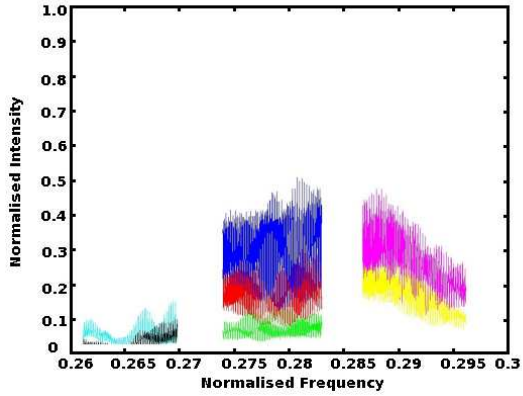

Figure 6.31: *Response of many bends in the frequency domain. The bandwidth is similar to that for the W1s, indicating the successful design and fabrication of a non-resonant bend.*

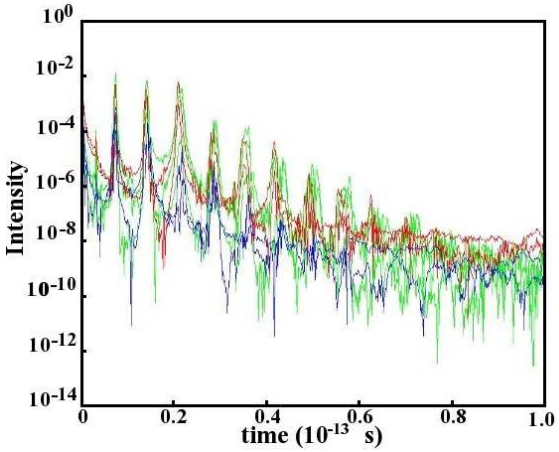

Figure 6.32: *Response of many bends in the Fourier domain. The secondary peaks are further enhanced, but the interfaces have not been changed. Because the crystal is so short relative to the die length, reflections due to the taper interfaces and the bends become lumped together.*

6.6.5 The complete device - taper, Y-splitter, and bends

For this final device, we have two output ports, and so double the number of measurement results. We consider the left and right arms separately.

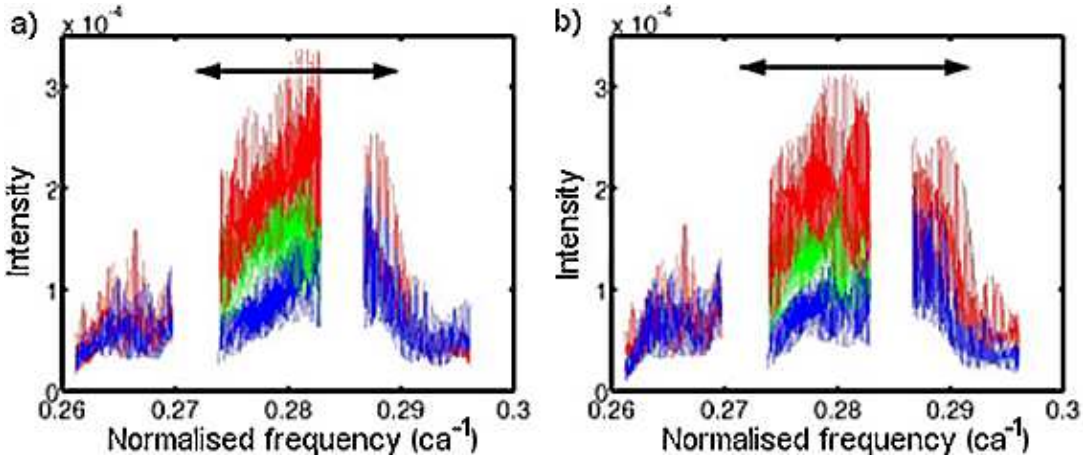

Figure 6.33: *Final device response in the frequency domain (shown previously). The left and right arm response are on the left and right respectively, and their similarity of shape shows a balanced response. The bandwidth of the W1 response is maintained, indicating that the Y-splitter is non-resonant, as designed. Further, at 110nm, it compares well with the 120nm bandwidth in the 3D simulation.*

Perhaps the most important result, is that whether we consider the spectral or Fourier space responses, the left and right arms of the Y show a very similar response. As discussed, this is not always the case for simple devices, simply to to the sensitivity to input coupling. The symmetry also shows that the fabrication is consistent across the whole device. Therefore, we consider the agreement in signal level and spectral response between the two arms to be a good sign of successful design and fabrication.

