NTNU – Trondheim
Norwegian University of
Science and Technology

Photonic Crystal Structures

Design and Fabrication

Marius Lorvik

Master of Science in Electronics
Submission date: July 2013
Supervisor: Astrid Aksnes, IET

Norwegian University of Science and Technology
Department of Electronics and Telecommunications

Problem Description

Microphotonics is a sub-branch of photonics oriented towards integrated systems on a single semiconductor chip. It involves the design, fabrication and testing of optical components to be fabricated at micro- or even nano-scale. Examples of these components are dielectric waveguides, microresonators, light modulators, couplers, periodic structures and more.

In this thesis, the focus will be on design and fabrication of a silicon photonic crystal waveguide structure. Simulations of non-ideal structures will be performed to investigate the effect of imperfections during fabrication. Both COMSOL Multiphysics and MIT Photonic Bands will be utilized.

The structures will be fabricated in the cleanroom at NTNU NanoLab. The fabrication processes will be based upon electron beam lithography, chemical vapor deposition, and plasma etching. Known fabrication processes in NTNU NanoLab will be drawn extensively upon in the preliminary stages. Recipes and processing parameters will be systematically changed to optimize the processing of the photonic crystal structures. Test structures will be characterized and evaluated. Much of the work is expected to be iterative, *e.g.*, processing parameters will be tuned on the basis of the characterization results.

The photonic crystal waveguide structure should be fabricated in a silicon on insulator (SOI) wafer. The possibilities of fabricating such a wafer in-house by chemical vapor deposition methods will be researched, but it might be more feasible to buy commercial SOIs. The silicon deposited by chemical vapor deposition will be amorphous, and it might suffer from high optical losses.

Cleanroom fabrication does not only rely upon well-established theories, but also a lot of 'know-how', which must be learned by experience. Fabrication is experimental work, and it is expected that a significant amount of time will be spent on learning the processes by hands-on experience. The thesis should seek to provide a contribution to the competence and common knowledge in the cleanroom, providing a foundation for future work on photonic crystals and silicon photonics.

Abstract

This thesis is devoted to the study of photonic crystal (PhC) design and fabrication. The design of a trigonal PhC waveguide is evaluated using finite element method (FEM) simulation software, and some considerations regarding non-ideal effects are discussed. The PhC of choice is a trigonal lattice PhC, with lattice constant and hole radius set to $a = 503.26\,\text{nm}$ and $0.4a$, respectively. Hole roughness is found to induce multiple scattering within the PhC waveguides, causing increased reflection coefficients. At high levels of roughness, a breakdown of the photonic bandgap is observed, causing light to leak out of the PhC structures.

The fabrication routines for the trigonal lattice PhC is researched, using electron beam lithography and inductively coupled plasma reactive ion etching. A silicon on insulator (SOI) wafer is created by plasma enhanced chemical vapor deposition of hydrogenated silicon dioxide and amorphous silicon (a-Si), but the resulting a-Si layer is found to be very rough. As this would induce optical losses, the fabrication is performed on single crystalline silicon wafers, but the results should be readily transferable to SOI wafers with a single crystalline top Si layer. The etched PhC structures are characterized by scanning electron microscope (SEM) inspection. Future work should seek to fabricate a PhC waveguide in SOI, using the routines outlined in this thesis, and measuring the transmission loss.

Two etch chemistries are investigated in this work. The first is based on sulfur hexafluoride (SF_6) and fluoroform (CHF_3), and the second is based on hydrogen bromide (HBr). The SF_6/CHF_3 etch recipe is found to provide near vertical etch profiles, while the HBr etch recipe has slightly positive profiles due to its reactive chemistry. Suggestions of improving the SF_6/CHF_3 recipe are outlined, and it is theorized that optimization will yield an etch recipe which is highly capable of defining PhC structures in Si.

A simple method of estimating the performance of fabricated structures is outlined, using enhanced SEM images to define the structures in FEM software. The results are limited by artifacts arising from image processing of the SEM images, but this should be easily mitigated given enough time.

Sammendrag

Denne oppgaven omhandler design og fabrikasjon av fotoniske krystaller (PhC). Det fotoniske krystallet som er undersøkt i oppgaven består av lufthull i silisium (Si), etset i et trigonalt gittermønster. Gitterkonstanten og hullradiusen er satt til henholdsvis 503.26 nm og $0.4a$, basert på simuleringer. Fotonisk krystall bølgeledere er simulert i finite element method (FEM) programvare, og noen ikke-ideelle effekter er beskrevet. Ujevnheter rundt lufthullene resulterer i spredning i bølgelederen, som medfører økte refleksjonskoeffisienter. Det fotoniske båndgapet er observert å bryte sammen som en følge av store ujevnheter, som resulterer i at lyset lekker gjennom PhC strukturene.

De fotoniske krystallene er fabrikkert ved hjelp av elektronstrålelitografi og plasmaetsing. En silisium på isolator (SOI) wafer er laget ved hjelp av plasmaassistert kjemisk deponering av hydrogenert silisiumdioksid og amorft silisium (a-Si) på en Si-wafer. Det er observert store ujevnheter i det amorfe silisium laget, noe som vil medføre store optiske tap. Derfor er krystallinske Si-wafere brukt for å undersøke fabrikasjonsmetodene, men metodene er overførbare til SOI wafere. Et sveipelektronmikroskop (SEM) er brukt til å karakterisere de etsede fotoniske krystallene. Fremtidige forsøk kan bruke fabrikasjonsmetodene som er beskrevet her for å lage fotonisk krystall bølgeledere i SOI, og måle transmisjonen gjennom de.

De fotoniske krystallene er fabrikkert ved hjelp av to forskjellige etsekjemier. Veldefinerte hull med nesten rette sidevegger ble etset ved hjelp av svovelheksafluorid (SF_6) og fluoroform (CHF_3). Etseoppskriften har en høy etserate, som gjør det problematisk å kontrollere etsedybde. En hydrogenbromid (HBr) oppskrift er også undersøkt, og den resulterer i positive profiler på grunn av en høy kjemisk etserate. Forslag til forbedringer av SF_6/CHF_3 oppskriften er beskrevet, og resultatene tyder på at forbedringer vil gjøre den godt egnet for PhC fabrikasjon i Si.

En metode for å estimere kvaliteten på fabrikkerte strukturer ved hjelp av FEM simuleringer er beskrevet. SEM bilder av strukturene kan brukes til å definere materialparametere i FEM, men bildestøy er funnet å være en feilkilde i simuleringene. Ved hjelp av bildebehandling kan dette unngås, og metoden kan nyttiggjøres.

Preface

This thesis is delivered in partial fulfillment for the degree of M.Sc. Electronics at the Department of Electronics and Telecommunications at the Norwegian University of Science and Technology, NTNU. The experimental work performed herein, and the text itself, is the individual work of the author, except for some important contributions by the people listed below.

Acknowledgments

First and foremost I would like to thank my supervisor, professor Astrid Aksnes. She has taken much time off from her busy schedule to offer helpful guidance during planning of my projects, interpretation of experimental and simulation results, and layout of this thesis. Her sincere interest in my work has motivated me throughout the last year, for which I am truly grateful.

Secondly I would like to thank Espen Rogstad, department engineer at NTNU NanoLab. He has been a tremendous resource during this work, mainly for sharing his insight on ICP etching. In addition, he has been very helpful in teaching me how to operate several of the instruments used extensively during my work in the cleanroom. It is safe to say that this work could not have been done without him.

Ph.D. student Oleg Żero has been very much a key person during my last year's project thesis and Master's thesis, taking the role as my co-supervisor. First in teaching me COMSOL Multiphysics and offering explanations to a few of the more intricate features of this program. More importantly, he has been responsible for pushing me to become an independent cleanroom user, and he has offered some much appreciated advice throughout this work.

Mark Chiappa, department engineer at NTNU NanoLab, has always offered tips and tricks in the cleanroom. In addition, he was very helpful during cleanroom training, answering many questions along the way.

I would also like to thank fellow M.Sc. student Leidulv Vigen for participating in some of the experiments, sharing his thoughts on my problems, and making some of the longest sessions in the cleanroom more cheerful.

Finally, I would like to thank my fellow students for engaging in many interesting conversations regarding their and my own Master's projects, and offering much needed breaks during busy times. A special thank you goes out to Per Øyvind Eid Fjuk for always providing a fresh pot of coffee!

Previous Work

During the fall of 2012 I worked on my project thesis, as a preparation for this Master's thesis. The project work included literature studies on photonic crystal theory, and finite element method simulations of photonic crystal wavegudies. The title of the thesis was "Investigating non-ideal effects in a 2D photonic crystal waveguide, using finite element method simulation software", and it provided a method of quantifying and interpreting hole roughness in photonic crystal waveguides. Some of the contents of the project thesis is reprinted in this thesis, and some of the simulation results are presented in a more suitable way.

The method of simulating photonic crystal waveguides is not thoroughly presented here, but some simulation results are. For a more in-depth presentation of the simulation method, the reader is referred to the project thesis.

Trondheim, July 2013
Marius Lorvik

Contents

List of Figures

List of Tables

List of Abbreviations

a-Si	Amorphous silicon
Au	Gold
BOX	Buried oxide layer in SOI
c-Si	Crystalline silicon
CCP	Capacitively coupled plasma
CHF_3	Fluoroform
CMOS	Complementary metal–oxide–semiconductor
DI water	Deionized water
EBL	Electron beam lithography
FEM	Finite element method
GaAs	Gallium arsenide
HBr	Hydrogen bromide
ICP	Inductively coupled plasma
IPA	Isopropyl alcohol, isopropanol
MPB	MIT Photonic Bands
N_2	Nitrogen gas
PBG	Photonic bandgap
PECVD	Plasma enhanced chemical vapor deposition
PhC	Photonic crystal
PML	Perfectly matched layer (COMSOL)
PMMA	Polymethyl methacrylate
PR	Photoresist
RF	Radio frequency

RIE	Reactive ion etching
RPM	Revolutions per minute
SEM	Scanning electron microscope
SF_6	Sulfur hexafluoride
Si	Silicon
SOI	Silicon on insulator

Chapter 1

Introduction

Photonic crystal structures are promising building blocks for various microphotonic applications, but they have not yet been successfully fabricated at NTNU. The goal of the work presented in this thesis is to investigate the possibilities of fabricating silicon based photonic crystal waveguides at NTNU NanoLab. Much of the work is transferable to other silicon based microphotonic devices, making it a contribution to the common knowledge at NTNU NanoLab.

1.1 Photonic Crystals

Photonic crystals (PhCs) were first conceived of by Eli Yablonovitch and Sajeev John in 1987 [1, 2]. A photonic crystal is an engineered structure with periodic dielectric constant, $\epsilon(\mathbf{r})$, where the periodicity can be exhibited as a 1D, 2D or 3D crystal structure. There is a range of wavelengths not allowed to reside within the PhC—a bandgap—which can be very accurately controlled by proper design of the crystal. If the PhC is designed with a defect, such as the lack of a row of holes in a 2D PhC, wavelengths within the bandgap are only allowed to reside within this defect—creating, *e.g.*, a PhC waveguide [3, 4].

Since a PhC is defined as a medium in which the dielectric constant changes periodically, it is infinite by definition. However, as will be shown later, the bandgap holds true even for finite structures.

The CMOS industry has been using silicon (Si) extensively, as it is found abundantly on Earth, and because of its excellent material properties. Because of the vastly successful history of this industry, micro- and nanoscale fabrication methods for Si are very well-developed.

Traditionally, photonics has been the domain of III-V semiconductors [5], but for incorporation of Si-CMOS and microphotonic devices, Si is becoming the new material of choice [5, 6, 7]. Other materials, such as gallium arsenide (GaAs), have been used for microphotonic devices, but the differing lattice constants of, *e.g.*, Si and GaAs, make monolithic incorporation of these materials difficult. As Si is transparent in the telecommunication wavelengths (1.3 µm–1.6 µm), it is an excellent waveguide material, and a poor detector material. However, research is being undertaken to resolve these difficulties, to open the possibilities of basing most microphotonic components on Si substrates [7].

1.2 Photonic Crystal Fabrication

Last year, initial studies surrounding PhC fabrication were undertaken at NTNU NanoLab [8]. This was also the first fabrication study on Si photonics at NTNU, and as such it provided a very important first step towards successfully demonstrating a working PhC structure at NTNU. However, the work yielded non-consistent results, and varying levels of roughness. This inspired the study of how to simulate non-ideal structures, and investigate how the roughness affects the propagation of light in a PhC waveguide [9].

The two dimensional (2D) trigonal lattice PhC, fabricated by etching air holes in a dielectric, has been extensively researched for various applications, such as waveguides with sharp bends [10]; out of plane emission of photons from point defects [11, 12]; and devices for slow light [13]. It serves as the PhC of choice for this work, and the focus is to successfully fabricate the hole structure. Once the fabrication routines for this structure are known, more application specific structures can be investigated, with their added complexity.

The fabrication process presented herein relies on several well-known fabrication routines, using electron beam lithography, plasma enhanced chemical vapor deposition and inductively coupled plasma reactive ion etching. Even though these processes are very well developed, the beginning cleanroom user faces challenges in learning the proper handling and usage of the instruments. In addition, most processes are dependent upon materials used, structure size and structural geometry. Hence, much time was spent during this work learning how to best utilize the resources available.

During the learning process there inevitably was some trying and failing. The lessons learned from failing are explained within this thesis. Hopefully, prospering students of PhC fabrication might use this work as a foundation to build upon, although it is likely that they must go through much the same process.

1.3 This Thesis

Chapter 2 outlines in some detail the theory behind PhCs. It is not a necessary prerequisite to understand the fabrication parts of this thesis, but it is necessary for discussions on PhC performance. PhC theory is readily available elsewhere, but the compact form of chapter 2 is thought to be useful for readers unfamiliar with PhC theory.

Chapter 3 surrounds the design of PhCs and is based on simulations and literature studies. The optimal geometry of the PhC for 1.55 µm operation is found, and some considerations regarding non-ideal structures are outlined.

Chapter 4 outlines the fabrication methods and background theory. It is designed to be a guide for starting cleanroom users and is divided up such as to follow the process flow in the cleanroom.

Chapter 5 contains both results and discussions. The two are intervened as the experimental results were continually evaluated to decide how the work should progress. In order not to disrupt the flow of information, chapter 5 is written more like a summary of the experimental work than a standard results chapter.

Chapter 2

Photonic Crystal Theory

Photonic crystal (PhC) theory is readily available from excellent textbooks and review articles (*e.g.*, [4, 14, 15]). A brief outline of the theory is included here for completeness, and to serve as a background for readers unfamiliar with PhCs. It is not strictly a necessary prerequisite for the fabrication parts of this thesis, but in designing the PhC and discussing fabrication results it is much needed.

2.1 Maxwell's Equations in a Periodic Medium

In studying any electrodynamic problem, the starting point is always the well known Maxwell's equations for the electric field, $\mathbf{E}(\mathbf{r},t)$, the magnetic field, $\mathbf{H}(\mathbf{r},t)$, the electric displacement field, $\mathbf{D}(\mathbf{r},t)$ and the magnetic flux density, $\mathbf{B}(\mathbf{r},t)$:

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \tag{2.1}$$

$$\nabla \times \mathbf{H} = \mathbf{J} + \frac{\partial \mathbf{D}}{\partial t} \tag{2.2}$$

$$\nabla \cdot \mathbf{D} = \rho \tag{2.3}$$

$$\nabla \cdot \mathbf{B} = 0, \tag{2.4}$$

where $\mathbf{J}$ and ρ denote the current density and the free charge density, respectively. In studying photonic crystals, it can be assumed that there are neither sources of current, nor any free charges in the material. Thus $\mathbf{J} = 0$ and $\rho = 0$, simplifying Maxwell's equations greatly.

Before Maxwell's equations can be made of use, a set of constitutive equations are needed, relating the electric field to the displacement field, and the magnetic field to the flux density. These can in general be very complicated, as described in [4, Ch. 2]. However, assumptions can be made that the materials of interest are linear, isotropic, lossless, and non-dispersive [16, Ch. 5.2.A]. These assumptions

give[1]

$$\mathbf{D}(\mathbf{r}) = \epsilon_0 \epsilon(\mathbf{r})\, \mathbf{E}(\mathbf{r}) \tag{2.5}$$

$$\mathbf{B}(\mathbf{r}) = \mu_0\, \mathbf{H}(\mathbf{r}), \tag{2.6}$$

where ϵ_0 and μ_0 are the electric permittivity and the magnetic permeability, respectively, of vacuum. $\epsilon(\mathbf{r})$ is the relative permittivity of the material—also known as the dielectric function. When a medium has a periodic dielectric function, it is called a *periodic medium*.

Since it is known from Fourier analysis that all signals can be thought of as a superposition of harmonic functions, the phasor notation will be used (see, *e.g.*, [16, Ch. 5.3]):

$$\mathbf{E}(\mathbf{r}, t) = \mathbf{E}(\mathbf{r})\, e^{-i\omega t} \tag{2.7}$$

$$\mathbf{H}(\mathbf{r}, t) = \mathbf{H}(\mathbf{r})\, e^{-i\omega t}, \tag{2.8}$$

where the physical fields are given as the real part of the time-varying fields.

Substituting (2.6) into (2.1), and utilizing (2.7) and (2.8), gives

$$\nabla \times \mathbf{E}(\mathbf{r}) - i\omega\mu_0\, \mathbf{H}(\mathbf{r}) = 0, \tag{2.9}$$

which is divided through by $e^{-i\omega t}$ in order to just look at the spatially varying field. Similarly, (2.2) simplifies to

$$\nabla \times \mathbf{H}(\mathbf{r}) + i\omega\epsilon_0 \epsilon(\mathbf{r})\, \mathbf{E}(\mathbf{r}) = 0. \tag{2.10}$$

As it always is when using Maxwell's equations, (2.9) and (2.10) should be combined in order to get one equation which fully describes the spatial field. Since $\mathbf{E}(\mathbf{r})$ is given by $\mathbf{H}(\mathbf{r})$, and vice versa, there is a choice in which field to calculate for. An equation for $\mathbf{H}(\mathbf{r})$ is chosen here, and justification of the choice is postponed for now. Dividing (2.10) by $\epsilon(\mathbf{r})$ and taking the curl gives

$$\nabla \times \frac{1}{\epsilon(\mathbf{r})}\, \nabla \times \mathbf{H}(\mathbf{r}) = -i\omega\epsilon_0\, \nabla \times \mathbf{E}(\mathbf{r})\,. \tag{2.11}$$

Substituting (2.9) for $\nabla \times \mathbf{E}(\mathbf{r})$ and remembering that the speed of light in vacuum is $c_0 = 1/\sqrt{\epsilon_0\mu_0}$, the equation for $\mathbf{H}(\mathbf{r})$ is found:

$$\boxed{\nabla \times \frac{1}{\epsilon(\mathbf{r})}\, \nabla \times \mathbf{H}(\mathbf{r}) = \left(\frac{\omega}{c_0}\right)^2 \mathbf{H}(\mathbf{r}),} \tag{2.12}$$

which is known as the master equation.

One interesting thing to note about (2.12) is that there is no fundamental length scale governing the equation (as it indeed is in, *e.g.*, quantum mechanics). As long as the system of interest is a macroscopic one, both $\mathbf{r}$ and $\epsilon(\mathbf{r})$ can be scaled to yield an unchanged equation, but with a scaled frequency, ω [4, pp. 20-21].

[1] Really, $\mathbf{B}(\mathbf{r}) = \mu_0 \mu(\mathbf{r})\, \mathbf{H}(\mathbf{r})$, but only cases where $\mu(\mathbf{r}) = 1$ will be considered.

2.2 Eigenvalue Problems and Hermitian Operators

The remaining theoretical analysis is largely based upon the derivations of [4, Ch. 2.3], and the reader is referred to this excellent text for a more in-depth analysis.

Upon investigating (2.12), it is seen that on the left side there is an operator operating on $\mathbf{H}(\mathbf{r})$, and on the right side there is a constant multiplied by $\mathbf{H}(\mathbf{r})$. This is known as an eigenvalue problem, where $\mathbf{H}(\mathbf{r})$ is the eigenvector.

The operator in (2.12) is defined as

$$\hat{\Theta} \triangleq \nabla\times \frac{1}{\epsilon(\mathbf{r})} \nabla\times . \tag{2.13}$$

In a given periodic medium, $\epsilon(\mathbf{r})$ is the periodic function in (2.13) which governs the resulting $\mathbf{H}(\mathbf{r})$ for a given frequency—much in analogy to the periodic potential in quantum mechanics.

By noting that $\hat{\Theta}$ is a linear operator, it is known that two solutions $\mathbf{H}_1$ and $\mathbf{H}_2$, both with frequency ω, can be combined linearly to form a third solution. *I.e.*, $\mathbf{H}_3 = a\mathbf{H}_1 + b\mathbf{H}_2$, where a and b are constants. Solutions $\mathbf{H}(\mathbf{r})$ with different frequencies are called harmonic modes and can be superposed to give every possible solution, as discussed earlier.

A brief digression is necessary in order to continue the analysis of the eigenvalue problem (2.12). The inner product of two functions is defined as

$$(\mathbf{F}, \mathbf{G}) \triangleq \int d^3\mathbf{r}\mathbf{F}^*(\mathbf{r}) \cdot \mathbf{G}(\mathbf{r}), \tag{2.14}$$

from which it is seen that for any operator $\hat{O}$, $(\hat{O}\mathbf{F}, \mathbf{F}) = (\mathbf{F}, \hat{O}\mathbf{F})^*$. It is simply stated here that if $(\mathbf{F}, \mathbf{G}) = 0$, the functions $\mathbf{F}$ and $\mathbf{G}$ are orthogonal. The reader is referred to [4, Ch. 2] for a more in-depth discussion.

The operator, $\hat{\Theta}$, is a Hermitian operator, which can be stated mathematically as[2]

$$(\hat{\Theta}\mathbf{H}_1, \mathbf{H}_2) = (\mathbf{H}_1, \hat{\Theta}\mathbf{H}_2). \tag{2.15}$$

Having this important feature for $\hat{\Theta}$ gives two important characteristics for the field $\mathbf{H}(\mathbf{r})$: 1) The eigenvalues ω^2/c_0^2 are real, and 2) two modes $\mathbf{H}_1$ and $\mathbf{H}_2$ with different frequencies are orthogonal.

The first feature can be proven by setting $(\hat{\Theta}\mathbf{H}, \mathbf{H}) = (\mathbf{H}, \hat{\Theta}\mathbf{H})^*$, which yields

$$\left(\frac{\omega^2}{c_0^2}\right)^* (\mathbf{H}, \mathbf{H}) = \left(\frac{\omega^2}{c_0^2}\right)(\mathbf{H}, \mathbf{H}), \tag{2.16}$$

so that $(\omega^2)^* = \omega^2$.

The second feature can be proven by using the mathematical definition of a Hermitian operator (2.15), where $\mathbf{H}_1$ has frequency ω_1, and $\mathbf{H}_2$ has frequency ω_2:

$$\frac{\omega_1^2}{c_0^2}(\mathbf{H}_2, \mathbf{H}_1) = \frac{\omega_2^2}{c_0^2}(\mathbf{H}_2, \mathbf{H}_1), \tag{2.17}$$

[2]The proof is left out here, but it is discussed in, *e.g.*, [4, Ch. 2].

which gives

$$(\omega_1^2 - \omega_2^2) \cdot (\mathbf{H}_2, \mathbf{H}_1) = 0, \tag{2.18}$$

from which it is seen that for $\omega_1 \neq \omega_2$, $(\mathbf{H}_2, \mathbf{H}_1) = 0$. Thus the two modes are orthogonal.

2.3 Energy Distribution in a Periodic Medium

Having established the equation for the modes of the magnetic field, the theoretical analysis is continued by finding the physical electromagnetic energy. As a result of the electromagnetic variational theorem [4, p. 14], the lowest frequency mode is given by the $\mathbf{H}(\mathbf{r})$ which minimizes

$$U_f(\mathbf{H}) \triangleq \frac{(\mathbf{H}, \hat{\Theta}\mathbf{H})}{(\mathbf{H}, \mathbf{H})}, \tag{2.19}$$

where U_f is known as the electromagnetic energy functional.

Rewriting the energy functional in terms of $\mathbf{E}(\mathbf{r})$, by use of (2.13), (2.9) and (2.10) along with the definition of the inner product (2.14), yields

$$U_f(\mathbf{H}) = \frac{\int d^3\mathbf{r} |\nabla \times \mathbf{E}(\mathbf{r})|^2}{\int d^3\mathbf{r} \epsilon(\mathbf{r}) |\mathbf{E}(\mathbf{r})|^2}. \tag{2.20}$$

From the above equation it is seen that in order to minimize the energy functional, the field must be concentrated within materials of high dielectric constant (to achieve a large denominator), and the amount of spatial oscillations must be minimized (to achieve a small numerator). It is important to remember that Gauss' law (2.3) must be satisfied.

From this it can be learned that the lowest order mode (lowest frequency) in a photonic crystal will distribute its energy in areas of high dielectric constant. Remembering that higher order modes must be orthogonal to the fundamental mode (and each other), the higher order modes are understood to have more of their fields in areas of low dielectric constant. This increases the energy functional and thus the energy, and is indeed what might give rise to a photonic bandgap (PBG), given the right geometry in the photonic crystal.

2.4 Symmetries in a Photonic Crystal

Because of the inherent symmetry in a photonic crystal structure, the electromagnetic fields within the photonic crystal will also display some interesting symmetries. The discussion of symmetries in the photonic crystal will draw extensively on the well-known symmetry considerations applied in solid-state physics. Readers unfamiliar with this area of physics might find it appropriate to consult the first chapters of [17].

A simple example will be utilized to describe translational symmetry. The structure of figure 2.1 is periodic in the y-direction with period a. That is, the structure has discrete translational symmetry in the y-direction. This can be stated mathematically with a translation operator for lattice vectors $\mathbf{R} = na\hat{\mathbf{y}}$,

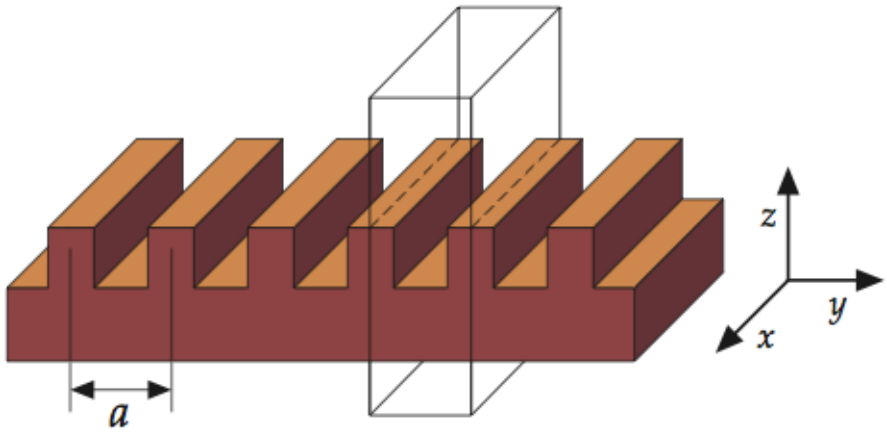

Figure 2.1: Dielectric structure with translational symmetry in y direction and infinite in x direction. Figure from [4].

where $n = 1, 2, 3, \ldots$. It should be obvious that there is also translational symmetry in the x-direction, although this symmetry is continuous: For any x vector, $d\hat{\mathbf{x}}$, the structure looks the same. The modes of $\hat{\Theta}$ are given as simultaneous eigenfunctions of the translation operators in x and y [4]:

$$\hat{T}_{d\hat{\mathbf{x}}} e^{ik_x x} = e^{ik_x(x-d)} = e^{-ik_x d} e^{ik_x x} \tag{2.21}$$

$$\hat{T}_{\hat{R}} e^{ik_y y} = e^{ik_y(y-na)} = e^{-ik_y na} e^{ik_y y}, \tag{2.22}$$

where k_x and k_y are the wave vector components in x- and y-direction, respectively. It can be seen quite easily from the above equations that two modes with their k_y vectors differing by $b = 2\pi/a$ have the same eigenvalues. The vector $\mathbf{b} = b\hat{\mathbf{y}}$ is the reciprocal lattice vector of the structure shown in figure 2.1, analogous to the reciprocal lattice of solid state physics (*cf.* [17, Ch. 2]).

From the above, it can be found [4] that the $\mathbf{H}$ field in this simple PhC is the incident field modulated by the periodicity of the crystal, which is known as Bloch's theorem

$$\mathbf{H}_{k_x,k_y}(\mathbf{r}) = e^{ik_x x} \cdot e^{ik_y y} \cdot \mathbf{u}_{k_y}(y, z). \tag{2.23}$$

Here, $\mathbf{u}(y + la, z) = \mathbf{u}(y, z)$, $l = \pm 0, 1, 2 \ldots$, *i.e.*, $\mathbf{u}$ is periodic in the y-direction with the same period as the crystal. Bloch's theorem states that all states with $k'_y = k_y + lb$ are identical. Thus, only states within $-\pi/a < k \leq \pi/a$ need to be considered to find the the complete y-dependence for the configuration in figure 2.1. This range of k_y's is referred to as the Brillouin zone [17, Ch. 2.3].

For the three-dimensional case, the lattice is described by its lattice vectors, $(\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3)$, and the lattice will have translational symmetry through all vectors described by $\mathbf{R} = l\mathbf{a}_1 + m\mathbf{a}_2 + n\mathbf{a}_3$, where (l, m, n) are arbitrary integers. The reciprocal lattice is described by its reciprocal lattice vectors, $(\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3)$, given by $\mathbf{a}_i \cdot \mathbf{b}_j = \delta_{ij}$. Bloch's theorem for three dimensional translational symmetry is given as

$$\mathbf{H}_{\mathbf{k}}(\mathbf{r}) = e^{i\mathbf{k}\cdot\mathbf{r}} \cdot \mathbf{u}_{\mathbf{k}}(\mathbf{r}), \tag{2.24}$$

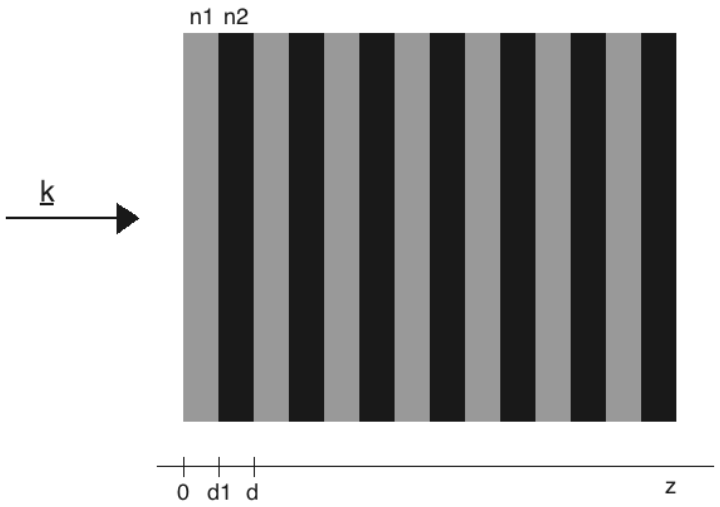

Figure 2.2: A 1D PhC of two alternating materials of high and low dielectric constant.

where $\mathbf{k} = k_1\mathbf{b}_1 + k_2\mathbf{b}_2 + k_3\mathbf{b}_3$ lies within the Brillouin zone and $\mathbf{u}_\mathbf{k}(\mathbf{r}) = \mathbf{u}_\mathbf{k}(\mathbf{r}+\mathbf{R})$.

The notion of Brillouin zones is extremely useful as the Brillouin zone contains enough information about the crystal to find its response everywhere. Some crystals can, however, exhibit more symmetries, such as rotational and mirror symmetry, yielding symmetries within the Brillouin zone itself. Thus it might not be needed to calculate the response inside the entire Brillouin zone. The smallest reciprocal lattice zone that can be utilized to calculate a crystal's band diagram, is known as the irreducible Brillouin zone [4, pp. 36,37].

2.5 An Example: Theoretical Treatment of a 1D PhC

In analyzing 2D-, 3D- or even complicated 1D photonic crystal structures, the use of simulation software is extensive. This is because the solution to the master equation often must be found numerically, and if it can be done analytically it is often highly complex.

Simulation software is extremely useful, but one does run the risk of not understanding the actual physics if the software is trusted blindly, and the results are not carefully investigated. In order to build a solid understanding of the origin of the bandgap in a photonic crystal, the most simple of all examples will be utilized: A 1D photonic crystal consisting of two infinite plates of high and low dielectric constant, repeated ad infinitum in the direction of periodicity (*cf.* figure 2.2).

