



NTNU – Trondheim
Norwegian University of
Science and Technology

Photonic Crystal Waveguide Fabrication

Jens Høvik

Master of Science in Electronics

Submission date: August 2012

Supervisor: Astrid Aksnes, IET

Co-supervisor: Oleg Zero, 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 types of components are dielectric waveguides, micro-resonators, light modulators, couplers, periodic structures and more.

In this thesis a silicon photonic crystal waveguide structure was simulated and optimized using various simulation software. The framework for constructing a PhC waveguide in the cleanroom was then developed.

Prior to this research, an in-depth study of photonic crystal theory was conducted, and a short summary is included within the thesis. Numerical simulations using the finite-element method in simulation software called COMSOL were utilized, and methods of studying photonic crystal phenomena with this software specifically was developed. Furthermore, another piece of simulation software called MIT Photonic Bands is utilized to study and, along with COMSOL, verify the results of the simulations. These two simulation methods together were used to not only simulate photonic crystals, but optimize a photonic crystal waveguiding device.

The main goal of the thesis is however, to lay the groundwork for fabricating such a device in the NTNU cleanroom. During the course of this research, a large amount of time was spent training to use the equipment facilitated by the NTNU NanoLab. The main challenge was understanding the principles of operation for the equipment, and developing methods to allow for the fabrication of a PhC waveguiding device. Methods of constructing a silicon-on-oxide wafer with the needed parameters in mind was developed by use of plasma-enhanced chemical vapor deposition, and

the results were subsequently characterized through refractometry, profilometry, and ellipsometry.

The pattern of the photonic crystal structure was defined with electron-beam lithography, using polymethyl methacrylate (PMMA) as a resist and an isopropanol solution as developer. The main etching was done by inductively coupled plasma reactive ion etching, and the characterization of the structures was done through optical microscopy, scanning electron microscopy (SEM), and scanning transmission electron microscopy (S(T)EM).

Other cleanroom instruments used throughout the process includes the plasma cleaner, RIE, MA56 mask aligner, scribe, and electron-beam evaporator.

Sammendrag

Denne oppgaven omhandler et teoretisk studium og deretter fabrikasjon av strukturer med periodisk dielektrisk konstant, også kjent som fotoniske krystaller. Disse strukturene viser interessante dispersjons-karakteristikker som gir dem bred rekkevidde med frekvenser som ikke kan propagere i krystallen. Denne egenskapen gjør de nyttige for et bredt utvalg av fotonikk-baserte komponenter.

En-dimensjonelle fotoniske krystaller er allerede kommersialiserte og blir hyppig brukt i for eksempel tynn-film optikk, og to-dimensjonale fotoniske krystaller er tilgjengelige som fotonisk krystall fibre. Mye arbeid er gjort rundt analyseringen av fotoniske krystaller og deres applikasjoner i forskjellige fotoniske komponenter. Når man skal designe en fotonisk krystall komponent må egenskapene til krystallen først være forstått og analysert.

Ved å bruke frekvens-domene simuleringsprogramvare har vi studert egenskapene til fotoniske krystaller og designet den ideelle strukturen for å begrense lys av enhver bølgelengde innen en tre-dimensjonell fotonisk krystallplate. Simuleringer har blitt benyttet for å finne den ideelle periodisiteten til den fotoniske krystallen, sammen med dens optimale geometri for bølgeledning.

Studien ble deretter utvidet til tre dimensjoner og den optimale tykkelsen til den fotoniske krystallplaten ble funnet. Ved å bruke disse verdiene kunne vi begrense lys inne i en defekt i den fotoniske krystallen. Dette fenomenet lot oss studere bruken av fotoniske krystaller til å lage bølgeledere.

Det neste steget i forskningen var å fabrikere den simulerte fotonisk krystall bølgelederen. Metoder for å lage en silisium-på-insulator skive

ved hjelp av 'plasma enhanced chemical vapor deposition' ble utviklet. Profilometri, refraktometri og ellipsometri ble brukt for å karakterisere kvaliteten til silisium-på-insulator skiven. En overflateruhet mellom 1.5-3.5 nm ble funnet. Tapene til det amorfe silisiumet ble forsøkt målt ved ellipsometri, men teknikken egner seg dårlig for tapsmålinger. Det er nødvendig å benytte andre teknikker for å måle tap på en pålitelig måte men dette lot seg ikke gjennomføre i løpet av masteroppgaven. En metode er å koble lys inn i bølgelederstrukturen og måle transmisjonstapet ved utgangen.

Den fotoniske krystall bølgelederstrukturen ble formet ved hjelp av elektronstråle litografi, og forskjellige tykkelser fotoresist ble testet og karakterisert. Videre ble forskjellige materialer forsøkt brukt som etsemaske. Polymethyl metakrylat (PMMA), silisium dioksid, og krom ble alle brukt, hvor bare den sistnevnte viste god nok selektivitet for å etse silisium. Dette førte til en kombinasjon av silisium dioksid og krom som den endelige etsemasken.

Flere forskjellige etsemetoder ble forsøkt. Både standard våtetsing, romtemperatur 'reactive ion etching' og kryogensk 'inductively coupled plasma reactive ion etching' ble brukt. Tørretsing viste seg å ha gode nok egenskaper, mens isotropien under våtetsing gjorde denne upraktisk for strukturer som kreves i en fotonisk krystall bølgeleder. Den endelige løsningen som tillot fabrikasjon av en fotonisk krystall var en flerlagsstruktur, som brukte både krom, silisiumdioksid, og PMMA igjennom etseprosessen.

PMMA ble brukt som en etsemaske til oksidet, oksidet ble brukt som en etsemaske til kromet, og krom ble igjen brukt som en etsemaske til silisiumet. Den kryogenske etseoppskriften ble karakterisert og vist til å være tilstrekkelig anisotropisk såvel som å ha en ganske høy maksimum etsedybde. Oppskriften viste en uønsket formasjon av såkalt silisiumgress, og en ganske høy sideveggruhet ble observert.

Til slutt ble en komplett fotonisk krystall bølgeleder fabrikert, med en gitterstruktur som smalner av for å koble lys inn i den fotoniske krystallen. Til tross for at de presenterte metodene fortsatt må optimaliseres, gir de grunnlaget for å fabrikere ikke bare fotoniske krystaller, men en mengde av optiske komponenter i renrom.

Abstract

This research is entirely devoted to the study and fabrication of structures with periodic dielectric constants, also known as photonic crystals (PhCs). These structures show interesting dispersion characteristics which give them a range of prohibited frequencies that are not allowed to propagate within the crystal. This property makes them suited for a wide array of photonic-based components. One-dimensional photonic crystals are already commercialized and are of widespread use in for example thin-film optics, and two-dimensional PhCs are available in the form of photonic crystal fibers.

Much work is being done on the analysis of photonic crystals and their application in various photonic components. In designing any photonic crystal component the properties of the crystal must first be understood and analyzed. By using frequency-domain simulation software we have studied the properties of bulk photonic crystals and designed the ideal structure for confining light of any design wavelength within a three-dimensional photonic crystal slab.

Through simulations, the ideal periodicity of the PhC has been found along with its optimal geometry for waveguiding purposes. The study was then extended to three dimensions and the optimal thickness of the PhC slab was found. By using these determined values we were able to properly confine light inside a defect in a photonic crystal structure. This phenomenon allowed us to briefly examine the use of a photonic crystal as a waveguide.

This research was then continued by attempting to fabricate such a device. Methods entailing how to create a silicon-on-insulator wafer through plasma enhanced chemical vapor deposition were developed. Profilometry, refractometry and ellipsometry were used to characterize

the quality of the SOI wafer. A surface roughness between 1.5-3.5 nm was found. The losses of the amorphous silicon was attempted to be measured through ellipsometry, but the ellipsometer was found incapable for loss measurements. Better testing methods must therefore be developed.

The PhC waveguiding structure is formed through electron-beam lithography, and various thicknesses of photoresist are tested and characterized. Furthermore, various materials were attempted as an etch mask. Polymethyl methacrylate, SiO_2 , and chromium were all utilized, where only the latter showed good enough selectivity for silicon etching. This led to using both chromium and SiO_2 as the mask of choice. Several etching methods were tested. Both standard wet etching, room-temperature reactive ion etching recipes as well as cryogenic inductively coupled plasma reactive ion etching recipes were employed. Dry etching was found to be of insufficient quality, while the complete isotropy of wet etching rendered it impractical for fine structures needed for a PhC waveguide.

The final solution which allowed the formation of the photonic crystal was a stacked structure, using both chromium, oxide, and PMMA throughout the etching process in order to achieve the wanted results. The PMMA was used as an etch mask to imprint the pattern in the oxide layer. The oxide layer was in turn used as an etch mask to imprint the pattern into the chromium, and the chromium served as the final mask to be used during cryogenic silicon etching.

The cryogenic silicon etching recipe was characterized and shown to be sufficiently anisotropic as well as having a rather high maximum etch depth. The recipe shows an unwanted formation of so-called silicon grass, and a high sidewall roughness is observed. Finally, a complete PhC waveguiding structure with both a grating, a taper, and a silicon waveguide is fabricated. Although they require more work to be optimized, the methods presented in this thesis provide the basis for fabricating a multitude of optical components, not only PhC's.

Acknowledgements

First and foremost I would like to thank my co-supervisor, Oleg Žero, for spending countless hours with me in the cleanroom, and for making the longest days more tolerable. His organizational skills contributed much to the work being done and it is safe to say that the work conducted in this thesis could not have been done without him.

I would also like to thank Astrid Aksnes, for abiding in the work in any way she could, and for allowing Oleg Žero to spend so much of his time helping in the lab. Another thank you goes out to Lars Martin Sandvik Aas, for allowing me to use the ellipsometer in realfagsbygget, and for helping with characterization of samples with ellipsometry.

Furthermore, I would like to thank Espen Rogstad, for being very flexible with instrument training times, for always being available whenever an instrument malfunctioned, and not to mention for taking time from his busy schedule to help in the cleanroom should it be needed.

Last, but not least, I would like to give a special thanks to David Barriet, whose expertise within the cleanroom was paramount for the success of this thesis. He was always available to answer questions about both cleanroom equipment and process methods, and he supplied us with some of his own personal etching recipes in which this thesis would not have been possible without.

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 1 |
| 2 | Theory | 5 |
| 2.1 | Calculating the band diagram of a one-dimensional photonic crystal | 5 |
| 2.2 | Symmetry considerations | 10 |
| 2.3 | Trigonal lattice | 16 |
| 3 | Method | 17 |
| 3.1 | Finding the ideal photonic crystal waveguide parameters | 17 |
| 3.2 | Fabrication equipment | 19 |
| 3.2.1 | Deposition methods | 19 |
| 3.2.2 | Pattern imprinting techniques | 22 |
| 3.2.3 | Etching techniques | 26 |
| 3.2.4 | Characterization techniques | 28 |
| 3.2.5 | Other cleanroom instruments | 32 |
| 3.3 | Coupling | 32 |
| 3.4 | The silicon-on-insulator wafer | 34 |
| 3.4.1 | The oxide layer | 34 |
| 3.4.2 | The amorphous silicon layer | 35 |
| 3.5 | Forming the pattern | 37 |
| 3.6 | The etching recipe | 39 |
| 3.6.1 | Finding the correct etching recipe and mask | 39 |
| 3.6.2 | Opening the mask | 45 |
| 4 | Results and discussions | 49 |
| 4.1 | The ideal photonic crystal waveguide | 49 |
| 4.2 | The silicon-on-insulator wafer | 56 |
| 4.2.1 | The oxide layer | 56 |
| 4.2.2 | The amorphous silicon layer | 61 |
| 4.3 | Forming the pattern | 66 |
| 4.4 | The etching recipe | 70 |

| | | |
|----------|---|------------|
| 4.4.1 | Finding the correct etching recipe and mask | 70 |
| 4.4.2 | Opening the mask | 76 |
| 4.4.3 | Characterizing the cryogenic etching recipe | 91 |
| 5 | Conclusion | 109 |
| 5.1 | The silicon-on-insulator wafer | 109 |
| 5.2 | Finding the good EBL parameters | 110 |
| 5.3 | The etching recipes | 112 |
| 5.4 | Future work | 114 |
| A | PECVD recipes | 117 |
| B | ICP-RIE recipes | 119 |
| C | Tutorial on MPB photonic bands | 121 |
| D | Simulating PhCs with COMSOL 4.2 | 125 |
| E | Sample parameters | 133 |

Chapter 1

Introduction

The manipulation of light has been subject to extensive studies over the past 100 years. From lenses to optic wires, one always seeks to further improve both our understanding and control of photonic phenomena. The photonic crystal is a new up-and-comer in this line of research, where a periodic change in refractive index has shown many promising features.

In essence, a PhC is simply a structure of periodic refractive index as one moves along an axis. As an example, figure 1.0.1 shows an incoming electromagnetic wave incident on a one-dimensional PhC. From the point of view of the incoming electromagnetic wave, as long as it is incident perpendicular to the left hand surface of the PhC, it will always experience a periodic refractive index. Depending on the design parameters of the PhC, such as the refractive index of the two materials of which it composes, the thickness of the slabs, as well as the frequency of the incoming electromagnetic wave, it will either be transmitted through the photonic crystal, or reflected back.

In other words, it serves as an optical filter. This is analogous to how an electron experiences a periodic potential as it moves along a crystal lattice, which leads to a bandgap for the electron, which means that only electrons of allowed energies will be able to propagate through the crystal [10]. In the same manner, this happens for photons trying to propagate through a photonic crystal, although the intuitive understanding is not as readily understood as for the electron. In short, the introduction of the photonic band gap means that a PhC will exhibit a band of prohibited frequencies in which no electromagnetic waves can propagate. This phenomenon is interesting for a wide array of photonic/based components. Much research has been done on the properties of a photonic crystal as a waveguide [2–4], and the use of photonic crystals as dielectric mirrors in many applications has also

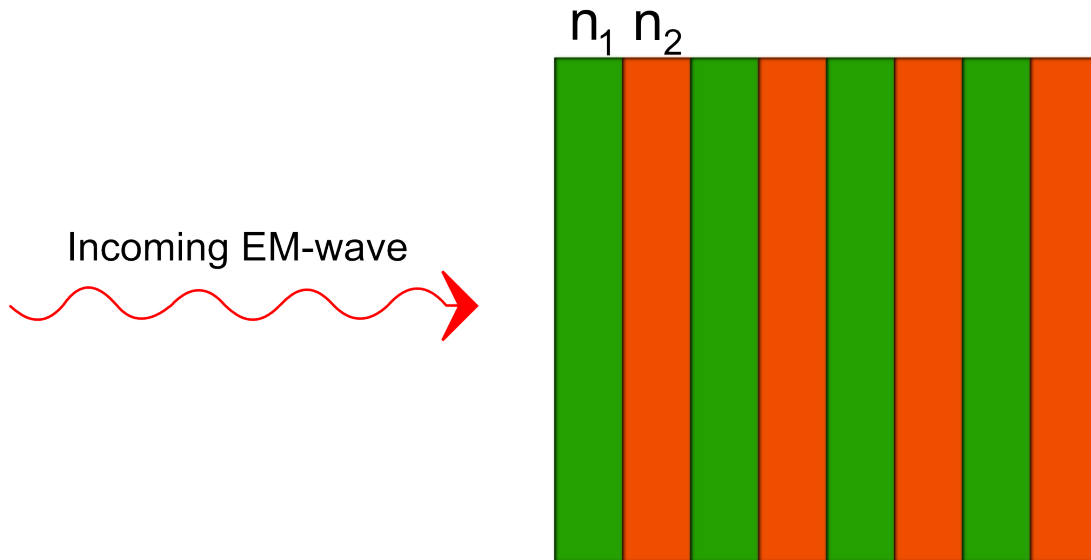


Figure 1.0.1: 1D photonic crystal

been shown [5–7]. The PhC has also found extensive use as a main component in sensing applications [12–15].

Unlike the periodic potential that causes the bandgap in a semiconductor, we can control the dimensions of the photonic crystal in a way that we never could for electronics. Our inherent control over the structural design of the PhC lets us introduce defects in forms of points or lines. This serves to further enhance the properties of our PhC, which in turn can be used to fabricate cavities [8, 9] or waveguides within the crystal.

Photonic crystals can be manufactured in one, two, or three dimensions and the dimensionality of the PhC has a direct influence on its properties. Most notably, the dimensionality of the structure defines for which directions of an incoming wave the properties of interest hold true. For example, an EM wave incident on a one-dimensional PhC will only perceive the band gap if its direction is perpendicular to the plane of periodicity. In other words, were we to create a one-dimensional PhC with a set of forbidden frequencies, these frequencies would only get filtered out if the direction of the incoming wave was in correlation with the periodicity of the structure.

In many applications there is a clear benefit to having your structure omni-directional. Only requiring periodicity in one dimension sensibly reduces manufacturing dif-

faculty, and creating higher-dimensional photonic crystals is always a challenge. Nonetheless, manufacturing methods are continuously being improved for both two- and three-dimensional PhC's [16–18]. As a result, PhC's of higher order dimensionalities are readily finding their way into photonic-based componenets.

Our component of interest is a waveguiding structure, and for this purpose a two-dimensional photonic crystal structure will suffice. Accordingly, most of our research will be on two-dimensional photonic crystals. One-dimensional photonic crystals will however also be considered and used for illustrative purposes. The main goal of this thesis is therefore to design the ideal PhC waveguiding structure, and then lay the groundwork for also fabricating such a structure.

It must be noted at this point that most simulations and calculations were done by the writer as a part of the project thesis [49]. Simulations finding the ideal parameters for a two-dimensional photonic crystal structure have already been conducted and are reiterated in this thesis to make it more complete. In other words, the correct design parameters for the photonic crystal waveguide have already been calculated and subsequently verified through simulations and mathematics. However, the issue of properly fabricating such a structure with the desired parameters is still at hand. In this thesis every step in the fabrication process is being done from scratch, apart from the construction of the actual silicon wafer. The PhC in itself is to be created on top of an Silicon-on-insulator (**SOI**) wafer. Many alternatives exist, but fabricating the PhC on an SOI wafer is useful since it is compatible with CMOS technology. Furthermore, signals in the infrared wavelength range, more specifically at $1.55\mu\text{m}$, experience very low losses when being guided through silicon. Equally important is the high refractive index contrast between silicon and silicon dioxide, which allows for good confinement of light.

SOI wafers can be purchased from experienced manufacturers, but are both expensive and, according to hearsay, of varying quality. The surface roughness of commercially purchased SOI wafers is a problem. Furthermore, a photonic crystal waveguide requires a very specific thickness of the amorphous silicon layer. As a result the decision was made to create the SOI wafer manually. This process was carried out through plasma enhanced chemical vapor deposition (**PECVD**). The imaging of the actual photonic crystal pattern along with the coupling section was done by electron-beam lithography (**EBL**), and the actual etching of the PhC was done by a dry-etch process called inductively coupled plasma reactive ion etching (**ICP-RIE**).

It is also important to point out that at the start of of this thesis, the author had next to no knowledge surrounding fabrication methods. As a result, much of the work being done in the first few months mostly revolved around taking courses to learn how to use the cleanroom equipment. Since the equipment had not previously been used to for the fabrication of optical components, the usefulness of cleanroom courses was limited. Methods to use the equipment available to fabricate optical components had to be developed, since no established method existed upon the beginning of this thesis.

In order to create a fully functioning structure in the nano- or micrometer regime, there is no room for erroneous process parameters and even the slightest of errors will lead to a waveguide not working or the coupling between the PhC and the waveguide to malfunction. In other words - to actually create a working structure, not only must every single piece of cleanroom equipment intended to be used be learned, it must be mastered.

As briefly touched upon, the NTNU cleanroom is a multi-user environment, but most progressed work at NTNU is not done with silicon, let alone within photonics. Even though the equipment in the cleanroom was fully capable of constructing PhC's, it had not yet been done before. Consequently, many issue arose and many questions had to be answered before such a problem could be solved. As a result, every single piece of work being done in the cleanroom has not only been done for the benefit of this thesis. Any knowledge gained will also be documented for the common knowledge of the cleanroom for future generations of cleanroom users, such that they will not make the same mistakes that were done during this research.

Chapter 2

Theory

2.1 Calculating the band diagram of a one-dimensional photonic crystal

Not much time will be devoted to cover the more complicated mathematical procedures involved in studying photonic crystals, as this is not the focal point of this thesis. Excellent literature regarding the study of photonic crystals exist [11], but in this thesis only a short section detailing how to calculate the band diagram of a one-dimensional photonic crystal is included. The band-diagram of a photonic crystal show the range of frequencies that are allowed to propagate within the crystal. The band-diagram therefore show the determining properties of a PhC. We acquire the band-diagram by finding the relation between frequencies and wave vectors for a one-dimensional PhC. Consider the PhC of figure 2.1.1.

From Maxwell's equations we know that the electric and magnetic fields are parallel to the boundaries, as well as continuous across them. Furthermore, as an analogy with electrons moving through a solid [10], we for now postulate that the field must be in the form of a Bloch wave $u(z)e^{ik_z z}$ where $u(z) = u(z + d)$.

We can then write the fields as

$$\mathbf{E}(z) = \vec{u}_e(k)e^{ik_z z} \text{ and } \mathbf{H}(z) = \vec{u}_h(k)e^{ik_z z}. \quad (2.1)$$

We also have that \mathbf{H} is perpendicular to \mathbf{E} , as well as proportional to the derivative of \mathbf{E} , so $\mathbf{H}(z) \propto \mathbf{E}'(z)$. We define the E-field to point in the x-direction and the H-field to point in the y-direction (ref figure 2.1.1). This allows us to drop the vector notation. From Maxwell's equations we then find the wave equation for the electric field




Figure 2.1.1: 1D photonic crystal with refractive indices n_1 and n_2 , and periodicity $d_1 + d_2 = d$. The electric and magnetic components of the incident wave are both parallel to all boundaries whereas the wave vector is perpendicular.

$$\frac{d^2 E(z)}{dz^2} - \frac{\omega^2}{c^2} \varepsilon(z) E(z) = 0. \quad (2.2)$$




Figure 2.1.2: 1D photonic crystal with refractive indices n_1 and n_2 , and periodicity $d_1 + d_2 = d$. The electric and magnetic components of the incident wave are both parallel to all boundaries whereas the wave vector is perpendicular.