Moreover, these spectral plots can show us the device bandwidth. Although the expected gaps in the spectrum obscure the response somewhat, we estimate

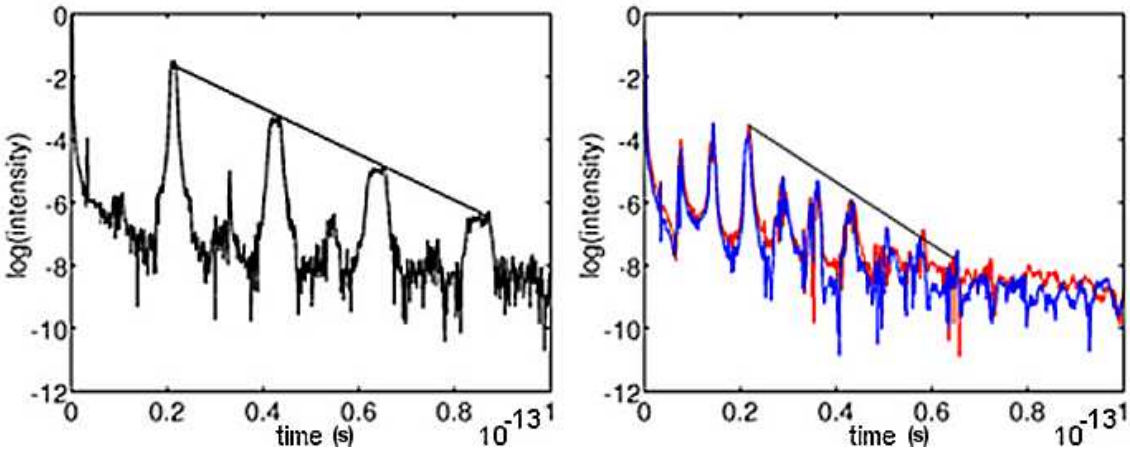

Figure 6.34: *Fourier transform response of the left and right arms of the published device on the right hand side. On the left, the reference spectrum (Fourier response of the ridge waveguide) is repeated for ease of comparison.*

the 3dB bandwidth of this device to be 110nm, in very good agreement with the 120nm found in the "verification" simulation (figure 6.18).

Finally, we can consider the Fabry-Perot response of this, our "complete" device - we find a transmission of more than 70%, relative to the 5μm reference waveguide! Again, this compares very favourably to the simulation, which has a peak transmission of 76%.

6.7 Discussion and conclusion

In this chapter, I have shown why PhC components can be a useful addition to the array of choices for optical routing, and explained the history and logic behind their design. Y-splitters and bends are considered in detail, and it has been shown that naive PhC designs of such devices have very narrow bandwidth and poor transmission. I have shown the incremental design modifications that have been discovered by various groups to improve either

transmission or bandwidth, and the trend towards the use of optimisation algorithms in the design of such devices. From this basis, I have shown the principles behind a numerical design methodology, which utilised realistic computing resources for the 2D optimisation and 3D verification of PhC Y-splitters and bends, as well as the short taper discussed in the previous chapter.

In 2D simulations, a composite device containing all of these elements was found to have a 3dB bandwidth over 100nm, and a peak transmission of 97%. In 3D, the device operating above the lightline in LOP AlGaAs heterostructure, has a 120nm 3dB bandwidth and a peak transmission of 76%. The difference in the transmission between 2D and 3D is readily explained by the intrinsic diffraction losses. We then fabricated this device in addition to a number of reference structures, taking care to accurately reproduce the optimised design, and perfect the etching to give minimal additional loss.