The analysis of the 1D PhC will start with Maxwell's equations (2.1–2.4). In much the same way as the master equation (2.12) was derived, an equation for $\mathbf{E}(\mathbf{r})$ can be found

$$\nabla\times\nabla\times\mathbf{E}(\mathbf{r}) = \left(\frac{\omega}{c}\right)^2 \epsilon(\mathbf{r})\,\mathbf{E}(\mathbf{r}), \tag{2.25}$$

which is no longer a Hermitian eigenvalue problem, but this will not be discussed here.[3]

In order to simplify (2.25), the incoming wave is assumed to have $\mathbf{k} = k \cdot \hat{\mathbf{z}}$, from which it follows that all electric fields (reflected and transmitted from each interface) are propagating orthogonal to the interfaces of the crystal. It also follows that the **E**-field will be polarized parallel to the boundaries. Once this is established, z is substituted for $\mathbf{r}$ in (2.25), and the vector notation is dropped.

From figure 2.2, the dielectric constant, $\epsilon(\mathbf{r})$, is seen to be periodic in the z-direction and is constant within each material. That is, $\epsilon(\text{material 1}) = \epsilon_1$, $\epsilon(\text{material 2}) = \epsilon_2$ and $\epsilon(z) = \epsilon(z + d)$.

The fact that all fields propagate parallel to the z-axis, and the periodicity of the dielectric function, yields a periodicity for the electric field, $|E(z + d)| = |E(z)|$. Using this periodic condition in (2.25) gives

$$E(z) = \mathrm{e}(z)e^{ik_z z}; \quad \mathrm{e}(z) = \mathrm{e}(z + d), \tag{2.26}$$

which is a Bloch wave, as discussed in chapter 2.4.

If the field in just one layer is inspected, the dielectric constant, ϵ_m, is constant, and the equation for the electric field is

$$\frac{d^2 E_m(z)}{dz^2} - \left(\frac{\omega}{c}\right)^2 \epsilon_m E_m(z) = 0, \tag{2.27}$$

where m denotes the material. The general solution to (2.27) is

$$E_m(z) = E_{+,m} e^{ik_m z} + E_{-,m} e^{-ik_m z}, \tag{2.28}$$

where $+$ and $-$ denote propagation in the positive- and negative z-direction, respectively.

Rewriting the constants of (2.28) as $E_{+,m} = ({}^{A_m}/_2 + {}^{B_m}/_{2i})$ and $E_{-,m} = ({}^{A_m}/_2 - {}^{B_m}/_{2i})$ gives

$$E_m(z) = A_m \cos(k_m z) + B_m \sin(k_m z), \tag{2.29}$$

and its derivative

$$E'_m(z) = -k_m A_m \sin(k_m z) + k_m B_m \cos(k_m z). \tag{2.30}$$

In order to find how the constants A_m and B_m relate to each other, some boundary conditions are needed. Knowing that the electric (and magnetic) fields, as well as the derivative of the electric field, must be continuous across all surfaces, the boundary conditions for boundary 0, 1 and 2 are denoted as[4]

$$\mathbf{E}_0 = \begin{bmatrix} E(0) \\ E'(0) \end{bmatrix}; \quad \mathbf{E}_1 = \begin{bmatrix} E(d_1) \\ E'(d_1) \end{bmatrix}; \quad \mathbf{E}_2 = \begin{bmatrix} E(d) \\ E'(d) \end{bmatrix}, \tag{2.31}$$

[3] It may seem strange not to use the master equation, after having spent time deriving it. As always, there is a free choice in what field to calculate for and for this specific problem (2.25) is most useful.

[4] Note that this requires the incoming field to be known, and as such might not seem any more informational than the original constants. However, this will prove most useful in order to find the dispersion relation.

where $d = d_1 + d_2$ in $\mathbf{E}_2$. The reason behind this notation will become apparent later.

By looking at (2.29) and (2.30) for material 1, it is seen that $E_1(0) = A_1$ and $E_1'(0) = k_1 B_1$, which must be equal to $E(0)$ and $E'(0)$, respectively. The electric field in material 1 can now be expressed in terms of the boundary condition $\mathbf{E}_0$ as

$$E_1(z) = E(0)\cos(k_1 z) + \frac{E'(0)}{k_1}\sin(k_1 z) \tag{2.32}$$

$$E_1'(z) = -k_1 E(0)\sin(k_1 z) + E'(0)\cos(k_1 z). \tag{2.33}$$

At the boundary between material 1 and 2, the field will be given by

$$E_1(d_1) = E(d_1) = E(0)\cos(k_1 d_1) + \frac{E'(0)}{k_1}\sin(k_1 d_1) \tag{2.34}$$

$$E_1'(d_1) = E'(d_1) = -k_1 E(0)\sin(k_1 d_1) + E'(0)\cos(k_1 d_1), \tag{2.35}$$

which can be written more compact in matrix form as

$$\mathbf{M}_1 \cdot \mathbf{E}_0 = \mathbf{E}_1, \tag{2.36}$$

where

$$\mathbf{M}_1 = \begin{bmatrix} \cos(k_1 d_1) & \frac{\sin(k_1 d_1)}{k_1} \\ -k_1 \sin(k_1 d_1) & \cos(k_1 d_1) \end{bmatrix}. \tag{2.37}$$

The electric field in material 2 will be found in much the same way as in material 1. In order to simplify the notation, the zero point on the z-axis is now moved to d_1. This is not a problem, since the boundary conditions at d_1 are known from $\mathbf{E}_1$. The constants A_2 and B_2 are found in the same way as for material 1, yielding

$$E_2(z) = E(d_1)\cos(k_2 z) + \frac{E'(d_1)}{k_2}\sin(k_2 z) \tag{2.38}$$

$$E_2'(z) = -k_2 E(d_1)\sin(k_2 z) + E'(d_1)\cos(k_2 z). \tag{2.39}$$

The field at the boundary $z = d_2$ is then found as

$$E_2(d_2) = E(d) = E(d_1)\cos(k_2 d_2) + \frac{E'(d_1)}{k_2}\sin(k_2 d_2) \tag{2.40}$$

$$E_2'(d_2) = E'(d) = -k_2 E(d_1)\sin(k_2 d_2) + E'(d_1)\cos(k_2 d_2). \tag{2.41}$$

Again, this can be written in the matrix form

$$\mathbf{M}_2 \cdot \mathbf{E}_1 = \mathbf{E}_2, \tag{2.42}$$

where

$$\mathbf{M}_2 = \begin{bmatrix} \cos(k_2 d_2) & \frac{\sin(k_2 d_2)}{k_2} \\ -k_2 \sin(k_2 d_2) & \cos(k_2 d_2) \end{bmatrix}. \tag{2.43}$$

From (2.26), $\mathbf{E}_2 = e^{ik_z d}\mathbf{E}_0$. Inserting this into (2.42) to find an expression for $\mathbf{E}_1$ and using this expression in (2.36) yields

$$\mathbf{M}_2 \cdot \mathbf{M}_1 \cdot \mathbf{E}_0 = \mathbf{M}_{tot} \cdot \mathbf{E}_0 = e^{ik_z d}\mathbf{E}_0, \tag{2.44}$$

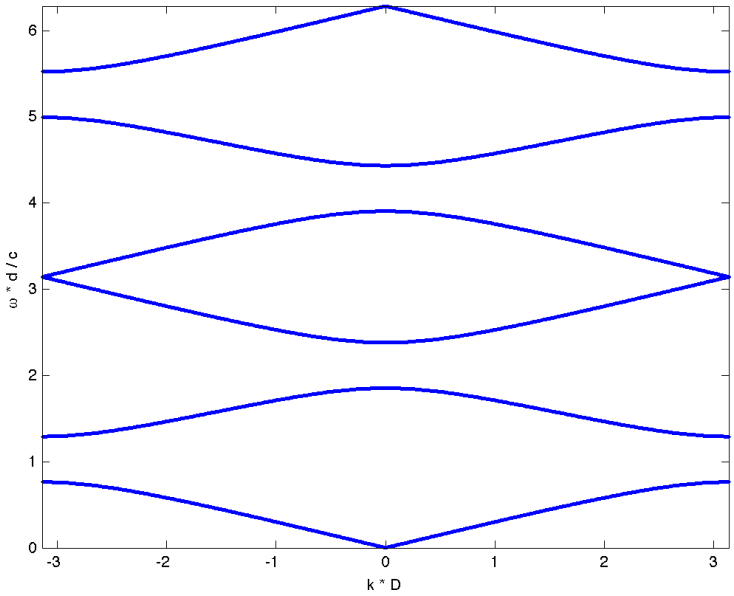

Figure 2.3: Band-diagram for a 1D PhC with $n_1 = 1$, $n_2 = 2$ and $d_1 = d_2 = D$. The x-axis is given as $k_z \cdot D$, and the y-axis is given as $\omega \cdot d/c$.

which is an eigenvalue problem. For a solution to exist it is required that $\det(\mathbf{M}_{tot} - \mathbf{I}e^{ik_z d}) = 0$, where $\mathbf{I}$ is the identity matrix. By setting the determinant equal to zero, an equation for the effective wavenumber, k_z, is found

$$\cos(k_z d) = \cos(k_1 d_1 + k_2 d_2) - \frac{(k_1 - k_2)^2}{k_1 k_2} \sin(k_1 d_1) \sin(k_2 d_2), \tag{2.45}$$

which can be plotted to yield the dispersion relation—or the band-diagram—of figure 2.3 [9].

2.6 The Trigonal Lattice, its Reciprocal Lattice and the Irreducible Brillouin Zone

Throughout the experimental work in this research, a trigonal lattice photonic crystal will be studied. Thus, the trigonal lattice will here be utilized to discuss some important theoretical concepts, which indeed will be needed in later chapters.

Figure 2.4 depicts the trigonal lattice and its basis vectors $\mathbf{a}_1$ and $\mathbf{a}_2$. The distance between neighboring lattice points is the lattice constant $a = |\mathbf{a}_1| = |\mathbf{a}_2|$,

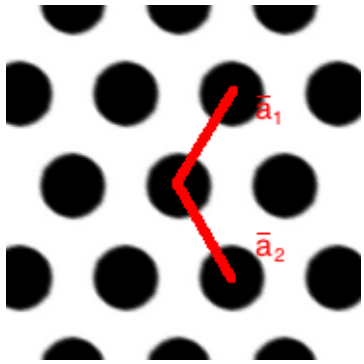

Figure 2.4: The trigonal lattice and its basis vectors.

and the angle between the basis vectors is 120°. The basis vectors are given by

$$\mathbf{a}_1 = \frac{a(\hat{\mathbf{x}} + \sqrt{3}\hat{\mathbf{y}})}{2} \tag{2.46}$$

$$\mathbf{a}_2 = \frac{a(\hat{\mathbf{x}} - \sqrt{3}\hat{\mathbf{y}})}{2}, \tag{2.47}$$

which is easy to verify by applying trigonometric considerations.

The reciprocal lattice, defined in section 2.4, is given by the reciprocal basis vectors $\mathbf{b}_1$ and $\mathbf{b}_2$ [17, p. 29]

$$\mathbf{b}_1 = 2\pi \frac{\mathbf{a}_2 \times \hat{\mathbf{z}}}{\mathbf{a}_1 \cdot \mathbf{a}_2 \times \hat{\mathbf{z}}} \tag{2.48}$$

$$\mathbf{b}_2 = 2\pi \frac{\hat{\mathbf{z}} \times \mathbf{a}_1}{\mathbf{a}_2 \cdot \hat{\mathbf{z}} \times \mathbf{a}_1}. \tag{2.49}$$

Using the expressions (2.46) and (2.47) in (2.48) and (2.49), the basis vectors for the reciprocal trigonal lattice are found

$$\begin{aligned} \mathbf{b}_1 &= \frac{4\pi}{a} \frac{(\hat{\mathbf{x}} - \sqrt{3}\hat{\mathbf{y}}) \times \hat{\mathbf{z}}}{(\hat{\mathbf{x}} + \sqrt{3}\hat{\mathbf{y}}) \cdot (\hat{\mathbf{x}} - \sqrt{3}\hat{\mathbf{y}}) \times \hat{\mathbf{z}}} \\ &= \underline{\frac{2\pi}{\sqrt{3}a} \cdot (\sqrt{3}\hat{\mathbf{x}} + \hat{\mathbf{y}}),} \end{aligned} \tag{2.50}$$

and

$$\begin{aligned} \mathbf{b}_2 &= \frac{4\pi}{a} \frac{\hat{\mathbf{z}} \times (\hat{\mathbf{x}} + \sqrt{3}\hat{\mathbf{y}})}{(\hat{\mathbf{x}} - \sqrt{3}\hat{\mathbf{y}}) \cdot \hat{\mathbf{z}} \times (\hat{\mathbf{x}} + \sqrt{3}\hat{\mathbf{y}})} \\ &= \underline{\frac{2\pi}{\sqrt{3}a} \cdot (-\sqrt{3}\hat{\mathbf{x}} + \hat{\mathbf{y}}),} \end{aligned} \tag{2.51}$$

which are shown in figure 2.5, along with the irreducible Brillouin zone.

As briefly discussed previously, the irreducible Brillouin zone is the smallest region by which the entire Brillouin zone can be reached by rotational and mirror operations. It should be clear from figure 2.5 that the crystal as a whole remains

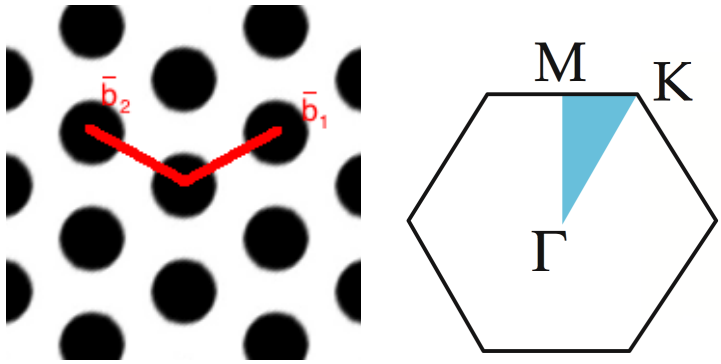

Figure 2.5: (a) The reciprocal lattice for the trigonal configuration and (b) its Brillouin zone with the irreducible Brillouin zone shaded [4, p. 76].

unchanged upon 60° rotations of the lattice. It is also seen to be mirror symmetry along all edges of the irreducible Brillouin zone. Thus it is understood that by calculating the band diagram in the irreducible Brillouin zone, the response of the entire crystal is found.

Furthermore, the maxima and minima of the band diagrams are usually found along the edges of the irreducible Brillouin zone.[5] For this reason, it is enough to calculate the response of the crystal along these edges to find the band-diagram and the crystals photonic bandgap. The band-diagram for the trigonal lattice PhC is plotted as a function of the in-plane wavevector, $\mathbf{k}_{||}$. This vector runs from the center point, Γ, to all the points along the edge of the irreducible Brillouin zone.

[5] At least this is the case for most structures, and indeed for the structure at hand. See [4].

Chapter 3

Design

In the previous chapter the underlying theory of photonic crystals (PhCs) was presented, and the 2D trigonal lattice was briefly discussed. This chapter will investigate the trigonal lattice PhC in more detail by use of simulation software. First—and most importantly—the optimal hole radius and lattice constant will be found for TE-like modes.[1] In addition, some methods of coupling light into and out of photonic structures will be described, and some considerations regarding non-ideal PhCs will be outlined.

3.1 Material and Structure

As discussed in chapter 1, the material of choice for eventual fabrication is silicon (Si) because of its excellent optical properties, well-known fabrication methods and possible integration with CMOS electronics. For single mode operation at 1.55 µm, slab waveguides in silicon-on-insulator (SOI) with a core thickness of around 220 nm and a core width of around 500 nm are often used [18]. The SOI wafer structure consists of a top single crystalline Si layer (often 220 nm thick) and a silicon dioxide (SiO_2, oxide) cladding, often referred to as a buried oxide (BOX), on top of a Si substrate. The BOX provides a high contrast in refractive index (1.528 in oxide *vs.* 3.477 for Si, at 1.55 µm wavelength), confining light in the top Si layer by total internal reflection. To ensure minimal coupling of stray fields to the substrate, the BOX needs to be thick enough (at least 1 µm) [18].

By fabricating a 2D PhC in the SOI wafer, it is possible to create a photonic waveguide by removing a line of holes—referred to as a W1 waveguide [18]. The photonic bandgap (PBG) of the 2D PhC provides confinement of light in the plane of the SOI, while total internal reflection provides confinement in the out of plane direction. However, the simple W1 waveguide is no longer single mode: The TE- and TM-like modes can couple, causing leakage by radiation away from the plane [4, pp. 144-147]. There are methods of improving the design, such as oxide removal (creating a Si membrane) [18], or reducing the optical volume of the waveguide to keep it single mode [19]. The relatively simple structure of the W1 waveguide does, however, make it well suited for initial studies.

[1] For an explanation on TE-/TM modes see, *e.g.*, [4].

The fabrication of PhC structures in SOI can be done by either deep or shallow etching. This corresponds to etching the holes of the PhC in both the top Si layer and the BOX, or in the top Si layer only. The deep etching has the advantage of increasing symmetry and refractive index contrast, but it often leads to an increase in surface roughness [18]. As this might severely increase scattering losses, the initial fabrication should focus on properly defining the PhC in the Si layer only. Once this is achieved, further investigations can explore the effect of etching into the BOX.

3.2 Finding the Band Diagram of the 2D Trigonal PhC

The simulations outlined in this section were performed using MIT Photonic Bands (MPB) [20], which is an open source simulation software designed to find the band diagram(s) of photonic structures.[2]

During simulations it is very helpful to normalize the radius and frequency by the lattice constant, a, thus solving the problem without considering absolute physical quantities. The normalized radius is given as

$$r_n = \frac{r}{a}, \tag{3.1}$$

and the normalized frequency is given as

$$\nu_n = \frac{\nu a}{c_0} = \frac{a}{\lambda}, \tag{3.2}$$

where c_0 is the speed of light in vacuum and λ is the wavelength of the incoming light. The reference frequency (for which $\nu_n = 1$) is seen to be at $\nu = c_0/a$, *i.e.*, with wavelength equal to the lattice constant.

When a photonic bandgap (PBG) is found from the simulations, the frequency at the center of the bandgap is taken to be the optimal frequency for the given geometry. The normalized center frequency is calculated as

$$\nu_{n,\text{center}} = \frac{\nu_{n,l} + \nu_{n,u}}{2}, \tag{3.3}$$

where $\nu_{n,l}$ and $\nu_{n,u}$ denote the normalized lower and upper frequencies of the bandgap, respectively. When the normalized center frequency for a given structure is found, the intended operational wavelength ($\lambda = 1.55\,\mu\text{m}$) can be put in to (3.2), yielding the optimal lattice constant of the structure.

Figure 3.1 contains the band diagrams for TE and TM modes, simulated with MPB for a hole radius of $r/a = 0.4$ in the trigonal lattice PhC. The PBGs are shaded in yellow, and the light cone is shaded dark purple.[3] The normalized lower and upper frequencies of the largest TE bandgap are 0.24418 and 0.40425, respectively, yielding a center frequency of 0.324685. From the center frequency, the optimal lattice constant for this configuration is found to be 503.26 nm, using

[2] nanoHUB.org [21] offer a free online version of MPB which can be used directly in any browser. It has been used extensively in this work.

[3] Modes that are inside the light cone are radiative, not guided modes [4, Ch. 7].

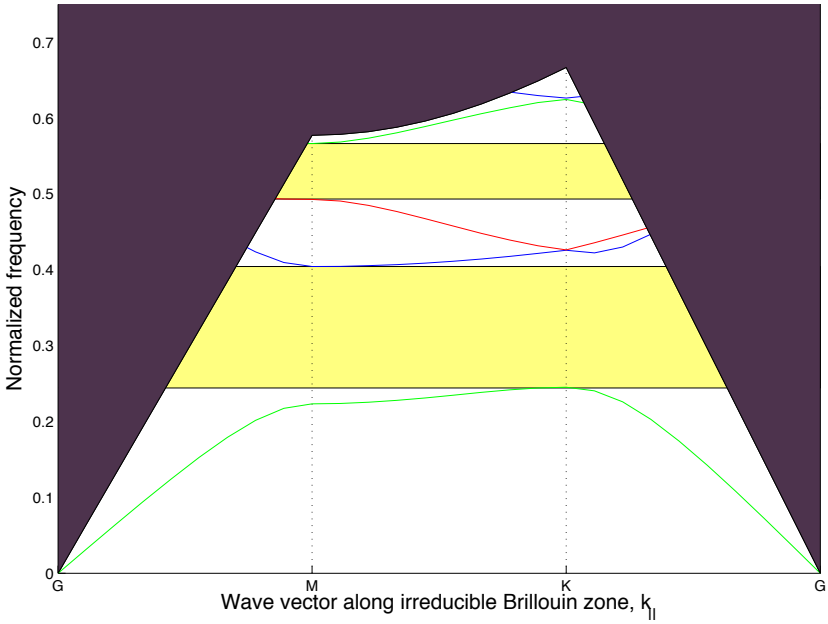

(a) Band diagram for TE modes.

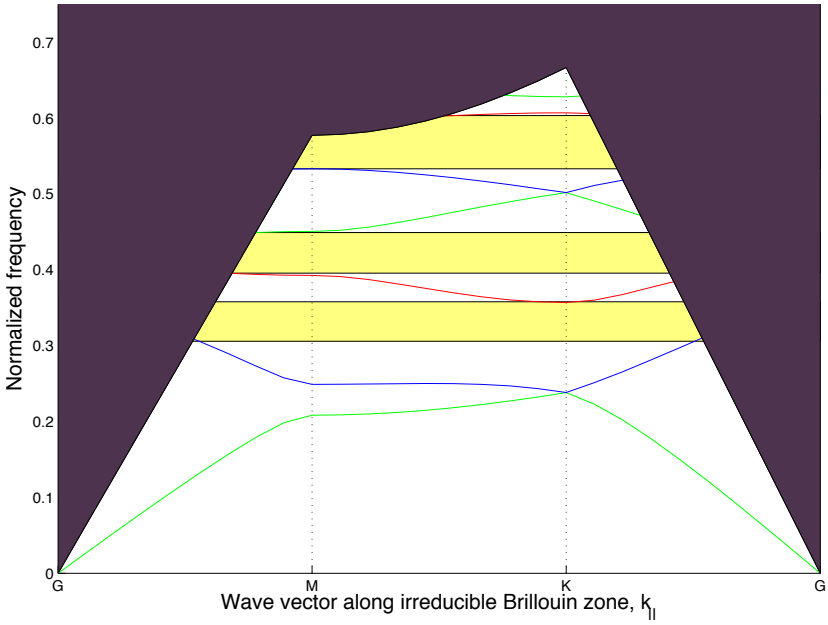

(b) Band diagram for TM modes.

Figure 3.1: Band diagram for (a) TE- and (b) TM modes for the trigonal lattice PhC with hole radius $r = 0.4a$. The y axis denotes the normalized frequency, given as $\nu_n = a/\lambda$, and the x axis denotes the in-plane wave vector along the irreducible Brillouin zone.

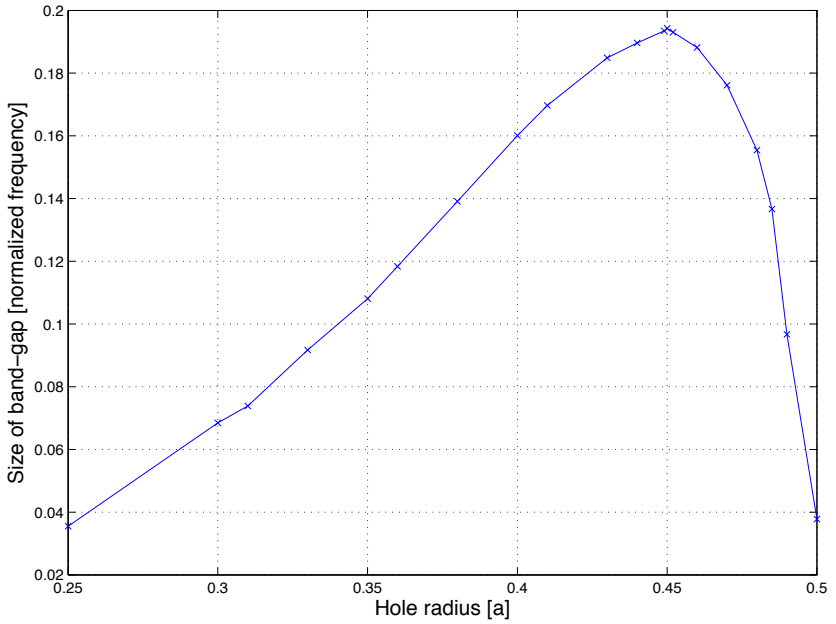

Figure 3.2: Size of largest bandgap (in units of normalized frequency) as a function of normalized hole radius in the 2D trigonal lattice PhC, for TE modes. There is a clear maximum for $r/a = 0.45$, which is the theoretically optimal hole radius.

(3.2). It can be seen from figure 3.1 that there is a complete bandgap for both the fundamental TE- and TM mode [4].

The bandgap simulations in MPB were performed for hole radii between $0.3a$ and $0.5a$, and the size of the fundamental bandgap along with the optimal lattice constant were found for each structure. They are plotted in figure 3.2 and 3.3, respectively. The largest bandgap for the TE modes in this configuration is found at $r = 0.45a$, and the corresponding optimal lattice constant is 611.6 nm. The MPB code is outlined in appendix A, along with the TE band diagrams for $r/a =$ 0.3, 0.35, 0.4 and 0.45.

In order to maximize the confinement of light in a future defect (*e.g.*, a waveguide), the PhC with the largest possible PBG should be chosen. However, there are practical limitations set forth by the processing tools used, such as minimum feature size attainable by electron beam lithography. If hole sizes larger than $0.25a$ are chosen, the minimum feature size will be the width of the Si region between nearest neighboring air holes. This width can be calculated from the previously found optimal lattice constant as $d = a - 2r$, and is plotted in figure 3.4 (in units of nm) as a function of normalized hole radius.

By use of figures 3.2, 3.3 and 3.4 it is possible to find the optimal PhC configuration, given the smallest feasible hole separation. The hole separation for $r = 0.4a$ is 100.7 nm, which is 64% larger than the hole separation for $r = 0.45a$, which is 61.25 nm. For the initial fabrication experiments, $r/a = 0.4$ is chosen as the hole radius, as this seems like an appropriate compromise between feature size and PBG size.

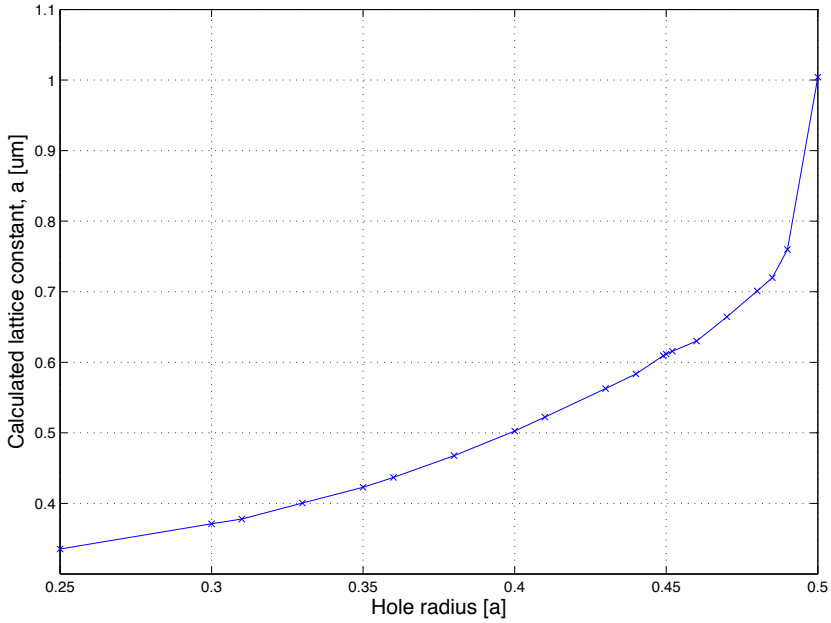

Figure 3.3: Optimal lattice constant for the PBGs found in figure 3.2, for operation at $\lambda = 1.55\,\mu\text{m}$, as a function of normalized hole radius. The optimal lattice constants for $r/a = 0.4$ and 0.45 are 503.26 nm and 611.60 nm, respectively.

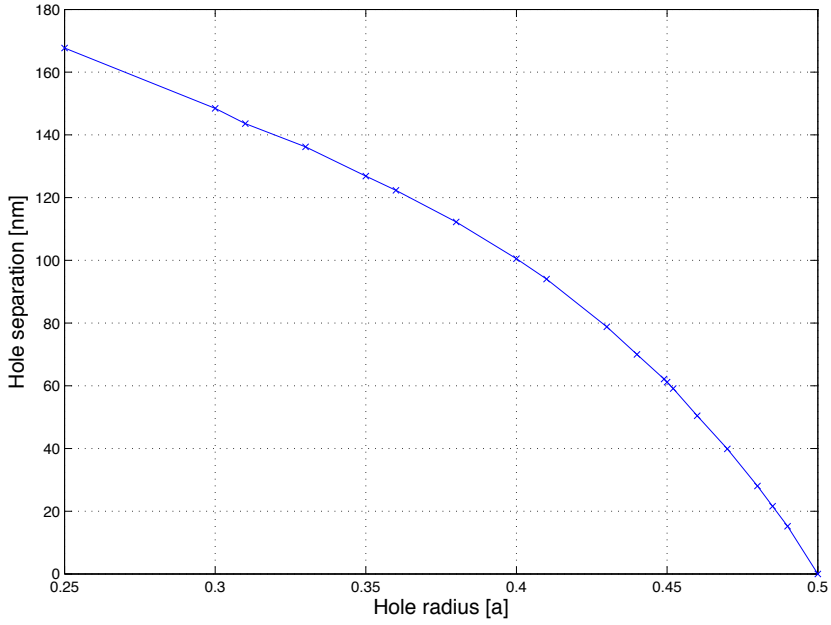

Figure 3.4: Distance between nearest neighboring holes in the trigonal lattice PhC as a function of hole radius. For holes larger than $0.25a$ this is also the minimum feature size, which might be limited by fabrication methods.

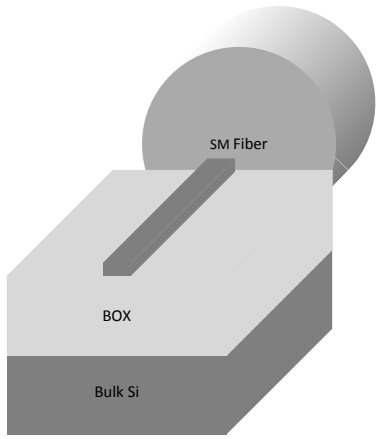

Figure 3.5: Illustration of butt-coupling between a single mode (SM) fiber and a slab waveguide in SOI. The SM fiber has a typical diameter of 8 µm, whereas the slab waveguide is 500 nm wide and 220 nm high.

3.3 Coupling

The ultimate performance test of a photonic waveguide is to measure the attenuation of light in the waveguide. To do so, light must be coupled into and out of the given waveguide, but this creates a problem: The coupling will inevitably introduce additional attenuation. Furthermore, the alignment of the input light (*e.g.* from an optical fiber) with the actual structures must be very precise, in order not to introduce even further attenuation which might vary between tests.

Several methods of coupling light into photonic structures exist, and some of them will be outlined below.

3.3.1 Coupling Light From a Slab Waveguide to a PhC

Depending upon application, it might be useful to couple light between slab waveguides and PhC structures. For instance, a possible test setup could consist of two slab waveguides placed in mirror symmetry on an SOI, where a PhC waveguide is placed in one of the two paths. This would create similar input and output conditions for the two slabs. Upon measuring the transmission through the individual paths, the difference would be the transmission loss in the PhC waveguide plus the coupling loss between PhC and slab waveguide. (This would of course require two identical slab waveguides, which in turn requires low variations in processing conditions.)

The coupling losses at the interface between a single mode slab waveguide and trigonal lattice PhC can be quite low. Sanchis *et al.* [22] used both simulation tools and analytic expressions to show that the coupling efficiency varies between 87% and almost 100%, depending on cut position of the PhC. The PhC of interest had a hole radius and a lattice constant of $r/a = 0.3$ and $a = 364.2\,\text{nm}$, respectively, at $\lambda = 1.55\,\mu\text{m}$.