Let the electric field in layer one and two be $E_1(z)$ and $E_2(z)$, respectively, with wave-vectors k_1 and k_2 . General solutions of equation (2.2) are then (ref figure 2.1.2)

$$E_1(z) = E_{+1} e^{ik_1 z} + E_{-1} e^{-ik_1 z} \quad (2.3)$$

$$E_2(z) = E_{+2}e^{ik_2z} + E_{-2}e^{-ik_2z}. \quad (2.4)$$

We rewrite equation (2.3) by letting $E_{+1} = (\frac{A_1}{2} + \frac{B_1}{2i})$ and $E_{-1} = (\frac{A_1}{2} - \frac{B_1}{2i})$, which gives us

$$E_1(z) = A_1 \cos(k_1z) + B_1 \sin(k_1z) \quad (2.5)$$

and its derivative then being

$$E_1'(z) = -k_1A_1 \sin(k_1z) + k_1B_1 \cos(k_1z). \quad (2.6)$$

We then introduce some boundary conditions for the system,

$$\mathbf{E}_0 = \begin{bmatrix} E_0 \\ E_0' \end{bmatrix}$$

and

$$\mathbf{E}_1 = \begin{bmatrix} E(d_1) \\ E'(d_1) \end{bmatrix}.$$

From these we find the constants A_1, A_2, B_1 and B_2 . We eliminate these and get

$$E_1(z) = E(0)\cos(k_1z) + \frac{E'(0)}{k_1}\sin(k_1z) \quad (2.7)$$

and

$$E_2(z) = E(d_1)\cos(k_2(z - d_1)) + \frac{E'(d_1)}{k_2}\sin(k_2(z - d_1)). \quad (2.8)$$

From Maxwell's equations we know that the fields are continuous across the boundaries, i.e. they are continuous at $z = 0, d_1, d_2$. As a result, we must have that $E_1(d_1) = E_2(d_1)$ and $E_1'(d_1) = E_2'(d_1)$. Equation (2.1) tells us that the field are Bloch waves. Hence, $E_2(d) = e^{ik_z d}E(0)$ and $E_2'(d) = e^{ik_z d}E'(0)$.

To simplify the equations we express these in matrix notation

$$\begin{bmatrix} \cos(k_1d_1) & \sin(k_1d_1)/k_1 \\ -k_1\sin(k_1d_1) & \cos(k_1d_1) \end{bmatrix} \begin{bmatrix} E(0) \\ E_0' \end{bmatrix} = \begin{bmatrix} E(d_1) \\ E'(d_1) \end{bmatrix}$$

$$\mathbf{M}_1\mathbf{E}_0 = \mathbf{E}_1 \quad (2.9)$$

$$\begin{bmatrix} \cos(k_2 d_2) & \sin(k_2 d_2)/k_2 \\ -k_2 \sin(k_2 d_2) & \cos(k_2 d_2) \end{bmatrix} \begin{bmatrix} E(d_1) \\ E'(d_1) \end{bmatrix} = e^{ik_z d} \begin{bmatrix} E(0) \\ E'(0) \end{bmatrix}$$

$$\mathbf{M}_2 \mathbf{E}_1 = e^{ik_z d} \mathbf{E}_0. \quad (2.10)$$

Inserting equation (2.9) into equation (2.10) we find

$$\mathbf{M}_{tot} \mathbf{E}_0 = e^{ik_z d} \mathbf{E}_0 \quad (2.11)$$

which is an eigenvalue problem, with solutions when $\det(\mathbf{M}_{tot} - I\lambda) = 0$. Solving the eigenvalue problem is rather straightforward. Since we are working with many terms and the solution process becomes very disorganized, we simplify by defining

$$\begin{aligned} c_{1,2} &= \cos(k_{1,2} d_{1,2}) \\ s_{1,2} &= \sin(k_{1,2} d_{1,2}). \end{aligned}$$

This gives

$$\mathbf{M}_{tot} = \mathbf{M}_1 \mathbf{M}_2 = \begin{bmatrix} c_1 & \frac{s_1}{k_1} \\ -k_1 s_1 & c_1 \end{bmatrix} \begin{bmatrix} c_2 & \frac{s_2}{k_2} \\ -k_2 s_2 & c_2 \end{bmatrix}$$

$$\mathbf{M}_{tot} = \begin{bmatrix} c_1 c_2 - \frac{s_1 s_2 k_1}{k_2} & \frac{c_2 s_1}{k_1} + \frac{c_1 s_2}{k_2} \\ -c_1 s_2 k_2 - c_2 s_1 k_1 & c_1 c_2 - \frac{s_1 s_2 k_2}{k_1} \end{bmatrix}.$$

The eigenvalue problem for a 2x2 matrix has the following form:

$$\det \begin{bmatrix} a_{11} - \lambda & a_{12} \\ a_{21} & a_{22} - \lambda \end{bmatrix} = 0$$

$$\lambda^2 - \lambda(a_{11} + a_{22}) + (a_{11}a_{22} - a_{12}a_{21}) = 0.$$

Inserting the values from the M_{tot} matrix and simplifying we obtain

$$\lambda^2 - \lambda \left[2\cos(k_1d_1 + k_2d_2) - \frac{(k_1 - k_2)^2}{k_1k_2} s_1s_2 \right] + 1 = 0 \quad (2.12)$$

with $e^{ik_zd} = \lambda$ inserted into (2.12) we obtain

$$e^{i2k_zd} - e^{ik_zd} \left[2\cos(k_1d_1 + k_2d_2) - \frac{(k_1 - k_2)^2}{k_1k_2} s_1s_2 \right] + 1 = 0 \quad (2.13)$$

$$e^{ik_zd} - \left[2\cos(k_1d_1 + k_2d_2) - \frac{(k_1 - k_2)^2}{k_1k_2} s_1s_2 \right] + \frac{1}{e^{ik_zd}} = 0 \quad (2.14)$$

$$(2.15)$$

$$\cos(k_1d_1 + k_2d_2) - \frac{(k_1 - k_2)^2}{2k_1k_2} s_1s_2 = \frac{e^{ik_zd}}{2} + \frac{1}{2e^{ik_zd}} \quad (2.16)$$

$$\cos(k_1d_1 + k_2d_2) - \frac{(k_1 - k_2)^2}{2k_1k_2} \sin(k_1d_1)\sin(k_2d_2) = \cos(k_zd) \quad (2.17)$$

Since $k_{1,2} = \frac{\omega}{c_{1,2}} = \frac{\omega n_{1,2}}{c_0}$ and introducing $k_{norm} = k_zd$ and $\omega_{norm} = \frac{\omega d}{c_0}$ we transform the above equation:

$$\cos(k_{norm}) = \cos\left(\omega_{norm} \left(\frac{n_1d_1}{d} + \frac{n_2d_2}{d}\right)\right) - \frac{(n_1 - n_2)^2}{2n_1n_2} \sin\left(\frac{\omega_{norm}n_1d_1}{d}\right) \sin\left(\frac{\omega_{norm}n_2d_2}{2}\right). \quad (2.18)$$

This is the dispersion relation for a one-dimensional photonic crystal. By using equation 2.18 the band diagram in figure 2.1.3 was obtained.

Evident from figure 2.1.3 are the gaps in allowed frequencies in the solution, as expected from the theory. Although they are rather straightforward, the band diagram




Figure 2.1.3: The band diagram of a one-dimensional PhC found by sweeping the dispersion relation in MATLAB. Its values are $n_1 = 1$, $n_2 = 3.48$ and $d_1 = d_2$.

calculations can be rather tedious, even for this one-dimensional case. For two and three dimensions they become exponentially more complicated. Fortunately, powerful simulation tools make easy the process of calculating and optimizing these. The simulations were performed through a simulation software called the MIT Photonic Bandgap Package (**MPB** [1]) and subsequently verified through simulations in a powerful simulation software package called COMSOL [19].

2.2 Symmetry considerations

In studying Photonic Crystals there are many simplifications that can be made based on symmetry arguments. As a result, we will in this section find the various symmetries of the PhC structure and see what simplifications can be made on account of these.

In the case of the one-dimensional PhC we only considered incoming waves that were perpendicular to the surface. This ensures that all waves experience the periodicity of the crystal. In contrast, any incoming wave that does not see a periodic refrac-

tive index does not adhere to the properties of the PhC. Spoken more plainly; the properties of the crystal only hold true for waves incident on the structure within a certain range of angles, severely limiting its properties. As a result, one-dimensional photonic crystals are not always of the greatest interest for several applications. This problem can be circumvented by extending the periodicity of the PhC into yet another dimension. There are numerous ways of doing this, such as having a periodic array of squares. However, the common way of extending photonic crystals into the second dimension is by creating either dielectric rods in air, or air holes in a dielectric. Considering a structure with such periodicity in two dimensions, it is self-evident that many simplifications can be made based on symmetry arguments.

Consider figure 2.2.1, showing two modes and their corresponding field distributions for a silicon rod in air. This particular structure shows several symmetries, most notably its mirror symmetry. It is evident that the structure geometry is invariant under a symmetry operation $\mathbf{r} \rightarrow -\mathbf{r}$. From this we can conclude that $\mathbf{H}(\mathbf{r})$ and $\mathbf{H}(-\mathbf{r})$ is the same mode, even though their field distributions might differ. This can also be shown by considering that any mode scaled with an arbitrary constant α is still the same mode, as its frequency would be left unchanged. If $\mathbf{H}(\mathbf{r}) = \mathbf{H}(-\mathbf{r})$ then $\mathbf{H}(-\mathbf{r})$ must simply be a multiple of $\mathbf{H}(\mathbf{r})$. Since inverting the mode distributions of figure 2.2.1 twice will return the modes to their exact initial field distributions, we see that $\mathbf{H}(-\mathbf{r}) = \alpha\mathbf{H}(\mathbf{r})$ where α is either equal to 1 or -1. This allows us to classify modes based on what happens during a single inversion. If the mode distribution is invariant under inversion, $\mathbf{H}(-\mathbf{r}) = \mathbf{H}(\mathbf{r})$, and we call the mode **even**. If it becomes its own opposite, $\mathbf{H}(-\mathbf{r}) = -\mathbf{H}(\mathbf{r}) \rightarrow -\mathbf{H}(-\mathbf{r}) = \mathbf{H}(\mathbf{r})$ we call it **odd**.




Figure 2.2.1: Two eigenfrequency solutions for a silicon rod in air.

We expand the example of a simple cell to a group of cells, and find that further

simplifications can be made on account of symmetries. Figure 2.2.2 shows a square lattice of silicon rods in air. Such a structure has no continuous symmetries, but do possess discrete ones in both the x- and y-directions. We define primitive lattice vectors $\mathbf{a}_1 = a\hat{\mathbf{x}}$ and $\mathbf{a}_2 = a\hat{\mathbf{y}}$ where the constant a is known as the lattice constant and is the length between two neighboring rods or holes. These vectors are defined such that from any position, moving any integer number of lattice vectors leaves the crystal invariant to the observer. This is known as translational symmetry. Evidently, position is only important within the confines of a single cell (since every cell is identical), and not within the confines of the entire crystal itself. From this we can conclude that for a crystal that is of infinite extent, we only need to consider a single cell as long as the crystal is periodic in all directions. In figure 2.2.2 this single cell is just one rod, and it is known as the unit cell.




Figure 2.2.2: Example structure of a square lattice of silicon rods in air, with given lattice constant a and radius r . Imagine the crystal to be infinite in extent in all dimensions.

The translational symmetries indicate that the operator $\hat{\Theta}$ must commute with all translation operators $\mathbf{R} = l\mathbf{a}_1 + k\mathbf{a}_2$. The modes of $\hat{\Theta}$ are thus eigenfunctions of the translation operators as well.

$$\begin{aligned}\hat{T}_{\mathbf{r}} e^{ik_y y} &= e^{ik_y(y-la)} = (e^{-ik_y la}) e^{ik_y y} \\ \hat{T}_{\mathbf{r}} e^{ik_x x} &= e^{ik_x(x-ma)} = (e^{-ik_x ma}) e^{ik_x x}\end{aligned}\tag{2.19}$$

We classify modes by first specifying the wavevectors k_y and k_x . According to equation (2.19), a mode with wave vector k_y and one with $k_y + 2\pi/a$ have the same eigenvalues. For both k_y and k_x all modes with wave vectors $k_y + m(2\pi/2)$ and $k_x + n(2\pi/2)$ form a degenerate set (where degeneracy is defined as two states or modes having the same energy). It follows that multiplying k_y or k_x by an integral multiple of $b = 2\pi/a$ leaves it unchanged. We thus call $\mathbf{b}_1 = b\hat{y}$ and $\mathbf{b}_2 = b\hat{x}$ the primitive reciprocal lattice vectors, since any linear combination of these degenerate eigenfunctions is yet another eigenfunction with the same eigenvalue.

The discrete periodicity in the x- and y-directions leads to a dependence for \mathbf{H} that is simply the product of a plane wave with a x- and y-periodic function. That is, we have a plane wave modulated by a periodic function resulting from the periodic lattice, much in analogy with solid state physics. The result from equation (2.20) is known as Bloch's theorem. The form of \mathbf{H} is in mechanics referred to as a Floquet mode and in solid state physics as a Bloch state, which is the conventional name utilized in PhC literature.

$$\begin{aligned}\mathbf{H}_{\mathbf{k}}(\mathbf{r}) &= \sum_m \mathbf{c}_{k_x, m}(z) e^{i(k_y + mb)y} \cdot \sum_n \mathbf{c}_{k_y, n}(z) e^{i(k_x + nb)x} \\ &= e^{ik_x x} e^{ik_y y} \cdot \sum_m \mathbf{c}_{k_x, m}(z) e^{imbx} \cdot \sum_n \mathbf{c}_{k_y, n} e^{inby} \\ &= e^{i\mathbf{k} \cdot \mathbf{r}} \mathbf{u}_{\mathbf{k}}(\mathbf{r})\end{aligned}\tag{2.20}$$

Resulting from the translational symmetries it is evident that a Bloch state with wave vectors k_x and k_y is equal to a Bloch state with wave vectors $k_x + mb$ and $k_y + nb$. Since there is no physical difference between the two, we can assume that the mode frequencies are also periodic in (k_x, k_y) : $w(k_x, k_y) = w(k_x + mb, k_y + nb)$. Thus we only need to consider k_y and k_x in the range $-\pi/a < k_y \leq \pi/a$ and $-\pi/a < k_x \leq \pi/a$. This region of wave vector values will be referred to as the Brillouin zone. The Brillouin zone is depicted in figure 2.2.3 and is defined as the smallest area of the crystal from which we can reconstruct the entire crystal by translation only. As we shall see, the Brillouin zone can be reduced further.




Figure 2.2.3: The left picture depicts the reciprocal lattice and the reciprocal lattice vectors. The right picture denotes the Brillouin zone of said lattice.

All the information surrounding the modes can be extracted by the wave vector \mathbf{k} and the periodic function $\mathbf{u}_{\mathbf{k}}(\mathbf{r})$. For every value of \mathbf{k} we expect to find an infinite set of modes with discretely spaced frequencies. We label every set of modes by a band index n . By inserting the Bloch state into the master equation [11], we get a new eigenvalue problem with a new Hermitian operator that is now dependent on \mathbf{k} :

$$\hat{\Theta}_{\mathbf{k}}\mathbf{u}_{\mathbf{k}}(\mathbf{r}) = (\omega(\mathbf{k})/c)^2\mathbf{u}_{\mathbf{k}}(\mathbf{r}) \quad (2.21)$$

where $\hat{\Theta}_{\mathbf{k}}$ is the new hermitian operator

$$\hat{\Theta}_{\mathbf{k}} \stackrel{\text{def}}{=} (i\mathbf{k} + \nabla) \times \frac{1}{\varepsilon(\mathbf{r})} (i\mathbf{k} + \nabla) \times . \quad (2.22)$$

Since \mathbf{k} now enters as a continuous parameter in $\hat{\Theta}_{\mathbf{k}}$, we can expect the frequency of each band n to vary continuously as the vector \mathbf{k} varies. This is in good agreement with the band diagram for the one-dimensional PhC, shown in figure 2.1.3.

We now consider the rotational symmetries of the PhC. For translational symmetries it was shown that any translation upon a given mode $\mathbf{H}(\mathbf{r})$ leaving the crystal and thereby also the mode invariant, would satisfy the master equation. The same can be shown for rotation, inversion, and mirror symmetries. We have

$$\omega_n(\mathfrak{R}\mathbf{k}) = \omega_n(\mathbf{k}) \quad (2.23)$$

where \mathfrak{R} is any rotational operation that leaves the crystal unchanged. Taking the rotational symmetry into account, it is evident that the Brillouin zone of figure 2.2.3 can be reduced further. This leads us to the irreducible Brillouin zone, as depicted in figure 2.2.4 for a simple square lattice.




Figure 2.2.4: The irreducible Brillouin zone of the square lattice shown in figure 2.2.2, the points Γ , χ and M are used to classify wave vectors.

Let us digress for a moment to once again briefly examine mirror symmetry. Most of the figures have been in two dimensions, though in reality they extend infinitely into the third dimension. It was briefly mentioned how mirror symmetries allowed us to separate the modes into **even** or **odd**. As we are working with electromagnetic waves, we can define these in terms of their field polarization. If we define the third dimension to be in the z -direction, we have that the operation $\hat{\mathbf{z}} \rightarrow -\hat{\mathbf{z}}$ is always a symmetry for the crystal as long as we are in the xy plane. It follows that $M_z \mathbf{k}_{\parallel} = \mathbf{k}_{\parallel}$ for all wave vectors \mathbf{k}_{\parallel} in the two-dimensional Brillouin zone.

Considering only two-dimensional PhC's we can then classify the modes into two distinct polarizations. The first case is when \mathbf{H} is perpendicular to the mirror plane (E_x, E_y, H_z), in which case the electric field is confined in the xy -plane and the mode is known as a transverse-electric (**TE**) mode. The other case has \mathbf{E} perpendicular to the mirror plane (H_x, H_y, E_z) and is known as a transverse-magnetic (**TM**) mode. On the topic of modes it is important to be able to distinguish between **TE** and **TM** modes for various areas of photonics. The definitions presented here are used in accordance with most photonic crystal literature, but are different in other areas of photonics.

2.3 Trigonal lattice

We diverge from the symmetry discussion to briefly describe another popular geometry for PhC's. The symmetry considerations thus far mostly concerned the square lattice of figure 2.2.2. The other geometry that will be studied is the Trigonal lattice, depicted in figure 2.3.1.




Figure 2.3.1: A PhC with a triangular lattice on the left, with its irreducible Brillouin zone depicted on the right.

This lattice geometry is of great importance, as it exhibits several distinct properties that its square counterpart does not, as will be demonstrated.

Chapter 3

Method

3.1 Finding the ideal photonic crystal waveguide parameters

As is evident from the very nature of photonic crystals, a one-dimensional PhC would only be useful for the construction of a straight waveguide, as its property of interest only holds true for one dimension. In a real scenario, one would like the waveguide to be able to bend, since its main region of interest is for integrated optics. To create a two-dimensional photonic crystal one needs only extend the periodicity of the refractive index into yet another dimension. This can be done in a number of ways. The most obvious would be to simply introduce yet another plane of periodicity, by forming squares of varying refractive index, as shown in figure 3.1.1.

This design is not ideal when it comes to fabrication. In order to fabricate such a PhC you would have to create very anisotropic square pillars, a feat which is far from easy. In contrast, it is much easier to obtain the periodic refractive index by simply "drilling" holes of air in a material of higher dielectric constant, for example silicon. This is why the photonic crystal of interest will look as illustrated in figure 2.3.1.

Since the ideal photonic crystal has a high contrast between the respective dielectrics, it is natural to assume that one of the materials will be air. Hence, two options exist: One can either design rods of high dielectric constant situated in air, or the inverse; holes of air situated in a dielectric substrate. Again, the fabrication of the photonic crystal will decide its design, and since rods are much harder to fabricate than holes, the air-holes in a substrate design is chosen.




Figure 3.1.1: Two-dimensional photonic crystal using square rods to achieve the needed periodicity.

Two pieces of software were used to analyze photonic crystals for waveguiding. A program called MPB (**MIT photonic bands**) was used to analyze the band diagrams for various configurations of photonic crystals. Meanwhile, a finite-element approach simulation program called COMSOL was used to verify the waveguiding properties of such a crystal. MIT photonic bands is a linux-based program developed by MIT to study namely photonic crystals. The program utilizes a scripting language called CTL to design PhCs of various shapes and sizes and runs simulations on them to verify how they would react to incoming electromagnetic waves. A short example on how to program and test a photonic crystal in MPB is given in appendix C.

After using the MPB to find the ideal band diagram for a PhC waveguide, the results were verified using a different simulation program, namely COMSOL. Unlike MPB, COMSOL has a graphical user interface which lends itself to verifying results and making them more easily presentable than in the case of MPB. An example of how to use COMSOL to simulate a photonic crystal is given in appendix D.

3.2 Fabrication equipment

3.2.1 Deposition methods

Two methods were employed to deposit thin films.

PECVD

The first method is PECVD. The PECVD in the NTNU cleanroom was the Plas-malab System 100-PECVD [28]. The PECVD is a powerful tool utilized to grow many different materials, most importantly SiO_2 and amorphous silicon (a-Si). The machine consists of a process chamber and a loadlock used to load a sample without contaminating the chamber. During the PECVD procedure, the process chamber is filled with various gases, the concentration of which is specified by the user. By use of an RF frequency generated within the chamber, these gases form a plasma that slowly deposits a layer on top of the sample inside the chamber. The material deposited as well as its quality is highly dependent on the gases in the reaction chamber, as well as their respective ratios. Many other parameters also play their part, such as the RF power that strikes the plasma, as well as the pressure and temperature within the chamber.

Not all recipes are compatible with the PECVD in the NTNU cleanroom. Various PECVD machinery use both different RF frequencies as well as having a different set of gases available. Furthermore - both gas flows and RF powers within the chamber are limited for the system. The available gases as well as the maximum gas flow rates for these gases in the PlasmaLab System 100 PECVD is given in table 3.2.1.

| Gas | Gas flow rate |
|----------------------------|---------------|
| SiH_4 | 50 sccm |
| NH_3 | 50 sccm |
| N_2O | 1000 sccm |
| N_2 | 2000 sccm |
| Ar | 2000 sccm |
| CF_4 | 500 sccm |
| 10% PH_3/Ar | 50 sccm |
| 0.1% $B_2H_6/0.23\%H_2/Ar$ | 50 sccm |

Table 3.2.1: The available gases as well as their respective maximum flow rates available in the NTNU cleanrooms PECVD system.

Another important characteristic of the PECVD is how films are deposited on a sample. The deposition is unfortunately not uniform, as illustrated in figure 3.2.1.




Figure 3.2.1: Cross section of a typical PECVD deposition on a structured wafer. The effects are exaggerated for illustrative purposes.

The nonuniformity of the PECVD deposition sets a minimum size for any sample. If the sample is too small, no area will have a constant thickness, which is never ideal for any sample. Instead, a large sample such as an entire wafer can first be processed with the PECVD, then diced into smaller samples afterwards.

Electron-beam evaporator

The other piece of equipment utilized for film deposition was the e-beam evaporator. The e-beam evaporator found in the NTNU cleanroom was the Vacuum Classic 500 [29].

The e-beam evaporator can, unlike the PECVD, deposit metals. Furthermore, the E-beam evaporator gives control of the film thickness down to the angstrom level while depositing. In contrast, the PECVD gives no reading of film thickness and the growth rate must be estimated and then subsequently verified after the process is complete. The e-beam evaporator takes wafers up to 4" diameters in size. For our smaller samples, specific methods of holding the samples in place must therefore be construed. In the process chamber the sample is held upside down over the evaporation reaction. A target anode is bombarded with electrons which causes atoms to shift from the solid phase to the gaseous phase. As the gas of a specific material rises upwards in the chamber, it hits the much cooler sample and deposits on it. The deposition rate can vary between 0.2 angstrom/second to 2.0 angstrom/second, which gives excellent control of the film thickness. Furthermore, the deposition method is uniform, meaning that the deposited edges are 'abrupt' with respect to the wafer surface, as shown in figure 3.2.3.