To characterise these devices, the Fabry-Perot method is introduced, whereby the measured spectra are Fourier transformed into the time domain and the loss evaluated from the round-trip response of the device. This gives the loss independent of input coupling, which can otherwise introduce significant errors. Using this, we have found that the short taper discussed in the previous chapter couples better into the W1 PhC waveguide than into the photonic wire it was designed with, which supports the argument presented in that chapter on the correct choice of optimisation results, as well as agreeing with simulations. We also see evidence of a remaining impedance mismatch between the taper and the devices, showing that further optimisation of interfaces would be advisable.

Using this method for the composite device, we find an experimental transmission of >70% and a 3dB bandwidth estimated to be 110nm. This agrees well with the simulation results, thereby validating the methods used

for optimisation and verification, and shows the high quality of the device fabrication. The results were published as [34], as the best PhC Y-splitter and bend devices at that time, and the only integrated composite device.

Chapter 7

Conclusions

In this thesis, I have presented my research into PhC devices, which has included modelling, design, fabrication and characterisation; the complete cycle!

To begin, I have discussed the background theory necessary to understand PhC devices, from which it is apparent that intuition is not enough - to design PhC devices detailed numerical modelling is necessary. I have also discussed the tools available for this modelling, their operating principles, and relative strengths and weaknesses. These tools have one major limitation, namely the huge computing resources they require, but there is no simple alternative. Instead, by carefully breaking a problem down into its constituent parts, I have shown that it is possible to completely model a PhC device, and even use numerical methods to optimise the design.

To design a useful device, it is not enough to have a thorough understanding of the available design tools, instead one must consider what is possible. To this end, I have also discussed fabrication techniques for the three material systems for the devices herein, and shown the constraints these provide on the design. The emphasis here is the practical understanding of the relevant

technologies for integrated optics devices.

By always keeping a fabrication perspective, I have shown that it is essential to consider the limitations of the available tools in device design, as it is much easier to change the design than the tools; in fact for current technologies I believe it is important to develop the fabrication techniques concurrently with the design, using a stepwise refinement process for both.

Having designed and fabricated devices, the next step is characterisation. There are two major measurement techniques used in this thesis. The first, analysis of the efficiency of surface grating couplers using the fibre to fibre coupling, is straightforward in principle, although the number of devices which must be fabricated successfully to complete this analysis is large, and hence demanding. The second is a more complex interferometric method, which yields useful data from a smaller set of devices. The end-to-end transmission spectrum is recorded at very high resolution, and then Fourier transformed to exploit the round-trip signature on the Fabry-Perot fringes. This allows us to calculate the loss due to the PhC, and also determine the points in the device where reflections occur, indicating impedance mismatches which must impair the device efficiency. Using these two methods, I have shown two separate cases - InP surface grating couplers and AlGaAs routing elements - that I have been involved in fabricating, operating very close to their designed levels, and thus validating the design.

Throughout the various chapters, the theme of coupling has emerged, both from fibre to waveguide, and from waveguide to PhC. This coupling is very important - I have shown that in some cases it provides the limiting factor on possible device performance. This is particularly important for "slow light" PhC designs, which have great potential for integration and miniaturisation of optical functionality, but have appeared to suffer from poor transmission

properties. I have shown that this is largely due to coupling problems, and have demonstrated a method which works in theory to alleviate this.

However, the most significant aspect of this work is that it addresses a complete story of PhC devices; as by starting at a fibre level we move closer to the ultimate goal of opto-electronic integration.

7.1 Specific results

In addition to the general conclusions offered above, I wish to address the specific results mentioned in the abstract, to emphasise the achievements in this thesis.

- Realisation of surface grating couplers in a novel material system demonstrating some of the highest reported fibre coupling efficiencies.

This comprises the work in chapter 4, where with my collaborators I have shown a successful design for an InP surface coupler, realised in an adhesive-bonded membrane system. Currently, coupling efficiencies are >50%, and simulation results show that for the final design polarisation independant coupling >80% can be expected. For these devices the lithography has been the particular challenge for me, as the required structures need both accuracy of reproduction and a great extent. I have developed the techniques to manage this.