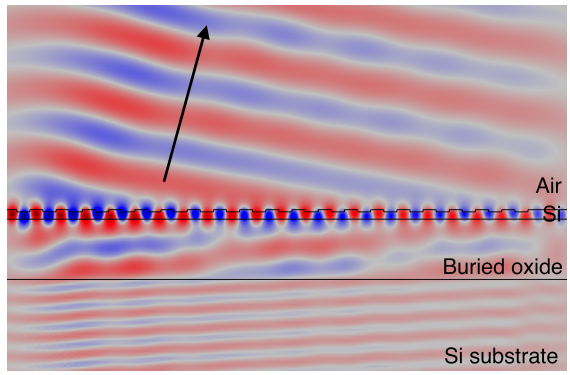

Figure 3.6: Electric field plot of the fundamental TE mode coupled from the SOI grating structure into air. Red and blue color depicts the maximum and minimum electric field, respectively (given in V/m). Notice the amount of light reflected towards the substrate. The simulation is based upon [25].

3.3.2 Butt-Coupling of Single Mode Fiber

Perhaps the most straightforward way of coupling light from a fiber to a photonic waveguide is a simple butt-coupling from a single mode fiber. A single mode fiber has a core diameter of about 8 µm for operation at the 1.55 µm telecommunications wavelength, whereas a slab waveguide in SOI has typical dimensions of 220 nm thickness and 500 nm width. The differing dimensions of the two waveguides leads to a mode mismatch and a high coupling loss of about 20 dB [23]. The setup is depicted in figure 3.5.

One method of lowering the coupling loss is to introduce a tapered waveguide segment between the fiber and the slab waveguide [24]. For long tapers (around 1 mm) it is possible to achieve very low radiation and coupling losses (less than 0.5 dB coupling loss for a 0.7 mm taper in [24]). However, for shorter tapers the radiation losses will increase.

3.3.3 Tapered Grating Coupler

A different approach to couple light into photonic structures is by near vertical coupling from a fiber to a grating. By use of a 10 µm × 10 µm square grating, it is possible to focus most of the light from the fiber onto the grating, significantly reducing the coupling losses due to size mismatch. The 10 µm wide grating can be tapered down to a 500 nm slab waveguide to further reduce the coupling loss, as mentioned above.

A Simple SOI Grating

A 2D simulation was performed in the finite element method (FEM) simulation software COMSOL Multiphysics (COMSOL) to find the electric field plot of a grating coupler, based upon the work of Van Laere *et al.* [25]. The grating has a period of 610 nm, etch depth of 50 nm and a 50% duty cycle, meaning that the width of the trenches is the same as the width of the plateaus. The number of

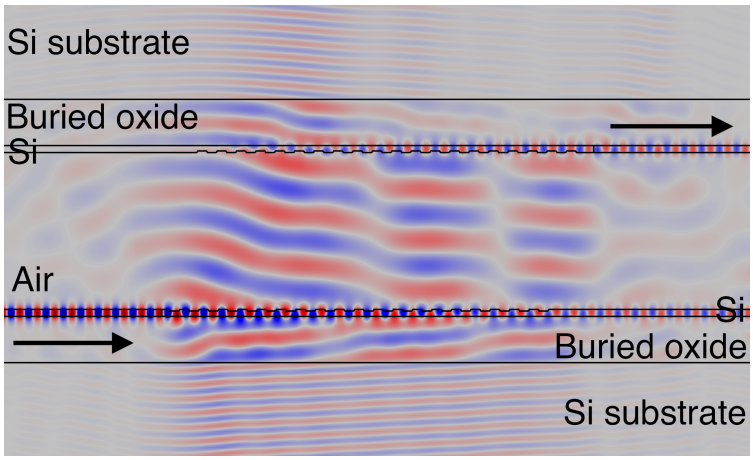

Figure 3.7: Electric field plot of the fundamental TE mode coupled from the bottom SOI grating structure into the top grating structure via air. 27.13% of the input light is coupled to the fundamental TE mode of the top slab waveguide.

periods is set to 20, resulting in a 12.2 µm long grating. The grating is created in SOI with a 220 nm top Si layer, a 1.45 µm buried oxide layer and a Si substrate which is terminated by using perfectly matched layers to attenuate any reflections from the simulation boundaries (*cf.* appendix B).

The resulting electric field plot is shown in figure 3.6, where the light is coupled to the fundamental TE mode in the left part of the slab waveguide. Light is seen to couple into air at 10° with respect to the normal of the SOI structure, while maintaining its fundamental TE mode. A substantial amount of light can be seen to reflect down towards the substrate, resulting in a lowering of the percentage of light available to couple to a fiber.

It is very important to note that these simulations are done in 2D, thereby assuming the structures to be infinite in the out of plane direction. This limits the accuracy of the results compared to a full 3D study, but the 2D simulations still provide useful information concerning the performance of the structures.

In order to find the coupling efficiency of the grating, two identical gratings placed at a distance of 10 µm are simulated. Light is coupled out of the bottom grating and into the top grating. The dimensions of the gratings are the same as for figure 3.6, and the resulting electric field plot is depicted in figure 3.7.

For the setup in figure 3.7, 27.13% of the incident light is found to couple to the fundamental TE mode of the top slab waveguide. Because the two SOIs are placed in close proximity (9.85 µm distance in the vertical direction) there might be multiple reflections between the Si surfaces. The slightly altered field distribution in the air region of figure 3.7, as compared to the field in figure 3.6, is probably due to such reflections. This might result in a higher coupling efficiency than what should be expected from coupling through a fiber, even if the mismatch between grating and fiber is accounted for. The backwards reflected light to the input port is found to be 3.50%, only slightly larger than for the single grating of figure 3.6 (2.80%). This might suggest minimal coupling due to reflections from

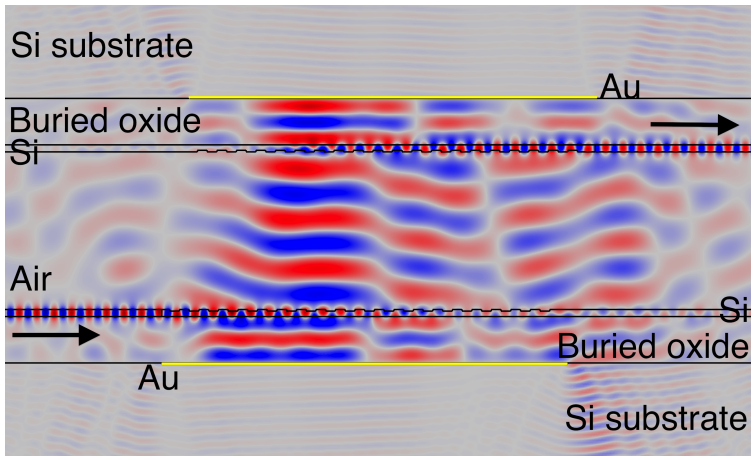

Figure 3.8: Electric field plot of the fundamental TE mode coupled from the bottom SOI grating structure into the top grating structure via air. 40 nm Au mirrors (highlighted in yellow) are added below the grating structures to reflect light scattered towards the substrate. 66.22% of the input light is found to couple to the fundamental TE mode of the top slab waveguide.

the SOI surface.

If the coupling efficiency in the double grating structure is taken to be correct, and by assuming the same efficiency for both outwards and inwards coupling, the efficiency for a single grating can be found. Taking the square root of the coupling efficiency of the double structure, the efficiency for the single grating is found as $\sqrt{0.2713} = 52.1\%$, which is in good agreement with [25]. Upon introducing the single mode fiber in the path of the outwards coupled light, the actual coupling efficiency for a grating-to-fiber coupler will probably be slightly less due to, *e.g.*, reflections at the air/fiber core interface. The coupling efficiency upon introducing the fiber is found to be above 45% in [25].

Adding a Bottom Mirror

Van Laere *et al.* utilized benzocyclobutene (BCB) in a complex oxide/Si/BCB/-Au/BCB structure (see [25]), in which the gold (Au) served as a bottom mirror to reflect light incident on the substrate back towards air. This was found to significantly increase the coupling efficiency of the grating structure.

Inspired by this work, a simple approach to enhance the coupling efficiency is described here. It is important to keep the SOI structure relatively unaltered in order not to disturb the performance of the other photonic structures; however, if the SOI is to be fabricated in-house by deposition methods, it might be feasible to add a gold mirror below the BOX layer in the part of the structure where the grating is to be placed. It should be noted that this adds a certain amount of processing steps, which might introduce higher uncertainties in fabrication.

The SOI and grating are kept at the same dimensions as before, but a 40 nm Au layer is added below the BOX layer in the area of the grating. The refractive index of Au is set to $0.55 + 11.5i$ for $\lambda = 1.55\,\mu\text{m}$ [26]. The resulting electric field

plot is presented in figure 3.8: Notice how the amount of light coupled to the substrate is significantly lower than for the electric field plot in figure 3.7.

Upon adding the Au mirror to both gratings, 66.22% of the light is coupled into the top slab waveguide. It is found that 8.31% is reflected back to the input, caused by the increased reflections in the system due to the mirrors. It is possible that the found value is artificially high because of the close proximity of the two SOIs. Looking apart from this, and employing the same logic as before, 81.38% of the light from the bottom waveguide is found to be successfully coupled into the top waveguide via the gratings. Upon comparison with [25] this value is much higher than what would be expected for coupling from a single grating towards a fiber. It is suspected that the Au mirrors cause multiple reflections between the SOIs, causing light which is not initially coupled into the grating to be coupled after some reflections. However, this result strongly suggests that introducing a gold mirror below the buried oxide will indeed increase the coupling efficiency. More thorough simulations must be employed to find the exact improvement, in order to weigh this against the disadvantages of introducing extra processing steps.

3.4 PhC Waveguide Simulations

3.4.1 The W1 Waveguide

Figure 3.9 depicts a COMSOL simulation of a W1 waveguide with butt-coupling to input and output ports. The light is excited from the left port, and the amount of transmitted and reflected light is measured at the right and left port, respectively. A transmission efficiency of 85.8% and a reflection of only 0.26% is found from the simulation. It is clear that the cut position of the PhC/slab waveguide interface is probably not optimal, so that most of the loss is attributed as coupling loss, not transmission loss.

The simulation presented in figure 3.9 serves as the reference point for the roughness considerations outlined in the next section. In the simulation model for the rough holes there are two columns of perfect holes at the ends of the PhC region. These are included to avoid any additional coupling losses, such that the additional loss upon adding roughness to the holes can be attributed mostly to transmission losses.

3.4.2 Introducing Hole Roughness

A simple method of quantifying hole roughness in 2D was developed in [9]. A brief summary of the method is given in appendix B.2 and some results are reprinted here. The parameter by which the roughness is quantified is the standard deviation, σ, of a normally distributed function. The average transmission and reflection coefficients for waveguides with σ between 0 and 0.30 are calculated, using 10 sample points at each value of σ. The resulting transmission and reflection coefficients are plotted in figure 3.10a and 3.10b, respectively, along with the sample standard deviation and a third degree polynomial fit of the experimental data.

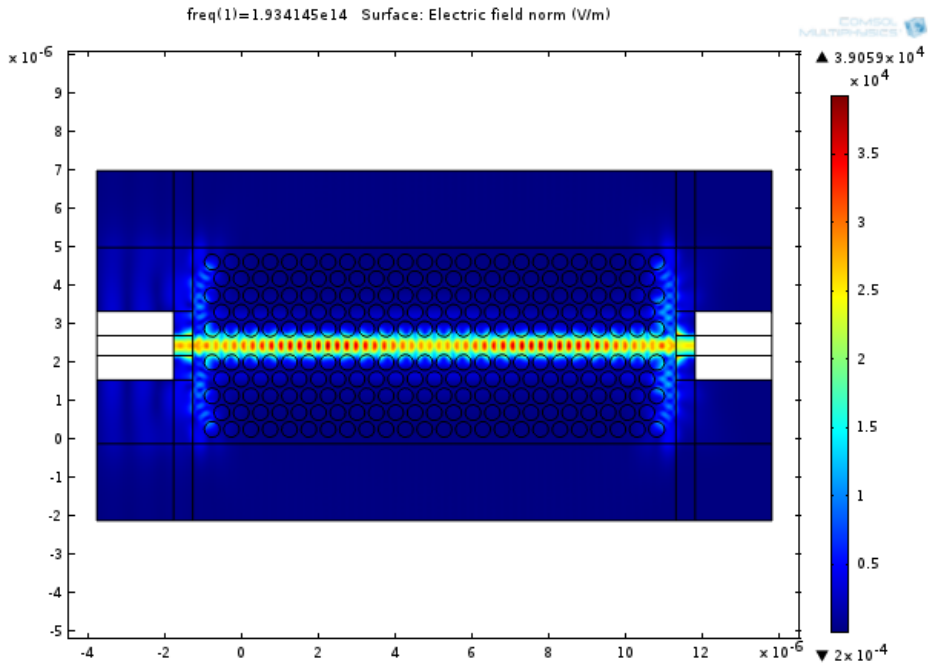

Figure 3.9: Normalized electric field plot in a W1 waveguide. The colorbar on the right defines the strength of the electric field, from zero (deep blue) to maximum (red). The numbers on the scale are given in units of V/m, but the absolute values are not important.

Figures 3.11–3.13 contain a few examples of the simulated PhCs with rough holes. Figure 3.11 depicts the electric field distribution in a structure with $\sigma = 0.08$. Here, the expected transmission efficiency should be about 70% from figure 3.10a, but for this specific structure the recorded value is only 0.64%. Of the incident light, 77.6% is reflected back to the input port. It is clear from figure 3.11 that a cavity is created by chance because of the roughness, and the light is reflected back to the input (left side). This demonstrates very well how relatively small fabrication uncertainties might have a severe impact on the performance of a waveguide. Even though light is not guided through this structure the bandgap is still intact, as can be seen from figure 3.11.

Figure 3.12 contains the electric field plot for a structure with $\sigma = 0.19$. The light is now seen to radiate slightly out of the structures, however, most of the hole-structure is seen to be a forbidden region for the light. Imperfections, such as slightly large holes or holes with sharp corners, are seen to cause multiple reflections back and forth in the structure. Eventually, 82.0% of the incident light is reflected back to the input, and only 0.0004% of the light is coupled to the output.

The electric field plot in figure 3.13 is obtained from a structure with $\sigma = 0.30$. The bandgap has now broken down completely in certain directions and light is seen to couple out of the PhC region (middle top of figure 3.13). In addition, some stray cavities within the hole-structure are excited by evanescent coupling.

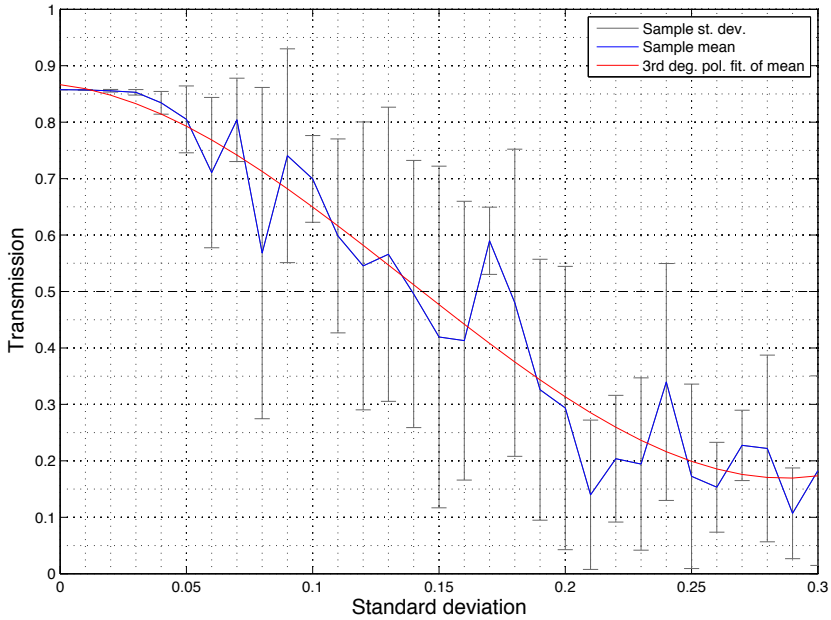

(a) Transmission coefficient

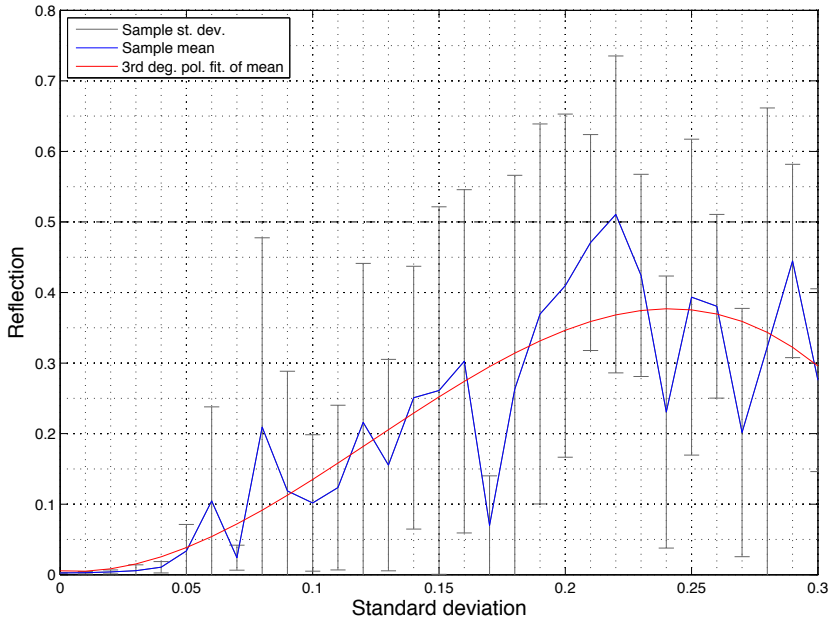

(b) Reflection coefficient

Figure 3.10: Experimentally observed average (a) transmission and (b) reflection coefficients (blue lines) in a W1 PhC waveguide for varying degrees of roughness. The sample standard deviation is depicted as gray bars, and a third degree polynomial fit of the average transmission and reflection is given as the red lines.

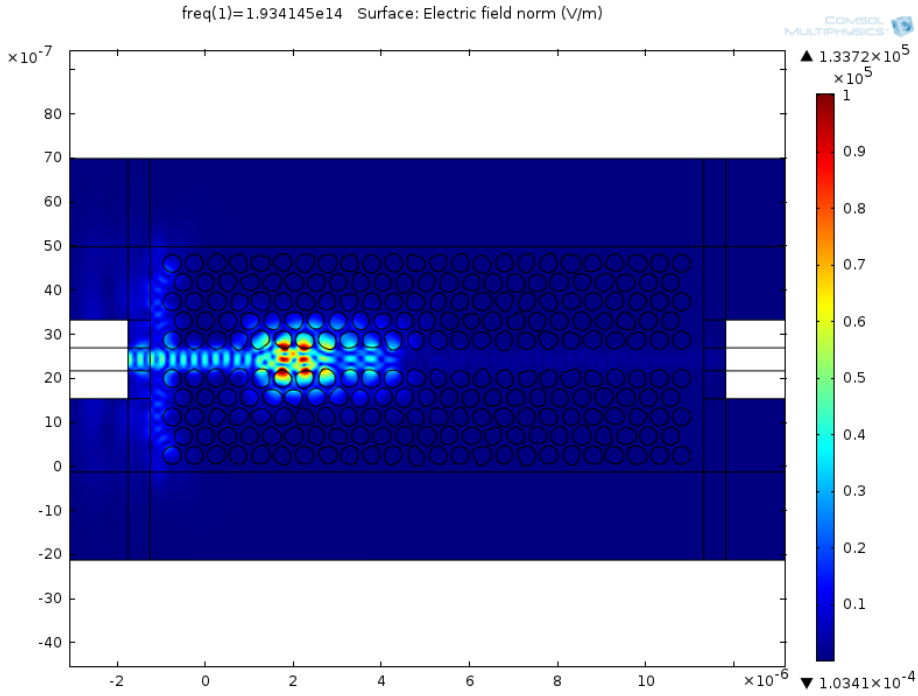

Figure 3.11: Normalized electric field plot for the PhC with $\sigma = 0.08$.

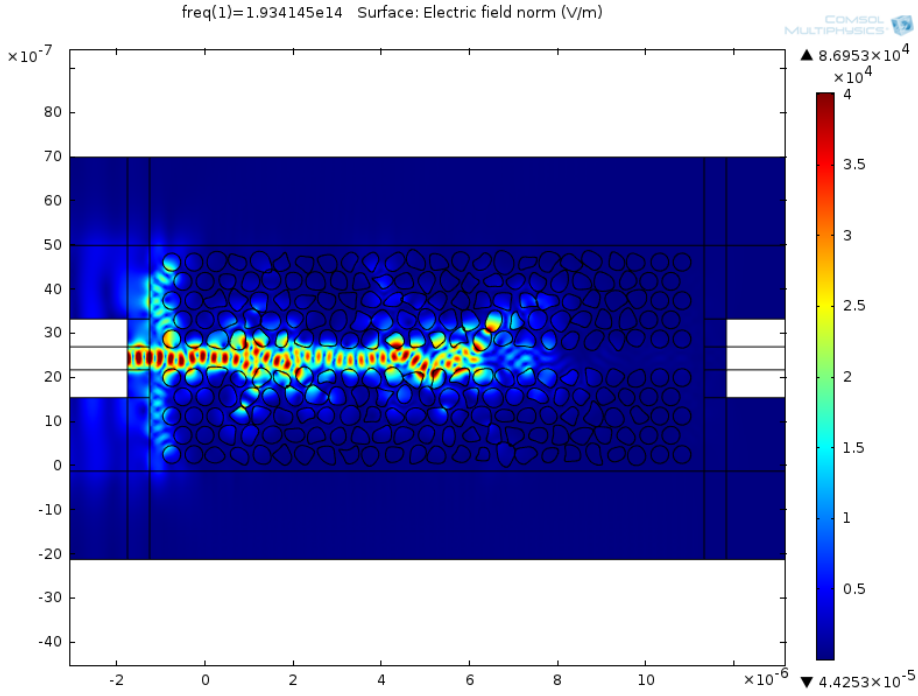

Figure 3.12: Normalized electric field plot for the PhC with $\sigma = 0.19$.

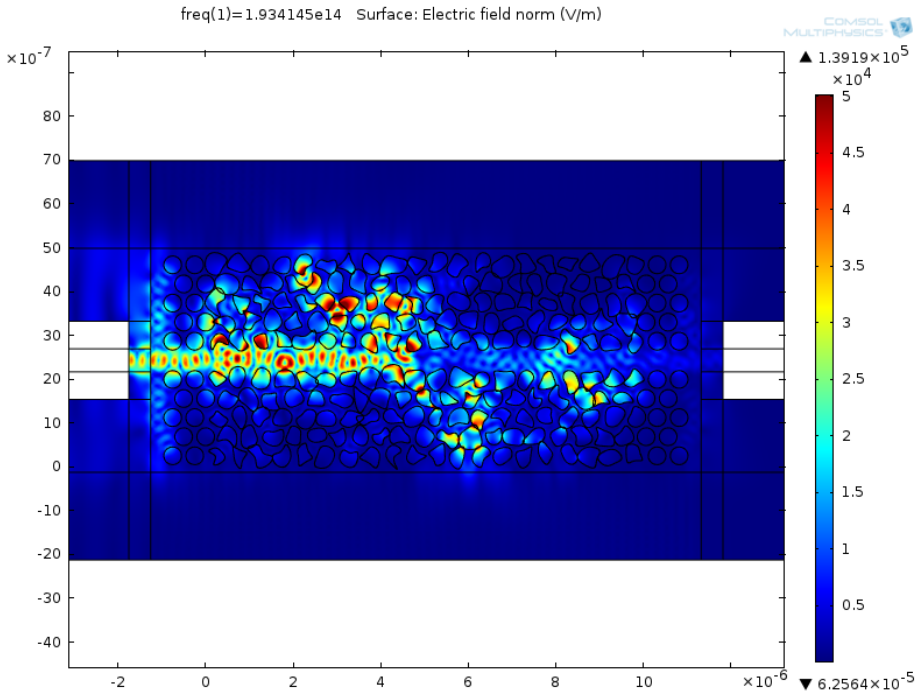

Figure 3.13: Normalized electric field plot for the PhC with $\sigma = 0.30$.

Light is seen to be damped very quickly upon hitting the outer regions of the model, this is because they are defined as perfectly matched layers in COMSOL, to avoid reflections from the model boundary (see appendix B.1).

The model presented here has a W1 waveguide segment which is only 12 µm long, thus limiting the losses of the waveguide. In applications it would probably be of interest to fabricate longer waveguides. Furthermore, the modeled roughness might not be directly comparable to that found in actual fabrication, but the resulting electric field plots are still of interest to discuss non-ideal effects. For a more thorough discussion see [9].

Chapter 4

Fabrication and Characterization

This chapter outlines the fabrication procedures for the work performed in the cleanroom. It also contains background information on the different processes used, in the hope that it will serve as a guide for the starting cleanroom user. It is divided up according to the process flow, outlining necessary background information immediately before each procedure is described.

4.1 Sample Handling

Proper sample handling is crucial in any fabrication process, in order to assure no contamination of the sample, or destruction or degradation of already fabricated structures. When dealing with fabrication of sub-micron structures, the sources of errors in fabrication become much more intricate than in fabrication of large scale structures: Even tiny particles are treated as contaminants, which is of course why this type of work is performed in a cleanroom.

4.1.1 Cleaning

Before most processes are performed on a sample, a thorough cleaning of the sample should be performed. This is done to remove any contaminants which might be a source of error in the subsequent process. The cleaning process used in this work is described here:

First, the sample should be rinsed in acetone, ethanol, isopropyl alcohol (IPA) and deionized water (DI water). Acetone is the strongest solvent, and is responsible for removing dust and particles from the sample. However, if the acetone is not rinsed off afterwards it will dry on the sample and leave particles. Therefore, the sample must be rinsed with gradually weaker solvents, and finally with DI water. It is important to rinse thoroughly with DI water to ensure that no solvents are allowed to dry on the sample. Finally, the sample should be blow-dried with nitrogen gas (N_2).

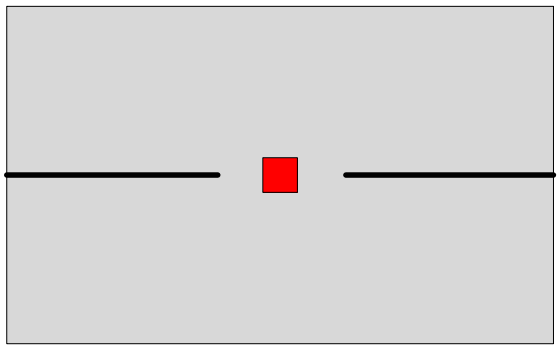

Figure 4.1: Scribe lines in sample before dicing the sample in two. The red area are the structures of interest, which should be cut approximately in the middle.

Between different processes the samples are stored in clean sample boxes, within N_2 purged (inert) cabinets. This ensures minimal contamination and oxidation of the samples.

4.1.2 Wafer and Sample Dicing

The Si used in this work is in the form of 250 µm thick 2 inch $\langle 1\,0\,0 \rangle$ oriented single crystalline Si wafers. The samples used for the experiments are 6 mm × 10 mm dies, which must be cut from the 2 inch wafers. The instrument used for this purpose is a Dynatex DX-III scriber and breaker, which uses a diamond tip to scribe straight lines on the surface of the sample where it should be diced. The scribing weakens the sample, allowing it to be broken along its crystal plane.

The subsequent breaking is done with an impulse bar breaker, which is impacted below the sample. If the alignment of the scribe line and the crystal plane of $\langle 1\,0\,0 \rangle$ Si is good, a very clean break can be achieved. The parameters used for scribing and breaking are given in appendix D.3.

After the final structures are etched into the samples, it is of interest to view the profile of the etched structures. In order to do so, the sample must be cut across the etched structures, exposing the cross section of the etched holes in the PhC. This is achieved by scribing the sample as shown in figure 4.1, allowing the subsequent break to be very clean at the area of interest, *i.e.*, the structures.

4.2 Silicon on Insulator Wafer: Plasma Enhanced Chemical Vapor Deposition

In order to guide light, the 2D PhC waveguide should be fabricated in a silicon on insulator (SOI) wafer, as discussed in chapter 3.1. The top Si layer should ideally be a 220 nm thick layer of single crystalline Si (sc-Si), and the oxide cladding (BOX) should be at least 1 µm. As commercially available SOIs are expensive to buy in small quanta, especially when the layer thicknesses must be strictly controlled, it is most convenient to fabricate the SOI in-house.

The SOI can be fabricated by depositing oxide and Si on a Si wafer, using Plasma enhanced chemical vapor deposition (PECVD). The major drawback of PECVD is its inability to deposit sc-Si; the deposited Si takes the form of hydrogenated amorphous Si (a-Si) [27]. a-Si generally has much higher optical losses than sc-Si, but recently there has been an increased interest in using a-Si for photonic applications, and low-loss devices have been demonstrated [27, 28]. Because of the relatively low deposition temperatures (around 300 °C), PECVD deposited oxide will generally have an increased amount of water impurities compared to oxide deposited at higher temperatures [29]. For this reason, PECVD deposited oxide is expected to etch faster than pure oxide, such as the oxide in high quality SOI wafers.

PECVD is a deposition method which utilizes a pressurized chamber in which the sample is exposed to reactive gases. The gases are ionized to form a plasma, much in the same way as for a reactive ion etching (RIE) system, *cf.* chapter 4.5. The reactive gases are driven towards the substrate by the same mechanism as in RIE, where they combine to form a solid film, while the by-products of the deposition reactions are carried away from the chamber by the exhaust. The resulting film is dependent upon the precursor gases and the conditions of the chamber.

PECVD offers a relatively low deposition rate, which make it possible to control the thicknesses of the resulting films accurately. In addition, the uniformity across the wafer or sample is generally quite good.

The PECVD system in NTNU NanoLab is an Oxford Instruments PlasmaLab System 100-PECVD. The instrument is fully automated: A connecting PC is used to input the parameters of the deposition, such as gas flow, temperature, pressure, RF power and time. The wafer is placed in a loadlock, which pumps down to a vacuum before inserting the wafer into the process chamber. This ensures that the process chamber is not contaminated by the outside environment, and vice versa.

4.3 Electron Beam Lithography

Electron beam lithography (EBL) or e-beam lithography is a direct-write (maskless) lithography technique. As with regular photolithography the sample is covered with a photoresist (PR, resist) which is activated by the incident electrons (or photons in the case of photolithography). For a positive (negative) resist, the exposed (unexposed) areas will be soluble in a developer, yielding the resulting pattern. By exposing the resist with electrons instead of light, the diffraction limit of regular photolithography is overcome, and nanometer-scale patterns can be achieved [30].

4.3.1 Principle of Electron Beam Lithography

In the EBL, the desired pattern is created by scanning the e-beam across the EBL's write-field using deflectors (electromagnetic lenses) [31, Ch. 2.2.4.2], exposing only the regions of interest. As such, it is a much slower technique than, *e.g.*, deep-UV lithography, where a physical mask between the light source and

the sample allows the entire sample to be exposed at the same time. EBL is not well suited for large scale fabrication because of its low throughput, but because of its excellent resolution it is much used in research [18].

An EBL setup consists of a scanning electron microscope (SEM) (see chapter 4.6.1), with a moving stage that can be controlled very accurately. The electron gun of the SEM is responsible for the actual exposure of the resist, and it also allows the user to view the sample before exposure, in order to align the stage with some orientation markers on the sample.

A beam-blanker allows the electron beam to be blanked between exposure areas without switching the beam off. If the blanking is too slow for a certain structure, the beam might expose areas which should not be exposed, possibly resulting in an elongation of the structures in the write direction [32, Ch. 1.2.2].

The electron gun of the EBL has a certain beam current and acceleration voltage, which can be set by the user. A higher acceleration voltage impacts more momentum to the electrons, and a higher beam current produces more electrons in the beam. A parameter called the exposure dose is often used to describe how many electrons are impacted on a certain area of the resist, causing the exposure. The dose is given as beam current times exposure time divided by area, *i.e.*, with units of $\mu\mathrm{A\,s\,cm^{-2}}$. For a given beam current and area, the exposure time must be increased to increase the dose. Thus it is clear that a higher beam current results in a faster exposure.

Throughout this work, the exposure dose is given as a unitless parameter. The unit of $\mu\mathrm{A\,s\,cm^{-2}}$ is implied when discussing exposure doses, and is therefore omitted for convenience.

The exposure is done at a certain magnification, yielding a corresponding write field size. Within this write field, the beam is deflected to create the pattern. The deflection of the beam needs to be calibrated by doing write field alignment, described in section 4.4.2. If the total area to be exposed is larger than the write field, the stage needs to move to subsequent write fields in order to write the entire mask. If the deflection of the beam is not calibrated, stitching errors presents as a misalignment of the exposed features at neighboring write fields. An example of a stitching error is presented in figure 4.2, from which it is obvious that the top write field has shifted downwards and rightwards (by approximately 200 nm) with respect to the bottom write field.