Figure 3.2.2: The plasmalab system 100 - PECVD in the NTNU cleanroom.




Figure 3.2.3: Deposition style of electron beam evaporation. The deposition is completely uniform, so the thickness of the deposited material is the same on the entire sample.




Figure 3.2.4:

The PFEIFFER VACUUM Classic 500 electron beam evaporator.

3.2.2 Pattern imprinting techniques

As with all fabrication of structures in the micro- or nanoregime, some form of imprinting technique must be employed. Two methods of imprinting patterns were utilized.

MA56 mask-aligner

The first was by standard lithography, performed by the MA56 mask-aligner [30]. A sample is coated with photoresist and loaded onto a stage in the MA56, and held in place by a vacuum. A hard mask featuring the pattern of interest is placed above the sample. The sample and the mask is then aligned so the pattern is placed in the correct position. The areas of the sample that are not covered by the mask are then exposed to light. After exposure, the sample is developed and, depending on the resist, the parts that were not exposed are either removed, or the parts that were exposed are removed.

The MA56 mask aligner is used to fabricate structures in the micrometer regime. It cannot expose samples with deep UV light, so diffraction severely limits the resolution of structures created with this method. It is therefore not applicable for

the fabrication of photonic crystal holes, but since a single MA56 exposure can cover an entire wafer with structures in a matter of minutes, it is very handy to create quick and disposable samples for etching tests.

Although standard photolithography does not possess the needed qualities to fabricate photonic crystals, it is possible through deep UV lithography. This requires a hard mask for the lithography process, which gives little to no possibility of changing the design parameters throughout the fabrication run, should this be needed. If it turns out the initial mask has too large photonic crystal holes in comparison with what is possible given the cleanroom equipment, for example, this would be disastrous with a hard mask. However, since deep UV lithography is able to expose the entire sample at once; subsequent samples would be processed very quickly instead of spending hours on every sample before taking it to developing, etching, and further processing. The decision not to use deep UV lithography was made easy, considering the NTNU cleanroom does not have the needed equipment accessible.

Focused Ion Beam

On the other side of the spectra was the helios NanoLab focused ion beam (**FIB** [31]). The focused ion beam requires in theory no mask at all, as the users themselves focus on specific spots on the sample to create the features required. In other words, one would simply be drilling the holes one by one. Although this method gives a lot of freedom with respect to the design parameters of the structure and allows the user to instantaneously change the structures should it be needed, it is tedious and takes a lot of time. Furthermore, this is a piece of labroom equipment which is highly desirable for a number of experiments and as such has a high amount of users and little availability.

Electron-beam lithography

The second piece of imprinting machinery used was therefore the Hitachi 4300 SEM electron-beam lithographer [32] (**EBL**). The EBL uses a mask defined through software. This makes the mask easy to modify between samples. Unlike deep UV lithography, the EBL does not expose the entire sample at once, but rather splits the sample into several smaller segments, or 'write fields'. For the intents and purposes of this thesis, using the EBL is the perfect tradeoff between effectiveness and flexibility.

The EBL imprints a pattern on a sample by first covering the sample with a photoresist, then exposing it to electrons. The electrons change the chemistry in the

photoresist in such a way that when the sample is later immersed in a solution known as the developer, the exposed areas are removed, while the unexposed areas remain intact. Figure 3.2.5 shows a picture of the EBL workstation.




Figure 3.2.5: Hitachi 4300 SEM with raith laser stage EBL at the NTNU cleanroom.

The EBL is one of the more complicated cleanroom machines to master. A sample is loaded into a process chamber, just like with previous machines. The sample is loaded on a stage, which is controlled by the user. Since the EBL uses an electron beam for patterning, it is also inherently a scanning electron microscope, which allows one to take pictures with very low resolution. Several parameters are important when writing the intended pattern with the EBL. First of all, the beam current must be decided. The beam current effectively decides the amount of electrons that hits the sample in a given amount of time, and is set by adjusting lenses in the optics of the EBL. In general, a high beam current gives a faster write speed but results in lower quality for fine features.

The next parameter that needs to be optimized is the acceleration voltage of the electrons. The acceleration voltage defines the speed at which the electrons hit the photoresist. A high acceleration voltage leads to more scattering of electrons on the sample, which reduces the quality of fine features. A low acceleration voltage on the

other hand might lead to electrons not properly piercing the photoresist. The effect of increased acceleration voltage in for example iron is shown in figure 3.2.6 [46].

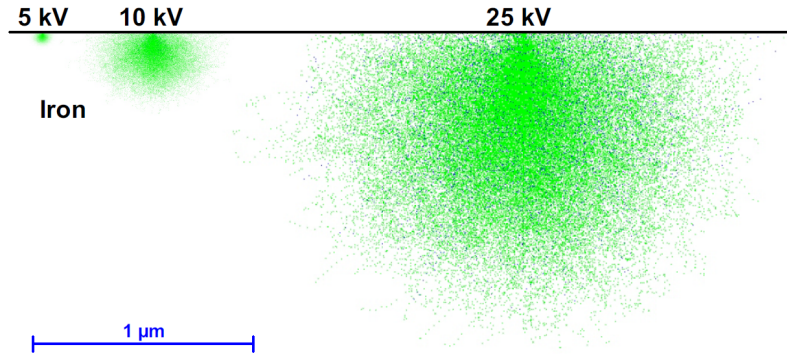


Figure 3.2.6: The effects of increasing the acceleration voltage during an EBL exposure. Higher acceleration voltages lead to more spread in the electrons which reduce the quality of finer features. On the other hand, too low acceleration voltages might prevent the resist from being properly activated.

The dose of current to hit the resist also must be selected before exposure. The dose is given in units of $\frac{\mu\text{As}}{\text{cm}^2}$ and decides to what degree the resist is activated. Choosing a dose which is too low leads to the resist not being activated and therefore the pattern not being formed. Choosing a dose which is too high leads to features in close proximity to each other to be exposed by stray electrons from each others exposures, known as proximity effects. This would lead to an overexposure of the photoresist and would typically cause the structures to collapse during development.

One of the downsides of electron-beam lithography comes from its inability to expose an entire sample at once. Because the EBL splits the sample into several write fields, it is imperative that when moving from one write field to the next, the two write fields are perfectly aligned. If they are not, a phenomenon known as stitching errors occur, an example of which is illustrated in figure 3.2.7. Fortunately, the stitching fields can be aligned in the EBL software, and if the job is done well enough, barely any stitching errors should be observed.

Photoresist

The available photoresist for use with the EBL in the NTNU cleanroom was poly-methyl methacrylate (**PMMA**). The thickness of the photoresist was decided by use of a spin-coater. PMMA is deposited on top of the wafer while it is held in place




Figure 3.2.7: Typical x and y stitching errors in a poorly aligned EBL.

by a vacuum chuck. It is then spun at a speed, time, and acceleration all specified by the user. The speed in which the sample is spinning decides the thickness of the remaining photoresist after the cycle is done. The thickness of the photoresist is important when considering the exposure conditions of the EBL. Furthermore, if the photoresist is being used as an etch mask, it must be thick enough to protect the sample from etching.

3.2.3 Etching techniques

Wet etch

A wet etch is performed by immersing a sample into an etching solution, typically an acid. The acid is selected such that it etches the target material and ideally nothing else. A wet etch has many benefits. For one, since the etch is entirely chemical it is possible to obtain very good selectivities during a wet etch. Wet etches are typically also quite quick, since the user can easily immerse a sample into the etching solution, remove it, and the sample will be ready for the next process step.

The main drawback of a wet etch is its complete isotropy. Isotropy in practice means that the sample is etched in all directions. When creating a structure, it is imperative that the etch is not too isotropic, or features that are close to each other will collapse. Due to isotropy, wet etching can not be employed to define structures that are very close to each other, as they will melt together. For very thin layers, however, a wet etch is in many cases applicable. Figure 3.2.8 shows the isotropy of a wet etch.




Figure 3.2.8: A perfectly isotropic etch. The substrate is etched in both x - and y -directions at the same rate, severely reducing the viability of etching fine structures with such a method. The gold-coloured layer is the etch mask. The blue layer is the target material.

Dry etch

The contrast to the wet etching process is a dry etch. Unlike the wet etch it is not completely chemical, as it etches by shooting ions at the surface of the sample. The gases used for the etching process are selected such that the target material has a higher etching rate than the mask used for the experiment, but since the surface is bombarded with ions, a portion of the mask is also etched. Resultingly, the selectivity is usually not as good as with the wet etch. However, in contrast with a wet etch the dry etching process is mostly anisotropic. In fact, with the correct parameters, an almost perfectly anisotropic etch can be achieved through dry etching processes. Figure 3.2.9 shows an example of a dry, perfect anisotropic etch.




Figure 3.2.9: A perfect anisotropic etch. The sidewalls are completely parallel with each other and vertical to the material surface.

The dry etch is done using the RIE or the ICP-RIE. Both machines were utilized and a large amount of recipes were used. The main difference between the two is that the ICP-RIE achieves higher etching rates and is more anisotropic due to its ability to cool down the process chamber to cryogenic temperatures.

ICP-RIE

During the course of this thesis, both the Plasmalab System 100 RIE [34] as well as the PlasmaLab System 100-ICP [33] were utilized. The ICP gives added power that improves the etching rate of any given recipe, an ability which often improves the quality of the etch. Furthermore, the ICP-RIE has the ability to cool down its chamber to cryogenic temperatures, and a cryogenic etch greatly increased anisotropy [47].

The design and operation of the ICP-RIE is similar to that of the PECVD, complete with a loadlock, a process chamber, and a user interface on a nearby computer. The gases required for the chemical interactions in the chamber, as well as the RF power used to strike the plasma along with the ICP power to guide the etching process are all specified in a recipe constructed by the user before the etching process begins. The recipe chosen is much like the PECVD limited by the available gases and their respective gas flow rates, given in table 3.2.2.

| Gas | Gas flow rate |
|---------|---------------|
| O_2 | 100sccm |
| Ar | 100sccm |
| SF_6 | 100sccm |
| BCl_3 | 50sccm |
| Cl_2 | 100sccm |
| HB_r | 100sccm |
| N_2 | 100sccm |
| CH_4 | 50sccm |
| H_2 | 50sccm |

Table 3.2.2: The available gases as well as their respective maximum gas flow rates available in the NTNU cleanrooms ICP-RIE system.

3.2.4 Characterization techniques

Several methods were used to characterize the samples.

Refractometer

The first was by use of the F20 filmetric refractometer [35], which served as the cornerstone in most experiments where the thickness of thin films needed to be pinpointed. Its basic principle is to use a technique called spectral reflectance, where

light is shone on a surface and the reflected light is analyzed to give film thickness and optical constants. Refractometry is a quick and easy method of characterizing samples, but has a few weaknesses. For one, this technique makes assumptions about the material it is trying to characterize, which means the user must input the theoretized thickness. If one does not have the faintest idea what the thickness of the layer to be analyzed is, finding the correct value might be difficult. Assumptions are also made regarding the refractive index of the layer to be characterized. This is yet another source of error. One must tell the refractometer to look for SiO_2 on Si for example. In this case, the results will only be completely correct if both layers have the same properties as those given by the instrument. This also means that if a substrate has multiple layers, the error becomes even larger, sometimes to the point where results can no longer be trusted. Nonetheless, the refractometer served as an important tool throughout the course of this thesis.

Profilometer

The next characterization tool was the DEKDAK 150 profilometer [36] whose main uses are for the quantification of the surface roughness of a sample. The process in which the surface roughness is identified constitutes of moving a diamond stylus across the surface for a specified distance with a specified contact force. As the stylus moves across the sample, small variations in its displacement is measured which produces data regarding the surface roughness. Depending on the quality of the equipment, a profilometer can measure variations on a sample surface as small as one nanometer. Equally important, however, is the profilometers horizontal resolution. This parameter depends on the radius of the stylus. The smaller the stylus radius, the smaller distance between separate measurable points. The radius of the DEKTAK 150 profilometer stylus is $12.5 \mu m$, which limits the profilometer resolution considerably when considering extremely small structures. In other words, analyzing the characteristics of constructions in the nanometer-regime is not possible with this profilometer, but it can still give us valuable information regarding surface roughness of various films grown on our samples. More importantly, it is a very useful piece of equipment for measuring step-sizes, an ability which is frequently exploited to measure etching rates of various etching recipes.

A picture of the dektak 150 as shown in figure 3.2.10.

Ellipsometer

The perhaps most powerful tool used for characterizing the dielectric properties of thin films was the RC2-ellipsometer. Ellipsometry is based on measuring the change in polarization upon reflection or transmission. As an incident wave is emitted




Figure 3.2.10: The dektak 150 profilometer found in the NTNU cleanroom

from one end of the ellipsometer (and subsequently linearly polarized through a polarizer), it hits the sample surface and changes polarization depending on the sample thickness as well as its refractive index. The signal then hits a detector which reads the data and compares it to some known model. Once the data is 'fitted' by the user, the correct thickness and refractive index values are found. This gives an accurate reading regarding both the thin films thickness as well as its surface roughness, both of which are qualities of profound importance for micro-photonics.

Scanning transmission electron microscope

The most potent way of studying fine features after an etch was through the S-5500 scanning transmission electron microscope [37] (**S(T)EM**). Although the EBL has a built-in SEM, its usefulness is ever decreasing as its spatial resolution is not good enough to discriminate features in the low nanometer regime. Considering the fact that some of the structures to be studied in this thesis enter the very low nanometer range, combined with the fact that studying the cross-section of photonic crystal holes is of great importance, the scanning transmission electron microscope was included in the wide variety of machinery used.

The S(T)EM combines the fundamentals of both the scanning electron microscope (**SEM**) as well as the transmission electron microscope (**TEM**), hence its name. A sharply focused beam of electrons scans the sample surface in a raster pattern. As the electrons hit the surface sample, some of them become backscattered and some are absorbed and re-emitted as secondary emitted electrons. Both of these incidents give valuable information about the sample. By looking at the backscattered electrons of the sample surface, valuable information regarding the materials in the surface is obtained. Heavy materials will for example shine brighter. Looking at the secondary emitted electrons on the other hand, gives information regarding the shape of the sample surface. When examining a sample with the S(T)EM, the acceleration voltage of the electrons can be chosen, just like with the EBL. The acceleration voltage of the electrons decide how far into the sample the electrons penetrate before being scattered. Using a very high acceleration voltage will illuminate heavier atoms, such as metals, and will to a certain degree make the sample 'transparent'. Using a low acceleration voltage will however mean that electrons simply hit the surface of the sample and scatter back. A low acceleration voltage is therefore more valuable when studying the surface of the sample.

Although the S(T)EM can also be used for transmission microscopy by using very thin samples, this was not of much importance in this thesis. However, the improved spatial resolution over the SEM became of great importance. Also, the S(T)EM

was the only machine in the cleanroom able to produce cross-sectional pictures of samples.

3.2.5 Other cleanroom instruments

A few cleanroom instruments does not fit into the preceding categories.

Scriber

There exists several methods of dicing a wafer into several smaller samples. This can either be done by hand with a pen-like instrument found in the cleanroom, it can be done with a dicing saw, or it can be done with a scriber. The Dynatex DXIII [38] scriber was utilized several times to dice large wafers into as many as 20 usable samples. The scriber utilizes a diamond tip scriber and an impulse bar breaker. The scriber reads a setup created by the user, telling it what force to use, what angle to hit the sample with and how deep it should scribe. After scribing, the sample is broken with the impulse bar breaker.

Plasma cleaner

The Diener Electronics Femto plasma cleaner [39] is a powerful method of cleaning wafer samples. It is in this thesis utilized mainly to clean samples before HF-etching. This is to make the samples more hydrophilic such that they attract the HF solution and therefore etch faster.

3.3 Coupling

Since the ultimate goal is to manufacture a photonic crystal waveguide, being able to verify if the structure is working is essential. The only way to test if such a structure works is, unequivocally, to couple light into it and see what exits on the other side. If light is coupled in, but nothing is measured on the other end, the structure obviously does not work. This in itself sounds easy, but presents a completely different problem. How is light coupled into the photonic crystal waveguide? Unlike electric circuits or optics in the macro range, this is extremely difficult.

There are two main methods available. The first is by simple "butt coupling" [26]. Butt coupling involves directing the laser beam directly to one end of the waveguide, as illustrated in figure 3.3.1. Butt coupling has very stringent requirements in terms of hitting the actual waveguide with the correct angle of incidence. Considering the cross section of the waveguide to be in the micrometer range, hitting it with a

laser is far from easy. Furthermore, butt coupling leads to heavy losses, and it is imperative that the end of the silicon waveguide is completely smooth in order to be successful with this method. In general, butt coupling is seldom used for complete optic components, but rather for characterization if no better option is available.




Figure 3.3.1:

Butt coupling in a photonic crystal waveguide

The other method used for coupling light into an optical circuit is through a grating and an adiabatic taper [27]. Using a grating leads not only to lower losses but is significantly easier to do given that the grating is designed and fabricated properly. The complete structure would then be as in figure 3.3.2. The electromagnetic waves are directed towards the grating with the correct angle such that the light is coupled into the silicon. The taper then directs the light so that it properly couples into the waveguide, and the reverse process occurs as light is coupled out of the structure.

Using a grating and a taper naturally leads to other complications. First and foremost, the grating has a different height than the waveguide and the photonic crystals, which means that a two-step etching process has to be employed. Furthermore, one cannot guarantee that process parameters that lead to good waveguides and photonic crystals will also be good for the gratings. The thesis aims towards optimizing fabrication parameters assuming that a grating will be used to couple the photonic crystal. The logic behind this decision is driven by the fact that creating both a grating and a photonic crystal during the same process run should not lead to a significant increase in work for every experiment. Exposing a grating on a sample only increases the exposure time by a minute or two. Meanwhile, since the entire sample has to go through development, etching and so on, the repercussions of including trenches in every single process run only costs minutes. Meanwhile, having defined trenches on the samples leads to a greater understanding of how the etching recipes




Figure 3.3.2:

Using a grating to couple light into a photonic crystal waveguide.

work, as well as allowing one to optimize for trenches simultaneously with holes.

3.4 The silicon-on-insulator wafer

3.4.1 The oxide layer

The first challenge was to create the actual SOI wafer on which the photonic crystal was to be placed. The plasmalab system 100 - PECVD comes with default recipes for both amorphous silicon as well as silicon dioxide. None of the recipes however mention the qualities of the film deposited or what the respective recipes were initially used for. There is a significant difference between amorphous silicon grown for optics compared to other applications. Not all applications require the amorphous silicon to have a low extinction coefficient and thus low absorption, a feature which is of grave importance for optics.

The first tests focused on characterizing the growth of SiO_2 , by using the recipe given in the cleanroom list of recipes. Since the refractometer in the NTNU cleanroom was made specifically to gauge oxide layers on silicon, the results were very reliable. Table A.0.2 of appendix A shows the standard recipe used for the first characterization, named "Standard SiO_2 recipe". Its growth rate is given as 50 nm/min, but this needs to be verified in order to be trusted.

For a photonic crystal the thickness of the oxide layer does not matter so long as it is thick enough to prevent leakage from the waveguide into the actual silicon wafer. However, the thickness of the oxide layer does matter for the coupling section. The coupling section as demonstrated in section 3.3 consists of a grating and a taper. It has been shown that if the thickness of the oxide layer is chosen carefully, some of the incident electromagnetic waves will reflect off the boundary between the silicon layer of the wafer and the silicon dioxide and couple back into the waveguide [48]. This is a feature that is not important until a functioning PhC waveguide can be manufactured, so at this point making sure the oxide layer is thick enough for the PhC is prioritized. An oxide thickness of $1\ \mu\text{m}$ was chosen as the target thickness. With a growth rate of $50\ \text{nm}/\text{min}$, such a thickness should be achieved by growing for 20 minutes. Several samples were processed around this range, to better estimate the real growth rate.

3.4.2 The amorphous silicon layer

As with the oxide, the cleanroom has a default recipe for amorphous silicon, given in table A.0.1 of appendix A. Finding the thickness of the amorphous silicon layer is more difficult than for the oxide. It is not possible to use the refractometer to measure amorphous silicon grown on top of crystalline silicon. At the very least - the quality of such a measurement would not have a high degree of certainty.

Methods of characterizing the silicon layer therefore had to be developed. The first method attempted was to simply grow silicon on top of the oxide layer already grown on previous samples, then use the refractometer to find the thickness of the silicon layer. Another method attempted in order to measure the growth rates of

the amorphous silicon recipe was to grow the amorphous silicon on glass plates available in the NTNU cleanroom. The benefit of using glass plates is that they are very cheap in comparison to silicon wafers. Where glass plates are practically free, silicon wafers can cost anywhere between 100,- NOK to 1000,- NOK, depending on quality. Furthermore, the hope was that using a glass plate would create a sufficient contrast between the two materials so that refractometry measurements of the amorphous silicon layer would be accurate.

A final approach to measuring the silicon layer thickness was made by using the profilometer. A glass plate was utilized and parts of it was covered during PECVD deposition, such that the profilometer could scan its stylus from the part of the glass with no silicon to the part of the glass with silicon, to find the stepsize. Picture

3.4.1 shows such a glass slide.



Figure 3.4.1: Glass plate gone through PECVD with part of it covered. It is evident that the growth profile is uneven.

These measurements mainly address one issue only - the thickness of the sample. The profilometer also gives the surface roughness, but more parameters need to be characterized, such as the absorption coefficient. However, the equipment needed for such measurements was not available in the cleanroom. Help was instead obtained from a different source with an ellipsometer at hand. Through ellipsometry, all previous measurements were made insignificant. Ellipsometry gives both very accurate readings of thin film thickness as well as surface roughness. In addition, ellipsometry gives an estimate of the silicon extinction coefficient.

In general, ellipsometry is quite accurate for analyzing thin films. However, since no ellipsometer exists in the NTNU cleanroom, the samples had to be removed from the cleanroom environment. This led to a contamination of the samples which reduced the accuracy of the tests. In addition, the samples taken out of the cleanroom were grown on pieces of glass instead of actual silicon wafers, further reducing the accuracy of the extinction coefficient tests.

3.5 Forming the pattern

As previously mentioned the formation of the pattern was done by electron beam lithography. The mask was created in a program called CleWin, and a typical mask pattern is given in figure 3.5.1. The blue area is the area that will be exposed during electron beam lithography. The developer used was a solution of isopropanol and de-ionized water. The ratio was one part DI water to nine parts IPA. This has generally been the standard concoction used to develop PMMA.

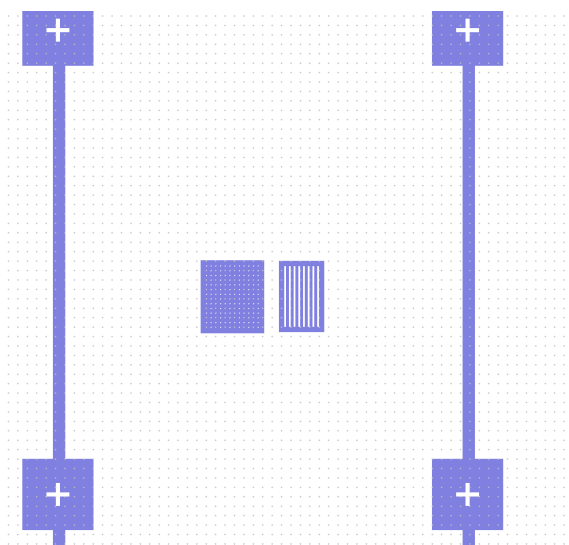


Figure 3.5.1: The mask used for the first process run of a photonic crystal and grating structure.