- Development of a short "injecting" taper for coupling into photonic crystal devices.

This short taper is discussed in chapter 5, and at around 9μm in length, represents a significant saving in device footprint - over this scale, standard

tapers are $\tilde{2}00\mu$m. There is no apparent compromise in performance. I have deconstructed the design given to me, to understand the operating principles of the device, as well as fabricating and characterising it. Again, this device works in practise as expected, and is a useful addition to PhC device designs.

- Optimisation and experimental validation of photonic crystal routing elements (Y-splitter and bend).

This PhC Y-splitter and bend, integrated with the short taper above, have been numerically optimised, fabricated in AlGaAs heterostructure and characterised using the Fabry-Perot technique discussed previously. This is in chapter 6. This device has a transmission in excess of 70%, published as [34], an efficiency that other groups report of an optimised Y-splitter alone.

- Exploration of interfaces and coupling for "slow light" photonic crystals.

This is presented in chapter 5, along with the 9μm injector. Here, I show that coupling is a major factor that must be resolved before PhC slow light can be usefully exploited. I discuss a formalism for exploring interfaces in a simplified manner, and show how the results relate to standard FDTD transmission spectra, for a fraction of the computational effort. Further, I propose a structure that improves the coupling into this slow light regime, and show that it works in simulations. Fabrication is in process.

7.2 Further work

.

There are many things discussed in this thesis that could be usefully further studied. The work towards a complete 2D polarisation-diversity surface

coupler is continuing, with the first 2D InP devices I have manufactured at St Andrews undergoing wafer bonding prior to testing at the time of writing. Of course, the work on coupling to PhCs presented in chapter 4 still requires much further effort, as all of the approaches that I have attempted to clearly demonstrate improved coupling to the slow light regime have been curtailed due to either software or computer memory limitations, although devices have been successfully fabricated and are currently awaiting transmission spectrum characterisation. However, with the development that I have presented in this thesis, I think that a complete integration of all the devices here is a realistic goal - with fibre coupling, short tapers, and efficient coupling into a slow-light photonic crystal device being demonstrated simultaneously!

Bibliography

[1] L. Pavesi and G. Guillot. Optoelectronic and photonic winter school, sardagna, italy. In *Optical Interconnects - the Silicon Approach*, 2005.

[2] M. Soljacic S.G. Johnson S. Fan M. Ibanescu E. Ippen and J.D. Joannopoulos. Photonic-crystal slow-light enhancement of nonlinearphase sensitivity. *J. Opt. Soc. Am. B.*, 19 (9):2052–9, 2002.

[3] Y.A. Vlasov M. O'Boyle H.F. Hamann and S.J. McNab. Active control of slow light on a chip with photonic crystal waveguides. *Nature Letters*, 438 (3):65–9, 2005.

[4] J. D. Joannopoulos R. D. Meade and J. N. Winn. *Photonic Crystals: Molding the Flow of Light*. Princeton, 1995.

[5] E. Yablonovitch. Photonic crystals. *JOURNAL OF MODERN OPTICS*, 41(2):173–194, 1994.

[6] E. Yablonovitch et al. 3-dimensional photonic band-structure. *OPT QUANT ELECTRON*, 24 (2):273–283, 1992.

[7] S. John. "strong localization of photons in certain disordered dielectric superlattices. *Phys. Rev. Lett*, 58:2486–9, 1987.

[8] E. Yablonovitch. Inhibited spontaneous emmision in solid state physics and electronics. *Phys. Rev. Lett.*, 58:2059–62, 1987.

[9] L.J. Wu M. Mazilu T. Karle et al. Superprism phenomena in planar photonic crystals. *J. Quant. Electron*, 38 (7):915–8, 2002.

[10] L.J. Wu M. Mazilu and T.F. Krauss. Beam steering in planar-photonic crystals: From superprism to supercollimator. *J. Lightwave Technol.*, 21(2):561–6, 2003.