4.3.2 EBL Resist

The EBL resists available in NTNU NanoLab are both based on PMMA (polymethyl methacrylate). They are MicroChem's 950PMMA A2 and 950PMMA A9. The "A" denotes the solvent anisole, which the PMMA is diluted in, and the following number refers to the percentage (by weight) of PMMA in the resist. A higher percentage of solids in the resist will result in a thicker resist layer, given the same coating parameters [33].

The resist is used as a mask for the etching step, either directly for the Si etching or as a mask for some intermediate hard layer, *e.g.*, chromium (Cr). The main problem with using PMMA as an etch mask for Si etching is a generally low selectivity between PMMA and Si, meaning that the PMMA will etch at a rate comparable to that of Si. For this reason it is necessary to have a PMMA layer

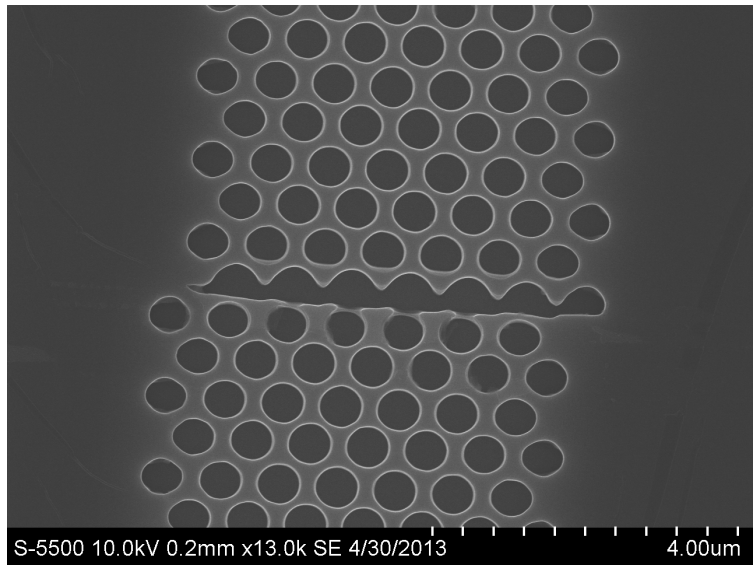

Figure 4.2: Example of stitching error during EBL in an etched PhC structure. Notice how the top half of the structure is shifted southeast with respect to the bottom half of the structure.

which is thick enough to withstand the etching, in order to successfully transfer the correct pattern to Si. If this is not possible, a thin layer of some metal can be deposited on top of Si before the resist is applied. The PMMA will then serve as a mask for etching this hard layer, and subsequent Si etching will utilize the metal as a mask. The downside of using a hard mask is—in addition to extra processing steps—that the etching of the hard layer might introduce roughness which is then inherited by the Si during Si etching.

After the sample has been through EBL and subsequent etching, the resist is no longer needed. Usually it is enough to soak the sample in acetone to remove all resist. If the resist still adheres to the sample surface, it can be placed in a beaker with acetone in an ultrasonic bath for 10 min–15 min.

4.3.3 Electron Scattering and Proximity Effects

Beam Widening

When the electrons impact the resist layer of the sample and travel through the resist, they experience forward scattering which results in a broadening of the electron beam [31, Ch. 2.3.1]. Although the forward scattering occurs at small angles, it might be a source of increased structural sizes and is best avoided by limiting the thickness of the resist [31, Ch. 2.3.1].

In addition to the forward scattering events, secondary electrons are created by the slowing of the primary electrons. The energy of the secondary electrons is much lower than the energy of the incident electrons, but they contribute in part to a widening of the beam. The secondary electrons are the main contributors to the resist exposure [31, Ch. 2.3.3].

Proximity Effects

As the primary electrons impact the substrate, they experience scattering events which might occur at much larger angles. These reflected electrons are known as backscattered electrons. For an acceleration voltage of 20 keV, the range of the backscattered electrons is on the order of 10 µm in Si and PMMA [31, Ch. 2.3.2].

When the backscattered electrons travel through the parts of the resist which are not exposed by the primary or forward scattered electrons, they will cause exposure of this resist. If several structures are placed close together—as indeed is the case for PhCs—the exposure caused by the backscattered electrons from neighboring structures might add up and render the resist soluble, even in the areas not exposed by the primary electrons. For this reason, small structures which are placed far from any other structures require a higher exposure dose than large structures, or structures which are placed in close proximity.

One possible approach to minimize the impact of proximity effects is to vary the exposure dose depending on how densely the PMMA is exposed. In regions where a high density of backscattered electrons are expected (*e.g.* the middle of a PhC), the dose can be decreased a bit with respect to regions where the density of backscattered electrons is less (such as the outer rows of holes in a PhC).

4.4 EBL in Practice

4.4.1 Sample Preparation

Proper sample handling before electron beam lithography is crucial. The sample must be cleaned thoroughly to remove any contaminants and to ensure good adhesion of the resist. This is important to achieve an evenly distributed resist layer.

After rinsing the sample, it should be pre-baked on a hotplate set at 200 °C for at least 20 min. This is done to evaporate any water from the surface to ensure good adhesion between the surface and the PMMA. The step is therefore often called a dehydration bake. If the sample is left on the hotplate for a longer time, less water will be left on the sample, but because of humidity in the air it is not possible to remove all water [34, Ch. 13]. Immediately after removing the sample from the hotplate, the spin-on process should commence.

Spin-On

The purpose of the spin-on step is to deposit a uniform layer of resist on the sample. The sample is placed on a vacuum chuck connected to rotating stage, and is sprayed with N_2 immediately before PMMA is deposited, to remove any dust which might have landed on the sample after cleaning. A drop of PMMA is deposited on the sample, using a pipette, and the spin-on is started. The spin-on recipe (taken from MicroChem's data sheet [33]) is given in table 4.1. The first step of the spin-on consists of five seconds of slow rotation to increase the uniformity. The next step is the actual spinning step, where the spin speed (x in table 4.1) decides the thickness of the resulting resist layer, within some thickness interval depending on the viscosity of the resist (given by the percentage of solids

Step	Time [s]	Speed [RPM]	Acceleration [RPM/s]
1	5	500	1000
2	45	x	1000

Table 4.1: Parameters for spin-on, collected from [33]. Note that the time for step 2 might be increased if a high final speed is needed, to compensate for the acceleration time.

in the resist). Normally, the spin speed is between 2000 and 4000 revolutions per minute (RPM). The acceleration of each step decides how quickly the speed is increased from the last step, and $1000\,\mathrm{RPM\,s^{-1}}$ is commonly used in NTNU NanoLab. If a high final spin speed is needed, the time of the final step should be increased by a few seconds due to the increased time necessary to accelerate to the final speed.

Soft Bake and Inspection

After spin-on, the resist is placed on a hotplate set at 180 °C for one minute. This is done to evaporate any solvents in the resist and to improve the adhesion between the sample surface and the resist. The timing of this step is not one of the most critical parameters, but for consistency it is kept constant at one minute.[1]

Before the exposure is commenced, the applied PMMA layer should be inspected in an optical microscope. It is important to check that the PMMA is evenly distributed in the area of exposure, and that no contaminants are present. Possible faults that might be detected during inspection are bubbles in the resist, uneven resist, or particles in the resist.

The resist will almost always be thicker near the edges of the sample. This is normal and hard to combat, but it is not a problem as long as the sample is relatively large in comparison with the exposure area. Once an area on the sample where the resist is evenly distributed and free of contaminants and bubbles is located, the sample should be scratched to define this area. Figure 4.3 is an example of such a scratch, defining the optimal exposure area as 1 mm to the right of the scratch. The reason for scratching the sample in this way is to have an orientation mark during EBL, as explained next.

4.4.2 Exposure

The EBL setup used in this work consist of a Hitachi S-4300SE SEM and a Raith EBL stage using laser interferometry. A computer with Raith Elphy Plus software provides origin and angle correction of the sample with respect to the stage (see later), write field alignment, and controls the exposure as defined by the mask file. A second computer is connected to the SEM and allows the operator to view the sample in order to calibrate the electron beam.

Once the sample is covered with resist and the optimal exposure area is defined by the scratch, the sample can be loaded onto the EBL stage. As for most other

[1] MicroChem suggests a soft bake at 180 °C for 60 s–90 s [33].

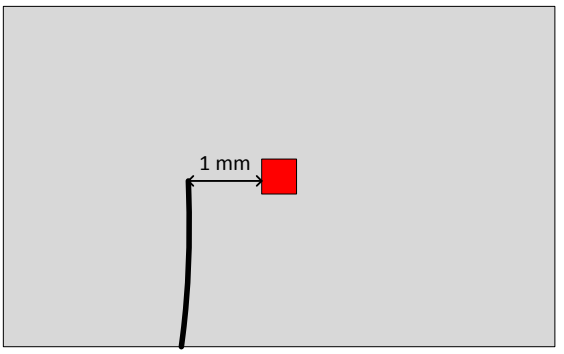

Figure 4.3: Scratch in sample defining the optimal exposure area (red) as 1 mm to the right of the endpoint of the scratch.

instruments, the sample is loaded through a loadlock to ensure no contamination of the EBL chamber. Once the sample is loaded, the acceleration voltage of the electron gun can be set and the voltage source is turned on. An acceleration voltage of 20 keV is commonly used in NTNU NanoLab, and serves as a good starting point for the following investigations.

Beam Setup

Before any exposure can be done, a series of calibrations must be performed. First of all the focus, stigma and aperture of the electron source must be calibrated, in order to get a decent electron beam. Next, the beam current must be measured and set to the correct value. The beam current will affect the exposure time, as a lower current will require a longer exposure time to achieve the correct dose, *cf.* chapter 4.3.1. A beam current of around −48.5 pA is often used in NanoLab. While calibrating the EBL it is very important to remember that the beam activates the resist. Thus it is important to keep away from the area which is to be exposed, in order not to cause any spurious exposure which might interfere with the intended patterns in the PMMA.

The coordinates of the sample are given in a local coordinate system, u-v, which needs to be related to the coordinates of the stage, x-y, *cf.* figure 4.4. The important parameters to adjust are the angle of the sample and the origin on the sample, both with respect to the x-y coordinate system. The angle correction is done by identifying the corners of the sample in the EBL, to calculate the angle difference between the global x-axis and the local u-axis. Once the orientation scratch in the sample is identified using the SEM, the x-y coordinates at the end of the scratch are set to be the origin of the u-v coordinate system.

In order to fine-tune the beam, contamination dots are burnt close to the exposure area. This is done by exposing the beam at only one spot for some time. If the beam is fine-tuned the dot will be round and sharp, but if the beam is unfocused or has some stigma the dot might be blurry and elliptical, and it will take a long time to burn it. By consecutively burning contamination dots and readjusting the beam using these dots, a well adjusted beam can be achieved. The routine should be repeated until a sharp dot can be created in less than 5 s. An

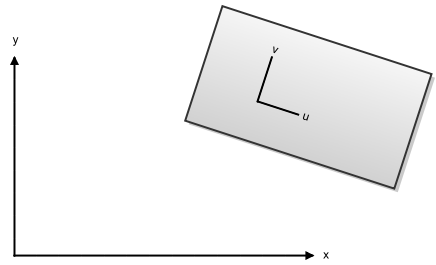

Figure 4.4: Example of misalignment between the global coordinate system of the EBL stage (x-y) and the local coordinate system of the sample (u-v). The sample is depicted by the gray rectangle. Before the sample can be exposed, the u-v coordinate system needs to be related to the x-y coordinate system by doing origin- and angle correction. The angle difference between the two coordinate systems is exaggerated for illustrative purposes.

example of contamination dots is depicted in figure 4.5, showing the incremental improvements in beam shape and focus. Notice also in this figure the highlighted rectangles in the resist: They are caused by the electron exposure from scanning the image during focusing on each dot.

Write Field Alignment

Once the beam is optimized, a write field alignment must be performed, to avoid stitching errors. The write field (WF) size is set to $50\,\mu m \times 50\,\mu m$ which corresponds to a magnification of 1000 on the EBL. The WF alignment is done by locating a sharply defined particle at this magnification and running a predefined WF alignment routine. During this routine, the stage moves by a small amount in certain directions and the beam is deflected to scan the area with the particle. At each position the particle should ideally be in the center of the image, if the WF alignment is correct. The user can correct the misalignment of the particle by specifying how far away from the center the particle is at each position. By doing subsequent WF alignments with smaller and smaller stage movements, the deflection of the beam can be very accurately aligned and stitching errors are circumvented.

Mask

The mask is given by a GDSII file, in which the areas to be exposed are specified with respect to the origin of the u-v coordinate system. Upon exposure, the EBL translates the u-v position list into exposure points in the x-y plane. The beam is then scanned across the write field, only exposing the given points.

The test mask that will be utilized to find the appropriate EBL and etching parameters should consist of holes with $r/a = 0.4$, where $a = 503.26\,nm$, in a trigonal array. As the sample should be scribed to expose the cross section of the structures, it is necessary to have structures that extend far enough to allow easy

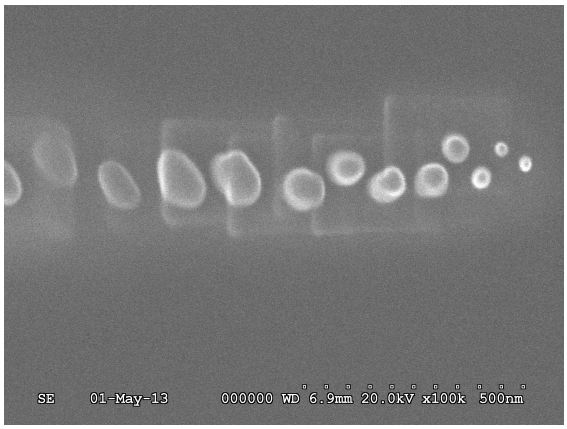

Figure 4.5: Contamination dots in PMMA during beam optimization for EBL. The dots are burnt consecutively from left to right, optimizing the beam at each dot until a very well defined dot could be achieved. The leftmost dot was burnt by a point exposure for 2 min, while the rightmost dot was burnt in only 5 s.

scribing. The array of holes will be approximately 5 µm × 100 µm. The mask is presented in figure 4.6, including a close-up of the hole structures. The highlighted green disks are the regions which will be exposed. Proximity corrections can easily be incorporated into the mask by designating the outer holes as separate layers in the GDSII file, which allows the user to specify different doses for the two (or more) layers during EBL, should it be needed.

4.4.3 Developing the Resist and Inspecting the Pattern

After the resist is exposed, the exposed areas should be dissolved in a developer. The developer of choice for this work is a mixture of nine parts IPA and one part DI water, as this is readily available in the cleanroom and found to provide good resolution. The developer is mixed in a 50 ml graduated cylinder, using pipettes when necessary. The developer is then poured into a beaker in which the sample is developed.

The timing of the development is crucial, and 30 s is found to be an appropriate time. The sample is held by tweezers and immersed in the developer while a timer is started. After 30 s, the sample is removed from the developer and soaked in DI water to stop the development process. The rinse-time is non-critical and a DI water rinse of about 10 s–20 s, with gentle agitation, is found to be enough. The sample is then blown dry with N_2, gently to avoid disturbing the patterns in the resist.

Conventional optical lithography processes usually include a postbake step to harden the resist. This step is optional for PMMA, and it will not be included in this process as the structures in the PMMA might round at temperatures around 125 °C [33].

The resulting patterns should be inspected with an optical microscope to ensure that the resist is dissolved. It is not possible to view the details of the actual structures in an optical microscope, as the resolution is too low. The

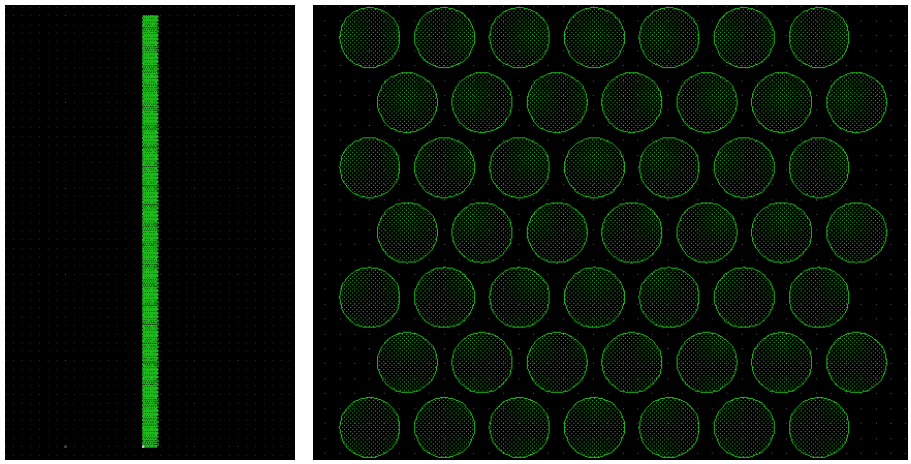

(a) **(b)**

Figure 4.6: (a) The test mask used to find the best EBL and etch parameters, and (b) close up of the hole structure of the mask. The green disks are the positions provided by the GDSII file to the EBL, defining the areas to be exposed.

resist patterns can be viewed in a SEM, however, the resist will be activated by the SEMs electron beam. Therefore, inspection in a SEM is—at least to some degree—destructive. Furthermore, the obtained images can not be trusted to be completely accurate, as some alteration of the resist will have occurred before the images can even by collected.

The best way of inspecting a resist pattern is therefore to look at the resulting structures from a known etch process. However, as both EBL and etching are to be studied in this work, the resulting structures must be discussed as a result of both EBL and etching processes.

4.4.4 Summary of EBL Process

As an overview, and for later reference, the EBL process is summarized in table 4.2.

Step	Name	Purpose
1	Clean	Remove contaminants
2	Dehydrate	Evaporate any water to enhance adhesion
3	Spin on	Apply resist evenly to the sample
4	Soft bake	Improve adhesion and evaporate solvents
5	Resist inspect	Check to see if the resist is evenly distributed
6	Exposure	Define the patterns in the resist
7	Develop	Dissolve the exposed resist

Table 4.2: Summary of EBL process flow.

4.5 Plasma Etching

4.5.1 Plasma Generation

A plasma is a neutral gas which has been energized such that some electrons in the gas are removed from their host atoms, effectively ionizing the gas. Neutral atoms and molecules, positive ions and free electrons are the constituents of a plasma. The generation of plasma naturally requires energy input, and this is often delivered by some external electromagnetic field which accelerates electrons in the neutral gas. As the electrons are accelerated away from the much heavier atoms and molecules, they leave behind positively charged ions thus creating the plasma. Collisions between the accelerated electrons and the relatively slow atoms and ions cause further ionization of the gas.

In a plasma etching system, the energy input is most often delivered by coupling a radio frequency (RF) field oscillating at 13.56 MHz to the gas in the chamber.[2] This gives the electrons a high mobility and the gas is ionized very effectively [35]. The coupling of the field is divided into two categories, as discussed next.

4.5.2 Reactive Ion Etching

Reactive ion etching (RIE) is a method of etching in which the process gas of the low pressure process chamber is imparted with energy to create reactive ions. The positively charged ions are then accelerated towards the sample to be etched, where they can remove the target material either by physical or chemical etching.

Capacitively Coupled Plasma

In a capacitively coupled plasma (CCP), the RF-source is coupled to an anode plate (often referred to as the table) through a series capacitance, while the cathode plate is grounded (see figure 4.7).

The series capacitance causes a buildup of negative charges on the table, resulting in a potential difference between the table and the charge neutral plasma, referred to as the self-bias, V_b. Positive ions are accelerated from the plasma to the table by the self-bias, thus allowing the ions to physically or chemically etch the sample.

Inductively Coupled Plasma

The inductively coupled plasma (ICP) is achieved by coupling the RF-source to a coil wrapped around the plasma region of the chamber. The induced electric field due to the inductive coupling results in a circulation of the plasma parallel to the table [35, p. 5].

Collisions between the accelerated electrons and the relatively slow ions ionize the gas, as with the CCP process. However, the ICP process do not cause ions to be driven towards the table. The difference between CCP and ICP is thus revealed: ICP controls (to a large degree) the density of ions in the plasma, while the CCP controls the incident energy of the ions delivered to the table. The

[2] Other frequencies are sometimes used, but 13.56 MHz is standard in NTNU NanoLab.

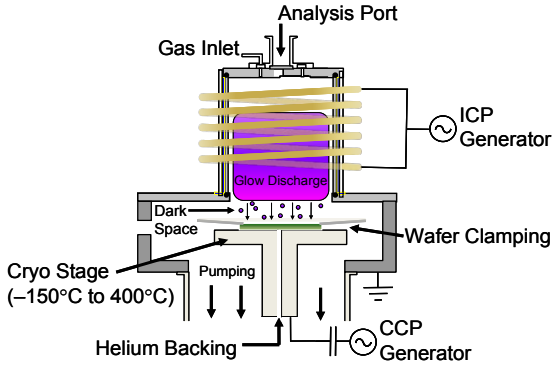

Figure 4.7: Cross-sectional view of ICP-RIE chamber. Figure reprinted with permission from [35].

ICP-RIE setup thus allows the two parameters, ion density and incident energy, to be decoupled. This provides a very flexible etching tool which can be adjusted to allow anisotropic etching of different substrates. The ICP-RIE chamber, with both CCP and ICP sources, is depicted in figure 4.7.

4.5.3 General Guidelines for ICP-RIE

The CCP power controls the strength of the electric field between the electrode and the table, and thus the energy of the ions in the chamber. An increase in the CCP power will give more energetic and thus faster ions. As a general guideline this will result in a more physical etch, which can provide very anisotropic etching, but at the expense of worse selectivity between mask material and target material. The physical etching is a result of ion bombardment which mechanically removes material from the sample, and the etch rate for the physical etching is thus determined by the hardness of the given materials.

The ICP power allows the ion density to be increased beyond what is possible with the CCP source alone, allowing the ion density to be controlled without much change in the incident energy. By increasing the ICP power, the number of reactive ions reaching the substrate will increase, which will increase the degree of chemical etching [35].

The chemistry of the etching recipe is indeed the most important parameter. In NTNU NanoLab the most used processes are based on either fluorine (F) or chlorine (Cl) gases, but hydrogen bromide (HBr) is also an available etchant. Parameters such as etch rate, selectivity and uniformity are highly dependent upon the etch chemistry and target materials. In addition to highly reactive gases used for removal of material, passivation gases might be used in a process. The passivation gas is chosen such that it deposits on the sample, effectively lowering the etch rate. If the mixture between passivation gas and etching gas is set correctly, the result is a deposition of a passivation layer on the sidewalls of the structures which is equal to the removal of said layer by the reactive species. This would result in an effective zero etch rate on the sidewalls, creating highly anisotropic etch profiles without any undercut.

In order for the plasma to ignite, the chamber pressure must be set correctly according to the gas chemistry used. However, it is possible to change the pressure once the plasma is stable. A reduction in the chamber pressure increases the electric field perpendicular to the table, so that the sidewalls become more vertical [35].

Generally, an increase in the temperature of the sample will make the chemical reactions at the surface more reactive. This results in an increased chemical etch rate, which creates a more isotropic etch. If cryogenic temperatures are used in conjunction with passivation gases, it is possible to achieve a very good sidewall passivation layer which becomes volatile—and is thus removed from the sample—once the temperature is increased [35].

4.5.4 ICP-RIE in Practice

The ICP-RIE system used in this work is an Oxford Instruments PlasmaLab System 100-ICP180. It consists of two independent 13.56 MHz RF sources connected to an ICP coil and CCP plates, as depicted in figure 4.7.

Before the sample is etched, the chamber should be preconditioned. This is achieved by running the intended recipe for at least 20 min on a dummy Si wafer. If the chamber is not preconditioned the plasma might become very unstable, or it might not ignite at all. By preconditioning the chamber before each etching run, the etching conditions are similar for all samples, providing comparable results.

The RF sources are connected to the coil (ICP) and the plates (CCP) through matching networks. These consist of two variable capacitors each, which are tuned to match the impedance of the plasma. They are controlled automatically, but if the plasma becomes very unstable the control unit might not be able to automatically match the impedance. If the impedance is not matched to the plasma, power will be reflected back to the source. This has two consequences: 1) The power delivered to the plasma will be less than intended, and 2) the process might automatically be aborted if the reflected power becomes too high. This is most easily avoided by preconditioning the chamber.

The sample to be etched is attached to a 2 inch Si wafer (referred to as the carrier) with Fomblin oil, to ensure a good thermal contact to the carrier. This is important to provide the correct temperature on the sample, and thus the right etching conditions. The timing of the recipe is crucial to achieve a desired etch depth. If the etch rate on the target material is known, the necessary time for the target depth is easy to calculate. However, due to fluctuations in the process conditions during the first few seconds of any recipe, the etch rate will not be exactly constant between runs. In order to minimize this impact, all recipes should be of at least one minute duration. For a target etch depth of 220 nm this would mean that the optimal recipe should have an etch rate less than $220\,\mathrm{nm\,min^{-1}}$.

The actual etching is an automated process: The parameters of the etch recipe are input in the system, and the carrier wafer with the attached sample is transfered from the load lock to the process chamber. As the gas chemistry for the etch might be highly toxic, the chamber and the load lock are pumped down to a high vacuum and purged with nitrogen (N_2) several times before and after the etching process.

End-Point Detection

The ICP-RIE is equipped with an end-point detection system, consisting of a 675 nm laser which can be used to measure the reflectance of the sample surface. Upon etching Si covered by PMMA, the reflectance will vary depending on the thickness of the PMMA layer. The wavelength of the laser light in the PMMA is $\lambda = \lambda_0/n_{\text{PMMA}}$, where $n_{\text{PMMA}} = 1.487$ is the refractive index of PMMA at $\lambda_0 = 675\,\text{nm}$. If the pathlength difference between the directly reflected light from the PMMA surface and the light reflected at the PMMA/Si interface is an integer multiple of λ, the reflectance is maximized by constructive interference. The pathlength difference is $2d$, where d is the thickness of the layer. It is clear that the difference in layer thickness between two reflectance maxima is $\lambda/2$, which is 227 nm for PMMA at $\lambda_0 = 675\,\text{nm}$.

The reflectance plot is a sine wave which is plotted as a function of etch time. The distance between to consecutive maxima in the plot is thus the time it takes to etch $\lambda/2$, which allows real-time monitoring of the etch rate. However, for the example above it is only possible to monitor the etch rate on PMMA, not Si. If the selectivity between Si and PMMA is known, this can however be used to approximately monitor the etch depth in Si.

After the etch recipe has finished, the recipe should be characterized. The important parameters to find are selectivity against mask material, etch rate and resulting profile.

4.6 Characterization

4.6.1 Scanning Electron Microscope

The scanning electron microscope (SEM) is already explained in part in chapter 4.3. The SEM at NTNU NanoLab is a Hitachi S-5500 S(T)EM with secondary electron and backscattered electron detectors, in addition to bright-/dark field transmission detectors. The secondary electron detector is the most commonly used detector, and as the name suggests it detects the electrons which are excited and emitted from the material by the primary electrons. The secondary electrons gives the highest contrast in surface geometry, whereas the backscattered electrons provides a high contrast between light and heavy materials. In addition to being capable of high resolution top-view images of the samples, the SEM has a cross section holder by which the samples can be viewed from the side. This is most useful in order to inspect how the profiles of etched structures appear.

The procedure for using the SEM is quite simple. The sample is attached to the sample holder by conductive tape if the normal (flat) holder is used. The cross-sectional holder has a slot, in which the sample can be fastened by screws. The slot is about 7 mm wide and 3 mm high, and the samples must thus be scribed to the correct size to fit in the holder. Once the sample is loaded into the SEM, the beam must be adjusted in the same way as for the EBL. The normal acceleration voltage and beam current for the SEM is 10 keV and 7 µA, respectively. For higher acceleration voltages, the beam will penetrate deeper into the material resulting in secondary emission originating deeper within the material. This might produce SEM images in which the surface looks smoother

than it actually is, but this is easy to check by investigating the sample using different acceleration voltages.

4.6.2 Reflectometer

The reflectometer is used to measure the thickness of partially transparent films, such as PMMA and oxide. It has a light source which emit light of several wavelengths, and the reflectance as a function of wavelength is recorded. The mechanism is much the same as for the end-point detection of ICP-RIE, but the multi-wavelength operation allows more accurate thicknesses to be found, if a suitable initial value is given.

As the reflectance is recorded, the software tries to match the reflectance spectrum to calculated spectra for the given material. The results are collected in a best-fit thickness and a corresponding goodness of fit parameter, denoted g. The goodness of fit is a number between 0 and 1, where 1 is a perfect match between calculated and measured reflectance spectra. Generally, g should be very close to 1, ideally within a few percent, for the recorded thickness to be trusted.

4.6.3 Profilometer

The profilometer at NTNU NanoLab is a Veeco Dektak 150, with a diamond contact stylus. The profilometer scans its stylus in a line across the sample surface, recording the relative height difference at each location. It is very useful for recording steps in the height profile on a sample, *e.g.*, the height difference between a partially covering layer and the substrate.

One possible application of the profilometer is to find the difference in etch rate between PMMA and Si. The etch rate on PMMA can of course be found using the reflectometer, but to find the etch rate on Si it is often necessary to inspect the sample in SEM from a cross section view. However, if the initial and final thicknesses of a PMMA layer are known, the profilometer can be used to measure the step between PMMA and Si, thus yielding a good estimate of the Si etch depth. Another possibility is to measure the step on the Si carrier after etching, between the exposed parts and the region in which the sample covered the carrier.

Chapter 5

Results & Discussions

This chapter outlines the experimental work performed in the cleanroom. The Results and the Discussions chapters are combined, as the work was iterative. This chapter is written in such a way as to allow the flow of information to be continuous. The bulk of the results are presented and discussed in sections 5.4 and 5.5, which surrounds the two etch chemistries. However, EBL parameters are an important part of the discussion also here.

5.1 Silicon on Insulator Wafer

The first activity in the cleanroom was the fabrication of a silicon on insulator (SOI) wafer. The PECVD recipe used and the measured layer thicknesses are given in appendix D.2. SEM pictures of the resulting wafer are shown in figure 5.1, from which it is clear that there is a lot of surface roughness in the amorphous silicon (a-Si) top layer.

The observed roughness in the a-Si layer would result in excessive scattering of light at the surfaces, which would deteriorate the transmission coefficient of the waveguide. In addition, a-Si has in general a larger absorption coefficient than single crystal Si (c-Si). For these reasons, the prospects of fabricating the SOI in-house were abandoned. There are several manufacturers of SOIs with excellent optical properties, and it was decided that the eventual fabrication of PhC structures should be done on such wafers.

Initial fabrication experiments were judged by the appearance of the structures, *i.e.*, the profile of the holes and the surface roughness, rather than on optical measurements. The fabrication experiments were therefore performed on c-Si wafers without any additional layers. The resulting profiles are expected to be very similar to those obtained by etching the c-Si top layer of prime-grade SOIs.

Upon leaving the idea of in-house fabrication of the SOI wafer, inquiries were sent to various SOI manufacturers to find a possible supplier of SOI wafers. Due to a combination of time, price and availability, it was not possible to obtain SOI wafers during the course of this work. However, the etching results achieved on $\langle 1\,0\,0 \rangle$ single crystalline Si wafers should be directly transferable to any high quality $\langle 1\,0\,0 \rangle$ SOI wafer.

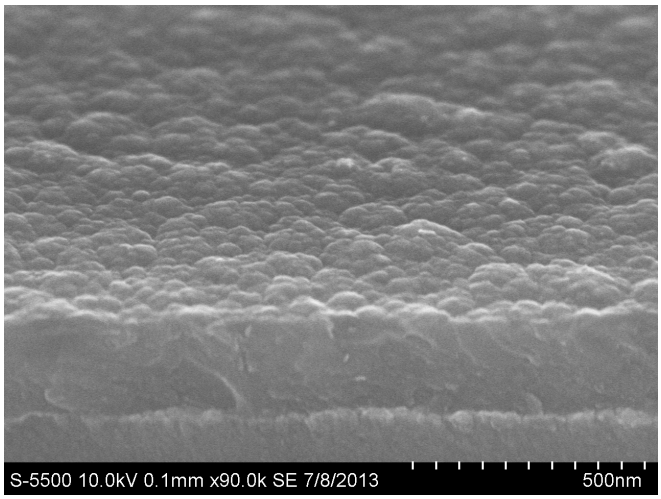

Figure 5.1: SEM micrograph of SOI structure fabricated by PECVD. It is obvious that the a-Si layer has a lot of surface roughness, which would cause severe scattering if it were to be used in photonic devices. The border between the top a-Si and the BOX is clearly visible.

5.2 Choice of Mask Material

During EBL exposure, the desired pattern is imprinted in the PMMA resist. In order to transfer this pattern to the underlying Si layer two approaches are possible. The first and easiest approach is to simply use the PMMA layer as the mask during etching. This puts a constraint on the PMMA layer in the form of a minimum thickness, as discussed in chapter 4.3.2. The second approach is to have some intermediate hard layer between the resist and Si layers. As the addition of a hard transfer layer highly complicates the processing, it was decided that initial investigations should use PMMA as the mask material.