The mask needs some explanation. It has the photonic crystal holes segregated from the grating structure. The idea was to test various exposure doses as well as various hole radii and groove thicknesses during the same exposure, to see how well the various parameters work after EBL. Figure 3.5.1 only shows one ninth of the entire mask, as it is repeated nine times along the plus-sign marks with different hole radii and groove thicknesses for every structure. The plus-sign marks were there to help locate the structure after exposure and etching, since it is likely too small to see with the naked eye.

Before EBL exposure, the wafer was first cleaned then dehydrated by baking at 200 C° for 20 minutes. A 100 nm thick layer of PMMA was applied by using a spin coater, the parameters of which are found in table 3.5.1.

The first step of 300 rounds per minute is to improve the uniformity of the PMMA.

| | Step1 | Step2 |
|---------------------|--------------|--------------|
| Rotation | 300 rpm | X rpm |
| Acceleration | 1000 rpm/s | 1000 rpm/s |
| Time | 12 s | 60 s |

Table 3.5.1: The standard setup for applying PMMA to the wafer. Vary X to vary the thickness. X=2000 gives a PMMA thickness of ~ 100 nm. X=5000 gives a PMMA thickness of ~ 70 nm.

The PMMA datasheet lists the minimum amount of time at full spin speed to be 45 seconds, which is why the total spin time of 60 seconds was chosen. To change the thickness of the PMMA, simply increase the spin speed of the second step. After the PMMA is applied the sample is post baked at 180 C° for 60 seconds. After the post bake the sample is ready to be taken to the electron beam lithography.

After the first few process runs it became apparent that the best way to find the ideal exposure parameters for the EBL was to hold the acceleration voltage and the beam current constant while varying the dose. Changing too many parameters at once defeats the purpose of changing them at all, since one would not be able to pinpoint the good and bad variables, respectively.

The final parameter that needs to be optimized is the actual electron beam of the EBL. Having a good focus on your sample and a clear beam is paramount to achieving a good exposure and therefore fine structures.

3.6 The etching recipe

Perhaps the most important and most difficult step of creating any structure in the micro- or nano-range is the etching step. Although the actual lithography is also of vital importance, it is in general a one-step process. In contrast, the etching process typically encompasses many steps all of which will likely go wrong during the development process, and sometimes it can be difficult to characterize which of the many steps need to be tweaked. Different types of etches and etching recipes were utilized throughout the course of this thesis. Many of them failed by either showing unwanted properties or poor selectivities with regards to the chosen mask material. This section will briefly explain the various etching methods utilized, how they work, and how they were characterized.

3.6.1 Finding the correct etching recipe and mask

NanoLab standard Si etching recipe and a PMMA mask

The first and easiest way of forming a pattern in silicon, is by simple lithography followed by an etching that utilizes the photoresist as an etching mask. Oxford Instruments, who designed the ICP-RIE and RIE in the cleanroom, supplies to purchasers a library of etching recipes. Most of the recipes used in this thesis are either taken from, or inspired by this library. The first etching method was by a PMMA mask, while employing the etching recipe named "Standard Si etch" found in appendix B. Figure 3.6.1 illustrates this simple etching process.

This method only works if the recipe shows a sufficient selectivity between the PMMA and the silicon in terms of etching rate. In general, the PMMA thickness may be somewhere between 70-120 nm, so to etch 300 nm, which is the bare minimum to create a decent photonic crystal, one needs a selectivity of around 1:3 in terms of PMMA:silicon. The etching rates of the recipe were found by coating a sample with PMMA, then cleaning half of the sample by dipping it in acetone to create a clear step between the PMMA and the silicon. The PMMA thickness was measured with a refractometer after every etch, while a profilometer measured the step between the silicon etch and the PMMA etch, as indicated in figure 3.6.2.

By measuring the PMMA thickness through refractometry, it is easy to find the etching rate of the silicon by measuring the step from the silicon to the PMMA layer. The sample was etched in intervals of 30 seconds and the process illustrated in figure 3.6.2 was undertaken each time.




Figure 3.6.1: Silicon etching by using the photoresist as an etching mask.

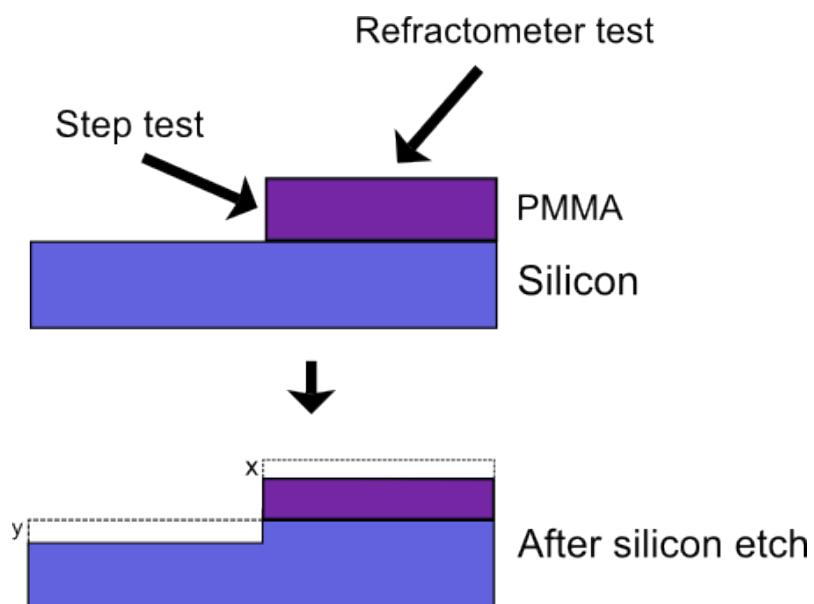


Figure 3.6.2: The characterization method used for finding the PMMA and silicon etching rate of the NanoLab standard Si etching recipe.

Cryogenic etching recipe with an oxide mask

On the other side of the spectrum to the standard Si etching recipe is a cryogenic ICP-RIE etching recipe. ICP cryogenic etching recipes are known for good anisotropy and otherwise good features during etching. The downside of such recipes is that they do not only etch chemically. The bombardment of ions on the sample surface means the mask will also etch quite fast. The recipe chosen was found in the cleanroom library of recipes, and was specifically designed with photonic crystal holes in mind. The recipe is given in appendix B and is labelled "Cryogenic silicon etching recipe". It is a two-step etching recipe that uses the first step to prevent so-called "Notching" [47], whereas the second step is used for the main etch. Notching occurs during an etch as it attempts to pierce from one layer to the next, illustrated in figure 3.6.3.




Figure 3.6.3: Example of notching during the etching process.

Since the recipe has a high ICP power, a PMMA mask will likely not withstand the etching. Another mask must therefore be found. The chosen material was SiO_2 , since the knowledge required to create such a mask was already possessed. To use an oxide mask, the process becomes more complicated than that indicated in figure 3.6.1. The PMMA must serve as an etch mask to define the pattern in the oxide, and the oxide in turn serves as an etch mask for the silicon. This poses a new problem, an etching recipe that etches the oxide without etching away the PMMA must be found. A wet etch was selected as the solution, since good selectivities can be obtained in this manner. The wet etch solution was a hydrofluoric (**HF**) acid mixed with NH_4F . 40% HF and 40% NH_4F in a 1:20 ratio. Every time a sample was etched with HF acid, it was first brought to the plasma cleaner and cleaned thoroughly.

By using a wet etch to open the oxide, the entire process is as shown in figure 3.6.4.




Figure 3.6.4: Using SiO_2 as an etching mask.

Since the wet etch is completely isotropic, this poses a great limitation in the thickness of the mask. If the mask is for example 30 nanometers, this would mean that in order to etch 30 nanometers downwards through the mask, one would also have to etch 30 nanometers in the outward direction. This poses a severe limitation. If two holes for example are within 60 nanometers of each other, a 30 nanometer wet etch would result in the holes completely opening and the mask most likely collapsing, indicated by figure 3.6.5.

In practice, this means that when performing a wet etch to form a mask, the mask should not be thicker than 30 nm. Since a silicon etch of at least ~ 300 nm is ideal, this means that the selectivity between the oxide and the silicon etching recipe chosen needs to be at least 1:10.

Chlorine-based silicon etch and an oxide mask

Most silicon etching recipes in this thesis are SF_6 based. To combat some of the downsides of the SF_6 based etch, a chlorine-based silicon etching recipe was also characterized. The characterization method utilized in section 3.6.1 is not entirely ideal because the photoresist is removed by hand, which makes the step quite uneven and therefore increases the uncertainties.




Figure 3.6.5: Illustration of an isotropic etch failing due to features being placed too close together.

A better method of researching etch rates was therefore developed for these measurements. The best method found was to simply spin a thick layer of photoresist on top of a silicon wafer, then use standard lithography to imprint a pattern in it. The photoresist used for this is called SPR700, and is a rather tough resist used for creating structures in the micrometer regime. Figure 3.6.6 shows the result of such a process.




Figure 3.6.6: Quarter of a 2-inch wafer with a pattern created by standard photolithography in SPR700 photoresist.

The photoresist in question is quite thick ($\sim 1\mu\text{m}$) and will withstand most etching

recipes so long as the etch time is not too long. The sample in question was then etched, typically for 30 or 60 seconds. The photoresist was then washed away with acetone, and the step size of the features was measured with the profilometer. Figure 3.6.7 shows an illustration of the process.




Figure 3.6.7: The figure shows the process of finding the etch rate of silicon with any arbitrary etching recipe. The sample is first etched, giving defined features in the silicon and an unknown decrease of the photoresist thickness. Typically, the etch rate of the photoresist is not known, so the photoresist is removed such that the actual dip in the silicon can be examined through profilometry.

This particular method was used to find the etching rate of the chlorine-based silicon etching recipe, and tended to give good results. The process of using a chlorine-

based recipe to etch the silicon was just as illustrated in figure 3.6.4, only the last silicon etching step was done with a different chemistry.

3.6.2 Opening the mask

The fourth way of defining the structures was by doing a cryogenic etch, like before, only this time with a chromium mask instead of an oxide mask. The reason chromium was not among the first choices as a mask material is because it required additional training in the cleanroom in order to be deposited, even though many groups utilize chromium masks for etching photonic crystals [40, 41]. Using a chromium mask further complicated the etching procedure. Since chromium is very hard, it is difficult to imprint a pattern in it. The solution used to create a chromium mask is shown in figure 3.6.8.

The PMMA layer on the top is used as a mask to imprint the pattern in the oxide. The oxide is in turn used as a mask to imprint the pattern in the chromium. The oxide layer underneath the chromium is used as a stopping layer for the chromium recipe. The recipe to etch the chromium was found in the library of cleanroom recipes, is given in appendix B and is labelled " $Cl_2 C_r$ etch". Since it is chlorine-based, the recipe will, once finished etching through the chromium, also etch the silicon. Since the recipe might not be ideal for silicon etching, the SiO_2 stopping layer serves to prevent the silicon layer being etched by the chlorine-based recipe.

The chromium layer was deposited through electron-beam evaporation. The e-beam evaporator has a crystal in the chamber which resonates at a certain frequency. As film is deposited in the chamber, the weight of the crystal changes and so does the crystals frequency. The frequency allows one to tell the amount of chromium deposited, down to the angstrom level. Ideally, the crystal should be 100 % accurate, but nothing is ever ideal in a cleanroom. As a result an alterior method of verifying the chromium thickness was utilized. Every time a sample was loaded into the e-beam evaporator chamber an auxillary test-sample with a thin piece of tape was brought to the chamber, shown in figure 3.6.9.

The figure shows an SOI wafer mounted on the standard e-beam evaporator sample-holder on the right side. To the left is a small auxillary sample with a thin piece of tape on it. After deposition, the piece of tape is removed and a clear step is defined. Since deposition through electron-beam evaporation is completely uniform, testing the step through profilometry will give very reliable results. The oxide layer used for the mask can either be deposited through PECVD or e-beam evaporation.




Figure 3.6.8: The final solution to etching silicon by using a chromium mask.




Figure 3.6.9: A two-inch wafer mounted onto the e-beam evaporator holder. An auxillary test-sample with a thin piece of tape is seen to the left.

Another important parameter that had to be tested was the resistance of the chromium during the cryogenic etch. Several samples were prepared with a lift-off technique to create chromium structures in silicon. The samples were then etched with different times to find if the chromium was still present.

Chapter 4

Results and discussions

A brief moment will be spared to explain how the results have been presented. Throughout the course of this thesis, the knowledge about the use of cleanroom equipment has steadily increased. Due to a lack of previous experience, there were inevitable failures, but knowledge was gained from every failure to improve the next process run. The development procedure was therefore iterative. Throughout the entire fabrication process, every experiment conducted was usually a consequence of some previously obtained results. To obtain a more apprehensible connection between the results and the experiments conducted, the discussion and results section have been combined, so that the reader can more easily understand the decisions made throughout the fabrication process.

Furthermore, the thesis can be separated into four separate sections, all of which are ultimately intertwined. The first section includes the design of the photonic crystal and simulations showing how such a structure can function as a waveguide. The second section includes the first process that had to be studied, which was the creation of an SOI wafer. The third section includes how to form the pattern of the photonic crystal and grating structures in PMMA on a wafer through electron-beam lithography, and the fourth section entails the etching recipes that ultimately form the structures in the actual silicon.

4.1 The ideal photonic crystal waveguide

Many parameters specifying the photonic crystal structures need to be optimized and a plethora of simulations were run. Among the most important parameters for a two-dimensional photonic crystal were the size of the photonic crystal holes, or more accurately their radius with respect to the fundamental lattice constant a . Furthermore, the actual formation of the lattice as well as the thickness of the

photonic crystal and the surrounding material was determined.

The ideal photonic crystal hole radius was found by simply varying the hole radius while simultaneously analyzing the band diagram. As previously mentioned a proper photonic crystal has a band gap in which no frequencies are allowed to propagate. Although our photonic crystal only really needs the band gap to be large enough to include the 1550 nm wavelength range, a larger band gap will lead to a higher localization of the electromagnetic wave and as such, less losses and a better waveguiding structure. As a result, all testing done aims to maximize the size of the band gap with the 1550 nm wavelength in the middle.

All bandgap testing was done using a program called MIT photonic bands (**MPB**) through a scripting language called CTL which allows one to design any photonic crystal and find its band diagram. Each simulation gives both the band diagram for the TE and TM mode, respectively. As previously mentioned the band diagrams are given as a function of the fundamental lattice constant \mathbf{a} . This way, the ideal photonic crystal can be designed around any given wavelength by simply changing the a parameter. The first test was run with a hole radius of $0.2a$ and the value was subsequently increased to find the ideal hole radius.

Since the ultimate goal of these simulations is to manufacture a photonic crystal waveguide, these simulations were brought to a halt at a maximum hole radius of $r=0.4a$, as beyond this point a further increase of hole radius would complicate the manufacturing process. The two extremes of these simulation results are shown in figures 4.1.1 and 4.1.2, respectively.




Figure 4.1.1: TE and TM band diagram of a square lattice of air holes in silicon with $r=0.2a$




Figure 4.1.2: TE and TM band diagram of a square lattice of air holes in silicon with $r=0.4a$

The graphs show that increasing the hole size also increases the size of the band gap. In other words one would ideally want the holes to be as large as possible. This may however not be feasible and the actual size of the holes will need to be determined once the limitations of the cleanroom equipment are properly gauged. After the ideal radius of the holes was found with respect to the band diagram, the actual geometry of the photonic crystal lattice was studied. As mentioned previously, two popular varieties of the photonic crystal lattice exist. These are the square and the triangular lattice. Figures 4.1.1 and 4.1.2 show results for the square lattice only. The exact same simulations as for the square lattice were carried out for the triangular lattice, and again it was shown that an increased hole size lead to a higher band diagram. More importantly, however, was the significant increase in band gap size for the triangular lattice, as shown in figure 4.1.3.




Figure 4.1.3: TE and TM band diagram of a triangular lattice of air holes in silicon with $r=0.4a$

We can deduce from the figure that there's a significantly sized TE gap for the triangular lattice type PhC at $r = 0.4a$. As a result, all further work was done with

a triangular lattice and a hole radius of $r = 0.4a$.

At this point the ideal conditions for the two-dimensional PhC were more or less determined, but only so long as one considers the PhC to be infinite in extent in the third dimension. Considering this is obviously not true, further considerations must be taken into account. The photonic crystal was thus modelled as a slab surrounded by another material, as shown in figure 4.1.4.




Figure 4.1.4: Three-dimensional model of a two-dimensional photonic crystal. The blue segments are silicon, the white are the air holes and the green segments are the materials surrounding the crystal.

The effect the slab thickness had on the actual band gap of the PhC was then tested by going from a very thin slab thickness to a very thick to see if any trends were observable. The results are shown in figure 4.1.5.

Easily discerned from figure 4.1.5 is the strong correlation between the slab thickness and the size of the bandgap. An ideal slab thickness is observed at around 0.6 of the fundamental lattice constant a . Finally, the ideal material to surround the photonic crystal slab is found. This was done by first assuming a material with $\epsilon = 1$, and then by incrementally increasing the dielectric constant to find its effects on the band diagram properties.

The results indicate that having the photonic crystal surrounded by a material of very low dielectric constant is ideal. In other words, the photonic crystal should be suspended in air. Again, certain limitations are being set by the fabrication limitations. Creating a suspended structure is no easy feat. In fact, it is much more difficult than simply creating a photonic crystal that is asymmetric. As a result, asymmetric photonic crystals were also put to the test. In our case, the asymmetry consisted of a photonic crystal slab situated on top of a layer of SiO_2 with air above it. Two situations were simulated, either with the photonic crystal being shallow etched, or a deep etch. The respective scenarios together with their results are given in figures 4.1.7 and 4.1.8, respectively.




Figure 4.1.5: Size of the band gap versus the slab thickness.




Figure 4.1.6: The band diagram of the PhC slab with a hole radius of $r = 0.4a$ in a triangular lattice surrounded by a material of $\varepsilon = 6$ to the left and $\varepsilon = 1$ to the right.




Figure 4.1.7: Cross-sectional illustration of a shallow etch in an asymmetric photonic crystal to the left. The corresponding band diagram to the right.

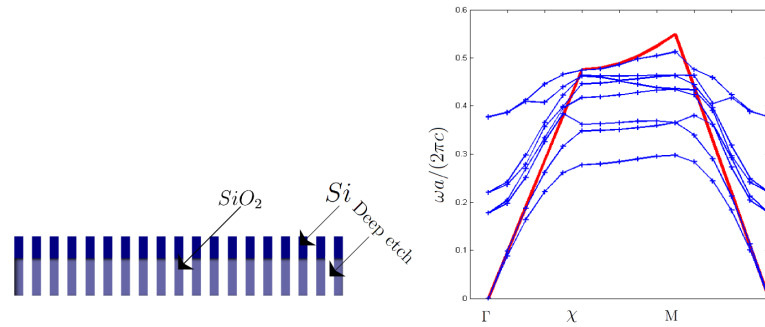


Figure 4.1.8: Cross-sectional illustration of a deep etch in an asymmetric photonic crystal to the left. The corresponding band diagram to the right.

Evident from the figures is the absolute loss of a complete photonic band gap as soon as the symmetry is broken. More notably, the deep etch is shown to be more ideal than the shallow etch, likely because the contrast in effective refractive indices are smaller for the deep etch. The question then remains; If the ideal case of the suspended photonic crystal structure can not be fabricated in the cleanroom in the allotted time, will the photonic crystal waveguide still function?

The results presented in this section were used to simulate a photonic crystal structure in COMSOL. For the wavelength 1550 nm, the fundamental lattice constant a of 465 nm is within the band gap. The height of the slab was given to be $0.6a$, since this was shown to be the ideal for maximizing the band gap. Finally, the holes were given a radius of 200 nm, to further optimize the structure. The structure was designed and tested by COMSOL simulations, which are shown in figure 4.1.9.

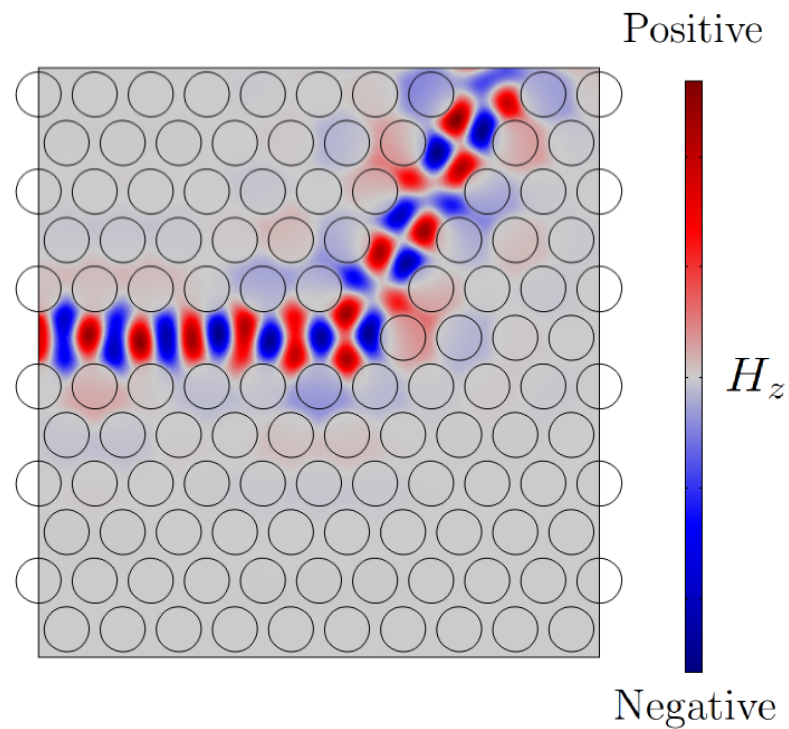


Figure 4.1.9: COMSOL simulations of a PhC functioning as a waveguide for light at 1550 nm wavelength. This particular wavelength is within the band gap of the PhC. Since the light cannot propagate in the photonic crystal, it is instead guided along the removed set of holes, effectively forming a waveguide.

4.2 The silicon-on-insulator wafer

4.2.1 The oxide layer

An SiO_2 and a-Si growth process with good control over the refractive index and growth rate is needed. The SiO_2 recipe used was found in the cleanroom library of recipes, and is given in table A.0.2 in appendix A. The recipe lists a growth rate of 50 nm/min. Since knowing the correct growth rate is of vital importance, this value was not entirely trusted and therefore retested. Since a 1 μ m oxide layer was desired the first sample underwent PECVD for 20 minutes. Subsequent attempts were altered based on the result of the previous ones. Table 4.2.1 show the results and figure 4.2.1 shows the transformation.

| Growth time | Thickness |
|-------------|-------------------|
| 20 min | 1230 \pm 0.5 nm |
| 17 min 10 s | 1060 \pm 0.5 nm |
| 15 min | 930 \pm 0.5 nm |
| 2 min 50 s | 170 \pm 0.5 nm |

Table 4.2.1: Growth times and the respective thickness obtained. The numbers are from repeated measurements at the very center of the wafer.