[11] A.S. Jugessur A. Bakhtazad A.G. Kirk et al. Compact and integrated 2-d photonic crystal super-prism filter-device for wavelength demultiplexing applications. *OPTICS EXPRESS*, 14 (4):1632–1642, 2006.

[12] R. Baets. Managing out-of-plane coupling in 2d photonic crystals. Two Dimensional Photonic Crystals, Monte Verita, Ascona, Switzerland, 2002.

[13] R. Wang and X.H. Wang. Effects of shapes and orientatoins of scatterers and lattice symmetries on the photonic bandgap in two-dimensional photonic crystals. *J. Appl. Phys*, 90 (9):4307–13, 2001.

[14] C. Kittel. *Introduction to solid state physics.* Wiley, 2005.

[15] T. J. Karle. *Dispersion engineering in planar photonic crystal waveguides.* PhD thesis, St Andrews University, 2006.

[16] J.M Lourtioz et al. *Photonic Crystals: Towards NanoScale Photonic Devices (English translation).* Springer, 2005.

[17] P. Bienstman. Modelling. In *ePIXnet Winter School Optoelectronic Integration -Technology and Applications*, 2006.

[18] P. Bienstman. *Rigorous and efficient modelling of wavelength scale photonic components*. PhD thesis, Universiteit Gent, 2001.

[19] S. G. Johnson and J. D. Joannopoulos. Block-iterative frequency-domain methods for maxwell's equations in a planewave basis. *Optics Express*, 8 (3):173–90, 2001.

[20] ab-initio.mit.edu.

[21] www.photond.com.

[22] camfr.sourceforge.net.

[23] J.P. Berenger. A perfectly matched layer for the absorbtion of electromagnetic waves. *J. Comput. Phys.*, 114:185–200, 1994.

[24] A. Taflove and S.C. Hagness. *Computational Electrodynamics*. Artech House, 2000.

[25] Min Qiu. Effective index method for heterostructure-slab-waveguide-based two-dimensional photonic crystals. *Appl. Phys. Lett.*, 81:1163–5, 2002.

[26] www.rsoftdesign.com, rsoft inc, ossining, n.y, usa.

[27] Liam OFaolain et al. Fabrication of photonic crystals using a spin coated hydrogen silsesquioxane (hsq) mask. *J. Vac. Sci. Tech. B.*, 24:336–9, 2006.

[28] W. Bogaerts. *Nanofotonische Golfgeleiders en Fotonische Kristallen in Silicium-op-Isolator, Nanophotonic Waveguides and Photonic Crystals in Silicon-on-Insulator*. PhD thesis, Universiteit Gent, 2004.

[29] D.H. Brown. *Tunable monolithic coupled cavity lasers*. PhD thesis, University of St Andrews, 2005.

[30] Y. Akahane et al. High-q photonic nanocavity in a two-dimensional photonic crystal. *Nature*, 425, 2003.

[31] G. Pagnotta et al. Shot-shifting for nanophotonic applications. *Proceedings of Micro- and Nano- Engineering (MNE)*, page unpublished, 2006.

[32] M. Suehiro et al. *J. Quantum. Electron.*, 29:2081–7, 1993.

[33] R. West et al. Limitations of proximity-effect corrections for electron-beam patterning of planar photonic crystals. *Opt. Eng.*, 44(4):043401.

[34] M. Ayre et al. Experimental verification of numerically optimised photonic crystal injector, y-splitter and bend. *IEEE Journal of selected areas in communications*, 23:1390–1395, 2005.

[35] M.V. Kotlyar et al. High-aspect-ratio chemically assisted ion-beam etching for photonic crystals using a high beam voltage-current ratio. *JVSTB*, 22:1788–91, 2004.

[36] M.V. Kotlyar et al. Low-loss photonic crystal defect waveguides in inp. *App. Phys. Lett.*, 84(18):1788–91, 2004.

[37] M.V. Kotlyar. *Photonic crystal waveguides and tunable filters in InP.* PhD thesis, University of St Andrews, 2005.