During previous work [8], the 950PMMA A2 was used as the resist. The resulting thickness was around 100 nm, which necessitated a Cr-etch mask. To simplify the process, a thicker resist layer was targeted. By mixing the two EBL resists available at NTNU NanoLab, 950PMMA A2 and 950PMMA A9, in a 1:1 ratio, a home-made 950PMMA A5.5 was made. According to MicroChem's data sheet [33], this percentage of anisole should provide thicknesses of around 400 nm–500 nm, at a spin speed of 2000 RPM.

The thickness as a function of spin speed was found by coating several test samples and measuring the resulting PMMA thickness. The resist was applied with the spin parameters given in table 4.1, with slightly longer step times for the higher spin speeds. In order to test the influence of the final step time, two samples were coated and spun at 3000 RPM for 45 s and 60 s. The resulting thicknesses were almost identical at 403.5 nm and 401.8 nm, respectively, indicating that the timing of the final step is non-critical—at least within some time interval. The thickness dependence on spin speed is given in figure 5.2.

It was important to have a PMMA layer thick enough such that it would not

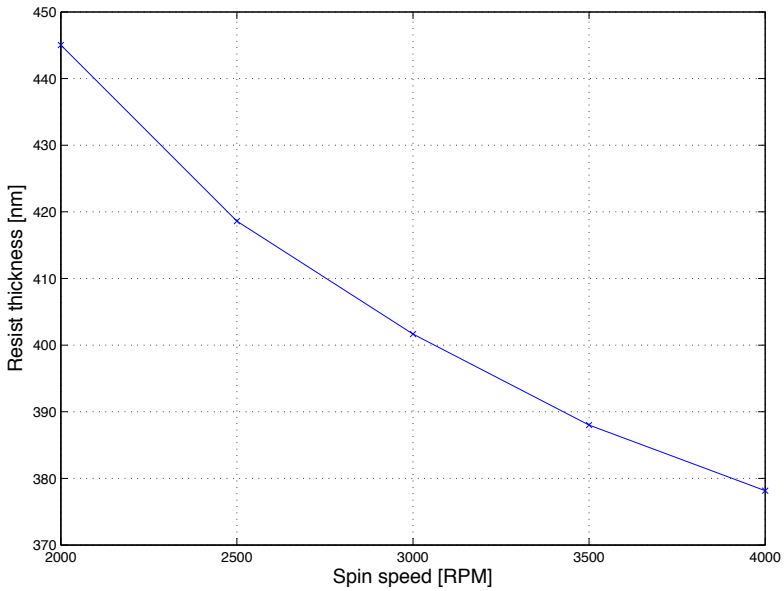

Figure 5.2: Thickness of 950PMMA A5.5 (in units of nm) as a function of final spin speed. The samples were coated using the recipe of table 4.1, with slightly longer spin times for the highest speeds. It is clear that this resist can be applied with a thickness of approximately 380 nm–445 nm.

Step	Time [s]	Speed [RPM]	Acceleration [RPM s^{-1}]
1	5	500	1000
2	45	2000	1000

Table 5.1: Spin-on recipe used to coat a 10 mm × 6 mm sample with 445 nm of 950PMMA A5.5. The first step spreads the drop of resist across the sample and the second step spins off excess resist, yielding the desired thickness.

etch away. The spin speed of 2000 RPM was chosen, as it gave a thickness of around 445 nm. The spin-on parameters necessary to achieve the 445 nm PMMA layer are repeated in table 5.1, for later reference. It is clear from figure 5.2 that PMMA thicknesses below 380 nm can be achieved with the 950PMMA A5.5. If thinner resist layers are needed, it is straightforward to mix a more diluted resist by increasing the amount of 950PMMA A2 in the mix.

5.3 Electron Beam Lithography

It is complicated to judge the parameters of EBL isolated from the etching parameters. Ultimately, the final structures etched into Si will be a function of resist type, exposure parameters and etching recipe. Of course all of these should be optimized, but it is possible—to some degree—to compensate, *e.g.*, a slight underexposure with increased lateral etching. For this reason, EBL parameters will

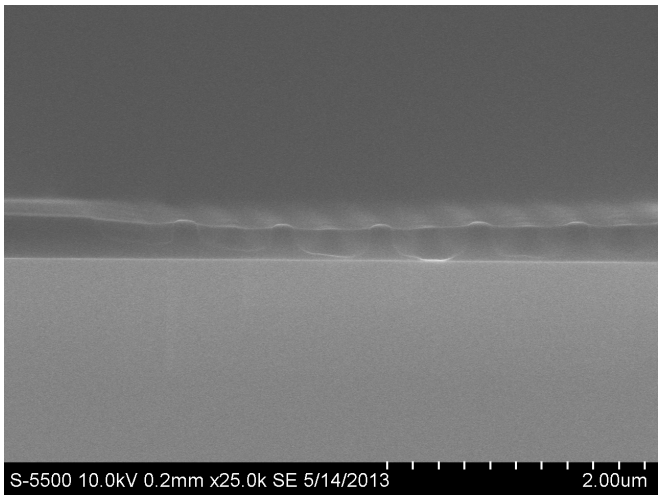

Figure 5.3: SEM micrograph of developed hole structures in PMMA after an EBL exposure dose of 100. The focus was not fine-tuned before capturing the image, due to quick alteration of the resist. However, it is seen that some resist remains at the bottom of the holes, indicating an insufficient exposure dose.

primarily be discussed and judged in combination with different etching recipes. However, some SEM micrographs were obtained of the developed resist patterns, primarily to inspect their profiles.

It should be noted that upon exposing the PMMA to the electrons in the SEM, it will quickly alter its shape as the resist is sensitive to electrons. For this reason, the micrographs of the developed resist patterns needed to be collected quickly, to ensure as little alteration of the patterns as possible. This inevitably means that there was no time to fine-tune the focus before each picture was taken. During SEM inspection the resist seemed to float out, thus rounding all sharp corners of the pattern. Because of this, the reader should note that the following figures are only an indication of the actual resist profile.

Figure 5.3 depicts the hole structure of figure 4.6 in the developed PMMA, after an exposure with dose 100. The profiles of the holes are not very straight, and it looks like the resist is not completely removed close to the surface of the Si wafer. This could be because of insufficient development time, but more likely the dose is just too low. The figure suggests that the resist is more developed in the center of the structure than near the edges, indicative of proximity effects.

The same pattern exposed with dose 120 is depicted in figure 5.4. From this figure it is clear that the increased dose resulted in complete removal of the resist in the holes. The profiles are very straight, and this seems like an appropriate exposure dose for the given structure and resist thickness. Some debris is seen in the left part of the image above the resist. This is believed to be caused by the scribing and breaking process, which naturally involves some particles. After most scribing processes it is normal to clean the sample, but this was not possible here as solvent cleaning would destroy or remove the resist.

In figure 5.5 the SEM micrograph of exposure dose 140 is presented. It is clear

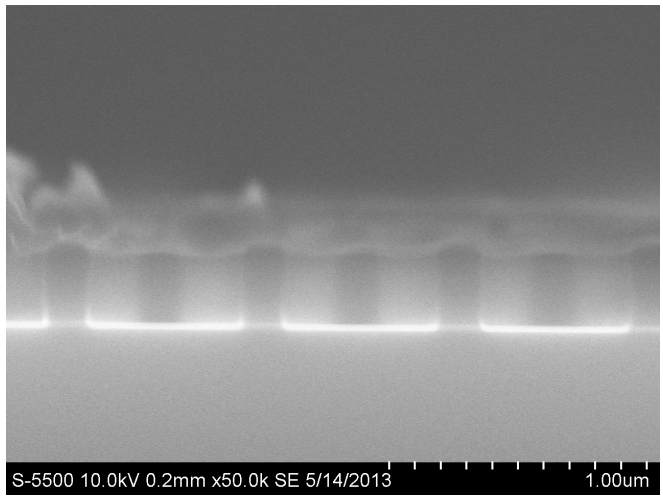

Figure 5.4: SEM micrograph of developed hole structures in PMMA after an EBL exposure dose of 120. The sidewalls are seen to be very straight, and the resist appears to be completely removed in the holes. The debris in the top left part of the figure is most likely particles deposited during scribing and breaking.

that this dose led to overexposing the smallest structures. The smallest features, *i.e.*, the lines between adjacent holes, are seen to be almost gone. This is probably due to proximity effects. The height of the resist is seen to be much less in the exposed area, but this can be partly caused by the electrons from the SEM. If indeed the resist thickness for the smallest features is too low, the etching step might remove all of the resist and start to etch the previously covered Si. It does seem, however, that the outermost holes are well defined with this exposure dose. The collapsed resist in figure 5.5 is thought to be caused by either the breaking step (after scribing), or by the electron exposure during SEM.

To sum up, dose 120 is clearly the best choice from this experiment, but it might be necessary to increase the exposure dose to, *e.g.*, 140 for the outermost rows of holes. However, the micrographs obtained of the PMMA should not be trusted blindly as the electrons from the SEM inevitably affect the PMMA.

5.4 Fluorine Based Etching

The first etching recipe that was investigated was based on sulfur hexafluoride (SF_6) as the reactive gas, and fluoroform (CHF_3) as a passivation gas—in order to combat the problem of lateral etching. The etching recipe was inspired by a standard Oxford Instruments recipe, and is given in table 5.2.

First, the recipe was run on two test samples (3.11 and 3.12) to find the etch rate on Si and the selectivity against PMMA. The samples were prepared by spin-on of 950PMMA A5.5 with the parameters given in table 5.1. The samples were then dipped in acetone to create a clear step between bare Si and Si covered by PMMA. Sample 3.12 was etched for one minute, after which all of the resist had etched away. Thus it was not possible to find the etch rate based on the step

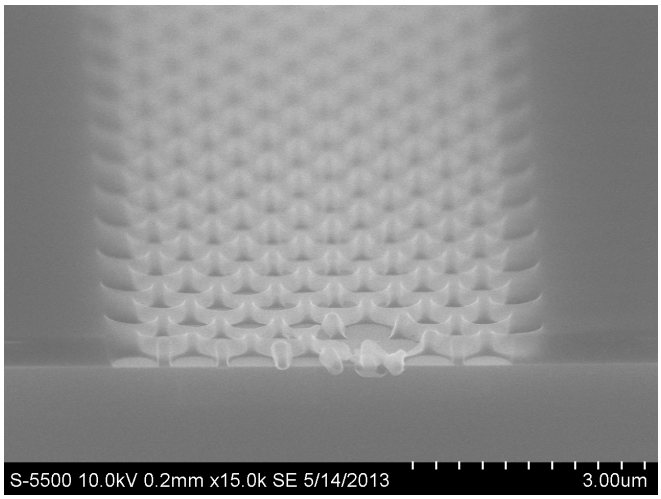

Figure 5.5: SEM micrograph of developed hole structures in PMMA after an EBL exposure dose of 140. Proximity effects are seen to have caused an overexposure of the smallest feature in the most densely exposed areas.

SF_6 flow rate	7.5 sccm
CHF_3 flow rate	50 sccm
Pressure	15 mTorr
ICP power	1200 W
CCP power	40 W
Temperature	20 °C
He backside	10 Torr

Table 5.2: Initial SF_6/CHF_3 recipe inspired by an Oxford Instruments recipe, but with adjusted flow rates.

between resist and Si. However, the step in the Si carrier between the etched Si and the Si covered by the sample was measured in the profilometer. The measurement is given in figure 5.6, from which it is clear that approximately 810 nm of Si was etched, yielding an etch rate of $810\,\mathrm{nm\,min^{-1}}$.

Sample 3.11 was etched with the same recipe, but the time was set to 20 s. The PMMA thicknesses before and after the etching were measured in the reflectometer. They were found to be 456.6 nm and 223.9 nm, respectively, resulting in an etch depth in PMMA of 232.7 nm. The step between the PMMA and the Si was measured in the profilometer, but it was clear that the step in the carrier (as above) gave a more accurate result. The profilometer measurement of the step in the carrier of sample 3.11 is presented in figure 5.7, and the etch depth in Si is now seen to be 280 nm.

For sample 3.11 the etch rate for Si and PMMA is then found to be approximately $840\,\mathrm{nm\,min^{-1}}$ and $698.1\,\mathrm{nm\,min^{-1}}$, respectively, yielding a selectivity of 1.20:1 between Si and PMMA. The results of the etching of samples 3.11 and

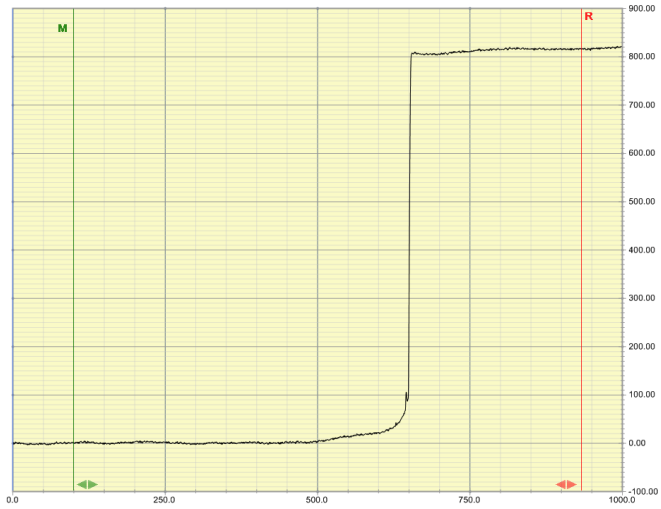

Figure 5.6: Profilometer measurement of the step between etched (left side) and not etched (right side) Si for sample 3.12, obtained by measuring the step in the Si carrier. The x-axis is the scanning distance along the wafer, given in units of µm, and the y-axis is the relative height difference, given in units of nm.

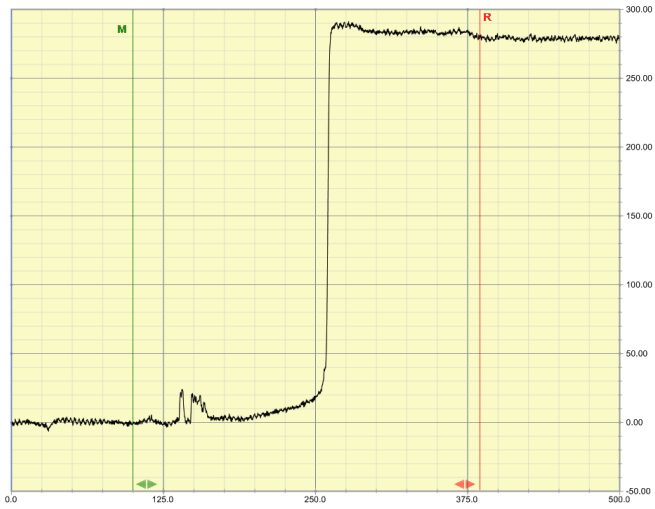

Figure 5.7: Profilometer measurement of the step between etched (left side) and not etched (right side) Si for sample 3.11, obtained by measuring the step in the Si carrier. The x-axis is the scanning distance along the wafer, given in units of µm, and the y-axis is the relative height difference, given in units of nm.

Sample	**3.11**	**3.12**
Recipe	Table 5.2	Table 5.2
Etch time	20 s	1 min
CCP reflected power	1 W	1 W
ICP reflected power	8 W	8 W
Plasma stabilization time	3 s	2 s
Si etch depth	280 nm	810 nm
Si etch rate	840 nm min^{-1}	810 nm min^{-1}
PMMA etch depth	232.7 nm	Entire layer
PMMA etch rate	698.1 nm min^{-1}	Unknown
Selectivity	1.20:1	Unknown

Table 5.3: Parameters collected from the etch tests of the initial SF_6 recipe.

Sample	**3.6**	**3.7**
Initial thickness	445.9 nm	427 nm
g	0.9881	0.8917
Final thickness	276.5 nm	281.4 nm
g	0.9919	0.8301
Etched PMMA	169.4 nm	145.6 nm

Table 5.4: Thickness of PMMA measured in the center of samples 3.6 and 3.7 before and after etching. g denotes the goodness of fit for the reflectometer measurements, which is explained in chapter 4.6.2.

3.12 are summarized in table 5.3.

It is clear that the given recipe etches very fast. From the etch test it is found that a total etch time of about 15 s is enough to etch the required 220 nm of Si. This poses a problem as such short etching times might result in difficulty of controlling the etch depth, *cf.* chapter 4.5.4. However, the RF powers in this recipe are quite high, and by decreasing them the etch rate should also decrease.

5.4.1 Etching the First Samples

In order to inspect the etch profile of the recipe of table 5.2, two samples were etched with this recipe before any adjustments were done. The samples were patterned by EBL, using the mask defined in figure 4.6, with varying exposure doses. The etch time was set to 15 s, which should give a resulting etch depth of around 220 nm. The PMMA thicknesses before and after etching were measured in the reflectometer and they are presented in table 5.4.

The measurements for sample 3.7 show a somewhat poor goodness of fit, g (*cf.* chapter 4.6.2), whereas the measurements for sample 3.6 are very accurate. It was not possible to improve the measurements, and the reason for this is unclear. The initial thickness of sample 3.7 is seen to be almost 20 nm less than the average PMMA thickness presented in figure 5.2. The bad g might indicate that some particles or bubbles are present in the resist, changing the properties of the layer

Sample	**3.6**	**3.7**
Recipe	Table 5.2	Table 5.2
Etch time	15 s	15 s
CCP reflected power	1 W	1 W
ICP reflected power	9 W	9 W
Plasma stabilization time	2 s–3 s	2 s
PMMA etch depth	169.4 nm	145.6 nm
Si etch depth	200 nm ± 2 nm	170 nm ± 5 nm
Si etch rate	800 nm min^{-1}	680 nm min^{-1}
Selectivity	1.18	1.17

Table 5.5: Observed etch parameters for samples 3.6 and 3.7. Notice how the etch rates are very different for the two samples, but the selectivities are almost identical. The difference in etch rates are thought to be due to instabilities during the first few seconds of the etches, yielding different effective etch times for the two samples.

such that the reflectometer did not record the expected reflections.

The observed etching parameters along with the resulting etch rates and selectivity, are given in table 5.5. The etch rate is seen to be higher for sample 3.6 than for sample 3.7, both on Si and on PMMA. The etch rate for Si is found to be about 13% less for sample 3.7 than for sample 3.6. Some difference is expected as the total etching time is only 15 s, but the observed difference is quite substantial. During the first few seconds of any recipe the plasma will be unstable, and the etch rate will be different than for the stabilized plasma. As the selectivity of the two etch runs are about the same, the discrepancy in etch rate is taken primarily as a strong indication that longer etch times—and thus lower etch rates—should be utilized.

The observed low selectivity between Si and PMMA might be an issue for two reasons: 1) Once the etching recipe is optimized and the structures are investigated more carefully, it might become a problem to have the relatively thick PMMA layer of 445 nm. If the PMMA thickness is reduced, the selectivity might be too low to allow the PMMA to be used as an etch mask. 2) It might be of interest at a later time to etch into the BOX layer of the SOI, to improve the performance of the PhC. If this is to be done there must be enough PMMA left to work as a mask for the oxide etch.

5.4.2 Cross-Sectional SEM View of the First Samples

Sample 3.6

In sample 3.6 the holes were exposed with doses 120, 140, 160 and 180. Dose 120 was barely visible and is not depicted here. Figure 5.8 depicts the etched structures for dose 140. In figure 5.8a, the shapes of the etched structures are seen to suffer from proximity effects, *i.e.*, the structures near the edge of the exposed area are not as well defined as the centered structures. Figure 5.8b is a close-up cross-sectional view of the holes in the middle of the exposed area. From this figure, the etch depth can be found to be about 200 nm ± 2 nm, in close

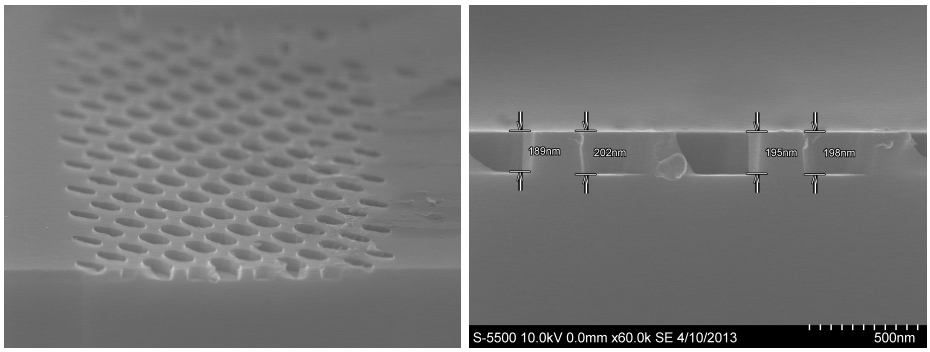

(a) Angled view **(b)** Cross section view

Figure 5.8: SEM micrographs of etched holes in sample 3.6 for exposure dose 140. Notice the apparent difference in etch depth and radius in the middle of the exposed area and near the edge.

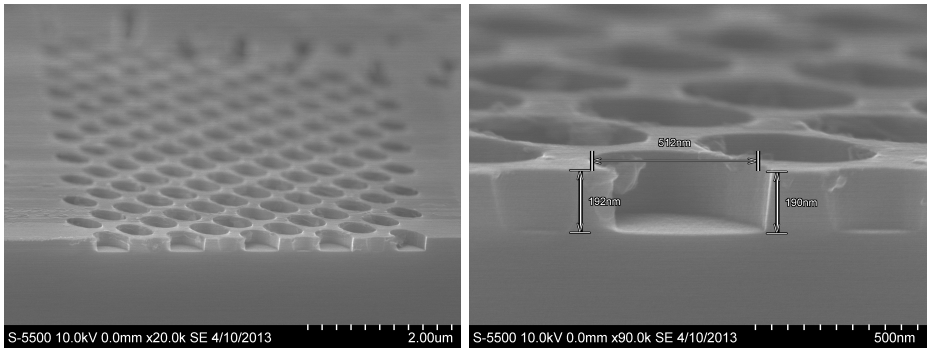

(a) Angled view **(b)** Close-up of the center holes

Figure 5.9: SEM micrographs of etched holes in sample 3.6 for exposure dose 160.

agreement to the previously found etch rate of $810\,\mathrm{nm\,min^{-1}}$–$840\,\mathrm{nm\,min^{-1}}$.

Figure 5.9 contains micrographs of the etched structures for dose 160. In figure 5.9a, the holes seem to be very well resolved. Except for slightly smaller holes near the edge of the exposure area, the holes look uniform with a sharp contrast between the exposed and unexposed areas. Based on this micrograph alone, 160 seems to be an appropriate exposure dose for all holes except the outer row(s). Figure 5.9b is a close-up of the centered hole in figure 5.9a. From this it is seen that the profile of the holes is slightly positive (as indicated by figure 5.8 as well). This might indicate an excess of sidewall passivation leading to decreased chemical etching in the lateral direction, or a slightly positive profile of the developed resist. The measurements shown in figure 5.9b are only used as an indication: As the sample is tilted, the actual depth is more. From a non-tilted micrograph, the depth is found to be the same as for dose 140.

The etched holes for dose 180 are presented in figure 5.10. From figure 5.10a, the smallest structures, *i.e.*, the regions between adjacent holes, are seen

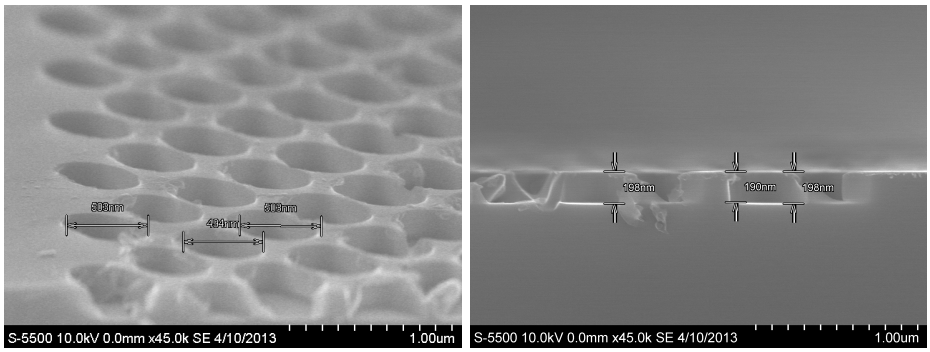

(a) Angled view **(b)** Cross section view

Figure 5.10: SEM micrographs of etched holes in sample 3.6 for exposure dose 180.

to be slightly over-etched. This is indicative of an overexposure, making the non-exposed areas soluble by an excessive amount of backscattered electrons. However, it is clear that the problem is most pronounced in the densely exposed regions; dose 180 might be appropriate for the outermost holes. Figure 5.10b is a cross-sectional view from which the previously found etch depth of $200\,\text{nm} \pm 2\,\text{nm}$ is found to be in good agreement. Please note how the measurements in figure 5.10b are slightly off (at least the 190 nm arrow). This is because the SEM picture shifts slightly over time due to drift effects.

Sample 3.7

In sample 3.7, the holes were exposed with doses 120, 140, 160 and 180. The resulting structures after etching for dose 120 are depicted in figure 5.11. From figure 5.11a it is clear that dose 120 was sufficient to properly define the structures of this sample. Indeed it seems like 120 is close to the optimal dose for this sample. This is in sharp contrast to sample 3.6, for which it was found that dose 160 was the best choice. If the reflectometer measurements for the PMMA thickness are correct, there was approximately 20 nm less PMMA on sample 3.7. This might be a possible explanation for the discrepancy in optimal dose, as a thicker resist layer will necessitate a higher exposure dose. However, the difference in optimal dose between samples 3.6 and 3.7 is found to be 25%, while the difference in resist thickness was only 4%. It is unclear at this point why the discrepancy is this pronounced. This will be further discussed in section 5.4.3.

It is seen from both figures 5.11a and 5.11b that there are some particles on the sample. It also seems from figure 5.11b that some residual PMMA is left between the holes, indicating that the cleaning procedure has not been thorough enough. The particles might have been deposited on the sample during SEM preparation, but the residual PMMA is left because of incomplete removal during acetone soaking. It should normally be enough to soak the sample in acetone to remove all PMMA, but this was not the case for the sample at hand.

Looking at figure 5.11 it is clear that the profile of the holes is again slightly positive. The etch depth is found to be about 170 nm–180 nm, but it is slightly difficult to assert from figure 5.11b as there seems to be some thin resist layer

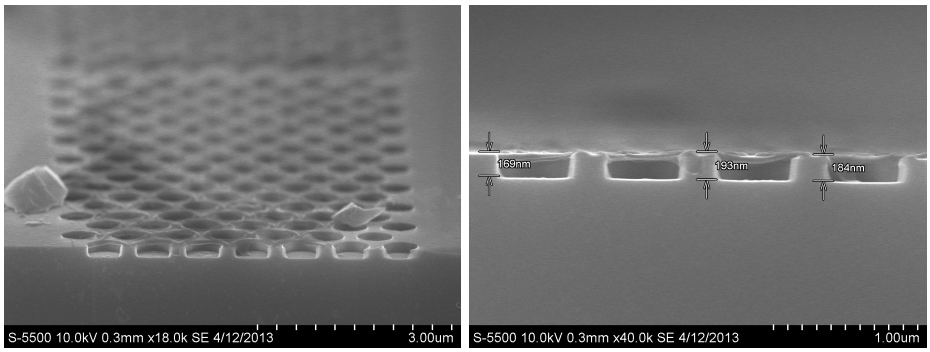

(a) Angled view **(b)** Cross section view

Figure 5.11: SEM micrographs of etched holes in sample 3.7 for exposure dose 120.

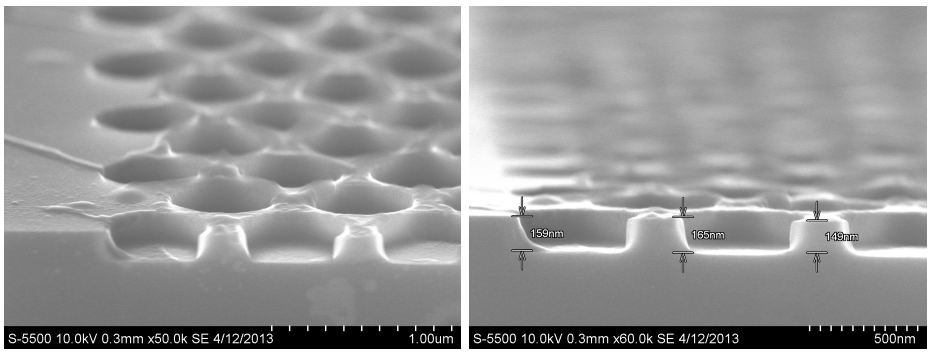

(a) **(b)**

Figure 5.12: SEM micrographs of etched holes in sample 3.7 for exposure dose 140.

left.

The holes of sample 3.7 exposed with dose 140 are depicted in figures 5.12a and 5.12b. From the figures it is clear that some resist remains on the sample. The holes appear to be well defined also for this exposure dose, but the profiles are still slightly positive. This is likely not caused by the EBL as it appears for all exposure doses. The outermost holes are seen to be a bit underexposed.

Figure 5.13 depicts the etched holes with exposure dose 160. It is clear that this dose is too high for the most densely exposed regions due to proximity effects. However, from figure 5.13b it can be seen that the outer row of holes is very well resolved, indicating that this is the correct dose at the periphery of the exposure area.

In figure 5.14, the etched holes for exposure dose 180 are depicted. In figure 5.14a the structures in the most densely exposed areas are seen to be completely etched away, indicating that the exposure dose is much too high. The outermost rows of holes are seen to be well resolved, however, but dose 160 was clearly sufficient even for the outer holes. From figure 5.14b, the etch depth is seen to be 165 nm–170 nm. Comparing figures 5.11b, 5.12b and 5.14b, an etch

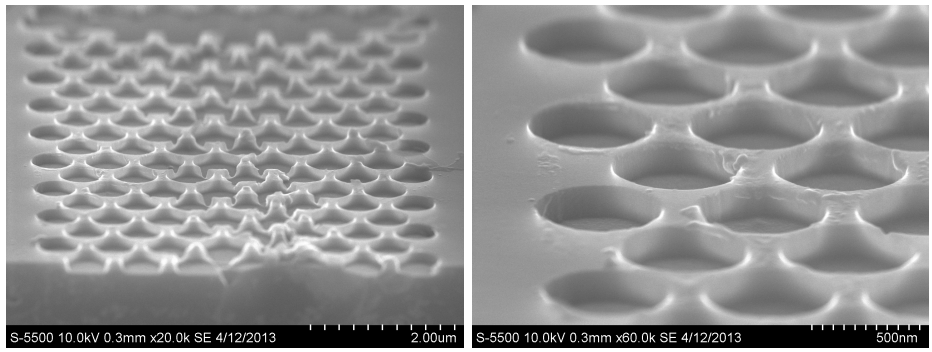

(a) **(b)**

Figure 5.13: SEM micrographs of etched holes in sample 3.7 for exposure dose 160.

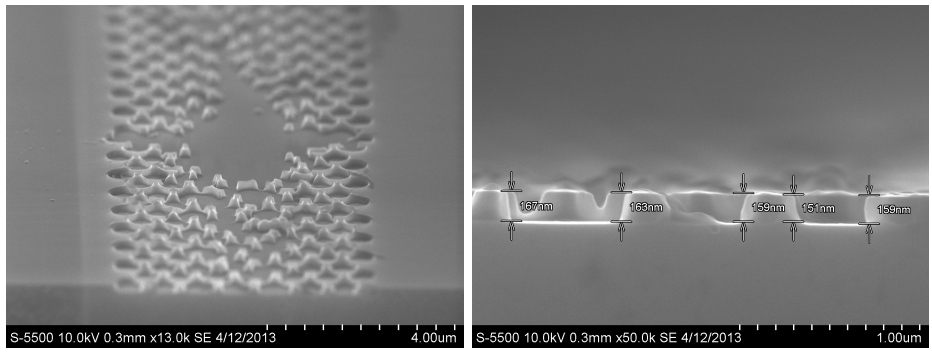

(a) **(b)**

Figure 5.14: SEM micrographs of etched holes in sample 3.7 for exposure dose 180.

depth of $170\,\text{nm} \pm 5\,\text{nm}$ is found from the experiment, yielding an etch rate of $680\,\text{nm}\,\text{min}^{-1} \pm 20\,\text{nm}\,\text{min}^{-1}$.

Comparing the Etching of Samples 3.6 and 3.7

Comparing the etch parameters of table 5.5, the etch rates for the two samples are seen to be substantially different. As discussed previously, this can be attributed to the short total etch time by considering the instability during the first few seconds of the etch. As the selectivity is found to be very similar, the actual etching conditions are probably equal. A difference in effective etch time between the two samples is likely the cause of the difference in etch depth. From this it should be concluded that lower etch rates are needed, in order to minimize the impact of the instability of the etch during the first few seconds.