The results indicate a growth rate much higher than 50 nm/min. After the first process run at 20 minutes which resulted in a thickness of 1230 nm, a hypothesis was made that the growth rate was actually closer to $\frac{1230nm}{20min} = 61.5nm/min$ which was the basis for the second run at 17min 10 seconds which should then result in 1056 nm. The other runs were made in an attempt to verify this growth rate and it was found that it was indeed around 60 nm. An inquiry was made to find why the growth rate found was so much higher than the documented value. According to the cleanroom staff, growth rates increases with time. This is typically caused by uncertainties within the chamber, such as the time it takes for the plasma to strike, and the chamber being conditioned to the gases used in the recipe. The documented growth rate was therefore likely to be used for thin layers of SiO_2 , and not for thick layers as the ones grown in this thesis.

After the growth rate of the SiO_2 was established, the next parameter to be researched was its uniformity. The first uniformity tests were done with the profilometer. The profilometer tip was scanned from one side of the wafer to the other. Based on the PECVD deposition pattern, a result similar to that of figure 4.2.2 was expected.




Figure 4.2.1: Quarter of a two-inch wafer on the left. The same wafer after $1 \mu\text{m}$ of SiO_2 is deposited on the right.




Figure 4.2.2: Based on what's known about PECVD deposition, the expected results of the profilometer measurements should be as above.

However, As the profilometer was utilized to scan across the entire wafer surface, the results shown in figure 4.2.3 were obtained.




Figure 4.2.3: The profilometer measurements obtained by scanning across the surface of a wafer with a $\sim 1\mu\text{m}$ layer of oxide on top.

At a first glance these results seem obscure. Were they to be trusted, they indicate the thickness of the oxide decreasing linearly towards the end of the sample. In an otherwise perfectly calibrated cleanroom this result is confusing. One must remember however that when measuring features in the micrometer range, the slightest particle or speck of dust underneath the wafer will slant it, and the entire measurement will be erroneous. Figure 4.2.4 illustrates how this could occur.




Figure 4.2.4: Typical source of error on the DEKTAK 150 profilometer.

Fortunately, there exists a solution to this problem. The program can 'level' the measurements by the user specifying two points on the sample that possess the

same y-value. With these two points forming the x-axis, the results should then be correct. After leveling, the results can be seen in figure 4.2.5.




Figure 4.2.5: The same measurement as for figure 4.2.3, only this time with the results being software-leveled.

These results indicate that the uniformity of the sample is absolutely horrible and the maximum thickness of the sample is at ~ 650 nm, even though the refractometer showed that the thickness should be ~ 1230 nm towards the middle of the sample. The problem in this case is that even though the sample has been levelled, there still exists too many uncertainties to get a good result. Were the two points chosen for the leveling perfectly level? If they were not, the results would again be erroneous. The profilometer is in general not intended to be used for larger sweeps, as painfully indicated by these results.

Another viable method of finding the film uniformity was by use of the refractometer. The thickness of the oxide layer was measured at many different points on the sample surface to see how much it varied from one point to the next. The results are shown in image 4.2.6.




Figure 4.2.6: The uniformity measurements taken with the refractometer. The left hand image shows an overview of the quarter of the wafer and the points that were measured. The table to the right show the respective thicknesses. All measurements had an uncertainty ~ 0.5 nm on the refractometer.

The refractometer measurements gave better results. Although the uniformity does not seem to be perfect, it still verifies a few things. First, it is evident that the expected PECVD deposition method is as previously assumed. The thickness at the edges (Points A, C, D and B) is higher than towards the middle (E, F and G). Although not many measurements were taken on this sample in particular, it is still evident that the thickness shouldn't vary much more than 5 nm from one point to the next.

The same measurements were carried out for the SiO_2 layers of ~ 1060 nm and ~ 920 nm respectively. The results are shown in figures 4.2.7 and 4.2.8.



Figure 4.2.7: Refractometer measurements for the SiO_2 layer with ~ 1060 nm thickness.

It looks as though the uniformity during these tests is also within acceptable boundaries, but not perfect.

The perhaps biggest downfall of growing the oxide through PECVD is that it is very hard to accurately control the thickness. When an oxide layer as thick as $1 \mu\text{m}$ is needed, this is not a problem. However, for very thin oxide layers such as

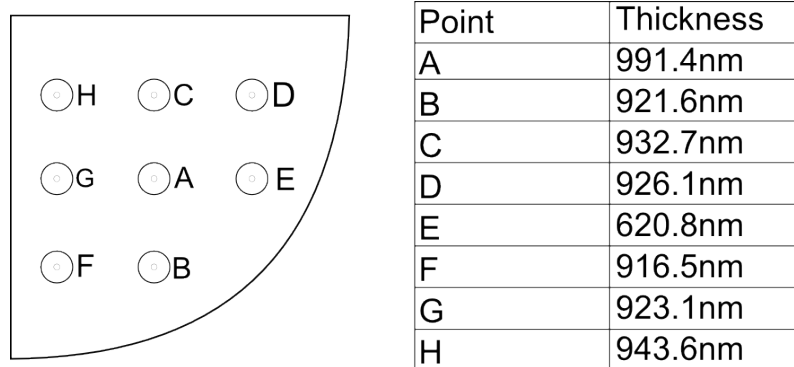


Figure 4.2.8: Refractometer uniformity measurements for the SiO_2 layer with ~ 930 nm.

the ones used for the etch mask, this can be of crucial importance. Since these layers are typically grown on something else than silicon, it more or less invalidates the viability of the refractometer as a method of checking the actual oxide layer thickness. This can be a rather severe problem when wanting to etch through a specific oxide layer, if the thickness is not accurately determined.

4.2.2 The amorphous silicon layer

With the SiO_2 recipe growth rate and uniformity sufficiently verified, the more difficult task of growing the amorphous silicon and characterizing it was undertaken. The PECVD was yet again utilized for the film growth. The first recipe attempted is given in table A.0.1 of appendix A, and was the NanoLab standard recipe. The growth rate given was 25 nm/min. With the ideal slab thickness being $0.6a$ and the fundamental lattice constant a at 1550 nm being ~ 465 nm the resulting slab thickness is $465nm \cdot 0.6 = \sim 280nm$. The target amorphous-silicon thickness was resultingly set to this value. The recipe was grown for 8, 10, and 12 minutes, respectively.

The first amorphous silicon growth was done on a previously processed sample, which already had a 1230 nm layer of SiO_2 grown on it. The silicon recipe was run for 12 minutes. After growth the sample was taken directly to the refractometer for thickness testing. The points measured along with the corresponding results are given in figure 4.2.9 and table 4.2.2, respectively.

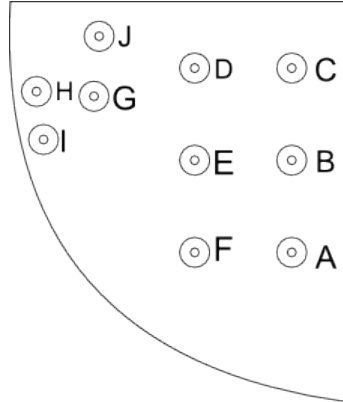


Figure 4.2.9: Refractometer uniformity measurements for the sample with a ~ 1230 nm SiO_2 layer with $a\text{-Si}$ grown for 12 minutes.

| Point | d(Si) | d(SiO_2) |
|-------|----------|---------------------|
| A | 330 nm | 1134 nm |
| B | 330 nm | 1134 nm |
| C | 330 nm | 1134 nm |
| D | 330 nm | 1134 nm |
| E | 330 nm | 1134 nm |
| F | 330 nm | 1134 nm |
| G | 330 nm | 1134 nm |
| H | 330 nm | 1300 nm |
| I | 329.3 nm | 1350 nm |
| J | 329.5 nm | 1134 nm |

Table 4.2.2: Results for the refractometer measurements of the sample with ~ 1230 nm after 12 minutes of $a\text{-Si}$ growth.

The results presented in table 4.2.2 are good, but do not look reliable. The refractometer shows a parameter called the "Goodness of fit" throughout all measurements that tells the user to what degree the measured results align with what the refractometer would expect with the given recipe. During previous measurements the goodness of fit was in the high end of 99%. During these measurements the goodness of fit was much lower, around 60%. This doesn't necessarily mean that something is wrong, but it is a good indication that the results can not be trusted. Considering the uniformity of the SiO_2 layer underneath the amorphous silicon layer (results of which are found in figure 4.2.6, it seems unlikely that the uniformity should improve as more layers are added. In other words, these results are likely incorrect. As a result, the refractometer was discarded as an accurate way of measuring multiple-layered samples. An alternate method would have to be found.

Silicon was therefore grown on glass plates, while covering parts of them to stop the growth, as shown in figure 3.4.1. The same sample could then be utilized for both profilometer and refractometer testing. From the picture, the thin film effect makes it evident that the thickness of the amorphous silicon is highly nonuniform. As a result, only one spot on the entire glass plate is viable for refractometer measurements - the middle spot. In other words, unlike the measurements done on previous samples, only one measurement will be conducted with the interest of finding the thickness at this spot.

Three samples were grown for these measurements, all three used the NanoLab standard amorphous silicon recipe found in table A.0.1. The first sample was grown for 8 minutes, the second sample was grown for 10 minutes whereas the last sample was grown for 12 minutes. The results from measuring the amorphous silicon layer on the three glass plates by refractometry can be seen in table 4.2.3 below.

| Growth time | Thickness |
|--------------------|--------------------|
| 8 minutes | 140 ± 0.1 nm |
| 10 minutes | 180.5 ± 0.1 nm |
| 12 minutes | 220.0 ± 0.1 nm |

Table 4.2.3: The thickness given by the refractometer for the amorphous silicon growth on three samples with the standard NanoLab amorphous silicon recipe

The last attempt to properly measure the thickness of the amorphous silicon layer was by using the profilometer to measure from the parts of the glass with no silicon,

to the middle of the glass which did have silicon, as shown in figure 3.4.1. The results are given in figure 4.2.10.




Figure 4.2.10: Profilometer measurements obtained by scanning the tip from the parts of the glass with no silicon on it to the middle of the sample, where the silicon should be at its most stable. The data has already been leveled.

Once again, the profilometer fails to assess the thickness of the sample. The samples were then taken out of the cleanroom to have their important properties investigated by ellipsometry. The sample which was grown for 12 minutes was mainly studied. The results of the studies are given in table 4.2.4.

From table 4.2.4 it looks like the growth rate is around 19 nm/minute, quite a bit lower than the given value of 25 nm/min. Furthermore it seems as though the refractometer, although slightly erroneous in its measurement, gives a good indication of the growth rate of amorphous silicon. The surface roughness of silicon layer is given to be around 3 nm. This value is a little bit high but can be further improved through electro-chemical polishing [24]. Although the extinction coefficient looks low enough, calculations indicate it is still too high and therefore this amorphous silicon is not applicable for optical components. Since the recipe is simply a standard recipe given in the cleanroom, the intended use of the silicon is not mentioned. It is therefore safe to assume it may not be used for optical applications and therefore not made with the extinction coefficient in mind.

| | |
|---|-----------------------|
| Thickness | 227.87 ± 0.053 nm |
| Surface roughness | 2.92 ± 0.009 nm |
| Extinction coefficient(at 1550.0 nm) | 0.00053899 |

Table 4.2.4: Ellipsometer measurements of the amorphous silicon sample grown for 12 minutes. The measurements for the films grown for 8 and 10 minutes respectively gave similar surface roughness and losses.

A literature study to find better recipes was conducted [20–22] and three recipes were then created and subsequently tested as a result of the findings of the literature study. The recipes are given in table A.0.1. Each recipe was grown for 8 minutes and the results are given in table 4.2.5.

| Variable | Test recipe 1 | Test recipe 2 | Test recipe 3 |
|-------------------------------|-----------------------|-----------------------|-------------------------|
| Thickness | 557.73 ± 0.128 nm | 229.58 ± 0.032 nm | 78.25 ± 0.014 nm |
| Surface roughness | 3.87 ± 0.01 nm | 2.51 ± 0.006 nm | 1.63 ± 0.007 nm |
| Extinction coefficient | 0.003 | 0.0197 | 0.01082 |

Table 4.2.5: Ellipsometer measurements of the other amorphous silicon recipes tested. The extinction coefficient given is for 1550 nm wavelength.

The results are not good. It seems as though none of the amorphous silicon recipes have losses that are acceptably low for optical applications, but time restraints did not allow for subsequent testing of this phenomenon.

The large gap between the reported losses and the ones encountered with these samples might be explained by different cleanroom conditions. Since the NTNU cleanroom is a multi-user cleanroom where every user uses a wide array of gases and materials, the conditions within the chamber can never be said to be the same from one process run to the next, even though one always attempts to make sure they are. In other words - it could be due to some taint in the process chamber that leads to the process not being as reliable as it should which could be the reason why one can not always copy a recipe and expect the same results.

On the other hand - such an error should be rather small. Even though the process chamber of the PECVD may be dirty or in other ways contaminated, one would not

expect several orders of magnitude difference in the process results. Some of the silicon recipes grown through PECVD used recipes which had reported losses below 10 dB/cm. The losses obtained from the extinction coefficients of the ellipsometry measurements show at least 10 times more than this. It is somewhat unlikely that a slight change in the process chamber could lead to such devastating variations. If this were the case - the growth of silicon for integrated optics would not be possible with this specific cleanroom equipment.

The other most likely case is then that the ellipsometry measurements can not necessarily be trusted for finding the loss of the silicon in itself. A spec of dust on the sample when it was removed from the cleanroom for example could invalidate the ellipsometry results. Furthermore, since all amorphous silicon tested by ellipsometry was grown on glass slides, there is a problem with reflections during the ellipsometry. The most precise way of finding the losses in silicon is to fabricate a waveguide in the cleanroom and find the actual losses caused by the silicon through thorough testing of this waveguide [25]. Of course - such a process would then have to be developed in order to verify the losses of the various layers of amorphous silicon. This is unfortunately beyond the scope of this thesis. As a result, any SOI structure to be grown would use an amorphous silicon recipe already verified and tested in literature, even though our tests can not verify that the losses are indeed low enough. At the very least, both the uniformity and growth rate of the recipes were confirmed and found to be within acceptable standards for our applications.

4.3 Forming the pattern

The first attempt at a structure utilized the mask of figure 3.5.1, which had many faults, to be explained later. The mask was therefore improved, and the mask in figure 4.3.1 was created.

This mask removed the alignment marks as it turned out they were not necessary, and added holes and trenches closer to each other. The holes have a radius of 100 nanometer and a horizontal separation of 800 nanometers. In other words, according to the proportions of an ideal photonic crystal, they serve only as a cautious starting point for fabricating holes with electron-beam lithography, but are far from ideal. The trenches all vary in size from 350 nm to 600 nm, with increments of 50 nm, to test what dimensions are feasible. The acceleration voltage of the electrons was for all samples presented here 20 kV, and the beam current was $\sim 40pA$.




Figure 4.3.1: More optimized mask for testing EBL exposure parameters.

Most of these tests were run parallel with fabrication of the SOI and before a proper etching recipe had been found. The main intent was to gain valuable experience with the EBL instrument, as well as determine the ideal doses for defining both holes and trenches.

Figure 4.3.2 shows the first developed sample using the mask of figure 4.3.1. The doses are as previously mentioned given in units of $\frac{\mu\text{As}}{\text{cm}^2}$, but for clarity the unit will be omitted. The doses used for this sample were 100, 150 and 200, respectively, but only the results from dose 200 are shown, as the other pictures gave no additional information.

For holes, even a dose of 200 seems too low to properly open the PMMA. By looking at the stitching error where the holes overlap, it is indicated that by doubling the dose the holes are closer to being properly open. As for the trenches, they seem to be sufficiently defined and opened at dose 200. It also looks as though the bigger trench has an easier time opening during exposure than the thin trench, but this could very well just be an illusion caused by the microscope, as more light would reflect from a thick trench than a thin one at this magnification.

The next sample had a wider array of doses exposed to better observe the changes as the doses are changed. The doses used were 250, 300, 350, and 400 respectively. The results are shown in figures 4.3.3. Only dose 400 is shown, as the other pictures gave no additional information.

It is fair to say the results look better than for the previous sample. Both holes and




Figure 4.3.2: Microscope pictures trenches and holes with dose 200 and a PMMA thickness of 100 nm.




Figure 4.3.3: Microscope pictures of trenches and holes at dose 400 with a PMMA thickness of 100 nm.

trenches now look much more defined, and the areas inside the stitching error look overexposed. The next sample employed a PMMA thickness reduced to 70 nm in order to assess whether or not a reduction in PMMA thickness would lead to finer features. Figure 4.3.4 show the results.




Figure 4.3.4: Microscope pictures of holes and trenches dose 100, 200, and 400.

It is difficult to assess the benefits of a lowered PMMA thickness through optical microscope pictures alone. It is evident however that much lower doses are required to obtain the same penetration in the PMMA. Lower doses means less exposure which should reduce unwanted issues like proximity effects. Furthermore it looks like around 100 is a good dose for both trenches and holes. A dose of 200 is slightly overexposed and at 400 all definition in structures seem gone. All in all, it seems beneficial to keep the PMMA thickness low as long as this is feasible with respect to the etching recipe. Furthermore, it is evident that holes need higher exposure doses than trenches.

4.4 The etching recipe

Different etching methods could be investigated once the ability to create an SOI wafer was in place, electron-beam lithography to form patterns were mastered, and the correct parameters for a photonic crystal waveguide had been found. Finding the correct etching recipe and etching mask is likely the most complicated procedure required to fabricate a structure, and therefore consists of the bulk part of the thesis. Considering the complexity of each of these steps, it seems redundant to include all process parameters for every sample before presenting the results after etching. Instead, only the parameters that are relevant for each result is presented. For example, presenting an in-depth explanation of the amorphous silicon layer, the oxide layer, the EBL-exposure, as well as the etching recipe if, for example, the etching recipe failed and ruined the sample is not beneficial. Instead, only the more important parameters are presented for each sample.

4.4.1 Finding the correct etching recipe and mask

NanoLab standard Si etching recipe and PMMA mask

The first recipe attempted was the NanoLab standard silicon etching recipe. It serves as a good starting point although the recipe is neither cryogenic nor applied ICP, i.e. it is not expected to be perfectly anisotropic or show good sidewall roughness. The etching rates of the recipe were first found, and are given in table 4.4.1.

| Etch time | PMMA thickness | Profilometer step |
|-----------|----------------|-------------------|
| 0 s | 161 ± 1 nm | 160 ± 5 nm |
| 30 s | 100 ± 1 nm | $75nm \pm 10$ nm |
| 60 s | 40 ± 5 nm | 90 ± 3 nm |
| 90 s | Volatile | 125 ± 5 nm |
| 90 s | 0 | 125 ± 5 nm |

Table 4.4.1: Results of the standard NanoLab silicon recipe etch test.

The profilometer measurements were somewhat volatile, perhaps because the resist is too soft for profilometry. Even with the uncertainties of the profilometer measurements, the results give the needed information. Step 4 of the measurements is listed as "volatile" since the refractometer gets confused if one tells it to look for a layer that is not there. This could either indicate that the layer has etched away, or it could indicate that it is so thin that the refractometer has difficulty measuring it accurately. At this point, the sample was cleaned and the tests redone to verify that the PMMA was in fact gone. The tests show that 160 nanometers of PMMA

lets one etch 125 nanometers of silicon, which is not enough for the purposes of this thesis. The recipe was still utilized for one process run to gain experience using the ICP-RIE.

A sample was prepared through electron-beam lithography with a 100 nm PMMA layer, and utilizing the mask depicted in figure 3.5.1. The sample was etched for 30 seconds using the standard Si etching recipe. The results are shown in figure 4.4.1.




Figure 4.4.1: Microscopy pictures of the very first sample processed through both EBL and RIE.

The lack of defined features is noticeable with the figure. Although both the alignment marks and the squares in which the structures are supposed to be placed can be seen, no holes or trenches are visible. The reasoning behind this is not evident. Perhaps the trenches and holes are simply too small to be seen with an optical microscope? The answer is readily found by further inspecting the mask used to define the structures. The mask is inverted, and instead of attempting to create holes, the mask has attempted to create cylinders. The cylinders will likely have collapsed and are therefore not visible. Why exactly the trenches are not visible, however, is more of a mystery. Furthermore, the stitching errors are clearly visible in the sample, which is not of much surprise since no stitching error corrections were made for this sample. To further enquire about the success of the etch tests, the sample was reviewed under a scanning electron microscope. Figure 4.4.2 shows the result.




Figure 4.4.2: SEM pictures of the trenches, holes, and stitching error supposedly seen in picture 4.4.1.

The SEM pictures verify that the structures are in fact not present. The trenches can be seen with the SEM, but they are very poorly defined.

Cryogenic etching recipe with an oxide mask

The etching rate of the buffered HF solution used to open the oxide mask was first found. A sample with a thick layer of SiO_2 was dipped in the solution for 30 second intervals. The oxide thickness was measured by refractometry after every etch. The results are given in table 4.4.2.

| Etch time | SiO_2 thickness | Etch rate/30 s |
|-----------|-------------------|----------------|
| 0 s | 1117 ± 0.5 nm | 0 nm |
| 30 s | 1044 ± 1 nm | 73 nm |
| 60 s | 989 ± 1 nm | 55 nm |
| 90 s | 949 ± 1 nm | 40 nm |
| 120 s | 902 ± 1 nm | 47 nm |
| 150 s | 858 ± 1 nm | 44 nm |
| 180 s | 817 ± 2 nm | 41 nm |
| 210 s | 762 ± 2 nm | 55 nm |
| 240 s | 718 ± 2 nm | 44 nm |
| 270 s | 674 ± 10 nm | 44 nm |
| 300 s | 629 ± 10 nm | 45 nm |
| 330 s | 578 ± 10 nm | 51 nm |
| 360 s | 531 ± 10 nm | 47 nm |

Table 4.4.2: Results of the HF etch test.

Apart from some irregularities during the first few etching intervals, it seems as though the HF etching rate is roughly ~ 44 nm per 30 seconds, or $\sim 1.4 \frac{nm}{s}$. A sample using an oxide mask was then quickly prepared with the parameters of table 4.4.3.

| | |
|------------------------|-------------------------|
| SiO_2 mask thickness | 25 ± 5 nm |
| HF etch time | 24 s |
| Expected HF etch | 33.6 nm |
| Cryo etch time | 1 min 45 s + 1 min 45 s |
| Expected Cryo Etch | $1.6 \mu m$ |

Table 4.4.3: Crucial parameters for the first cryogenic etch sample.




Figure 4.4.3: The sample after ICP-RIE cryo etching on the right. On the left; the carrier wafer on which it was placed.

The cryogenic recipe indicates a $\sim 1.6 \mu m$ etch when both steps are set to 1 minute 45 seconds. This is roughly 5 times more than necessary for our intents and purposes. If this etching rate is indeed correct this is good news, as long as the oxide mask withstands the etching. Figure 4.4.3 shows the sample after the cryo etch. Without going into too much detail surrounding these results it is evident they do not bode well. All but a smoking crater remains in the region where the structures are supposed to be, and even the carrier wafer was completely destroyed in the areas exposed to the actual etch. Bringing the samples under the microscope gives more disturbing information, as seen in figure 4.4.4.