[38] L. O'Faolain. *The fabrication of microstructures for ultrashort pulse generation in monolithic modelocked semiconductor lasers.* PhD thesis, University of St Andrews, 2005.

[39] C.J.M. Smith et al. Use of polymethylmethacrylate as an initial pattern transfer layer in fluorine- and chlorine-based reactive-ion etching. *JVSTB*, 17(1):113–7, 1999.

[40] D. Taillaert. *Grating couplers as Interface between Optical Fibres and Nanophotonic Waveguides*. PhD thesis, Universiteit Gent, 2005-5.

[41] J. Doylend. Efficient, manufacturable fibre-soi waveguide optical coupler. The Rank Prize Funds Mini-Symposium on Si-Based Optoelectronics, 2006.

[42] M. Notomi. Photonic crystal waveguides. EpixNET Winter School, Pontressina, Switzerland, 2006.

[43] N. Moll S. McNab and Y.A. Vlasov. Ultra-low loss photonic integrated circuit with membrane-type photonic crystal waveguides. *Optics Express*, 11:2927–39, 2003.

[44] P. Bienstman. *Rigorous and efficient modelling of wavelength scale photonic components*. PhD thesis, Universiteit Gent, 2001.

[45] F. van Laere et al. Compact and efficient fiber-to-waveguide grating couplers in inp-membrane. *Proceedings, Photonics Europe*, 2006.

[46] I. Christiaens et al. Thin film devices fabricated with bcb wafer bonding. *IEEE Journal of Lightwave Technology*, 23:517–23, 2005.

[47] F. van Laere et al. Compact grating couplers between optical fibers and silicon-on-insulator photonic wire waveguides with 69 *OFC*, page PDP15, 2006.

[48] F. van Laere et al. Compact and efficient fiber-to-waveguide grating couplers in inp-membrane. *Electronics Letters*, 42(6):343–5, 2006.

[49] F. van Laere et al. Compact and highly efficient grating couplers between optical fiber and nanophotonic waveguides in bonded inp-membranes. *Proceedings, ECOC*, 2006.

[50] J. Leijtens. Photonic integrated circuits. In *ePIXnet Winter School Optoelectronic Integration -Technology and Applications*, 2006.

[51] S. Boscolo et al. Numerical analysis of propagation and impedance matching in 2-d photonic crystal waveguides with finite length. *Journal of Lightwave Technology*, 20:304–10, 2002.

[52] R. Biswas et al. Impedance of photonic crystals and photonic crystal waveguides. *App. Phys. Lett.*, 84(8):1254–6, 2004.

[53] S. Boscolo and M. Midrio. *Opt. and Quant. Electr.*, 2004.

[54] J.M. Heaton. *App. Phys. Lett*, 61 (15).

[55] T. Shibata et al. Silica-based waveguide-type 16*16 optical switch module incorporating driving circuits. *IEEE Phot. Techn. Lett.*, 15:1300, 2003.

[56] D. L. Lee. *Electromagnetic principles of integrated optics*. Wiley, 1986.

[57] Y.A. Vlasov and S.J. McNab. Losses in single mode silicon-on-insulator strip waveguides and bends. *Opt Express*, 12:1622, 2004.

[58] S. Noda et al. Full three-dimensional photonic bandgap crystals at near-infrared wavelengths. *Science*, 289:605, 2000.

[59] A. Chutinan and S. John. 3 + 1 dimensional integrated optics with localized light in a photonic band gap. *Opt. Express*, page 1266, 2006.

[60] A. Mekis et al. High transmission through sharp bends in photonic crystal waveguides. *Phys. Rev. Lett.*, 77:3787, 1996.

[61] S. Olivier et al. Resonant and nonresonant transmission through waveguide bends in a planar photonic crystal. *Applied Physics Letters*, 79:2514–6, 2001.

[62] T. Baba et al. Observation of light propagation in photonic crystal optical waveguides with bends. *ELECTRONICS LETTERS*, 35:(8), 1999.

[63] S. Olivier et al. Improved 60 degree bend transmission of submicron-width waveguides defined in two-dimensional photonic crystals. *Journal of Lightwave Technology*, 20:1198–1203, 2002.