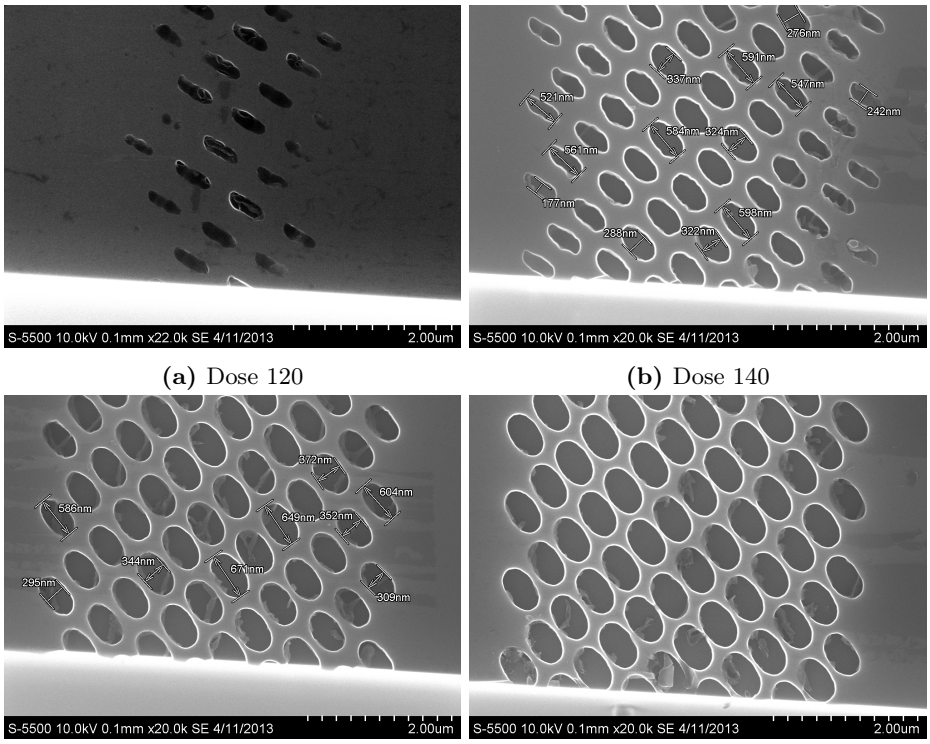

(a) Dose 120 **(b)** Dose 140

(c) Dose 160 **(d)** Dose 180

Figure 5.15: Top view SEM micrographs of etched holes in sample 3.6 with exposure doses 120, 140, 160 and 180.

5.4.3 SEM Top View of the First Samples

Sample 3.6

Sample 3.6 was inspected from the top in the SEM as well, to investigate the shape of the holes in the plane. Figures 5.15a, 5.15b, 5.15c and 5.15d depict the etched holes for doses 120, 140, 160 and 180, respectively. Immediately upon inspection of the figures it is clear that the exposure is not correct: The holes do not form the trigonal lattice pattern as given by the mask (figure 4.6). Rather, the entire mask seems to have been stretched along one direction and compressed along another direction, resulting in elliptical holes. As the holes were defined by perfect circles in the mask software the fault does not lie there, rather this must be an error during EBL. It almost seems like the u-v coordinate system of the EBL is not perpendicular: The local coordinate system in which the exposures are defined seems to have shifted its basis vectors such that the mask is transformed into the observed structures. This is likely due to improper write field alignment.

Putting aside the faults of the exposure for a while, it is still interesting to discuss the results at hand. Attention is called to figures 5.15a and 5.15b: Note how the circumferences of the holes have a substantial amount of roughness, *i.e.*, they are not smooth curves. It is quite obvious that both doses 120 and 140 were

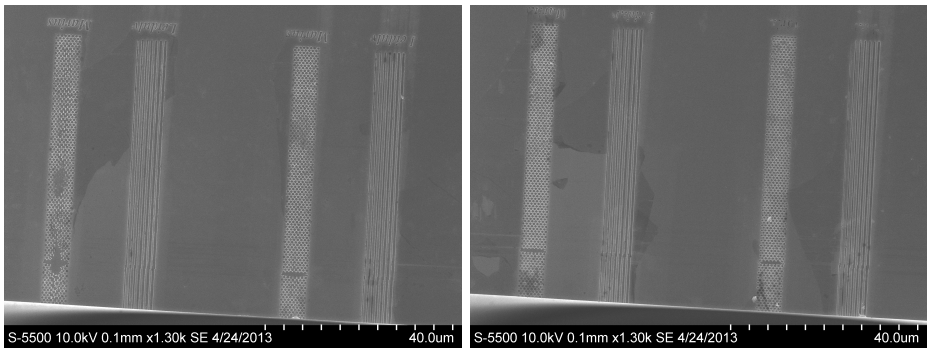

(a) **(b)**

Figure 5.16: Overview of exposed structures in sample 3.7.1, with exposure dose increasing from right to left. Notice the dark areas on the sample. There are both holes (h) and trenches (t) in the sample, and the exposure doses are **(a)**: 180 (h), 220 (t), 160 (h), 200 (t) and **(b)**: 140 (h), 180 (t), 120 (h), and 160 (t).

not sufficient to properly define the structures. In figure 5.15a, only the most centered holes are etched slightly. In figure 5.15b, it is clear that the outer holes are very underexposed, but the most centered holes seem to be properly defined as supported by figure 5.8.

Figure 5.15c depicts the top view of the holes exposed with dose 160. This micrograph supports the previous conclusion that 160 was the optimal dose for sample 3.6: The holes are seen to be very smooth and well defined apart from the outermost rows, which are slightly smaller in size and have some observable degree of roughness.

The holes exposed with dose 180 are presented in figure 5.15d, and it is clear that the outermost holes are now well defined. However, the most densely exposed areas are overexposed, as supported by the angled SEM view of figure 5.10a.

Sample 3.7

Figure 5.16 is an overview of the exposed structures in sample 3.7. Some strange residue is seen to remain on the sample in the form of the dark areas on Si. Between etching and SEM inspection this sample accidentally stuck to the top of the sample box. This probably smeared the PMMA, resulting in better adhesion to the surface and incomplete removal. Note that this did not affect the structures in the Si as this happened after etching.

The etched holes in sample 3.7 are presented in top view SEM micrographs in figure 5.17. It is clear from figure 5.17a that 120 is a too low dose near the edges, but in the most densely exposed areas 120 is high enough, as previously found in figure 5.11. Overall, dose 140 (figure 5.17b) is the optimal dose, however, it is still a bit low at the periphery of the exposure area. The best combination would be dose 120 in the very center of the exposure area, dose 140 close to the edges, and dose 160 (or slightly lower) near the corners of the exposure area. Dose 180 is depicted in figure 5.17d. This is a much too high dose, as the smallest features are now completely etched away in the middle of the structure.

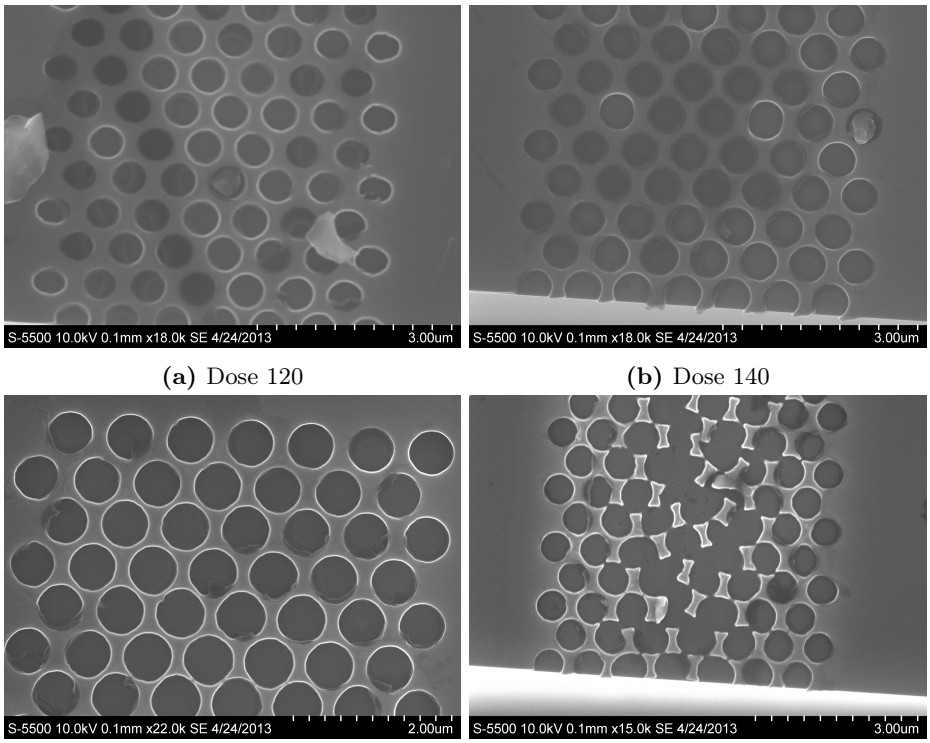

(a) Dose 120 **(b)** Dose 140

(c) Dose 160 **(d)** Dose 180

Figure 5.17: Top view SEM micrographs of etched holes in sample 3.7, with exposure doses 120, 140, 160 and 180.

Some particles can be seen in figures 5.17a and 5.17b, as previously seen in the cross-sectional micrographs. The sample could have been cleaned and new micrographs obtained, but as the particles did not hinder the evaluation of the structures in the sample this was not done.

In the last section, it was stated that dose 120 seemed like the best dose for sample 3.7, based on the cross-sectional micrographs. After this latest discussion, the best dose is found to be 140, although the dose should be varied over the exposure area for an optimal result. Thus the difference between the best dose for samples 3.6 and 3.7 is seen to be only 12.5%. This difference can be explained by considering 1) the apparent difference in resist thickness (4%), and 2) the fact that sample 3.6 was exposed with narrower structures because of the error in EBL adjustment. As discussed in chapter 4.3.3, narrower structures in the resist will require a higher dose to become soluble.

5.4.4 Adjusting the Recipe

After the first round of etching, a plan was made for the next etch tests. The slightly positive profile of the etched structures is probably due to an excess of sidewall passivation. This can be overcome by increasing the SF_6/CHF_3 ratio to

SF_6 flow rate	15 sccm
CHF_3 flow rate	50 sccm
Pressure	15 mTorr
ICP power	800 W
CCP power	20 W
Temperature	20 °C
He backside	10 Torr

Table 5.6: Adjusted SF_6 recipe, based on the results of the first etch tests. The decreased RF powers will lower the etch rate, and the increased SF_6 flow will hopefully combat excessive passivation, yielding vertical profiles.

reduce the percentage of passivation gas. In addition the ratio of ICP to CCP power should be increased, to increase the chemical etch rate with respect to the physical etch rate.

The second problem with the initial SF_6 recipe is the high etch rate. The etch rate is dependent upon the amount of reactive species in the chamber—as controlled by the ICP power—and by the ion bombardment—as controlled by the CCP power. A reduction of both RF powers should thus decrease the etch rate. An adjusted SF_6 recipe based on the discussion above is given in table 5.6.

5.5 Hydrogen Bromide Based Etching

After the first few samples were etched using the previous SF_6/CHF_3 recipe, the ICP-RIE for fluorine broke down. Fortunately there exist two almost identical ICP machines in the cleanroom, one for fluorine based recipes and one for chlorine based recipes. However, the ICP for chlorine did not have CHF_3 supplied, so it was not possible to follow the previous plan of improving the recipe.

As a result, a hydrogen bromide (HBr) based recipe was investigated. One main drawback of HBr etching of Si is a generally worse selectivity against PMMA, but for a PMMA thickness of 445 nm it should not be a problem to etch 220 nm Si. A possible advantage of HBr etching is a relatively good selectivity towards SiO_2 [36]. Upon etching of an actual SOI, it might be beneficial to use the buried oxide (BOX) as a stop layer for the Si etch, resulting in easier control over etch depth.

5.5.1 The Starting Point

The HBr based Si etch given in table 5.7 is taken from Oxford Instruments' catalog of recipes. It has a quoted low etch rate and relatively high selectivity against photoresist (PR), in addition to a vertical profile, making it ideal for PhC etching.

The recipe of table 5.7 quotes an etch rate of $62\,\mathrm{nm\,min^{-1}}$, resulting in an expected etch time of

$$t_{\mathrm{etch}} = \frac{220\,\mathrm{nm}}{62\,\mathrm{nm\,min^{-1}}} = 3.55\,\mathrm{min}. \tag{5.1}$$

HBr flow rate	50 sccm
Pressure	7 mTorr
ICP power	600 W
CCP power	50 W
Temperature	50 °C
He backside	10 Torr
Si etch rate (quoted)	62 nm min^{-1}
Selectivity Si:PR (quoted)	3:1

Table 5.7: HBr based Si etch recipe from Oxford Instruments. In addition to the quoted low etch rate and decent selectivity, the recipe should provide vertical profiles.

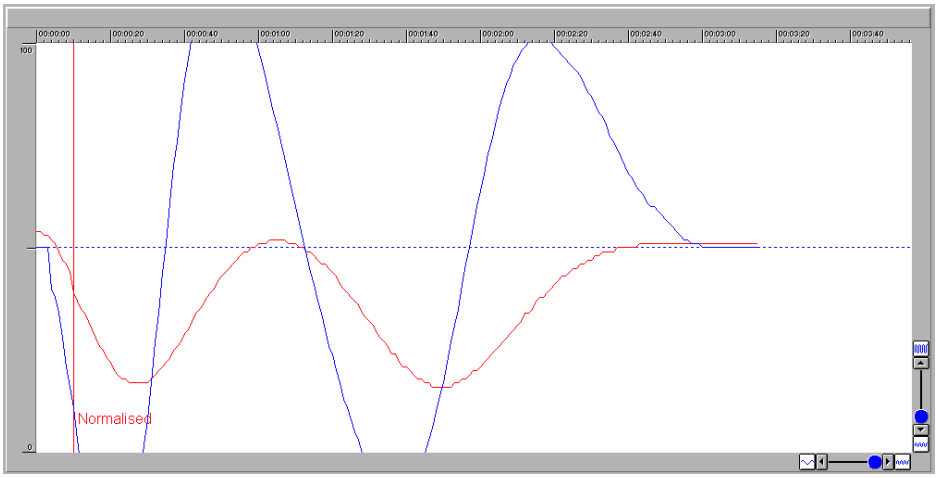

Figure 5.18: Reflectance plot from the end-point detection during etching of sample 3.8. It is clear that the PMMA layer is completely etched away after 3 min.

The recipe was run on sample 3.8 while the thickness of the PMMA layer was monitored using end-point detection, as described in chapter 4.5.4. After three minutes and twenty seconds of etching it became clear that all of the resist was etched away in the middle of the sample, and the recipe was then terminated.

The reflectance plot from the end-point detection is included in figure 5.18, from which it is clear that the resist layer was etched away after approximately 3 min. The initial thickness of the resist layer of sample 3.8 was measured to be 445 nm, yielding an etch rate on PMMA of

$$r_{\text{PMMA}} = \frac{445\,\text{nm}}{3\,\text{min}} = 148.3\,\text{nm}\,\text{min}^{-1}. \tag{5.2}$$

The resulting structures can be seen in figures 5.19–5.22. It is clear from the figures that the recipe causes a lot of trenching, as well as positive profiles. The etch depth is approximately 600 nm, with some of the trenching extending deeper. Figure 5.19 depicts the resulting structures after exposure with dose 100. In figure 5.19a spikes are seen to form between the holes, which can be understood

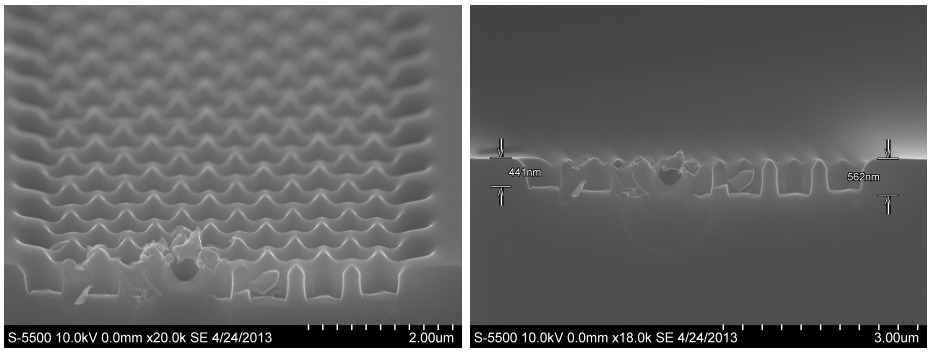

(a) **(b)**

Figure 5.19: SEM micrographs of etched holes in sample 3.8 with exposure dose 100.

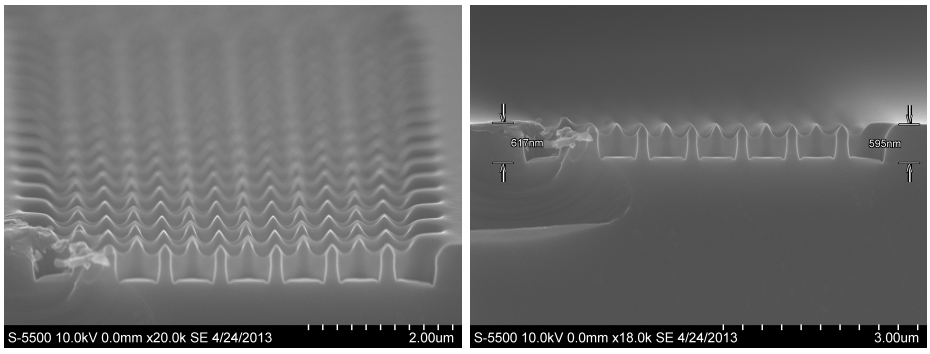

(a) **(b)**

Figure 5.20: SEM micrographs of etched holes in sample 3.8 with exposure dose 120.

considering overexposure of the PMMA: It is clear that smaller features will be more affected by proximity effects than larger structures, as more backscattered electrons will activate the resist in the smallest structures. Smaller features will thus have a thinner PMMA layer covering them than the bulk of the covered Si. Because of the thinner PMMA layer, the Si underneath will be exposed to the etching gas before the bulk of the PMMA is removed, leading to the observed shapes in figure 5.19a.

From figure 5.19b it is seen that the etch depth in Si is more than 560 nm and probably closer to 580 nm. Please note that as the PMMA was etched away during the first three minutes, the etch time in Si to get the observed depth was also only 3 min. During the last 20 s of the etching time all of the resist was gone, and thus the Si was etched both in the previously covered and uncovered areas. It is important to note that as the etch depth is almost three times that of the optimal PhC (220 nm), the etched holes would not look like the holes presented in these micrographs if the etch depth had been correct.

Figure 5.20 contains the micrographs for dose 120. The same shapes are observed in the Si, but the spikes are now narrower and sharper. This can be

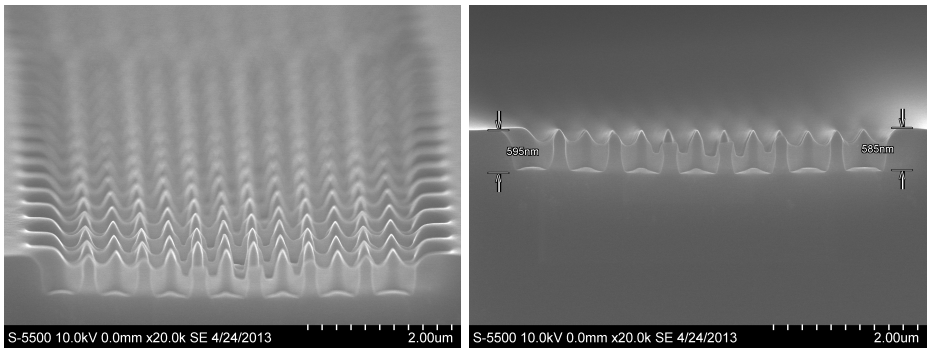

(a) **(b)**

Figure 5.21: SEM micrographs of etched holes in sample 3.8 with exposure dose 140.

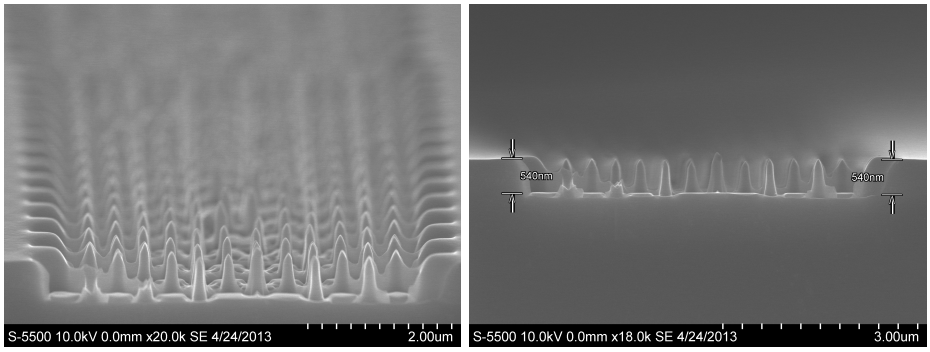

(a) **(b)**

Figure 5.22: SEM micrographs of etched holes in sample 3.8 with exposure dose 160.

understood to be caused by the increased dose, leading to a smaller region in which the PMMA has the thickness of the bulk. It is clear from figure 5.20b that the profile is only slightly positive, and the average etch depth is seen to be less than the recorded 595 nm at the right corner. (There is excessive trenching near the corner, giving an increased etch depth as compared with the holes in the center.) Again the etch depth is found to be approximately 580 nm.

The etched structures resulting from exposure doses 140 and 160 are presented in figures 5.21 and 5.22, respectively. Again, the increase in exposure dose leads to a blurring of the smallest features. The etch depth for dose 140 is observed to be 585 nm. In figure 5.22 the sample is seen to be cut at a place where there is a small height between the holes and the edge of the sample. This results in a recorded etch depth of 540 nm which is less than the actual etch depth.

Discussion of Etch Parameters

By evaluating all the cross-sectional SEM micrographs above, the etch depth in Si is approximated as $580\,\text{nm} \pm 5\,\text{nm}$. By noting that all PMMA was gone after

Sample	**3.8**
Recipe	Table 5.7
Etch time	3 min, 20 s
PMMA removed after	3 min
CCP reflected power	1 W
ICP reflected power	1 W
Plasma stabilization time	3 s
PMMA etch depth	445 nm
PMMA etch rate	148.3 nm min^{-1}
Si etch depth	580 nm ± 5 nm
Si etch rate	193.3 nm min^{-1}
Selectivity	1.30:1

Table 5.8: Etch parameters determined from the first HBr etch.

three minutes, this gives an etch rate in Si of

$$r_{\mathrm{Si}} = \frac{580\,\mathrm{nm}}{3\,\mathrm{min}} = 193.3\,\mathrm{nm\,min^{-1}}, \tag{5.3}$$

which results in an etch time of 1.14 min, or 68 s, for a target depth of 220 nm.

The resulting selectivity between Si and PMMA is found from equations (5.2) and (5.3) as

$$S_{\mathrm{Si:PMMA}} = \frac{r_{\mathrm{Si}}}{r_{\mathrm{PMMA}}} = \frac{193.3\,\mathrm{nm\,min^{-1}}}{148.3\,\mathrm{nm\,min^{-1}}} = 1.30{:}1, \tag{5.4}$$

which is marginally better than for the SF_6/CHF_3 etch recipe of chapter 5.4 (1.18:1). The observed etch parameters of this first HBr etch are summarized in table 5.8.

5.5.2 Increasing the CCP Power

Because of the excessive trenching for the original recipe, the CCP power was doubled from 50 W to 100 W. This was done to increase the DC bias and the physical etching, to combat the reactive ions from trenching. During etching, the thickness of the PMMA layer was measured using the end-point detection of the ICP. The selectivity was previously found to be 1.30:1, but the increased CCP power was assumed to decrease the selectivity closer to 1:1. By terminating the etching once approximately 200 nm of PMMA had been etched, the resulting etch depth in Si should be close to the desired 220 nm.

The resulting hole structures for exposure doses 100, 120, 140 and 160 are presented in figure 5.23. There is some trenching for the highest dose, but for the lower doses the trenching is not present. Dose 100 is found to be an appropriate dose in the middle of the structure, this can be understood from the increased physical etch rate: Even though the resist has probably not been completely removed in the holes from this low dose (*cf.* figure 5.3), the increased ion bombardment from the doubled CCP power will likely remove the remaining

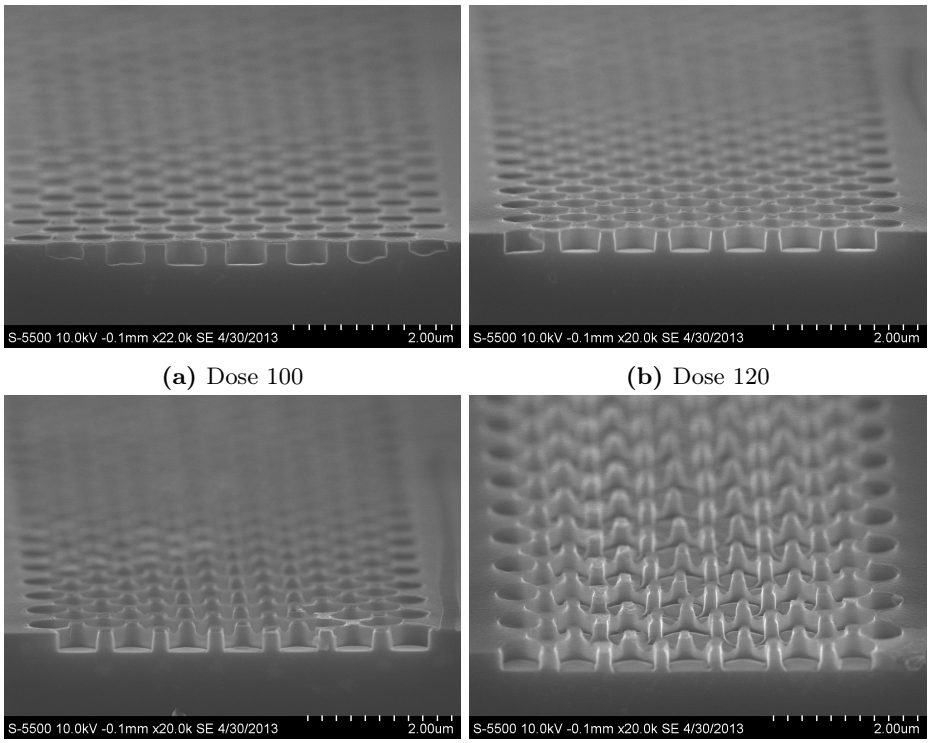

(a) Dose 100 **(b)** Dose 120

(c) Dose 140 **(d)** Dose 160

Figure 5.23: SEM micrographs of etched holes in sample 3.9, using the etch recipe of table 5.7 with CCP power set to 100 W. The EBL exposure doses are 100, 120, 140 and 160.

thin PMMA layer quite quickly, exposing the underlying Si. Already at dose 140 there is seen to be excessive etching of the thinnest structures, which is also likely to be caused by the high CCP power.

The positive profiles of figures 5.23b - 5.23d are thought to be the result of a slightly thinner resist layer around the edges of the holes, caused by proximity effects during EBL. As the etch recipe is run, the resist near the edges is gradually removed exposing more and more of the Si, resulting in gradually wider holes. It is clear that this problem is not present for dose 100, nor at the outermost holes of doses 120 and 140. From the figures and the discussion above, the optimal exposure dose is 100 or slightly higher in the middle, 120 at the edges and 140 near the corners of the exposure area.

The etch depth was found to be $280\,\text{nm} \pm 5\,\text{nm}$, and the difference in resist thickness before and after etching was found to be 256.9 nm. The etch parameters of this run are collected in table 5.9. The increased etch rate and the decreased selectivity are expected results of increasing the CCP power. The increased ion bombardment might also heat up the resist, causing its etch rate for PMMA to increase further.

Sample	**3.9**
Recipe	Table 5.7 (CCP = 100 W)
Etch time	58 s
CCP reflected power	1 W
ICP reflected power	1 W
Plasma stabilization time	2 s
PMMA etch depth	256.9 nm
PMMA etch rate	265.8 nm min^{-1}
Si etch depth	280 nm $\pm$ 5 nm
Si etch rate	289.7 nm min^{-1}
Selectivity	1.09:1

Table 5.9: Etch parameters determined from the second HBr etch, with increased CCP power.

HBr flow rate	50 sccm
Pressure	7 mTorr
ICP power	600 W
CCP power	100 W
Temperature	20 °C
He backside	10 Torr
Etch time	4×13 s
Pauses (N_2 purge)	3 min

Table 5.10: Adjusted HBr etch recipe, based on the recipe of table 5.7. The recipe is run in four cycles of 13 s with three minutes of N_2 purged pauses between the cycles.

5.5.3 Decreasing the Temperature

In order to make the etching recipe less reactive the temperature was decreased from 50 °C to 20 °C, which should also decrease the etch rate. In addition, pauses consisting of three minutes of N_2 purging steps were added. This might give the resist time to cool down if indeed the low selectivity is caused by resist heating due to ion bombardment.

The adjusted etch recipe is given in table 5.10. Assuming an etch rate somewhere between those found in the last two sections, the total etch time necessary to remove 220 nm of Si is expected to be around 50 s–60 s. The etch time for each cycle was set to 13 s, giving a total etch time of 52 s. Due to initial instability, the effective etch time might be somewhat lower.

The etching results of this run are summarized in table 5.11. From the table it is seen that the Si etch rate did in fact decrease to 78% of the previous etch rate (table 5.9). However, the etch rate on PMMA increased by 24%, yielding a selectivity of only 0.75:1 between Si and PMMA. Thus the added pauses did not serve their intention, and this low selectivity will likely be a problem.

SEM micrographs were obtained as for the other samples, but the results for doses 100 and 120 were very similar to those presented in figure 5.23. For dose

Sample	**3.10**
Recipe	Table 5.10
Etch time	$4 \times 13\,\mathrm{s}$
CCP reflected power	$1\,\mathrm{W}$
ICP reflected power	$1\,\mathrm{W}$
Plasma stabilization time	$2\,\mathrm{s}$
PMMA etch depth	$259\,\mathrm{nm}$
PMMA etch rate	$298.8\,\mathrm{nm\,min^{-1}}$
Si etch depth	$195\,\mathrm{nm} \pm 5\,\mathrm{nm}$
Si etch rate	$225\,\mathrm{nm\,min^{-1}}$
Selectivity	0.75:1

Table 5.11: Etch parameters determined from the third HBr etch, with decreased temperature and added pauses. Note how the selectivity is now below 1:1.

140 and higher, the increased etch rate of PMMA resulted in earlier removal of the resist near the holes, yielding more positive profiles.

PMMA After Etching

Sample 3.10 was not stripped of PMMA before the SEM micrographs were obtained. This was done intentionally, in order to inspect the pattern of the PMMA after etching. Figure 5.24a and 5.24b are close-up SEM micrographs of the remaining PMMA layer around the holes for dose 140 and 160, respectively. These figures contain a very important piece of information: Notice how the PMMA layer between the neighboring holes in the right part of figure 5.24a is the same height as the bulk PMMA, but very thin. From this it can be understood that the smallest features in the PMMA do not etch faster than the bulk PMMA in the vertical direction. Instead, the reactive species remove the PMMA from the side by lateral etching, causing the apparent higher etch rate for the smallest structures.

In figure 5.24b there is only a small area between the holes still covered by PMMA, as the rest of the PMMA has been etched away laterally, shrinking the smallest features. This is indeed what has given rise to the spike-like structures observed at high doses for the HBr recipes. Notice also the continued trenching in figure 5.24.

To conclude, the higher doses give rise to proximity effects, increasing the radius of the holes defined in the PMMA by backscattered electrons.

5.5.4 Comparing the HBr Based Recipes

Upon comparing the three different HBr based etch recipes, the original recipe of table 5.7 seems to be the best choice. The increased CCP power, as discussed in section 5.5.2, was found to decrease the trenching and as such it was a good improvement. However, the decreased selectivity is a problem.