Although remnants of the structures can still be observed, they seem to be more or less lost. However - knowing what exactly had happened during this etch is hard to say without closer inspection through the SEM. The sample was therefore loaded into the scanning electron microscope and the results are seen in figure 4.4.5. The SEM pictures do not clarify much as to what has happened, but all pictures together verify that the results are not good enough.




Figure 4.4.4: After ICP-RIE cryo etching shown on the right. Top left picture gives an overview. Top right is a zoom-in on the trenches. Bottom left and right pictures are of trenches and holes respectively, both with 250 dose.




Figure 4.4.5: SEM pictures of what remains of the structures after the first ICP-RIE cryo etch. The two top pictures show the trenches. The bottom pictures show the holes.

Chlorine based silicon etch

Following the failure of attempting to utilize an oxide mask to do cryogenic etching, two options were left available. Either, a better mask had to be employed. This was problematic at this point in time since training in additional cleanroom instruments had to be undertaken to create a harder mask. The other option was to simply find another recipe which was able to achieve good selectivity between SiO_2 and $a-Si$. Chlorine-based recipes in general can somehow show good selectivities. Many chlorine-based silicon etching recipe were found in the NTNU library of recipes. The recipe that boasted the best selectivity between oxide and silicon was chosen for testing. The recipe chosen is given in appendix B, labelled "Cl₂ Si etch".

In order to verify whether or not a chlorine based silicon etch would be viable, the etching rates of PMMA, SiO_2 and $a-Si$ all had to be tested. In general, only the selectivity between the SiO_2 and Si is of relevance as the actual mask in the oxide is defined through buffered HF etching, however, should the PMMA be somewhat resistant to the etching, it might help the oxide mask withstand some extra etching should the selectivity in general be too poor. The resulting etching rates are given in table 4.4.4.

| Material | Etch rate |
|----------|--------------|
| PMMA | 240 nm/min |
| SiO_2 | 25 nm/min |
| Si | 80-90 nm/min |

Table 4.4.4: Etching tests of the chlorine based silicon etching recipe. The recipe itself boasted a selectivity between oxide and silicon to be 1:10.

The total selectivity between the SiO_2 and Si is then less than 1:4. This means that with a 30 nm thick layer of oxide one would not be able to etch more than ~ 120 nm of silicon. With a ~ 100 nm layer of PMMA one would gain an extra $\frac{100}{240} \cdot 85 = 35$ nm of silicon etch for a maximum total of $120 + 35 = 155$ nm. This is still far below the requirement to etch around 280 nm of silicon.

4.4.2 Opening the mask

Since the chlorine-based silicon etch showed subpar qualities, a new mask had to be utilized. With a chromium mask in place instead of an oxide mask, the cryogenic etching recipe could once again be employed. A number of samples were processed while trying to find the correct parameters needed to open the mask. Tables listing

all these samples and their relevant parameters are given in appendix E, labelled table E.0.1, and E.0.2. First, of course, the selectivity of the cryogenic etch recipe with respects to chromium had to be found, and the chromium etching rate of the chlorine-based chromium etching recipe had to be tested.

The chromium etching rate was found to be somewhere between 6-10 nm/min, but these values were very hard to verify due to uncertainties with the profilometer. However, most measurements lead us to believe it was much closer to 10 nm than 6 nm. The same measurements executed with SiO_2 proved difficult. Barely any step-size at all was measured. This should in practice mean that the oxide barely etches with this recipe, which is good news indeed.

Several samples with chromium structures were then loaded to the ICP-RIE process chamber and underwent cryogenic etching to see if the chromium would resist. Figure 4.4.6 shows a sample before and after an etch test with both etching steps set to 1 min 45 s





Figure 4.4.6: Photographs of the sample used for testing the chromium resistance against the cryogenic etching recipe. The left picture is before the etching, the right is after.

Simply by looking at the sample it is obvious that the chromium withstood the etching, but once again the silicon does not look good. Whether this is a macroscopic effect that occurs since the entire wafer is etched at once, or an effect that will also occur once a proper mask is in place is unknown, but will have to be studied at some point. Measuring the step between the silicon and chromium part gave a step size of ~ 350 nm. It must also be mentioned that this particular sample had a rather thick chromium layer of ~ 50 nm, which is thicker than what the final mask will be.

For the purposes of this thesis, being able to etch 350 nanometers or more before the mask etches away is perfectly fine, although this gives an etching rate that is not even 25 % of the etching rate boasted by the recipe. Still, for our intents and purposes, this is adequate. The question remains whether a thinner chromium layer would be able to withstand the etching. To test this, another sample from the same wafer was etched for 3 minutes. The results were more or less the same - the chromium layer was still clearly visible, meaning that a 30 nm layer should be able to withstand 1 minute 45 seconds of etching. Although these tests do not give the exact etching rate of this recipe on chromium, they supply sufficient information by confirming that a 30 nm layer of chromium will be sufficient to etch the needed depth.

At this point, all recipes had been sufficiently characterized and samples were ready to be processed. The first sample to go through cryogenic etching while using a chromium mask was sample 1, whose parameters are given in table E.0.1. A microscope picture after the cryogenic etch is shown in figure 4.4.7.




After cryo etch. Dose 200


Figure 4.4.7: Sample 1 after cryogenic etching. Dark spots have formed on the sample surface.

Dark "spots" have formed on the wafer, likely due to a fault somewhere in the etching process, leading to 'dents' in the chromium. Perhaps the oxide layer serving as a mask for the chromium was too thin, such that an uneven etching of the chromium began once the oxide was gone. Another explanation could be that the cryogenic etching recipe has this effect on the chromium layer. The sample was loaded into the S(T)EM to further research this phenomenon. The results are shown in figure 4.4.8.


The cross-sectional pictures answer a lot of questions while raising a few new ones. The picture of the top layer indicates that there are 75 nanometers of some substrate




Top Layer



Trenches dose 150



Top layer again



Top layer BSE

Figure 4.4.8: Cross-sectional $S(T)EM$ pictures of sample 1. The top left image shows a layer of some material resting on top of the silicon substrate. The top right image shows the trenches, which do not look like they are well defined. The bottom pictures show the top layer from a better angle, indicating that the top layer consists of two different materials.

on top of the silicon. At first glance, this looks like it could be chromium, as it appears much brighter than the silicon substrate underneath. If this layer is indeed as thick as 75 nanometers, it is no surprise that the structure has not been properly formed, as the cross-sectional picture of the trenches clearly show. The third and fourth pictures of figure 4.4.8 however show a slightly different story. The leftmost image indicates a 30 nm layer of chromium as well as a 30 nm layer of oxide underneath, something which corresponds better with previous measurements. Furthermore, the bottom right picture of figure 4.4.8, taken with a back-scattered electron detector clearly shows there being two different layers on top of the sample. It is reasonable to assume that the top one is chromium, whereas the bottom one is the stopping-layer oxide.

The question then remains to find out why the structures did not form. If the chromium layer is indeed 30 nm, then 4 minutes of chromium etching with an etchrate of 10 nm/min should be sufficient to open it. The cryo etching recipe has already been shown to etch SiO_2 , so the stopping layer is not the cause either. Attention is therefore turned to the top oxide layer. It is highly possible that it was not properly opened during the HF etch. If this is indeed the case, the chromium etch would simply be stopped by the unopened oxide layer, and the cryogenic etching would in turn be stopped by the chromium layer.

Before the next sample was processed, a few previously stated conclusions were questioned. Since it appeared that the oxide layer did not properly open during the HF etch, the HF etching rate was retested. Doing refractometry measurements on the same sample multiple times can be difficult if the sample surface is not sufficiently uniform. The last time the HF etching rate was characterized, a small sample with a very nonuniform surface was utilized. It was therefore imperative that every refractometer measurement was conducted on the same location on the sample, which is difficult when the sample is constantly being removed for HF etching. This time, a large sample with a uniform surface area was employed, so that the uncertainties were kept to a minimum. The results of the HF etching test are given in table 4.4.5.

The etching rate is now closer to ~ 30 nm per 30 seconds or $1\frac{nm}{s}$. A result which is much lower than the previously tested value. Since these tests had much lower uncertainties, they are more trusted than the previously found value. Furthermore - such a low etching rate would explain why the previous sample failed at some point during the etching. Most likely - the oxide layer had not been opened when the chromium etch began.

| Etch time | SiO ₂ thickness | Etch rate/30s |
|-----------|----------------------------|---------------|
| 0s | 755.5±0.1 nm | 0 nm |
| 30s | 722.8±0.1 nm | 32.7 nm |
| 60s | 694.2±0.1 nm | 28.6 nm |
| 90s | 665.5±0.1 nm | 28.7 nm |
| 120s | 635.3±0.3 nm | 30.2 nm |

Table 4.4.5: Results of the second HF etch test.

With the new HF etching parameters found, sample 2 was ready for processing. Since the previous sample did not open the mask at all, the next sample was purposely overetched to see if the fault indeed was due to the HF etch.


Figure 4.4.9 shows the results of sample 2 after etching.

There are two things of interest in this figure. First is a slight change in the EBL mask. An arrow has been added to indicate the direction in which the dose increases. This makes it easier to orient oneself when looking at the samples in the S(T)EM. It is evident from the figure that the long HF etch has ruined the structures, but a single structure still stands. The etch step is far from good, but it is difficult to say whether this is caused by failures when breaking the sample or by the etching recipe. An unwanted formation of silicon grass is also evident, which most likely explains the effects the etch has had on silicon samples. The grass is the reason the samples turn dark after etching.


To try and combat the overetching evident in the previous sample, the HF etching time was reduced for sample 3, in hopes that the HF etch would still pierce the oxide but leave the structures unscathed. The results of sample 3 after etching are shown in figure 4.4.10.

The structures still collapsed during the etch. A few are still standing, but their quality are subpar. As with the previous sample, a large amount of silicon grass is forming in the substrate. Sample 4 therefore experienced a much shorter HF etching time. The results are shown in figure 4.4.11.


This time the structures seem well defined in the microscopy pictures. However, S(T)EM pictures indicate otherwise. The silicon etch was stopped by the mask. If the chromium mask had been opened, there should be an etch step of at least 300 nm. Instead, only a 50 nm etch step is visible. The structures can be seen, but the light surface of the structures give the impression that they are simply imprinted in chromium. This poses a difficult problem. 50 seconds of HF etching opened the mask, but ruined the structures. On the other hand, 40 seconds did not ruin the




140 dose



Etch step



Arrow 100 dose



Standing structure 120 dose

Figure 4.4.9: Microscope and S(T)EM pictures of sample 2 after silicon etching. The optical microscope picture clearly show that both trenches and holes have been overetched. The S(T)EM images confirm this.




Figure 4.4.10: Microscopy and $S(T)$ EM pictures of sample 3 after silicon etching. Reducing the HF etching time was clearly sufficient to still open the mask, but once again the structures have collapsed.


structures, but was not enough to open the mask. To circumvent this, the next few samples were immersed in HF before electron-beam lithography to make the oxide layer thinner, so a shorter HF etch could be employed to prevent the oxide layer being overetched.

Sample 5 was immersed in HF before electron-beam lithography. The results after silicon etching are shown in figure 4.4.12.


The figure shows no indication of any holes, even though they are supposed to be located by the tip of the arrow. The reason why is difficult to gauge. The only thing that changed between this sample and the previous one is that it supposedly has a thinner oxide layer. Perhaps the holes had somehow completely etched away. Alternatively, something could have gone wrong during electron-beam lithography.

In taking the sample to the cross-sectional $S(T)$ EM, the structures could not be found even after extensive searching. The only likely explanation is that somehow the structures had not been formed. Why this is is difficult to say, as one would think with the thinner oxide a 40 second etch should be enough to open the mask. If nothing else, it should be enough to see an outline of the structures. With sample 5 being unsuccessful, the HF etching time was increased for sample 6. The results after silicon etching are shown in figure 4.4.13.


The microscopy pictures show good definition of both trenches and holes even after cryogenic etching, which is good news indeed. However, when attempting to study the structures under the cross-sectional $S(T)$ EM, they could not be found, just




Dose 100



Etching step



Holes 100 dose



Trenches 100 dose

Figure 4.4.11: Microscopy and $S(T)EM$ images of sample 4 after silicon etching. The optical microscope image shows that the trenches are still properly defined. In other words, the structures have not collapsed. The $S(T)EM$ images however clearly show that the mask has not been properly opened.




Figure 4.4.12: Microscopy pictures of sample 5 after silicon etching. The trenches look well defined but the holes are somehow not visible.




Figure 4.4.13: Microscopy pictures of sample 6 after cryogenic etching. The trenches look well defined, but the holes seem to have collapsed.

like for sample 5. Understanding why the structure could be so visible under the microscope and not at all visible under cross-sectional S(T)EM figures is a difficult task. In general, any surface roughness would quickly lead to a high contrast under the microscope, whereas a surface roughness during S(T)EM would be seen as only that - surface roughness. The structures are probably there, but with a complete lack of contrast, very difficult to find with the S(T)EM. In any case, not finding the structure with the S(T)EM means that at some point, a layer has not been opened, such that the cryogenic silicon etch did nothing. Once more, finding the exact cause of this is not easy.

After careful consideration and some analysis regarding the results presented for previous samples, the decision was made to abandon all attempts at a wet etch, and instead try to find a proper dry etching recipe that etches silicon dioxide while retaining a good selectivity between the photoresist PMMA and the oxide underneath. Fortunately, another cleanroom user had toiled with this problem for quite an extensive amount of time and supplied a recipe that did exactly this. The recipe in question is given in table B.0.3 of appendix B.

The gas used in the recipe, CHF_3 is not available for the ICP-RIE and therefore the standard RIE was used. Without being able to cool down the chamber, or use the ICP-power, the recipe is not completely anisotropic, but still much more so than a wet etch. With oxide layers as thin as 20-30 nanometers, the slight isotropy of the etching recipe will only amount to a few nanometers or so, in other words completely within acceptable boundaries.

The etching rate of the oxide as well as the etching rate of the PMMA was tested quite easily. By using one sample with an oxide layer on top as well as one sample with only PMMA on top, a short 30 second etch followed by refractometer testing gave a rough estimate of the etching rates. The results are given in table 4.4.6.

| | Before etch | After etch | Etch rate/30 s |
|----------------------------------|-------------|------------|----------------|
| PMMA thickness | 103.6 nm | 60 nm | 43.6 nm |
| SiO₂ thickness | 912 nm | 863 nm | 49 nm |

Table 4.4.6: The SiO₂ recipe used for the SOI wafer.

The selectivity of this recipe is clearly good enough for our intents and purposes, so no further testing had to be conducted.

Three samples were processed in parallel to gauge the usefulness of the dry etching recipe to open the oxide. All three were etched down to a 20 nanometer layer for the top oxide. They then had their masks defined through EBL and the oxide mask defined with a dry RIE etching process. The RIE etching times were 30, 35, and 40 seconds respectively. The chromium and cryo silicon etch were kept constant at 5 minutes and 2+2 minutes respectively. 2+2 here meaning that the cryogenic etching recipe was run for 2 minutes on both steps. In all three samples, the structures were not visible. This led to a rather large amount of questions, and to some surprising conclusions.

First of all - if the oxide layer is only 20 nanometers thick, even the sample who was dry-etched for 30 seconds should at the very least have some structures visible. With the values presented, the oxide mask had to open. As with the previous samples, the chromium etch lasted for 5 minutes. It seems as though the main protagonist of our experiments was not the oxide etch afterall, but rather the chromium etch. With regards to the results presented earlier, the chromium etch rate was not sufficiently pinpointed. An etch rate of 6-10 nanometers per minute were found, with the 10 nanometer value being taken as the most accurate one.

The reason a 5 minute chromium etching time had been utilized throughout the last few tests is because when the HF was overetched, the mask did indeed open. This would leave one to believe that the oxide etch had to be the perpetrator. However, one option was not considered. During the HF overetch, the buffered HF will have pierced all of the oxide rather early in the etching process. There is a good chance that the buffered HF would then start to etch the chromium layer. Prior to the HF etching a quick test to acquire the HF etching rate of chromium was undertaken and the results were underwhelming. An etching rate of not more than 5 nanometers per 30 seconds was found. However, this could be just enough to tip the chromium etch of 5 minutes in our favour.

At this point, there was logically no reason why the oxide should not have opened, so the only possible explanation is that the HF overetch had somewhat helped the chromium layer to open. As a result, the next few samples used a 6 minute chromium etching time instead.

Before these tests were conducted, the mask was modified. The setup of the trenches has been changed. The size of the trenches no longer vary. Instead, every trench is 50 nanometers wide, but the distance between subsequent trenches is no longer constant. The distance between the first two trenches is 50 nanometers, then 100,

then 200, then 300, and finally 400. This was done so one could gauge the minimum distance between two trenches before the rib in between would collapse. Also, the size of the holes has been increased significantly to better match an ideal photonic crystal. The diameter of the holes is now 200 nanometers, and their horizontal separation is 400 nanometers, from center to center.

Following the results found from previous samples, sample 7 was prepared with the dry oxide etching recipe. The results are given in figure 4.4.14.




Figure 4.4.14: Cross-sectional S(T)EM pictures of sample 7. The trenches are now very well defined. The anisotropy of the etch is quite good, but an unwanted formation of silicon grass can be seen.

Finally, it can be seen that the trenches are well defined. However, it is apparent from the figure that the holes have once again collapsed. The dose used for the holes lead to good results after EBL exposure for the previous sample, but for this sample they have collapsed. This is most likely caused by proximity effects, given that the holes are both larger and closer to each other in this mask. The right way to correct this would be to reduce the dose for the holes. A few holes are still visible at the edge, where proximity effects are the least severe. Otherwise, it looks like the masks have properly opened between every etching step. The S(T)EM pictures also show that for dose 100 the trenches are quite well defined, even for really small separations. The pictures also clearly show the chromium and stopping oxide layers. The anisotropy is also quite good, although not perfect. This is probably attributed to the fact that the pressure was increased from its standard value due to a problem getting the plasma to strike. The most concerning detail in the figures is the vast amount of grass still forming in the silicon.

The only value that changed for the sample 8 was the length of the oxide etching recipe. The length of the oxide etching recipe was reduced since the value utilized for sample 7 was much too high when regarding the etching rates of the dry oxide etching recipe. Figure 4.4.15 shows the result of sample 8 after silicon etching.




Figure 4.4.15: Cross-sectional $S(T)EM$ pictures of sample 8. Trenches are well defined. Holes are now also visible, but look far from perfect.

A few things have changed between this and the previous sample. The trenches of 100 dose are not properly opened. This is likely caused by the reduced oxide etch. 100 dose is probably too low to properly open trenches, something a long oxide etch might make up for by first etching through the PMMA, and then etching the oxide. Since this sample had a reduced oxide etching time, it makes sense that the trenches might not open if the exposure dose was too low to begin with. Otherwise, 120 and 140 dose was enough to properly open the oxide such that the trenches are again nicely defined. Most importantly is the occurrence of holes. Even though the holes are far from nicely defined, they are present. Whether this is caused by the reduced etching time of the oxide etch, simply a matter of chance, or due to a slightly thinner photoresist is hard to say. Nonetheless, the holes need a lower dose during electron-beam lithography.

Sample 9 was processed with a thicker PMMA layer, and a thinner stopping layer was used to see if this would have any effects on the sample. Furthermore, both the oxide and chromium layer were slightly thinner. The results of sample 9 after silicon are shown in figure 4.4.21.

There are a few things of interest in this picture. First of all, a severe undercut is visible under the mask. This also happened with previous samples but not as




Figure 4.4.16: Cross-sectional $S(T)EM$ pictures of sample 9. Undercut is visible as the silicon has etched under the mask.

severely. It is hard to attribute this underetch to any specific parameter, but since both the oxide etch and chromium etch have not changed much, the most likely cause is the reduced thickness of the stopping layer. If this is the case it would mean that the chromium etch quickly etched through the chromium layer and started etching the oxide, which seems unlikely given the chromium recipes etching rates.

For sample 10, the chromium etching time was reduced from 5 minutes to 4 minutes, to see if this would help with the severe underetch observed in the previous sample. The results are shown in figure 4.4.17.




Figure 4.4.17: Cross-sectional $S(T)EM$ pictures of sample 10. Undercut is still visible as with sample 9. The right image shows stitching errors in the trenches.

Changing the chromium etching recipe from 5 minutes to 4 minutes does not seem to have helped tremendously, indicating that the chromium etch did not cause the underetching. Perhaps the undercut is resulting from step1 of the cryogenic etching recipe, and the stopping oxide layer helps prevent this. One more sample is processed with the mask in mind, sample 11. Sample 11 attempted to research whether the formation of silicon grass was caused by the oxide stopping layer. Not knowing how the cryogenic silicon etching recipe etches oxide, the sample had its oxide removed by an RIE oxide etch before beginning the cryogenic etch. The results are shown in figure 4.4.18.




Figure 4.4.18: Cross-sectional S(T)EM pictures of sample 11. Due to the sample breaking poorly during scribing, the image instead shows the stitching error. The surface of the trenches is very dirty, likely due to the poor breaking.

Unfortunately, the breaking of the sample failed so no proper cross-sectional pictures are available, and the pictures were instead taken of the stitching-error to give some insight. As a result, it is difficult to gauge whether or not the dry oxide etching before cryogenic etching had any impact on the undercut. The figures do however conclude that a dry oxide etch before the cryogenic etch does not decrease the formation of silicon grass.

4.4.3 Characterizing the cryogenic etching recipe

With the tests executed so far, enough information has been gathered regarding the opening of the mask. The next tests therefore attempt to characterize the cryogenic etching recipe. The thickness of the layers as well as the the recipes used to open the mask were kept constant throughout these tests, so that only the cryogenic etching recipe was changed from one sample to the next. The relevant parameters of the

samples used to characterize the cryogenic etching recipe are given in appendix E in table E.0.3 and E.0.4.

The two steps

The first thing to be characterized is the difference between the first and the second etching step. Sample 12 was the first to undergo such testing. For this sample, only the first etching step was utilized. The results are shown in figure 4.4.19.

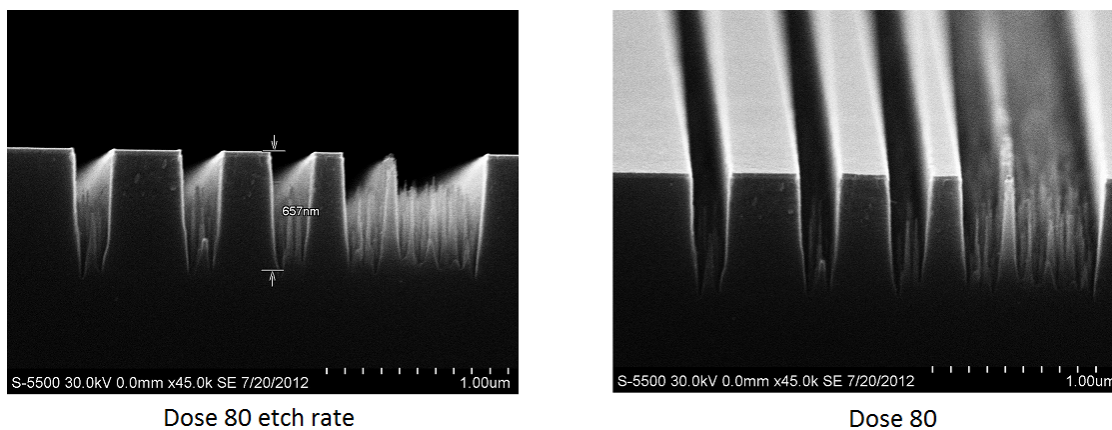


Figure 4.4.19: Cross-sectional $S(T)EM$ pictures of sample 12. The trenches are nicely defined and there is not much undercut. The profile does however not look perfectly anisotropic.