[64] A. Talneau et al. Photonic-crystal ultrashort bends with improved transmission and low reflection at 1.55 microns. *Applied Physics Letters*, 80:547–9, 2002.

[65] H. Benisty et al. Models and measurements for the transmission of submicron-width waveguide bends defined in two-dimensional photonic crystals. *IEEE journal of quantum electronics*, 38:770–85, 2002.

[66] E. Chow et al. Quantitative analysis of bending efficiency in photonic-crystal waveguide bends at lambda = 1.55 microns wavelengths. *Optics Letters*, 26:286–8, 2001.

[67] Y. Sugimoto et al. Two-dimensional semiconductor-based photonic crystal slab waveguides for ultra-fast optical signal processing devices. *IEICE Trans. Electron.*, E87-C:316–27, 2004.

[68] Y. Sugimoto et al. Fabrication and characterisation of photonic crystal based symmetric mach-zehnder (pc-smz) structures based on gaas membrane slab waveguides. *IEEE Journal of selected areas in communications*, 23, 2005.

[69] B. Miao et al. High-efficiency broad-band transmission through a double-60 degree bend in a planar photonic crystal single-line defect waveguide. *IEEE Photonics Technology Letters*, 16(11):2469–71, 2004.

[70] L. H. Frandsen. Ultralow-loss 3-db photonic crystal waveguide splitter. *Optics Letters*, 29(14):1623–5, 2004.

[71] A. Tetu et al. Broadband topology-optimized photonic crystal components for both te and tm polarisations. *Opt. Express*, 13:8608, 2005.

[72] A. Chutinan and S. Noda. Waveguides and waveguide bends in two-dimensional photonic crystal slabs. *Physical Review B*, 62:4488–92, 2000.

[73] A. Chutinan et al. Wider bandwidth with high transmission through waveguide bends in two dimensional photonic crystal slabs. *Applied Physics Letters*, 1698-1700, 2002.

[74] T. Sondergaard and K.H. Dridi. Energy flow in photonic crystal waveguides. *Physical Review B*, 61:15688–96, 2000.

[75] S. Fan et al. Waveguide branches in photonic crystals. *J. Opt. Soc. Am. B*, 18:162–165, 2001.

[76] I. Park et al. Photonic crystal power-splitter based on directional coupling. *Optics Express*, 12:3599–604, 2004.

[77] S.Y. Lin. Low-loss, wide-angle y splitter at 1.6micron wavelengths built with a two-dimensional photonic crystal. *Optics Leters*, 27:1400–2, 2002.

[78] M. Bayinder et al. Photonic crystal based beam splitters. *Applied Physics Letters*, 77:3902–4, 2000.

[79] S. Boscolo et al. Y junctions in photonic crystal channel waveguides: high transmission and impedance matching. *Optics Letters*, 27:1001–3, 2002.

[80] Y. Sugimoto et al. Light-propagation characteristics of y-branch defect waveguides on algaas-based air-bridge-type two-dimensional photonic crystal slabs. *Optics Letters*, 27:388–90, 2002.

[81] R. Wilson. Efficient photonic crystal y-junctions. *Journal of Optics A*, 5:S76–80, 2003.

[82] D. Hofstetter and R. L. Thornton. Loss measurements on semiconductor lasers by fourier analysis of the emission spectra. *Appl. Phys. Lett.*, 72:404–6, 1998.

[83] D. Hofstetter and J. Faist. Measurement of semiconductor laser gain and dispersion curves utilizing fourier transforms of the emission spectra. *IEEE Photonics Technology Letters*, 11:1372–4, 1999.

[84] D. Hofstetter and R. L. Thornton. Theory of loss measurements on fabry-perot resonators by fourier analysis of the transmission spectra. *Optics Letters*, 22:1831–3, 1997.

[85] A. Talneau et al. Compound cavity measurement of transmission and reflection of a tapered single-line photonic-crystal waveguide. *Appl. Phys. Lett.*, 82:2577–9, 2003.