Having a higher selectivity is judged to be more critical than decreasing trenching for two reasons: 1) If the etch rate on oxide is sufficiently low, the

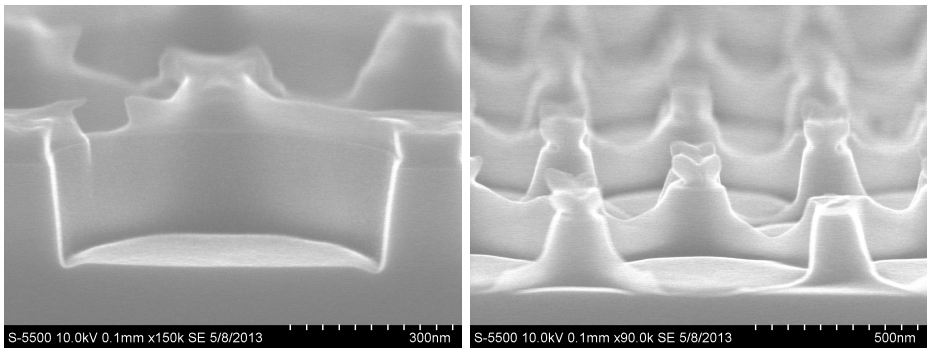

(a) Dose 140 **(b)** Dose 160

Figure 5.24: Close up SEM micrographs of remaining PMMA after etching of sample 3.10. (a) Notice how the resist layer between the holes is very thin, and almost gone for dose 140. (b) The Si has started to etch away between the holes for dose 160.

Location	**A**	**B**	**C**	**D**
Initial thickness [nm]	133.9	131.2	129.2	129.4
g (initial)	0.9959	0.9955	0.996	0.9967
Final thickness [nm]	112.8	110.1	108.8	109.7
g (final)	0.9999	0.9999	0.9999	0.9999
Etch depth [nm]	21.1	21.1	20.4	19.7

Table 5.12: Thickness of oxide layer before and after HBr etching. The measurements are collected using the reflectometer where g gives the goodness of fit for each measurement. The average etch depth is found to be 20.6 nm.

BOX layer of the SOI can be used as a stop layer for the HBr based Si etch. The etch recipe can then be run for a slightly longer time than necessary, giving an equal etch depth across the holes. 2) It is probable that the relatively thick PMMA layer might be identified as a source of roughness, once the process is optimized. This would require the selectivity against PMMA to be sufficiently high, in order to avoid adding processing steps by utilizing a harder mask (*e.g.*, chromium). In addition, the original recipe had an acceptably low etch rate, resulting in a target etch time of more than one minute.

5.5.5 Finding the Etch Rate against SiO_2

SiO_2 was deposited on a test sample by PECVD to find the etch rate against oxide for the original HBr recipe of table 5.7. The etch recipe was run for one minute on the sample, and the oxide thickness was measured in the reflectometer at four different locations on the sample before and after the etching. The measurements and the resulting etch depth are collected in table 5.12.

The average etch depth was found to be 20.6 nm, resulting in an etch rate of 20.6 nm min^{-1} and a selectivity of 9.4:1. The observed etch parameters are collected in table 5.13, along with the previously found etch rate against Si. The

Sample	**M.test.1**
Recipe	Table 5.7
Etch time	60 s
CCP reflected power	1 W
ICP reflected power	1 W
Plasma stabilization time	2 s
SiO_2 etch depth	20.6 nm
SiO_2 etch rate	20.6 nm min^{-1}
Si etch rate	193.3 nm min^{-1}
Selectivity (Si:SiO_2)	9.4:1

Table 5.13: Etch parameters observed during oxide etching of sample M.test.1. The etch rate for Si is copied from table 5.8 and is used to calculate the selectivity.

selectivity is much less than what should be expected from HBr based recipes [36], and it is uncertain whether it is high enough for the oxide to be useful as an etch stop layer. It should be noted that the oxide deposited by PECVD contains some hydrogen, which might make it etch faster than what should be expected from pure oxide in prime grade SOIs (*cf.* chapter 4.2).

5.5.6 Correcting for Proximity Effects

Based upon the previous discussions of proximity effects, a new mask was created in which the holes were divided into different layers depending on their placement within the structure. An extract of the mask is depicted in figure 5.25b. Sample 4.7 and 4.16 were processed through EBL with various combinations of exposure doses, and subsequently etched using the HBr recipe of table 5.7. The etch time was set to 68 s, based on the previously found etch rate in table 5.8 (193.3 nm min^{-1}). The observed etch parameters are summarized in table 5.14.

Figure 5.25a depicts the best result from the etch test, which were the hole structures of sample 4.7 exposed with doses of 100, 120 and 140 as the low, medium and high dose, respectively. It is clear that the division of the structure was not very successful: The holes defined by the highest dose (lower part of figure 5.25a) are seen to be quite well resolved, but the fourth row of holes from the bottom is clearly overexposed. This could be corrected by only letting the three bottom rows be exposed with the highest dose. The holes exposed with the lowest dose are seen to be well resolved in the middle of the structure, but a slightly higher dose is needed near the edges. Possibly a dose which is halfway between the low and the medium doses.

The holes along the left edge of the structures are seen to be elongated along the vertical direction, giving them an elliptical form. This might be due to proximity effects from the neighboring holes with the same dose, or an error during EBL. It is strange, however, that the same is not observed at the right side of the structure, as the exposure pattern is symmetric along the center of the hole structure.

Sample	**4.7**	**4.16**
Recipe	Table 5.7	Table 5.7
Etch time	68 s	68 s
CCP reflected power	1 W	1 W
ICP reflected power	1 W	1 W
Plasma stabilization time	7 s	2 s
PMMA etch depth	171.3 nm	200.1 nm
PMMA etch rate	151.1 nm min^{-1}	176.6 nm min^{-1}
Si etch depth	210 nm $\pm$ 5 nm	220 nm $\pm$ 5 nm
Si etch rate	185.3 nm min^{-1}	194.1 nm min^{-1}
Selectivity (Si:PMMA)	1.23:1	1.10:1

Table 5.14: Observed etch parameters from the HBr etching of samples 4.7 and 4.16. Note the difference in selectivity against PMMA. This might be due to a long instability time of the plasma for sample 4.7.

The holes in the middle of the structure are also seen to be slightly elliptical, but in the perpendicular direction with respect to the holes along the edge. One possible explanation is that the beam blanking between exposed and unexposed areas (*cf.* chapter 4.3.1) has been too slow, resulting in the observed shapes. As the holes along the edge are defined in a single layer, the beam is deflected much faster in the vertical direction for these structures. This might be a possible explanation for the observed structures, but again it is unexpected that the elongation is only present along one of the edges.

5.6 Simulating the Performance of a PhC Waveguide

Due to time limitations and the inability to procure SOI wafers for fabrication, it was not possible to fabricate any functioning PhC waveguides and measure their performance. As an alternative, and an approximation of actual performance, a method of simulating the performance of fabricated PhC structures in COMSOL was developed. The method is briefly outlined in appendix B.3.

The PhC structure of sample 4.16 exposed with doses 100, 120 and 140 (in the same way as for figure 5.25a) was used for the simulation. A portion of a top-view SEM micrograph was used to define the structures, and it is shown in figure 5.26a. In the obtained micrograph the axes of the PhC structures were not aligned with the axes of the image, so the image had to be rotated to do this alignment. Upon loading the SEM micrograph into COMSOL, a problem was quickly detected: The edges of the holes light up in the micrograph, as there is increased secondary electron emission from these edges. This leads the colormapping in COMSOL to assign higher values to the edges of the holes, resulting in a higher index of refraction along the edges.[1] To combat this problem, the SEM micrograph of figure 5.26a was converted to a binary image based on a threshold value.

[1]For more details, see appendix B.3.

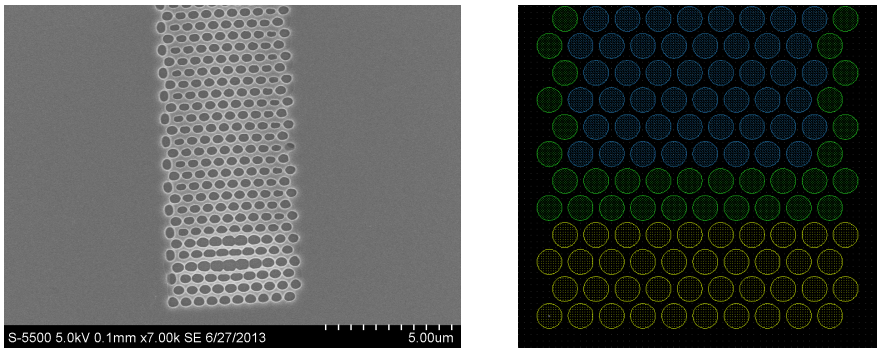

(a) **(b)**

Figure 5.25: (a) SEM micrograph of etched holes in sample 4.7 with proximity corrections, and (b) corresponding EBL mask. In the mask the yellow, green and blue layers correspond to high, medium and low dose, respectively. Notice how the medium dose seems to be too high along the edges, and the low dose seems to be too low near the edges in (a).

The threshold value was set such that the most resemblance with the perceived structures was obtained. The adjusted image is shown in figure 5.26b.

It is clear that the hole-structure of figure 5.26 is far from optimal, and the bandgap might not even be present. The holes are elliptical in shape, and the proximity correction was not very accurate. This makes it, however, well suited to study the simulation method, as interesting results are expected. In order to create a W1 PhC waveguide from the structure of figure 5.26b, the image was used to define the holes above and below a row of missing holes, and the spacing was set according to the target lattice constant of 503.26 nm. The lattice constant in units of pixels was found from the image by averaging the distance between neighboring holes in units of pixels. Note that the structure of figure 5.26b was used both above and below the missing row, but without mirroring the image, *i.e.*, the structure is not symmetric about the missing row. If it had been, the symmetry would cause symmetric scattering, which should not be expected in actual non-ideal PhC waveguides.

The simulation of the W1 waveguide created from figure 5.26b resulted in the electric field plot depicted in figure 5.27. It is clear from this figure that the structure does not work well as a waveguide, and much light is seen to leak through the PhC structures. From the simulation, 29.2% of the incident light is found to couple to the output, and 22.1% is reflected back to the input. Considering that this waveguide is only about 5 µm long, this is a a very high transmission loss. Upon creating a model of the same size with perfect holes, the transmission and reflection are found to be 92.3% and 3.2%, respectively.

The refractive index of the model is plotted in figure 5.28. It is clear that there are a lot of artifacts arising from the image processing: Even though most of the total area is well divided into red (Si) or blue (air), there are several bright spots where the refractive index is substantially different from the surroundings. This will of course introduce effects which should not be expected in actual fabrication. In addition, the periphery of the holes are seen to have a refractive index

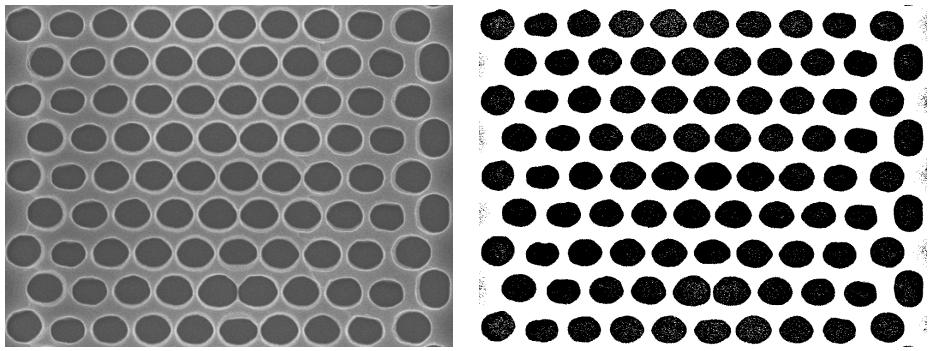

(a) **(b)**

Figure 5.26: (a) Original SEM micrograph of hole structures in sample 4.16 rotated and cropped to align the image with the PhC axes. (b) Enhanced version of (a) where the black and white levels has been set such that the white area is Si and the black areas are air. (b) was used in COMSOL to simulate waveguide performance.

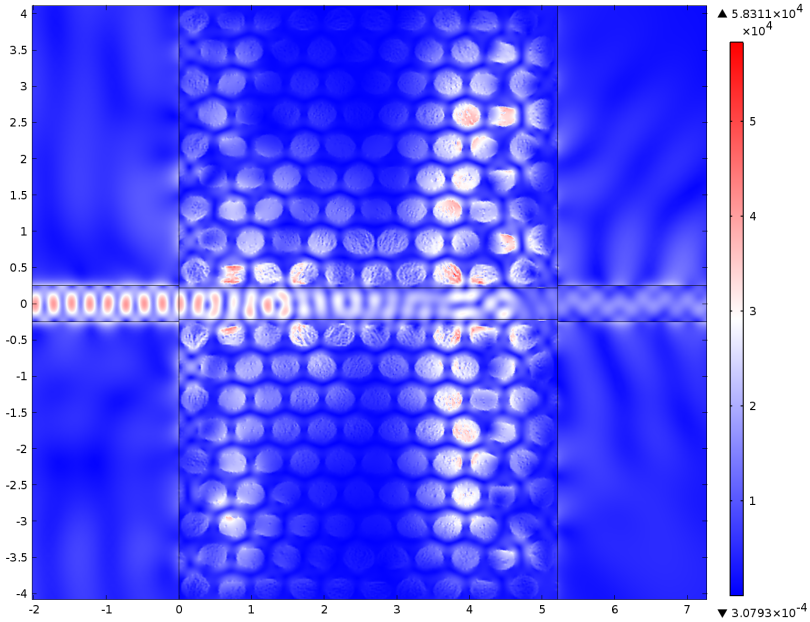

Figure 5.27: Electric field plot of the W1 waveguide, simulated by using the SEM micrograph of a trigonal lattice PhC to define the refractive index above and below the defect line. The strength of the field is given by the color-bar on the right side. The axes of the plot gives the size of the structures in units of µm.

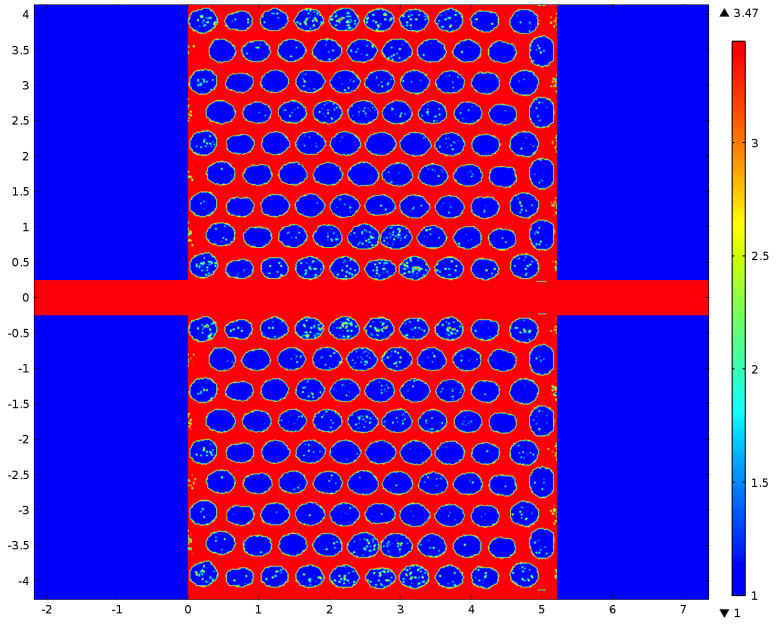

Figure 5.28: Refractive index of the model in figure 5.27. The refractive index is given by the color-bar on the right; it is seen that red corresponds to Si ($n = 3.47$) and blue corresponds to air ($n = 1$).

somewhat lower than Si.

In FEM simulations, the model is divided up in small elements which constitute the mesh, as discussed in appendix B.1. Normally, the boundaries between different materials are defined by geometrical elements in the model, and the mesh is set so that its edges line up with the material boundaries. In the model presented here, however, the boundaries between air holes and Si substrate are not explicitly defined. The meshing thus needs to be extremely fine in order to resolve the holes well enough, and a mesh size of 26 nm was chosen (one sixtieth of the vacuum wavelength). This is much smaller than what is normally needed in FEM simulations, resulting in a simulation time of approximately 45 min for the structure of figure 5.27.[2]

The simulation software assumes identical conditions within each element, yielding regions along the hole edges which have refractive index somewhere between that of Si and air

From comparing figures 5.26 and 5.28, it is clear that the size of the mesh introduces additional roughness around the holes. The size of the mesh is essentially a quantization of space, which is the cause of this additional roughness. It is clear that a denser mesh is needed to provide a more accurate solution, but that is not possible with the limited computing power available.

[2]From experience, this is about the longest simulation time the given hardware can successfully perform. More computing power is needed to simulate denser meshes or larger structures.

Chapter 6

Conclusions

6.1 Fabrication

6.1.1 Etching

Two etch chemistries have been investigated in this research and both have been demonstrated to be capable of PhC fabrication, although further optimization is necessary. The SF_6 based recipe could not be investigated much due to system malfunction, but it is believed to have a lot of potential for improvement.

The fluorine based recipe provides almost vertical profiles, which is very important in PhCs. As the passivation gas/reactive gas ratio and the CCP/ICP ratio can be adjusted freely, it is believed that optimization will lead to vertical profiles. The HBr based recipes generally seem to be more reactive with significant lateral (chemical) etching, yielding positively sloped profiles. This could be compensated for by increasing the CCP/ICP ratio, but as demonstrated in chapter 5.5.2 this also results in a decreased selectivity against PMMA which might become problematic. Trenching was observed when the CCP power was 50 W, but disappeared upon increasing the CCP power to 100 W. This was, however, at the expense of decreased selectivity and increased etch rate.

The selectivity between Si and PMMA was found to be quite low for all recipes. For the SF_6/CHF_3 recipe the selectivity was approximately 1.2:1, but upon decreasing the CCP/ICP ratio the selectivity will increase. The selectivity for the original HBr recipe was found to vary between 1.1:1 in chapter 5.5.6 and 1.3:1 in chapter 5.5.1. This suggests that the HBr etch chemistry is more unstable and subject to small fluctuations in chamber conditions, which obviously is not good for reproducible results. In addition, there is a trade-off between selectivity and anisotropy as discussed above.

Even though it was not investigated here, fluorine based etch chemistries generally have a low selectivity between Si and oxide. For the original HBr recipe, a decent selectivity of 9.4:1 was demonstrated for PECVD deposited oxide. As discussed in chapter 4.2, the selectivity against pure oxide will be higher. Whether or not this selectivity is important depends on the target PhC structure (*cf.* chapter 3.1): For a shallow etched PhC it is important to have a high selectivity, such that the BOX of the SOI can be used as a stop layer.

It is important not to have too high etch rates, as the recipes ideally should be of minimum 1 min duration to ensure good control of etch depth and reproducible results. This was observed during the SF_6 etching of samples 3.6 and 3.7 (chapter 5.4.1). The SF_6 recipe was run for 15 s on both samples and the resulting etch rates were 800 nm min^{-1} and 680 nm min^{-1}, respectively. The etch rate for the adjusted SF_6 recipe will be lower than this, but it is uncertain how much it is possible to reduce it. If a shallow etched PhC is targeted it is critical to have good control of the etch rate to achieve the target 220 nm etch depth. For the HBr based recipes, the etch rate was found to be much lower (193 nm min^{-1} for the recipe of table 5.7). Due to the expected high selectivity against pure oxide it would also be possible to overetch slightly, to ensure complete removal of the top Si layer within the holes.

To sum up, the SF_6 recipe clearly offer more well defined holes with vertical profiles, but due to the high etch rate it is difficult to achieve control over etch depth. Optimization of the recipe will hopefully yield both vertical profiles, relatively low etch rate, and higher selectivity against PMMA.

6.1.2 EBL

From the first SEM investigations of the exposed resist, dose 120 was found to provide vertical resist profiles. The impact of changing the acceleration voltage was not studied in this work, mainly because of time limitations. The voltage of 20 keV was chosen because it is commonly used in NTNU NanoLab. This might seem arbitrary, but when the optimal parameter is not known it is wise to base the choice upon the body of knowledge developed by previous users. Based upon the almost vertical profiles of the structures in Si etched with the SF_6 recipe, it is concluded that the acceleration voltage is appropriate, although more thorough investigations should research the effect of varying voltages.

The only parameter which was varied during EBL was the exposure dose. Based upon the SEM micrographs of the exposed resist and the etched holes in sample 3.7, dose 120 is found to be the best choice for the most densely exposed areas. For the HBr based etch recipes, the radius of the holes are seen to increased from the intended values as this chemistry etches more laterally. This should be combated by decreasing the chemical etch rate as described in the last section, probably yielding 120 as an appropriate dose also here.

Lower doses are typically seen to decrease the hole radius and etch depth, and introduce roughness around the edges of the holes. This is thought to be due to remaining resist in the exposed holes (as seen in figure 5.3) and insufficient exposure along the edges of the holes leading to the observed rough borders. As the undeveloped PMMA in the holes is etched away, the Si substrate is exposed to the reactants. If the remaining PMMA is slightly uneven, the resulting structures in Si will inherit the uneven shape of the PMMA—resulting in unevenly etched structures.

Conversely, an overexposure of the PMMA typically leads to very smooth structures, but at the expense of increased hole radius. The increased exposure dose allows complete removal of the resist, even near the edges of the holes, but

because of an increased amount of backscattered and secondary electrons there is a widening of the exposed structures. For all the etched samples, it is clear that proximity effects result in different optimal doses depending on the density of the exposure pattern.

Proximity corrections were made for the etching of samples 4.7 and 4.16, but the division of the mask into different doses was not accurate enough. However, the method is explained and it is straightforward to investigate how a certain PhC should be divided up in layers of lower and higher exposure doses. The best result would obviously have a continuously decreasing dose from the edge of the pattern towards the center, but decent results can probably be achieved with only three different layers if the borders between high and low doses are set at the right place in the PhC.

6.2 Simulations

Simulations were performed in MIT Photonic Bands, yielding $r/a = 0.45$ as the optimal hole radius for the trigonal lattice PhC. However, fabrication limitations were briefly discussed and it was decided that $r/a = 0.4$ was a good compromise between a large bandgap and a manageable minimum feature sizes. The fabrication results suggests that it might be possible to create larger holes, *i.e.*, smaller regions of Si between the holes.

FEM simulation software was used to simulate the electric field in W1 PhC waveguides, with and without rough holes. A degradation of the bandgap was observed at increasing levels of roughness, leading to radiation of light through the hole-structures. At lower levels of roughness, the bandgap was found to be intact, but the imperfections in the waveguide were observed to reflect much light towards the input side. The rough holes are seen to be sources of scattering, which in the 2D case only cause reflections within the plane. In the 3D case, however, scattering from rough holes are likely to scatter light away from the SOI plane; *i.e.*, sidewall roughness causes leakage out of the plane and is thus a source of loss.

A method of simulating the performance of fabricated structures was developed, and the electric field in a W1 PhC waveguide was simulated. The SEM micrograph used to define the structures was converted into a binary image, in which white (black) corresponded to Si (air). However, there were some bright spots in the black holes and some dark spots on the substrate. This was interpreted by the FEM software as small regions of Si in the air holes and small pinholes of air in the Si substrate, both of which are unlikely to occur in fabrication.

The finite size of the mesh was found to introduce additional roughness along the edges of the holes, as some of the mesh elements comprised regions of different refractive index. This must be resolved by increasing the mesh size, or possibly defining the structures explicitly. It is obvious that the SEM micrographs need to be enhanced better before the simulation results can bear meaning. However, this is an exercise in image processing which is beyond the scope of this thesis.

6.3 Test Setup and Coupling

Some considerations regarding coupling were outlined in chapter 3.3. Based on these discussions, a test setup for measuring the transmission through a PhC waveguide should probably utilize a grating for input and output coupling, connected to tapered waveguide segments. The main drawback with this approach is that the measured transmission efficiency will be a combination of 1) alignment error between grating and fiber, 2) coupling efficiency of the gratings, 3) transmission efficiency of the tapers, 4) coupling efficiency between slab waveguide and PhC waveguide, and 5) transmission efficiency of the PhC waveguide. One possible solution would be to fabricate one structure which only consists of gratings and a slab waveguide, from which all losses except the coupling loss between slab waveguide and PhC waveguide could be found. This would of course assume identically processed structures and proper alignment of the fiber with respect to the grating.

A main concern with fabrication of a grating coupler is the differing etch depth with respect to the other structures. For a shallow etch of the taper and the PhC, 220 nm of Si should be etched, whereas the grating outlined in chapter 3.3.3 has an etch depth of 50 nm. The structures would most likely need to be fabricated individually: First the PhC, taper and the outline of the grating should be defined by EBL and subsequently etched. After these structures are fabricated in the Si, the trenches that make up the grating will need to be defined perfectly aligned to the already defined structures. This naturally requires a very accurate alignment during EBL processing.

Even though it is difficult to find the exact transmission efficiency of a PhC waveguide, this might not be a big concern in the first measurement setup. The first goal will probably be to successfully couple light in and out of a PhC waveguide, and any substantial transmission will imply that light is in fact guided through the waveguide—yielding the experiment successful.

6.4 Future Work

It is the authors hope and belief that this thesis will serve as a platform to build upon for new students wishing to explore PhC fabrication and design. A focus throughout this thesis has been the limitations and possibilities of all methods and results. As such, there are several hints herein towards future work, but the most important points are listed below:

- Optimize the SF_6/CHF_3 etch recipe to yield higher selectivity, lower etch rate, and vertical profiles.
- Optimize the proximity correction to achieve identical holes across the PhC.
- Investigate the possibility of overexposing the resist slightly to yield smooth holes, compensated by slightly smaller holes defined by the mask.
- Procure commercial SOIs, or research the possibilities of depositing low-loss a-Si by PECVD to fabricate SOI in-house.
- Enhance the image processing of SEM micrographs.
- Explore the possibilities of using cross-sectional SEM micrographs to define the profile of a PhC in FEM simulations.

Appendix A

MIT Photonic Bands

A.1 Calculating the Bandgap

The control sequence for a simulation in MPB is defined by a `.ctl` (control file) script, which is described by the example of the trigonal lattice photonic crystal of figure 2.4. The band-gap for TE- and TM-polarized light in a crystal, consisting of air holes formed in a trigonal lattice in a silicon substrate, is calculated.

First, the parameters of the model need to be specified. These include the radius of the holes `r`, the dielectric constant `eps`, and the number of k-points to calculate for between the Γ-M-K-Γ points (*cf.* Ch. 2.6):

```
(define-param r 0.4) ;normalized radius of the holes
(define-param eps 12.1104) ;dielectric constant of Si
(define-param k-interp 8) ;# of points between k-points
(define Si (make dielectric (epsilon eps))) ;defines Si
```

A few words about the syntax are in order: `(define-param a 1)` defines the variable `a` to have the constant value of 1, whereas `(define A (obj_type (obj_var b)))` defines `A` to be an object of `obj_type` (*e.g.*, a vector or a material) which is set to have the property defined by `obj_var` equal to `b`. Notice also that "`;`" marks the start of a comment.

The crystal lattice is created:

```
;create the trigonal lattice:
(set! geometry-lattice
                (make lattice (size 1 1 no-size)
                (basis1 (/ (sqrt 3) 2) 0.5) ;basis vectors...
                (basis2 (/ (sqrt 3) 2) -0.5))) ;of the lattice
(set! default-material Si) ;sets Si as default
;create the holes in the lattice:
(set! geometry (list (make cylinder ;cylindric holes
                (center 0) (material air) ;holes=air
                ;assume infinity in z-dir. and set radius=r
                (radius r) (height infinity))))
```

The `(set! ...)` command is used to change predefined entities to those appropriate for the model under study. Here, the lattice of the crystal is set to be a 2D lattice, with basis vectors equal to those of a trigonal lattice. Notice that MPB uses Polish notation for mathematical calculations, *e.g.*, `(/ 1 2)` is ½ and `(sqrt 3)` is $\sqrt{3}$. The second `set!` command defines the default material to be Si, *i.e.*, all volumes not specified to be of another material will be Si. The final command creates the cylindrical holes centered in the origin, consisting of air, with radius `r` and infinite in z-direction.

The reciprocal lattice points where the band diagrams should be calculated need to be set. These k-points lie within the irreducible Brillouin zone, as discussed in chapter 2.4.

```
;the k's are set to be at the Gamma, M, K, Gamma points,
(set! k-points (list (vector3 0 0 0)
                (vector3 0 0.5 0)
                (vector3 (/ -3) (/ 3) 0)
                (vector3 0 0 0)))
;...but we also need calculation points in between (interpolate)
(set! k-points (interpolate k-interp k-points))
```

First, the k-points are set to be only at the corners of the irreducible Brillouin zone, specified by the vectors. Some points in between these corners should also be calculated, to yield the dispersion relation. The built-in function `interpolate` is used to get `k-interp` number of calculation points (which was defined previously to be 4) between the corners of the irreducible Brillouin zone, defined in `k-points`.

The final steps are to set the resolution of the lattice cell (which should be a power of 2, or a power of a small prime, for computational efficiency [20]), and to decide how many bands to calculate.

```
(set-param! resolution 32) ;how to discretize the cell
(set-param! num-bands 8) ;number of bands to calculate

        (run-te) ;find TE band-diagram
        (run-tm) ;find TM band-diagram
```

Notice that the `set-param!` function differs from the `define-param` function, in that `set-param!` changes—or assigns a value to—a parameter which is already defined by default. (The default value for `resolution` is 10.) Finally, `run-te` and `run-tm` tell MPB to calculate the band diagrams for TE- and TM-modes, respectively. After having run the analysis, one can extract the calculated data to MATLAB for better plotting possibilities.

A.2 Band Diagrams for TE Modes

Figures A.1–A.4 contains the band diagrams for $r/a = 0.3$, 0.35, 0.4 and 0.45 to show how the size of the bandgap depends on hole radius.

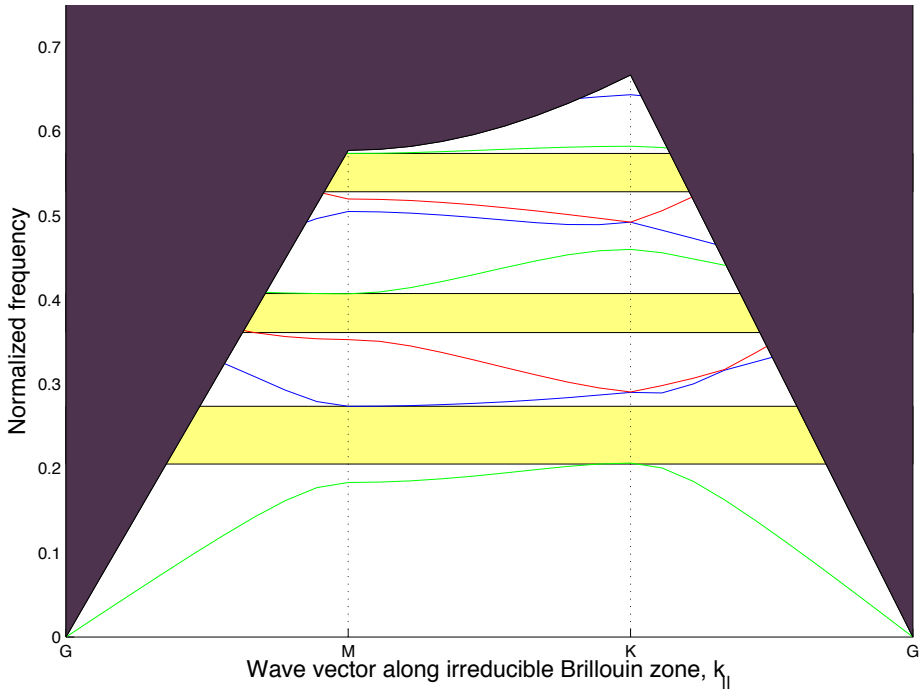

Figure A.1: Band diagram for TE-modes in the trigonal lattice at $r = 0.3a$.

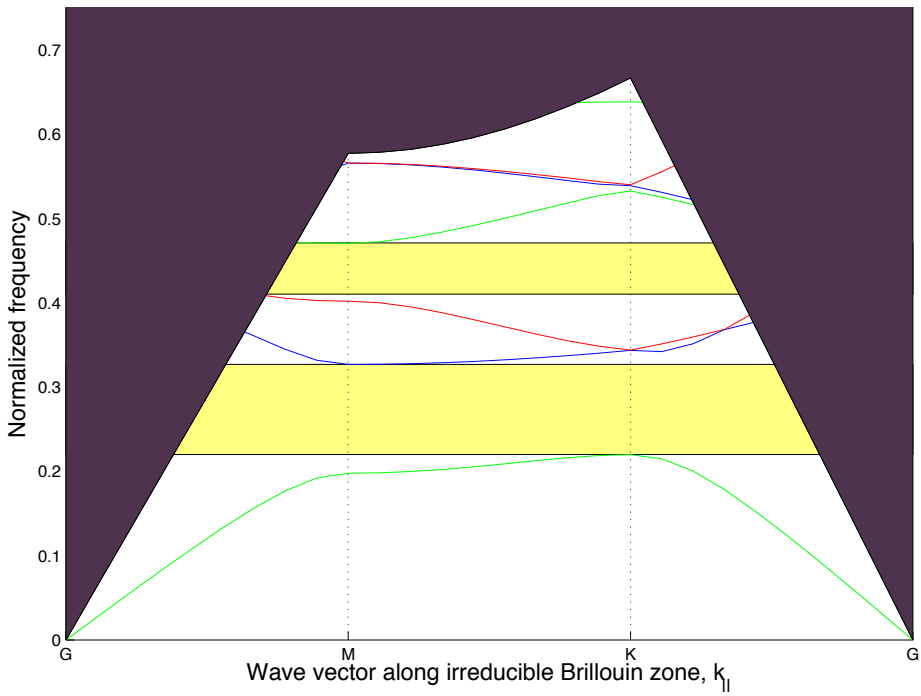

Figure A.2: Band diagram for TE-modes in the trigonal lattice at $r = 0.35a$.