Surprisingly, not much difference is observed between using only step1 and using both steps. For sample 13, step 1 was removed and only step 2 was run. The results are shown in figure 4.4.20.

Once again, it seems like there is barely any difference between using only step 1 and using only step 2 of the etching recipe. A result which is somewhat confusing, given the high difference between the two steps in terms of gas flows as well as the RF power. At the very least it seems as though the silicon grass is not related to either one of the two steps.

Gas flows

The effects of changing the gas flows of the recipe were then tested. The oxygen flow was the first value to be changed, since silicon grass is usually a result of a high oxygen flow in the chamber. Sample 14 was therefore processed with both etching steps, but step2 had its oxygen flow drastically reduced. The results are shown in figure 4.4.21.




Figure 4.4.20: Cross-sectional $S(T)EM$ pictures of sample 13. The trenches are nicely defined and not much undercut can be seen. The etch is however not completely anisotropic.




Figure 4.4.21: Cross-sectional $S(T)EM$ pictures of sample 14. The trenches look poorly defined in this figure. The formation of silicon grass does not appear to have been reduced.

It seems as though reducing the oxygen flow slightly increases the etching rate, as this sample has etched as much as 800nm. Furthermore, the reduction of the oxide does not look like it has sufficiently decreased the formation of silicon grass. It is evident that the anisotropy has become much worse when the oxygen flow was reduced. Even though the reduction of oxygen did not help, another sample is prepared with an even lower oxygen flow, so that one can conclude that the oxygen flow does not affect the silicon grass to any great extent. The results of sample 15 can be seen in figure 4.4.22.

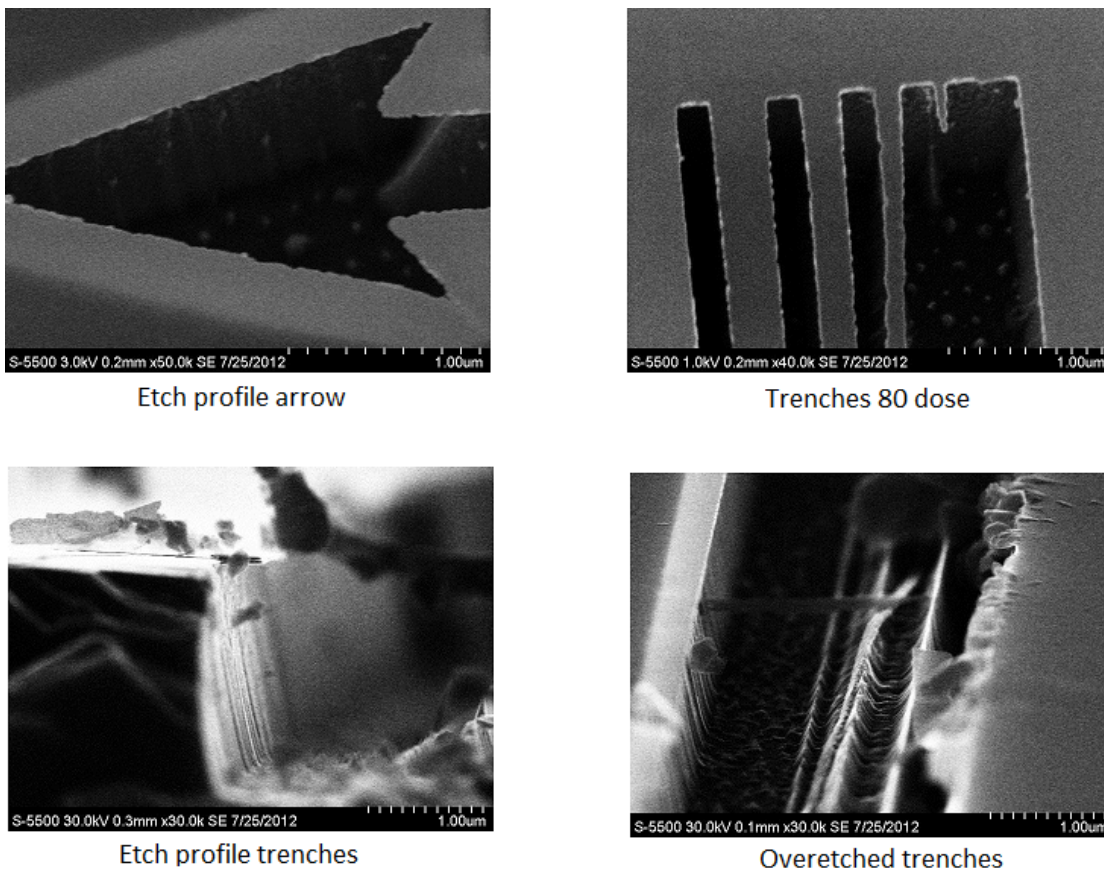


Figure 4.4.22: Cross-sectional $S(T)EM$ pictures of sample 15. The top left image shows the arrow. The lighter parts are caused by the silicon underneath the chromium missing. On the right side of the image of the arrow the slope of the silicon is visible. The same effect can be seen for the other images.

The pictures show a severe underetch. Look for example at the arrow, it is greatly underetched and it is highly fortunate that the chromium is still floating above the rest of the structure. The same can be seen in the pictures of the etch profile of the trenches. The recipe in this state is in other words not applicable for photonic

components. However, reducing the oxygen to this extent did overcome the problem with the silicon grass. Due to the underetch, however, such a recipe is not applicable.

The temperature

The next parameter to be investigated, was the temperature. Sample 16 had a slightly decreased temperature in comparison with previous samples. The results are shown in figure 4.4.23.

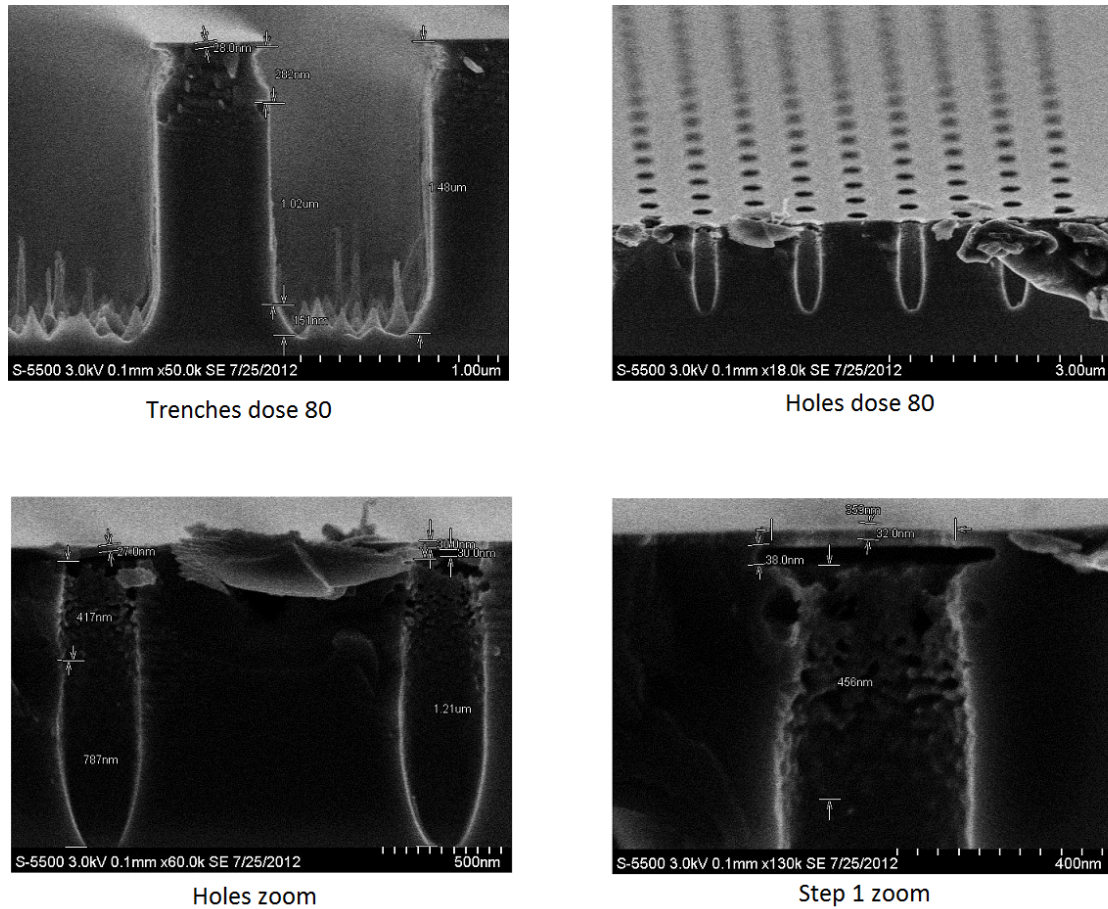


Figure 4.4.23: Cross-sectional S(T)EM pictures of sample 16. The top left image shows that step 1 of the recipe undercuts and etches underneath the mask. Step 2 is very anisotropic. Holes are also defined and are well defined, with a very deep etch.

The pictures indicate that this is the best sample yet. Although the silicon grass is back, the sidewalls are nearly 90 degrees vertical to the surface. Furthermore, the holes look well defined this time but are, unfortunately, not entirely anisotropic.

This could be the result of a higher temperature, as lower temperatures should in theory improve the anisotropy. Furthermore, the difference between step1 and step2 is discernable from these pictures. It is evident that step1 of the etching recipe is a much more "rough" etch, as it is intended to be a burst step only to prevent notching.

The sidewall roughness created by step1 of the etching recipe, makes this step unviable for serious etching procedures. Furthermore, a rather long step 1 has been utilized thus far. However, running step 1 for a long amount of time should not be necessary when the mask is so thin. For projects utilizing for example a standard photoresist mask of up to 1 micron, a long burst step to prevent notching is needed, but when using hard masks, it should not be necessary. In fact, by looking at sample 13 which only used step 2 of the recipe, no notching was observed, so there is a good chance that step 1 is in fact not necessary.

Sample 17 was processed with an even higher temperature. The results are shown in figure 4.4.24.




Figure 4.4.24: Cross-sectional S(T)EM pictures of sample 17. The two etching steps are even more different when the oxygen flow has been further reduced. Neither of the steps seem to be at all anisotropic, and some of the ribs are barely still standing.

Evident from this sample is how much of a difference the temperature actually makes during the etch. The trenches are severely overetched and are barely still standing. The effects of increasing the temperature highly reduces the anisotropy. The holes look far from good enough for use in a photonic crystal waveguiding device. In other words, -110C is much too high, while -115C showed good results. The next sample is therefore processed at -120C, and the results are shown in figure 4.4.25.




Figure 4.4.25: Cross-sectional $S(T)EM$ pictures of sample 18. Both the image of the arrow and of the holes clearly indicate the difference between step 1 and step 2 in terms of sidewall roughness. Step 1 gives a high amount of roughness.

The picture of the arrow clearly shows the difference between the first and the second etching step. Although the second step is not perfect, it performs much better than the first step, particularly when it comes to sidewall roughness. It looks as though the roughness obtained during the first etching step affects the quality of the second etching step. The roughness protruding out from the sidewall obtained by the first etching step seems to increase the roughness of subsequent etching. The holes look sufficiently anisotropic during the second etching step and are deeply etched.

Finding a good combination of parameters

Sample 19 further investigated the effects of the oxygen flow, but now utilized that a temperature of -120 degrees gave good results. The process was a three-step process where every step changed the oxygen flow. The results are shown in figure 4.4.26.

The figure gives a lot of additional information. The difference from each step is clearly visible in the pictures. The second step, when the oxygen flow was set to 15 sccm seems to be the best. It shows both the best sidewall roughness as well as the best anisotropy. The first step when the oxygen flow was set to 20 sccm shows half-decent sidewall roughness but a high amount of undercut in the etch. When the oxide is set to 10 sccm on the other hand, the sidewall roughness is very large and the etch is not at all anisotropic.




Figure 4.4.26: Cross-sectional $S(T)EM$ pictures of sample 19. The three different oxygen flows utilized clearly show different qualities in terms of both surface roughness and anisotropy.


Structures on an SOI wafer

Up until this point all processed samples have been fabricated directly on a standard silicon wafer. Since the main goal was to characterize the etching recipes, it was not found necessary to grow the entire SOI stack for every sample processed. One of the more evident ways of reducing the silicon grass was to actually create the entire stack of layers needed for a final fabrication run. Since in theory, the amorphous silicon is to be situated on top of a $\sim 1\mu m$ layer of silicon dioxide, this should help overcome the formation of silicon grass during the etch. The oxide layer was aimed to be around $1\mu m$ and the amorphous silicon layer was aimed to be around 240 nm.


Sample 20 was the first sample which was a complete SOI sample. Since only 240 nm needed to be etched on this sample, it was not etched for long. The results are shown in figure 4.4.27.

The results are interesting. It seems as though the inclusion of an oxide layer beneath the amorphous silicon does indeed prohibit the formation of silicon grass. However, a lot of roughness is still observed. It must be remembered that the chromium layer has served as an etching mask and naturally has a high surface roughness. However, the chromium layer is intended to be stripped at some point, so this is not an issue. Still, quite a bit of roughness is observed in the silicon as well.


The next sample was the last of the old design used for testing holes and trenches before creating the entire tapered structure. Both etching steps were used. A long etching time for this sample was attempted to verify whether a deep etching of




Dose 80 trench



Dose 80 etch step



Dose 80 holes



Dose 120 trenches

Figure 4.4.27: Cross-sectional $S(T)EM$ pictures of sample 20. The inclusion of a buried oxide layer clearly rectifies the silicon-grass problem. The images do however show a very high surface roughness on the chromium mask, something which might not be a problem after stripping the mask.

photonic crystal holes was possible. The results are shown in figure 4.4.28.

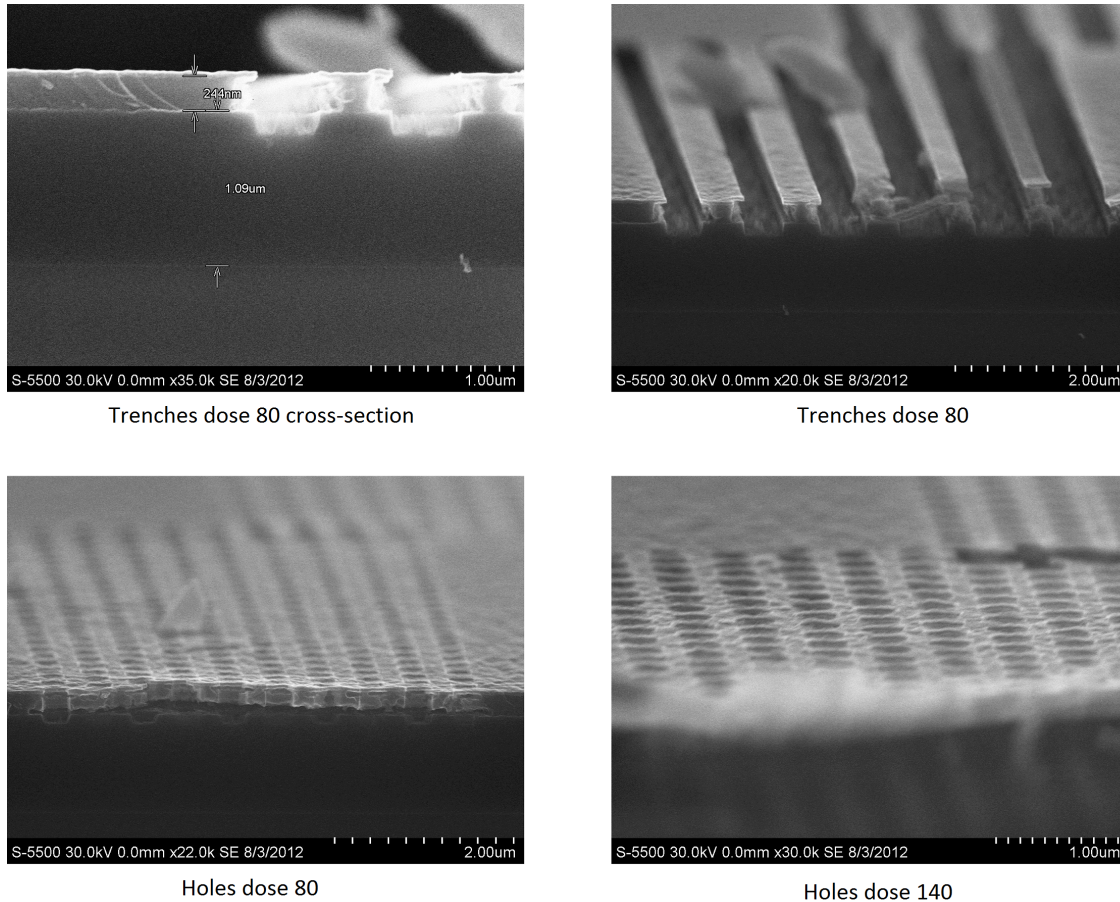


Figure 4.4.28: Cross-sectional $S(T)EM$ pictures of sample 21. The structures are not at all anisotropically etched. The difference between the etching profile of the amorphous silicon and the oxide is clearly visible. The holes also look well defined.

The figures are somewhat interesting. The first picture of figure 4.4.28 clearly shows both the amorphous silicon layer as well as the oxide layer underneath it. The oxide layer is quite close to one micrometer, but a little bit on the thicker side. The amorphous silicon layer is 240 nm, just as expected. Furthermore, the picture clearly shows the difference in the etching recipe between etching amorphous silicon and silicon dioxide. The silicon etch seems to have an inward slope, whereas the oxide etches with an outward slope. Furthermore, a 3 minutes and 30 second etch is much too long, as it greatly reduces the anisotropy of the silicon etch. The holes also look more or less well defined, although their sidewall profiles are hard to gauge from this sample.

The next sample, sample 22, was the first of a new breed. For the first time, the entire structure was to be fabricated, including the grating and taper for coupling light in and out of a photonic crystal waveguide. Pictures of the structure after electron-beam lithography are shown in figure 4.4.29.




Figure 4.4.29: Microscopy pictures of the first sample constructing the entire structure. The top left image shows the mask exposed five times at five different doses. The top right image shows the entire mask.

After electron-beam lithography the structure looks well defined, apart from the holes. The holes had a realistic proximity to each other in terms of a photonic crystal waveguide, so they have been severely overexposed due to proximity effects. The mask employed also needs some explanation. There are four structures in the mask. The first structure employs two tapered structures used to couple light into - and out of - a photonic crystal waveguide. The three other structures do not have a photonic crystal waveguide but are intended only for testing the silicon waveguides. There are three different waveguides, each of a different length. The three lengths

attempted are 30, 50, and 80 micrometers.

For such a large sample, it made little to no sense to try and take a cross-sectional image, unless one could cut it along the waveguide, a feat which is most likely not possible. As a result, the pictures are always taken from a birds-eye view, sometimes with a tilt of around 30 degrees. The results of sample 22 after silicon etching are shown in figure 4.4.30.




Figure 4.4.30: Cross-sectional S(T)EM pictures of sample 22. The grating looks well defined initially, but upon further inspection it is clear that the sidewalls and surface is very rough. The PhC waveguide has collapsed due to proximity effects.

The first picture shows an overview of the grating at the lowest dose. The grating looks well-defined from a top-down view but when rotating the sample to look at the sidewall, it is once again obvious that the etching recipe has a problem with sidewall roughness. Furthermore, a portion of the trenches was not properly defined,

probably due to some photoresist remaining on the sample after the EBL exposure. The holes are overexposed and have therefore not survived the fabrication process.

For the next sample the dose used to expose the holes was reduced. Furthermore, the etching recipe continued to be changed to try and reduce the sidewall roughness. The sample was processed with only one step, but with greatly reduced gas flows. The results are shown in figure 4.4.31.

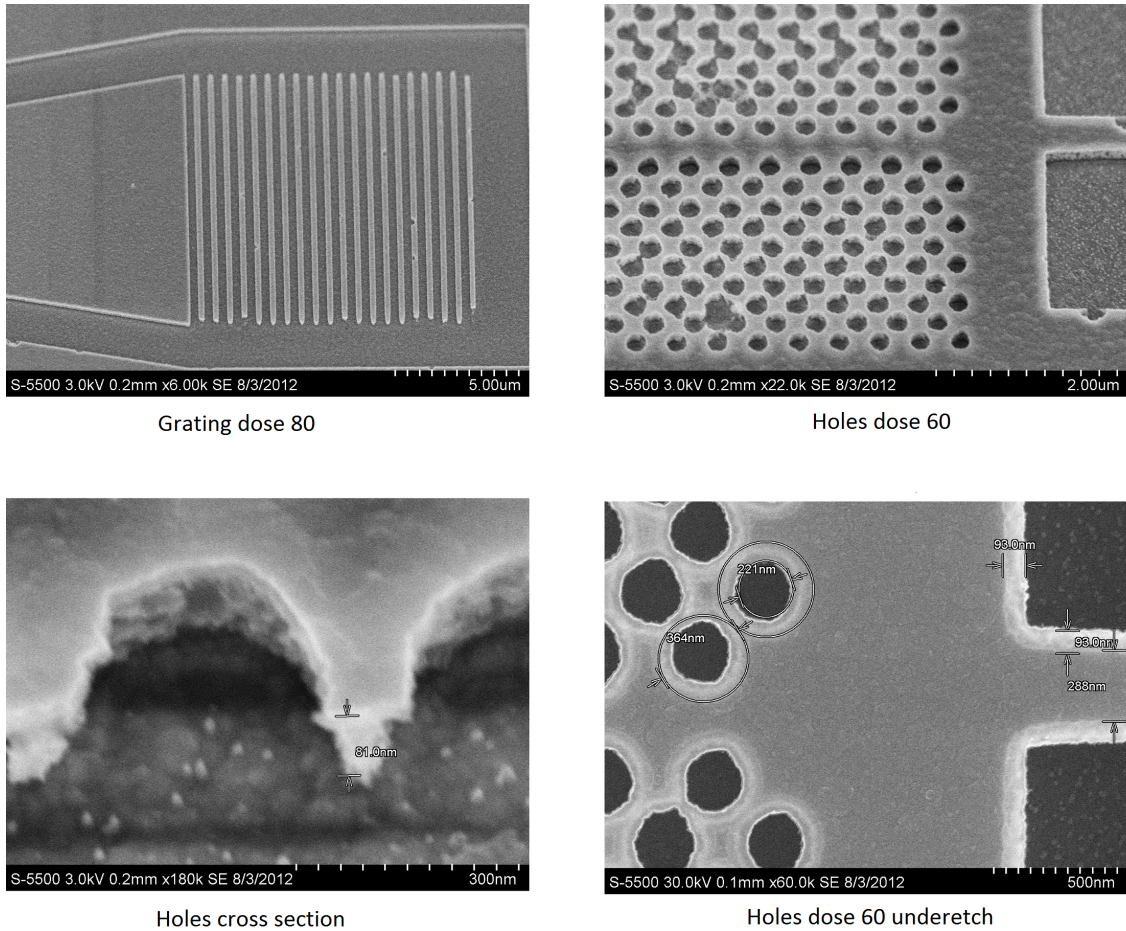


Figure 4.4.31: Cross-sectional $S(T)$ EM pictures of sample 23. As before the grating looks well defined. The holes now also look much better defined, but proximity effects are still visible near the center of the array of holes. Close-up images of the holes show the etch to perform poorly, as it is poking holes in the amorphous silicon.