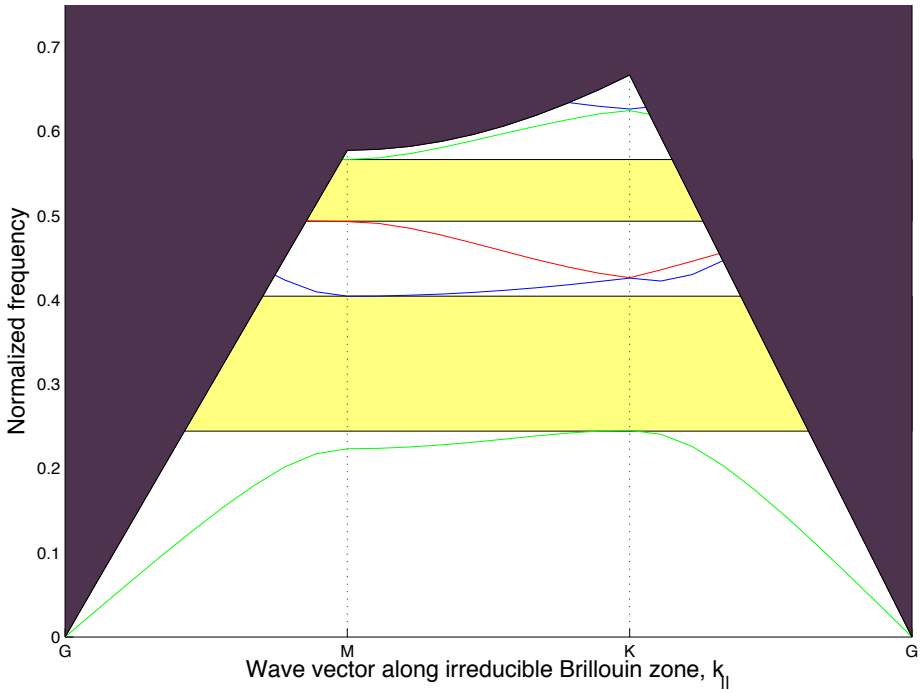

Figure A.3: Band diagram for TE-modes in the trigonal lattice at $r = 0.4a$.

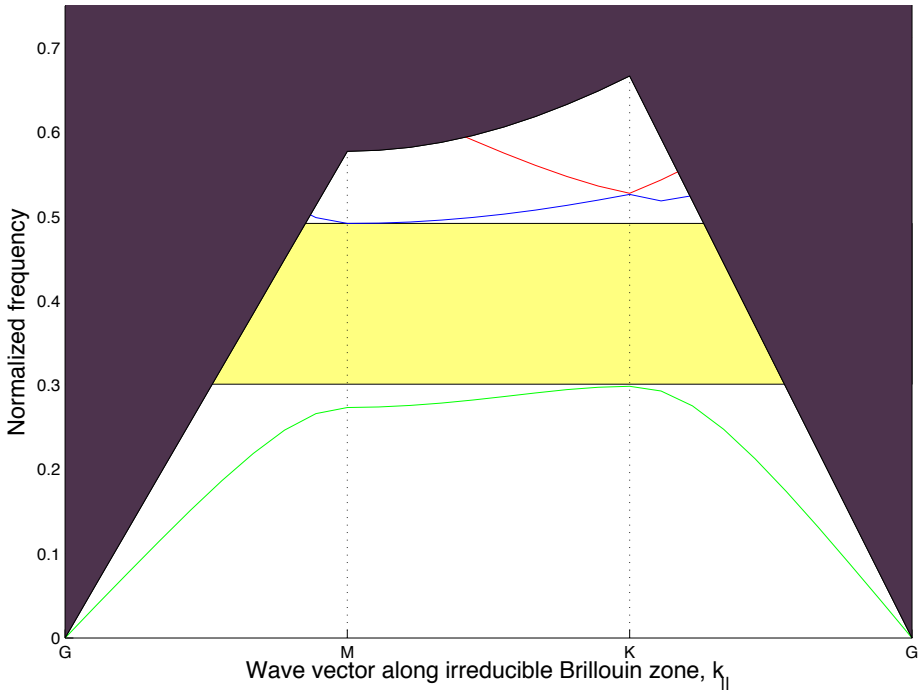

Figure A.4: Band diagram for TE-modes in the trigonal lattice at $r = 0.45a$.

Appendix B

COMSOL Multiphysics

B.1 Finite Element Method

The finite element method (FEM) divides the simulation space up in small regions (finite elements), which can be linked together by the governing equations. Rather than a full solution with infinitesimal regions, the FEM provides a numerical solution that calculates how the electric field (or other physical measurable quantity) varies between the elements. The simulation space needs to be divided up into small elements by a mesh, and the maximum size of this mesh should be chosen small enough to get an accurate solution. A rule of thumb for electromagnetic problems is to choose the maximum mesh size to be no larger than one tenth of the wavelength.

It should be obvious that a finer mesh results in a more accurate solution. On the other hand, longer computation time will result when the number of finite elements increases. In addition, a larger simulation space will increase the demand for computational power, unless the mesh size is increased.

In any simulation, some assumptions have to be made. In studying the band diagram of the trigonal PhC, it was assumed that the crystal was infinite in all directions. An actual PhC will of course not be infinite, but the band-gap holds true even for finite structures. In the COMSOL simulations, the model needs to be confined to a relatively small space, in order to minimize the computational time. The boundaries of the model need to impose the proper boundary conditions on the field, to avoid spurious reflections of light. A perfectly matched layer (PML) can be added around the boundaries of the model, to absorb any scattered waves. This is an excellent way of limiting the simulation space: If an actual photonic device were to be placed in an environment without any interfering objects (*e.g.*, in air or vacuum), the scattered waves would leave the device without being reflected back to it. This is much the same situation as a PML: The scattered waves are not reflected back into the PhC. The PML should be at least one wavelength thick, in order to absorb most of the scattered waves.

B.2 Roughness Considerations

The following text is adapted from [9], and provides some background information on the roughness simulations of chapter 3.4.2. For a more thorough walk-through, including a step-by-step guide for COMSOL simulations of a PhC waveguide, please see the original thesis [9].

B.2.1 Modeling Non-Ideal Photonic Crystals

Once the 2D photonic crystal is known in terms of its band-diagram, and how it works as a waveguide, a study of more practical interest was undertaken: How will a PhC waveguide perform once imperfections are taken into account? The performance metric that made the most sense to find during a simulation the transmission of light through the waveguide. Thus, a COMSOL model had to be created, in which it will be possible to compare PhC waveguides with different degrees of hole roughness to that of a perfect PhC waveguide.

Please note that as the roughness experiments herein are done in 2D, their results can not be coupled directly to any actual fabrication processes. Neither was this at any time the intention of this research as this is a demanding task which can not be performed in the limited time available. This is merely a first-pass into the study of roughness, with its main goal being to get acquainted with the possibilities of FEM simulations on non-ideal structures. This research should be followed up with a more in depth study of how an actual 3D structure will respond to imperfections in fabrication.

The first thing to consider before performing any simulations is how to model the roughness. Several important factors play into this decision: Obviously, the model should be based upon the known limitations of processing equipment used in processing of PhCs. This is highly dependent upon the processing steps, materials, the individual machines used in the fabrication, and even the fabrication facility itself. However, the more accurate the model becomes, the harder it will be to describe it mathematically. Even more important is the effects the roughness will have on the computation time, as resources such as time and computational power is severely limited during this research. Rough structures with high frequency noise will have a lot of small sharp features that will require a very fine mesh, *i.e.*, there will be a great increase in the degrees of freedom, thus increasing the computation time.

In order to find a good way to model the roughness, it is useful to look at how roughness often presents itself in cleanroom fabrication. During resist development, and later etching of the structures, the quickest variations (the high frequency noise) will most likely be smoothed out. Thus it is probably a good assumption that the amplitude of any high frequency noise contributions will be minimal.

During an e-beam lithography process on an array of holes, it is expected that the inconsistencies will come from three main contributions:[1] 1) The periodicity of the holes might have small variations throughout the crystal. This is the most crucial inconsistency—intuitively understandable as a PhC should be perfectly

[1]Thanks to PhD students Magnus Breivik and Jon Olav Grepstad for useful insight on the processing of PhC's.

periodic—but also the parameter which is easiest to control. 2) The radius of each hole will vary throughout the crystal, depending on variations in material parameters. 3) The holes will not be perfectly round—there will be a roughness component superimposed around the circumference of each hole. Often the holes will appear slightly elliptical.

As a first-pass study of non-ideal PhCs, it is not necessary to create a model that incorporates all possible forms of inconsistencies. Because of limited simulation time available, the structures that are to be studied will be small. Thus there is probably little value in investigating the effect of the periodicity, as it is not likely that slow variations in periodicity will yield interesting results in a small model. The obvious starting point for this research is the effect of noise contributions around each circle, *i.e.*, the effect of non-ideal—or *rough*—circles upon the transmission of light through a PhC. This roughness will be superimposed upon an ideal radius, r_0, that will be constant for each circle. It is possible to later look at the effect of variations in the ideal radius throughout the crystal in addition to variations in each single hole, but it is best to start the study with just one changing variable.

The radius of each hole will be quantified by the function

$$r(\phi) = r_0 \cdot (1 + \mathcal{N}(0, \sigma^2)), \tag{B.1}$$

where r_0 is the ideal radius, and $\mathcal{N}$ denotes the normal distribution, here with mean 0 and standard deviation σ. For each value of ϕ—the angle around the circle—a new random number will be drawn from the normal distribution, thus creating a normally distributed $r(\phi)$ centered around the ideal radius value.

In addition to the amplitude of the roughness at each point ϕ, given by the standard deviation in (B.1), the frequency of the noise will contribute greatly to the final shape of the holes. This is determined by the number of points around the circle where r will be calculated: Only one point will give no variation in radius, whereas a large number of points will give very rapid variations around the hole. A few different configurations are tried, and remembering that the high frequency noise will probably be evened out, the final decision is to use eight points around the circle. These points will be interpolated by COMSOL to yield the shape of the holes.[2]

B.2.2 Simulating the Roughness

The transmission and reflection coefficients are calculated automatically in COMSOL with each simulation run, when ports are used to excite the input of a waveguide and to absorb the transmitted waves [37]. The roughness study will be done by varying the standard deviation, σ, of equation (B.1), and plotting the resulting transmission and reflection. Because of the normally distributed roughness around each hole, it should be expected that some of the created structures will have much better or worse performance than what should be expected from

[2]The decision to interpolate exactly eight points in the creation of the holes might seem arbitrary. Indeed to some degree it is, but it was found that this produced the hole shapes with most resemblance of previously fabricated PhC structures [9]. A different number of interpolation points would yield different results. A more thorough model should have several contributions of different frequencies, and the different components should be compared.

the value of σ: Even at the highest levels of uncertainties in a fabrication, some structures will perform excellent!

In order to have statistically significant data, the simulations of transmission and reflection coefficients at each value of σ need to be repeated a number of times. The most interesting part of the study will be around the standard deviations most likely to occur in actual fabrication, and this is where it is most critical to have a good prediction of how much loss will be expected.

The field distributions of all the simulations will be studied, and compared to the corresponding transmission and reflection coefficients, and hopefully some trends will emerge. It is expected that increasing hole roughness will result in non-uniform reflections, and worse waveguiding. Some stray defects within the PhC—*i.e.*, regions where the light is no longer forbidden—might appear if the hole shapes are rough enough. It would be interesting to reveal some correspondence between scattering patterns and transmission-/reflection coefficients.

B.3 Simulating Fabricated PhC Structures

This very short guide will explain how to load images in COMSOL to map material properties, specifically refractive index, based on a SEM image. It will assume familiarity with other aspects of COMSOL, which can be gained from the step-by-step guide on PhC waveguide simulations in [9, App. B] or the COMSOL user guide [37].

It is possible to import images into COMSOL, mapping the colors (or brightness) of the image to numeric values. These values can in turn be used in various ways, *e.g.*, to set material properties. However, in order for the properties to be set in a meaningful way, the mapping must done correctly. For SEM images there often are brightness variations due to, *e.g.*, higher emission of secondary electrons from edges. Figure B.1 depicts how the mapping was done upon loading the unprocessed SEM micrograph of figure 5.26a into COMSOL.

Even though the air holes seem to be properly defined, the Si region is seen to have great variations in assigned value. This would lead to difficulties in defining the refractive index correctly. For this reason, the image used in the simulation of chapter 5.6 was transformed into the binary image of figure B.2.

The image is loaded into COMSOL by right clicking *Global Definitions, Functions, Image*. The size of the image must be specified in units of meter. To do this, the lattice constant can be found from the image in units of pixels. If, *e.g.*, the distance between two nearest neighboring holes is 50 pixels and the image is 500 pixels by 400 pixels, the absolute size of the structures contained within the image should be $10a$ by $8a$.

Upon loading the image of figure B.2 into COMSOL, a plot representing the assigned value of each pixel is created. The plot is depicted in figure B.3, from which it is obvious that the value of 1 is assigned to Si and 0 to air. Also obvious are the artifacts arising from improper image processing.

The function created by the image, $im(x, y)$, should now ideally contain only values of 1 and 0 for Si and air, respectively. To use this function to specify the PhC region above a missing row of holes in a W1 waveguide, create a rectangle in COMSOL of the same size as the structures in the image. Now, the material

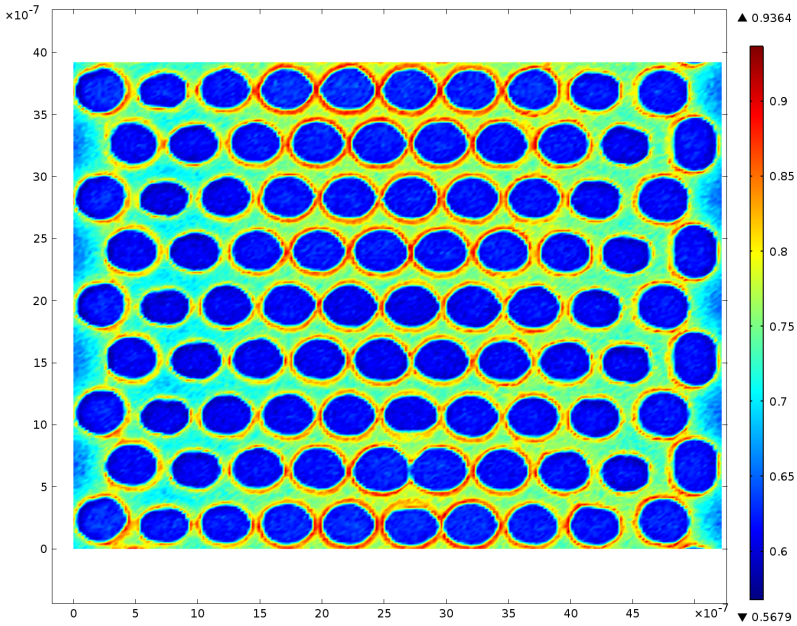

Figure B.1: 2D plot of numeric value assigned to the function $im(x, y)$ upon loading the unprocessed SEM micrograph of figure 5.26a into COMSOL.

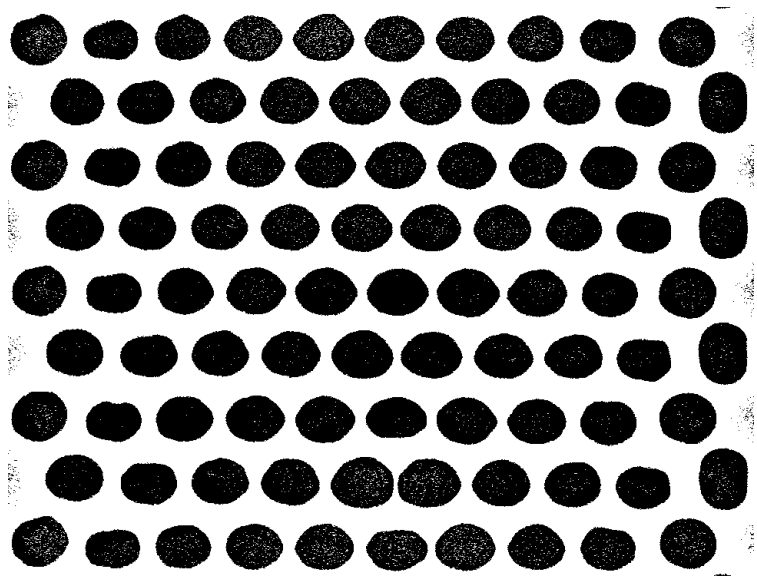

Figure B.2: Binary version of SEM micrograph of PhC structure, used to define refractive index in COMSOL. Same image as in figure 5.26b.

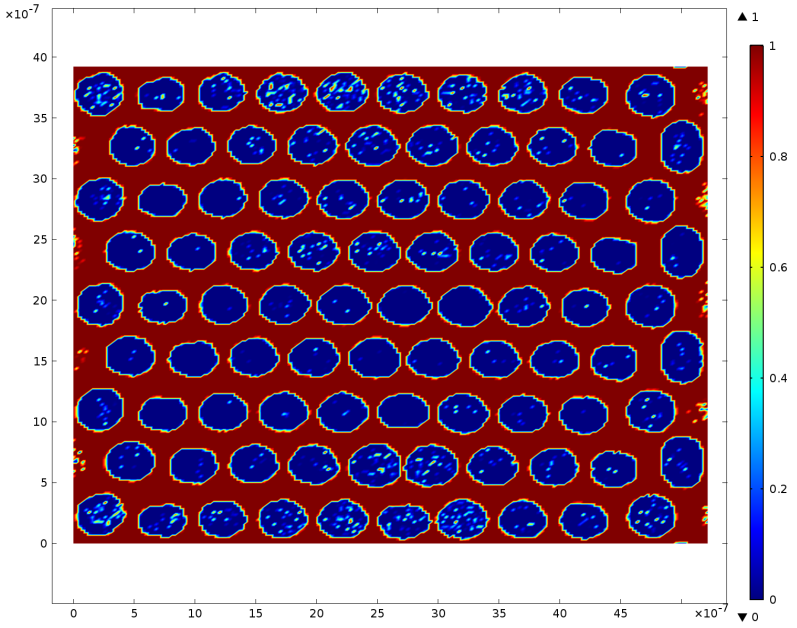

Figure B.3: 2D plot of numeric value assigned by the binary SEM micrograph of figure B.2 to the function $im(x, y)$.

parameters for this rectangle can be set by calling the function $im(x, y)$. Assume the rectangle representing the structures is placed with its base corner in the origin. In the *Materials* tab in COMSOL add a new material, assign it to the rectangle, and set its refractive index to be dependent upon $im(x, y)$. For the example at hand, the refractive index should be set to

$$n = 1 + 2.47 \cdot im1(x, y). \tag{B.2}$$

For air $im = 0$ and thus $n = 1$, whereas for Si $im = 1$ and $n = 3.47$.

If the rectangle had not been placed at the origin, the function im should be called with relative coordinates, *e.g.*, $im(x - dx, y - dy)$ if the rectangle is based at (dx, dy).

Appendix C

List of Instruments

Tool	**Manufacturer**	**Model**	**ID#**
PECVD	Oxford Instr.	PlasmaLab System 100-PECVD	1225
EBPVD	Pfeiffer	Vacuum Classic 500	1205
EBL	Hitachi	4300 SEM	1103
		With Raith electronics	
ICP-RIE (F)	Oxford Instr.	Plasmalab System 100 ICP-RIE 180	1207
ICP-RIE (Cl)	Oxford Instr.	Plasmalab System 100 ICP-RIE 180	1230
Scriber	Dynatex	DXIII	3021
SEM	Hitachi	S-5500 S(T)EM	1512
Reflectometer	Filmetrics	F20	1500
Profilometer	Veeco	Dektak 150	1720

Table C.1: List of instruments used in this work. The ID# refers to NTNU NanoLab's internal ID system.

Appendix D

Process Flow and Fabrication Recipes

D.1 Process Flow

As an overview of the process flow, a step-by-step guide to the fabrication process is listed below. The details of each step are presented in chapter 4. The process begins with a 2 inch wafer.

- Scribing and breaking:
  - Scribe the wafer into smaller samples using the recipe of appendix D.3.
  - Clean all samples to remove the particles created by the scribing and breaking process.
- Electron beam lithography:
  - Clean the sample with acetone, ethanol, IPA, and DI-water. Blow-dry with N_2.
  - Dehydrate on hotplate at 200 °C for at least 20 min.
  - Blow with N_2 and apply a drop of 950PMMA A5.5.
  - Spin-on 950PMMA A5.5 at 2000 RPM for a target resist thickness of 445 nm.
  - Pre-bake on hotplate at 180 °C for 60 s to evaporate solvents and improve adhesion.
  - Inspect sample in microscope to detect any contaminants or bubbles in the resist.
  - Measure resist thickness in reflectometer.
  - Scratch sample to define the optimal exposure area.
  - Expose the sample in EBL.
  - Develop the exposed resist in a mixture of nine parts IPA and one part DI-water for 30 s, stop the development by soaking in DI-water. Gently blow-dry with N_2.
  - Inspect the sample in microscope, to ensure proper development.
- Plasma etching:

  - Precondition the plasma chamber by running the desired recipe for 20 min on a dummy Si wafer.
  - Attach the sample to a 2 inch Si carrier wafer by a drop of Fomblin oil.
  - Set the desired etch time in the ICP-RIE system and load the sample into the loadlock.
  - Begin the automated process and monitor the stability of the plasma to ensure proper etching conditions.
  - Unload the sample.

- Remove resist (optional):
  - Soak the sample in acetone.
  - Place the beaker of acetone, with the sample, in an ultrasonic bath for 10 min–15 min at moderate strength (optional).
  - Clean the sample with acetone, ethanol, IPA, and DI-water. Blow-dry with N_2.

- Scribe sample to allow cross-sectional inspection (optional). Follow the same routine as when scribing the initial wafer, but make sure the sample is broken across the structures of interest. The scribing should not be across the entire wafer, as this will destroy the structures: Scribe a line on each side of the structures to allow a clean break at the area of interest.

- Inspect sample in SEM:
  - Clean the sample thoroughly.
  - Attach the sample with a conductive tab to the sample holder. OR:
  - Fasten the sample with screws in the cross-sectional holder.
  - Set the desired acceleration voltage and beam current in the SEM and adjust the beam.

D.2 PECVD Fabricated SOI Wafer

The PECVD recipes used to deposit the oxide and a-Si layer for the SOI wafer of chapter 5.1 are given in table D.1 and D.2, respectively. They are collected from Oxford Instruments' catalog of recipes and are commonly used in NTNU NanoLab.

SiH_4 flow rate	8.5 sccm
N_2O flow rate	710 sccm
N_2 flow rate	161.5 sccm
Temperature	300 °C
RF power	20 W
Pressure	1000 mTorr

Table D.1: Oxford Instruments' standard PECVD recipe for SiO_2 deposition.

SiH_4	25 sccm
Ar flow rate	475 sccm
Temperature	250 °C
RF power	7 W
Pressure	1000 mTorr

Table D.2: Oxford Instruments' standard PECVD recipe for a-Si deposition.

The observed deposition parameters, and the resulting layer thicknesses are given in table D.3. The layer thicknesses were measured by cross-sectional SEM micrographs of the sample.

	Oxide layer	**a-Si layer**
Recipe	Table D.1	Table D.2
Deposition time	21 min, 30 s	8 min, 48 s
Reflected power	0 W	1 W
Plasma stable after	1 s	2 s
Resulting layer thickness	1.63 µm	240 nm
Deposition rate	75.8 nm min^{-1}	27.3 nm min^{-1}

Table D.3: Observed deposition parameters during PECVD fabrication of SOI wafer.

D.3 Scribing Parameters

Tables D.4 and D.5 contain the parameters for scribing and breaking a 250 µm thick 2 inch Si wafer into 10 mm × 6 mm samples. The same parameters were used to dice the samples before cross-sectional SEM investigations, but the scriber and breaker was set to interactive mode and the scribe lines were chosen manually.

Scribe parameters	**X**	**Y**
Methdod	Continuous	
Step size	10 mm	6 mm
Impulse bar height	0 µm	0 µm
Extension	90 µm	90 µm
Scribe force	1600 cnts	1600 cnts
Scribe angle	36°	36°
Scribe speed	6 mm s^{-1}	6 mm s^{-1}
Approach speed	1 mm s^{-1}	1 mm s^{-1}

Table D.4: Scribe parameters used during scribing and breaking of a 2 inch Si wafer into 10 mm × 6 mm dies.

Break parameters	**X**	**Y**
Method	Anvil	
Air pressure	65 kPa	65 kPa
Anvil height	0.426 mm	0.426 mm
Gap	0.376 mm	0.376 mm
Dwell time	0.1 s	0.1 s
Cycle time	0.2 s	0.2 s

Table D.5: Break parameters used during scribing and breaking of a 2 inch Si wafer into 10 mm × 6 mm dies.

Bibliography

[1] Eli Yablonovitch. Inhibited spontaneous emission in solid-state physics and electronics. *Phys. Rev. Lett.*, 58:2059–2062, May 1987.

[2] Sajeev John. Strong localization of photons in certain disordered dielectric superlattices. *Phys. Rev. Lett.*, 58:2486–2489, Jun 1987.

[3] J.D. Joannopoulos. The almost magical world of photonic crystals. In *LEOS '99. IEEE Lasers and Electro-Optics Society 1999 12th Annual Meeting*, volume 1, pages 232 –233 vol.1, 1999.

[4] John D. Joannopoulos, Steven G. Johnson, Joshua N. Winn, and Robert D. Meade. *Photonic Crystals Molding the Flow of Light*. Princeton University Press, Priceton, NJ, 2nd edition, 2008.

[5] B. Jalali and S. Fathpour. Silicon photonics. *Lightwave Technology, Journal of*, 24(12):4600 –4615, dec. 2006.

[6] Chen Xia, Li Chao, and Tsang Hon K. Device engineering for silicon photonics. *NPG Asia Mater*, 3:34–40, 01 2011.

[7] Zhou Fang and Ce Zhou Zhao. Recent progress in silicon photonics: A review. *ISRN Optics*, 2012(428690), 2012.

[8] Jens Høvik. Photonic crystal waveguide fabrication. Master's thesis, Norwegian University of Science and Technology, Aug 2012.

[9] Marius Lorvik. Investigating non-ideal effects in a 2d photonic crystal waveguide, using finite element method simulation software. Project thesis, Norwegian University of Science and Technology, December 2012.

[10] Masatoshi Tokushima, Hideo Kosaka, Akihisa Tomita, and Hirohito Yamada. Lightwave propagation through a 120[degree] sharply bent single-line-defect photonic crystal waveguide. *Applied Physics Letters*, 76(8):952–954, 2000.

[11] Susumu Noda, Alongkarn Chutinan, and Masahiro Imada. Trapping and emission of photons by a single defect in a photonic bandgap structure. *Nature*, 407:608–610, 2000.

[12] Bong-Shik Song, Susumu Noda, and Takashi Asano. Photonic devices based on in-plane hetero photonic crystals. *Science*, 300(5625):1537, 2003.

[13] Toshihiko Baba. Slow light in photonic crystals. *Nat Photon*, 2(8):465–473, 08 2008.

[14] John D. Joannopoulos, Pierre R. Villeneuve, and Shanhui Fan. Photonic crystals: putting a new twist on light. *Nature*, 386:143–149, 1997.

[15] Rajesh V. Nair and R. Vijaya. Photonic crystal sensors: An overview. *Progress in Quantum Electronics*, 34(3):89 – 134, 2010.

[16] Bahaa E.A. Saleh and Malvin C. Teich. *Fundamentals of Photonics*. John Wiley and Sons, Inc., Hoboken, NJ, 2nd edition, 2007.

[17] Charles Kittel. *Introduction to Solid State Physics*. John Wiley and Sons, Inc., Hoboken, NJ, 8th edition, 2005.

[18] W. Bogaerts, R. Baets, P. Dumon, V. Wiaux, S. Beckx, D. Taillaert, B. Luyssaert, J. Van Campenhout, P. Bienstman, and D. Van Thourhout. Nanophotonic waveguides in silicon-on-insulator fabricated with cmos technology. *Lightwave Technology, Journal of*, 23(1):401–412, 2005.

[19] M. Notomi, A. Shinya, S. Mitsugi, E. Kuramochi, and H. Ryu. Waveguides, resonators and their coupled elements in photonic crystal slabs. *Opt. Express*, 12(8):1551–1561, Apr 2004.

[20] MIT Photonic Bands. MPB webpage. `http://ab-initio.mit.edu/wiki/index.php/MIT_Photonic_Bands`.

[21] nanoHUB.org. NanoHUB's MPB client. `http://nanohub.org/tools/mpb`.

[22] P. Sanchis, P. Bienstman, B. Luyssaert, R. Baets, and J. Marti. Analysis of butt coupling in photonic crystals. *Quantum Electronics, IEEE Journal of*, 40(5):541–550, 2004.

[23] J.V. Galan, P. Sanchis, and J. Marti. Low-loss coupling technique between SOI waveguides and standard single-mode fibers. In *European Conference on Integrated Optics*, April 2007.

[24] O. Mitomi, K. Kasaya, and H. Miyazawa. Design of a single-mode tapered waveguide for low-loss chip-to-fiber coupling. *Quantum Electronics, IEEE Journal of*, 30(8):1787–1793, 1994.

[25] F. Van Laere, G. Roelkens, M. Ayre, Jonathan Schrauwen, D. Taillaert, D. Van Thourhout, T.F. Krauss, and R. Baets. Compact and highly efficient grating couplers between optical fiber and nanophotonic waveguides. *Lightwave Technology, Journal of*, 25(1):151–156, 2007.

[26] Edward D. Palik. *Handbook of Optical Constants of Solids*. Academic Press, Boston, 1985.

[27] T. Lipka, O. Horn, J. Amthor, and J. Müller. Low-loss multilayer compatible a-si:h optical thin films for photonic applications. *Journal of the European Optical Society - Rapid publications*, 7(0), 2012.

[28] Shiyang Zhu, G. Q. Lo, and D. L. Kwong. Low-loss amorphous silicon wire waveguide for integrated photonics: effect of fabrication process and the thermal stability. *Opt. Express*, 18(24):25283–25291, Nov 2010.

[29] Jr. M. F. Ceiler, P. A. Kohl, and S. A. Bidstrup. Plasma-enhanced chemical vapor deposition of silicon dioxide deposited at low temperatures. *J. Electrochem. Soc.*, 142(6):2067–2071, 1995.

[30] C. Vieu, F. Carcenac, A. Pépin, Y. Chen, M. Mejias, A. Lebib, L. Manin-Ferlazzo, L. Couraud, and H. Launois. Electron beam lithography: resolution limits and applications. *Applied Surface Science*, 164(1–4):111 – 117, 2000.

[31] Harry J. Levinson, Mark A. McCord, Franco Cerrina, Robert D. Allen, John G. Skinner, Andrew R. Neureuther, Martin C. Peckerar, F. Keith Perkins, and Michael J. Rooks. *Handbook of Microlithography, Micromachining, and Microfabrication. Volume 1: Microlithography.* Society of Photo-Optical Instrumentation Engineers, 1997.

[32] Sanja Hadzialic. *Photonic Crystal Slabs for Optical Micro Electromechanical Systems (OMEMS).* PhD thesis, Universiy of Oslo, Nov. 2009.

[33] MicroChem. PMMA Data Sheet. `http://microchem.com/pdf/PMMA_Data_Sheet.pdf`.

[34] Michael Quirk and Julian Serda. *Semiconductor Manufacturing Technology.* Prentice-Hall, Upper Saddle River, NJ, 2001.

[35] Michael David Henry. *ICP Etching of Silicon for Micro and Nanoscale Devices.* PhD thesis, California Institute of Technology, 2010.

[36] C.C. Welch, A.L. Goodyear, G. Ditmer, and G. Tan. Choice of silicon etch processes for opto- and microelectronic device fabrication using inductively coupled plasmas. In *Optoelectronic and Microelectronic Materials and Devices, 2004 Conference on*, pages 13–16, 2004.

[37] COMSOL. Rf module user's guide, May 2012.