The figures are of great interest. The holes are much better defined this time around, undoubtedly a result of the lowered dose used when exposing the holes. However,

it is apparent that the etching recipe is not ideal. The sidewalls now have a very "sponge-like" surface with many smaller holes, as visible in the closeup figure of the holes. Furthermore, the holes are somewhat underetched, as indicated by the bottom right picture in figure 4.4.31. The white circles around the actual holes indicate electrons passing through the top chromium layer, hitting the oxide substrate and being reflected back up, as if there is no silicon layer there. In other words, a quite serious underetch.

The next sample had the same exposures as the previous sample, but the etching recipe was once again changed. The RF power was severely increased to see if a quick etch would help with the sponge-like surface indicated by the previous sample. The results are shown in figure 4.4.32.

It is evident the holes near the middle of the structure become overexposed due to proximity effects. A method to overcome this must be developed. Furthermore, although the sidewalls are better than they were with previous samples, they are still far from good enough. After such a vast amount of different tests performed on the cryogenic etching recipe, the question arises whether or not the surface roughness is caused by this recipe or somehow caused by one of the previous recipes.




Figure 4.4.32: Cross-sectional $S(T)EM$ pictures of sample 24. The waveguide sidewall profile still shows a lot of roughness, but not as much as for the previous sample. One of the gratings show a stitching error, indicating a poor correction of stitching errors. The holes look rather good, but a slight underetch is still visible.

Figure 4.4.33 shows an illustration of a bad EBL exposure leading to a curvy surface on the PMMA. During the oxide etch, the same curvy surface will be imprinted in the oxide, and so on. In other words, any one of the 4 steps in the process may be at fault for the sidewall roughness observed. All four of the process steps must therefore be investigated with the same tenacity as for the cryogenic etching recipe, if the proper cause of the surface roughness is to be found.




Figure 4.4.33: Example of how a bad exposure can lead to a high sidewall roughness for the end silicon.

Ideal parameters and comments

The next two pages give a summary of the ideal parameters found through the tests.

| Parameter | Value | Comment |
|--------------------------|--|---|
| SOI oxide layer | $\sim 1\mu m$ | A $1\mu m$ oxide layer prevents substrate leakage. Might not be optimal for coupling with a grating. |
| SOI a-Si layer | 280 nm | A 280 nm silicon layer is optimal for PhC waveguide. |
| SiO_2 stopping layer | ~ 30 nm | A 30 nm oxide stopping layer seems to reduce undercut during etching. Needs to be stripped after etching. A wet etch to remove the oxide layer will also begin etching the SOI oxide layer. |
| Cr mask layer | 30 nm | A 30 nm chromium layer seems to be enough to etch the needed ~ 280 nm amorphous silicon layer. Can probably be even thinner. Also needs to be stripped after etching. |
| Oxide mask layer | 20-25 nm | Needs to be thick enough to function as a mask for the chromium. Selectivities should be good enough to allow the oxide layer to be thinner than the chromium layer. The dry oxide etch is not completely anisotropic. Resultingly, the thinner the oxide layer, the less this effect will be noticeable. |
| PMMA | 70 nm. 30s development in 1:9 deionized water: Isopropanol | a 70 nm PMMA will allow for very fine resolution during electron-beam lithography. The layer only needs to be thick enough to prevent the oxide from etching during the dry oxide etch. Selectivities are good enough to allow a 70 nm oxide layer to suffice. |
| EBL acceleration voltage | To be decided | A low acceleration voltage gives less scattering and finer features. A high acceleration voltage gives less undercut. |
| EBL current | 40pA | Seems to give fine structures. |
| EBL trench exposure dose | 80 | Seems to give good resolution for the needed trench size. |
| EBL hole exposure dose | Around 60 | 60 is fine as long as proximity effect correction is not needed. Mask will however have to be optimized to take this into account. |
| Dry oxide etch | 30s | Enough to open the oxide layer without etching away the PMMA. Recipe is however isotropic. |
| Chromium etch | 6min | Enough to open the chromium layer without etching away the oxide. Recipe is not optimized. |

Table 4.4.7: Best process parameters found for creating a PhC waveguide.

| Cryo etching recipe | Value | Comment |
|---------------------|----------|---|
| Step 1 | 0min | Does not seem necessary. |
| Step 2 | 1min | 1 min seems to be roughly the correct etch time to etch 280 nm of silicon. |
| O_2 | 15 sccm | 15 sccm showed the least sidewall roughness while maintaining good anisotropy during the etch. |
| SF_6 | 60 sccm | Effects of changing SF_6 value were not studied sufficiently. SF_6 gas flow should affect etching rate. |
| RF Power | 8 W | Might not be optimal but has functioned well so far. |
| ICP Power | 600 W | Might not be optimal but has functioned well so far. |
| Temperature | -115 C° | Best results seemed to be obtained at -115 C°. |
| Pressure | 10 mTorr | Pressure should probably be lower. Initial recipe lists pressure as 7.5 mTorr, but this value was too low for the plasma to properly strike. Methods can be developed to circumvent this, such as using a striking pressure. An increased pressure will give less selectivity and an etch which is more sensitive to the oxygen flow. |

Table 4.4.8: Best cryogenic etching recipe parameters found for creating a PhC waveguide.

Chapter 5

Conclusion

5.1 The silicon-on-insulator wafer

The results obtained from processing the SOI wafer are a mixture of positive and negative. The growth rate of the silicon dioxide recipe is roughly 60 nm/minute. This is quite fast, but is not a problem for oxide layers as thick as wanted in this thesis. The results indicate that growing an oxide layer of any target thickness should not be a problem, depending on the accuracy needed. In attempting to grow a 1 μm thick oxide layer, S(T)EM pictures reveal the layer to be roughly 1.08 μm . For photonic crystals this accuracy is well within acceptable boundaries. It might be that the uncertainty is too high to be ideal for the grating section, but in a worst case scenario the oxygen layer can have its thickness reduced through HF-etching.

Results concerning the amorphous silicon growth indicate the same. In attempting to grow a ~ 240 nm layer of amorphous silicon, S(T)EM pictures show that 244 nm were grown. By changing the thickness of the amorphous silicon layer with ± 5 nm, barely any change is noticeable in the photonic crystal properties. It is therefore safe to conclude that this is also within acceptable boundaries. Controlling the thickness of both the amorphous silicon layer as well as the oxide layer is therefore a completely feasible task. Perhaps the most difficult part of this process is verifying the actual thickness of the layers. Using a refractometer is not viable for measuring the thickness of amorphous silicon layers, only to measure the thickness of oxide layers situated on top of silicon. Furthermore, when utilizing S(T)EM to inspect the thickness of the layers, the results might not be correct. The sample might break in such a way that when looking at its cross-section, one or more of the layers may appear thicker than they actually are. This means that the oxide layer thickness might in fact be lower than 1.08 μm and the amorphous silicon layer might be thinner than 244 nm, but since both values are slightly above what's expected, this

is not a problem.

The surface roughness of layers created with the PECVD is found to be between 1.5-3.5 nm, depending on the recipe employed. 3.5 nm surface roughness might be pushing it in terms of scattering losses, but 1.5 nm should be well within acceptable boundaries [23]. The surface roughness is not ideal, and can still be improved through methods such as chemical polishing [24]. The downpart of the results are given by the losses in the amorphous silicon. With extinction coefficients measured to give losses of over 100 dB/cm, they are far outside of adequate boundaries. In other words, were these results to be trusted, no optics would be applicable using the amorphous silicon in question. The silver lining is how the ellipsometry measurements showed the same results for every single amorphous silicon layer grown, even for layers grown following established recipes by experienced groups [20–22]. These groups in question report losses between 2-3 dB/cm, which is two orders of magnitude less than the results obtained in this thesis. Although only one of the three sources cited utilized parameters that were compatible with the PECVD in the cleanroom [22], this recipe in particular was duplicated down to the smallest detail. According to ellipsometry however, the results still showed losses above 100 dB/cm. This raises the question of whether or not the ellipsometry measurements can be trusted, since such a large deviation in results seems unlikely.

As was mentioned, to properly gauge the losses of the amorphous silicon a waveguide would have to be fabricated and subsequently tested [25]. In other words, the chapter regarding fabrication of the SOI wafer can not be concluded as closed. More work characterizing the amorphous silicon layer must be undertaken, before an acceptable SOI wafer can be constructed.

5.2 Finding the good EBL parameters

Many tests surrounding good EBL parameters were done through the thesis. Testing started early and, admittedly, many of the tests were not particularly beneficial until a proper method of etching the samples were in place. Although the experiments conducted in section 4.3 gave helpful training with the electron-beam lithography, the microscopy pictures are not sufficient to conclude whether a certain exposure is optimal or not. However, they give a clear indication of whether or not a certain exposure is in the correct range. Only after etching can the results of changing the EBL parameters be sufficiently characterized.

Furthermore, the ideal parameters of the electron-beam lithography were heavily

reliant on the etching method used. If the thickness of the PMMA layer changes, so do the ideal doses, and most likely also the ideal acceleration voltage and current, evident from the results of section 4.3, which clearly showed that a thinner PMMA layer needed much lower doses to be sufficiently exposed. In practice this means that thorough testing of EBL exposures is not advocated until such a time when the crucial parameters surrounding the etching recipe are in place. The time spent trying to optimize the EBL-exposures was however not in vain, as valuable knowledge surrounding one of the more complicated pieces of equipment was gained.

Two structures of main importance are created in these experiments. Trenches which form a grating used for coupling light into a structure, and holes used to pattern the photonic crystal itself. Throughout the testing performed, it is apparent how the trenches are more rigorous in terms of exposures. So long as the trenches have a specific separation and sufficient stitching-error corrections are made, the trenches are typically well defined.

The holes on the other hand show a serious problem with proximity effects. Methods to overcome the overexposure of holes must be developed. Reducing the dose may lead to holes in close proximity to the center of the array of holes to be properly defined, whereas the holes close to the rim might be underexposed. In other words, one must either find the very core of the proximity effects and try to reduce it, or design the mask taking into account that proximity effects will occur. The second option is likely the easiest. For example, one can make a mask where the array of holes are defined in different layers that are exposed with different doses depending on the amount of holes surrounding the layer in question, as in figure 5.2.1. The less holes that need to be exposed in the near vicinity, the higher the dose.

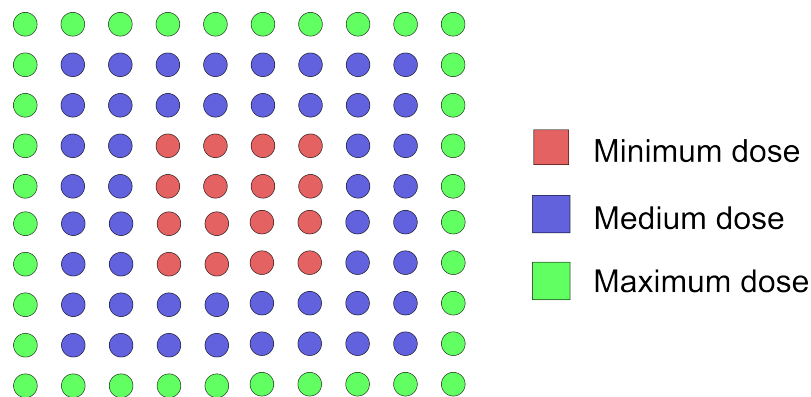


Figure 5.2.1: One theory of how to correct for proximity effects when exposing holes.

Many more parameters might also need to be investigated more thoroughly, but it

is too early to conclude at this point. There is a good chance that the developer used after electron-beam lithography is not ideal for our intents and purposes. The solution utilized was the same for every sample: one part de-ionized water mixed with nine parts isopropanol with a development time of 30 seconds. For one sample, a significantly longer developer time was attempted, with no visible change occurring in the sample. The development time was therefore assumed to be of sufficient quality. Perhaps changing the acceleration voltage, either by increasing it or decreasing it might help to decrease sidewall roughness experienced in most samples after etching. It is evident that holes in general require a lower exposure dose than trenches, and that for a PMMA thickness of around 70 nm, the good hole exposure dose is somewhere between 40-60 whereas a good dose for trenches should be around 80-120.

5.3 The etching recipes

Many different configurations were attempted to optimize the definition of the patterns. The first attempts using a standard silicon etching recipe were quickly discarded. The reason behind this was the lack of established research groups using such recipes for photonic components. The effects of increasing the temperature from -125C to -110C showed a complete disappearance of anisotropy and to heavy undercut during the etching. It is evident that using room-temperatures will not give good results. A lot of this knowledge was not present when the first etching tests were executed, which is why the recipe wasn't immediately discarded.

Another difficult task was to find a suitable mask for the etching process. Once the proper cryogenic etching recipe was found, defining the proper mask was of crucial importance. PMMA was the only photoresist available in the NTNU cleanroom, and it was found to be a very soft photoresist that etches quickly. Many established research groups use harder photoresists, such as ZEP520 [44, 45]. ZEP520 was unfortunately not available at NTNU, and was not available for purchase due to funding considerations. Had ZEP520 been an option, it could have been used as the etching mask during this entire process, and months of work might have been circumvented. Not only could work finding a proper etching mask be avoided, the possible causes of error would be greatly reduced, making it easier to optimise the fabrication process. Hard masks generally perform better during etching and although finding the right mask and etching recipe is tedious work, it is likely worth it in the long run.

The first attempt was to use an oxide mask defined with an HF etch. Unfortunately the oxide etched much too quickly to be used as a mask. Spending enough time optimizing the etching recipe might have given a suitable selectivity between the

oxide and the silicon. This is a process that could take anywhere between a week and a month. In other words, an oxide mask is not the ideal solution when etching with SF_6 and a high ICP power. Instead, chromium was utilized.

Using chromium posed a different problem, which was how to define the pattern in the chromium itself. The solution was given through chemistry. Finding a recipe that etches silicon but not silicon dioxide is difficult because the two substrates are very similar. Chromium and silicon dioxide on the other hand are two completely different substrates. As a result, recipes that etch chromium but not oxide are commonplace. The same goes for recipes that etch silicon but not chromium. A stacked structure is therefore employed to imprint the pattern in the chromium. PMMA is used as a mask for the oxide, the oxide is used as a mask for the chromium, and, finally, the chromium is used as a mask for the silicon.

Attempts at using a wet etch to define the structures were discarded in the end, even though it's not entirely impossible that such a process could work. For it to work however, it is imperative that no layer is thicker than needed, or the anisotropy during the wet etch becomes too great. The dry etching recipes had faults of their own however. The oxide etching recipe, for example, was not anisotropic due to a lack of the required gas in the ICP-RIE. The chromium etching recipe was not likely to be optimized either, and still needs to be properly characterized.

Finally, a substantial amount of research regarding the main silicon etching recipe was conducted. Trying to find the optimal anisotropic etching as well as the one that gave the least sidewall roughness is no easy task. Removing the silicon grass was not feasible, even when changing almost every parameter in the recipe. Reducing the oxygen flow greatly reduced the amount of grass substantially, but this also led to very poor anisotropy and a severe undercut in the etching. Fortunately, the formation of silicon grass was not of utmost importance since the final structure is to be placed on an SOI wafer so there will be no silicon underneath the structure in which the silicon grass can be formed.

Decreasing the gas flows of the recipe served no purpose. Instead of creating a smoother etch, it resulted in holes in the silicon, giving it a sponge-like sidewall. The optimal oxygen value, obtained from the testing in figure 4.4.26 is 15 sccm. This gives both the best anisotropy as well as the best sidewall roughness. Some etching tests also showed that -115 degrees is a good value for the etching, but for longer etches the isotropy of this increased temperature becomes apparent. This can be seen in figure 4.4.28, where a 3 minute and 30 seconds long etch completely

undercut the grating. It is also possible however that this is actually a notching effect, since such a long etch will quickly etch through the silicon and begin etching the oxide. This is a typical situation in which notching will occur. This indicates that a different etching recipe if be utilized should a deep etching of the photonic crystal holes is desired.

In conclusion, the testing performed on the cryogenic etching recipe indicates that previous recipes are likely the culprits of the poor sidewall roughness. By examining, for example, figure 4.4.27 it is apparent that the sidewalls of the chromium layer are also poor. If the sidewall of the mask is poor, then so are naturally the sidewalls of the substrate underneath the mask, which again leads to sidewall roughness. The question then arises whether this is caused by the chromium etching recipe, or by the oxide etching recipe. It is equally plausible however that the original sidewall roughness originated from the PMMA itself, and that a part of the EBL-exposures were not good enough.

In conclusion - the results presented give a good method of fabricating a photonic crystal waveguide. PECVD is utilized to create grow the SOI wafer, and EBL is used to form the patterns in the photoresist. A chromium mask combined with cryogenic ICP-RIE etching defines the structures in the silicon. Although the methods presented can construct an entire photonic crystal waveguide with a coupling section, much work is left to optimize the developed methods to create such a structure which also possesses acceptable losses.

5.4 Future work

Much work is left in order to improve the structures sufficiently such that a photonic crystal waveguide can be rigorously characterized.

Growing low-loss amorphous silicon

The growth of low-loss amorphous silicon is paramount to the functionality of any silicon-based optic component. The work done on this subject in this thesis using ellipsometry as a measurement technique show high losses for the amorphous silicon grown by PECVD. This issue needs to be address by first establishing a more reliable method of identifying the losses, such as fabricating a waveguide and finding its losses [25]. Once a method for identifying the losses of the amorphous silicon has been found, the PECVD recipes can be properly characterized and the optimal amorphous silicon recipe can be found.

Minimizing sidewall roughness

The sidewall roughness present for every sample is a serious problem. Much research must be done in order to minimize this effect. Every etching step should be characterized in the same way that the cryogenic etching recipe was. Furthermore, the EBL exposure parameters should also be studied under resolutions obtained with S(T)EM, to verify that the sidewall roughness or undercut shown for many samples are not attributed to the lithography step. It is plausible that different amorphous silicon recipes act differently during etch. Some recipes might lead to more sidewall roughness than others, which is yet another factor that needs to be analyzed.

Making a multi-step etching process

In the work presented in this thesis, only one silicon etching step was utilized, and every feature is etched to the same extent. The grating, the silicon waveguide, and the photonic crystal holes are all etched to the same depth. This is not sufficient, as the grating is supposed to be much more shallow than the waveguide. The ideal photonic crystal has holes that are etched through the oxide, so the PhC holes should be etched much more than the waveguide. In practice this means that a three-step etching procedure is needed to create the optimal structure. There are several methods of doing this, eg. etching once and then covering the grating with another mask before the next etch. This requires a perfect re-alignment of the mask, which is no trivial ordeal. Furthermore, to deeply etch photonic crystals through the buried oxide, another etching recipe aimed towards etching SiO_2 must likely be found.

Minimizing surface roughness

A high amount of surface roughness leads to considerable losses as a result of scattering. The surface roughness found for the amorphous silicon recipes tested was given to be somewhere between 1.5 - 3.5 nm. These values are far from ideal and in a worst case scenario they would lead to losses that are much too high to make the structure applicable. These surface roughness values were measured for amorphous silicon by ellipsometry before any form of etching had taken place. The results of the etching tests however show that a considerable amount of surface roughness is observed after the etching, so in reality the surface roughness on the structures are probably much higher than those reported by ellipsometry measurements for pure amorphous silicon. Several methods exist to reduce surface roughness, such as electro-chemical polishing [24].

Stripping

Symmetry is one of the cornerstones of a good photonic crystal. If the structure

is not sufficiently symmetric, it will not function. Both oxide and chromium is deposited on top of the SOI wafer, as a result of the etching process. After the masks have been used and the structures defined, these layers need to be removed, so that only the SOI structure remains. There are many methods of doing this, such as a wet chemical etch that removes oxide or chromium, but for this to be applicable the etching must not affect the amorphous silicon at all, or the etch will destroy the structures. Furthermore, removing the oxide with a wet etch is not ideal either since the entire structure is situated on top of SiO_2 , which means the wet etch would start etching away the buried oxide, which could very well lead to structures collapsing.

Creating a functional silicon waveguide

The focal point of this thesis is to fabricate a photonic crystal waveguide. Before such a structure can be characterized, a standard silicon waveguide must be created and confirmed working. Without having a functioning method of guiding light, it is difficult to fully characterize a photonic crystal waveguide. The first step should therefore be to create a standard silicon waveguide and making sure this works, before guiding it into a PhC waveguide.

Optimizing the coupling between the PhC and the silicon waveguide

To be able to properly couple light from the silicon waveguide into the PhC waveguide, it is of great importance that the particular mode that propagates in the silicon waveguide is matched with the mode that will propagate in the PhC waveguide. The two therefore need to be designed such that they are compatible. Furthermore, the coupling section between the PhC and the silicon waveguide must be optimized. Several methods exist, such as simply coupling via a matched waveguide or a Bragg mirror [42] or by evanescent coupling [43], although the latter is probably not applicable for the intents of this thesis.

In conclusion, the groundwork for fabricating photonic components using an SOI platform has been performed. However, there is yet work to be done to optimize the fabrication process and recipes to match the strict requirements for photonic crystal structures or photonic circuits combining different elements.

Appendix A

PECVD recipes

| | NanoLab standard | Test recipe 1 | Test recipe 2 | Test recipe 3 |
|------------------|------------------|---------------|---------------|---------------|
| SiH ₄ | 25 sccm | 50 sccm | 50 sccm | 20 sccm |
| Ar | 475 sccm | 150 sccm | 150 sccm | 0 sccm |
| Pressure | 1000 mTorr | 1000 mtorr | 1000 mtorr | 412 mtorr |
| RF power | 7 W | 50 W | 10 W | 7 W |
| RF frequency | 13.56 Mhz | 13.56 Mhz | 13.56 Mhz | 13.56 mhz |
| Temperature | 300 C° | 300 C° | 300 C° | 300 C° |

Table A.0.1: The characterized Si recipes used to fabricate the SOI wafer.

| Parameter | Standard SiO ₂ recipe |
|------------------|----------------------------------|
| SiH ₄ | 8.5 sccm |
| N ₂ O | 710.0 sccm |
| N ₂ | 161.5 sccm |
| RF power | 20 W |
| RF frequency | 13.56 Mhz |
| Temperature | 300 C° |

Table A.0.2: The SiO₂ recipe used for the SOI wafer.

Appendix B

ICP-RIE recipes

| Parameter | Cl ₂ Cr etch | Standard Si etch | Cl ₂ Si etch |
|------------------|-------------------------|------------------|-------------------------|
| O ₂ | 3.0 sccm | 2.5 sccm | 0sccm |
| SF ₆ | 0 sccm | 12.5 sccm | 0sccm |
| Ar | 0 sccm | 10.0 sccm | 0sccm |
| Cl ₂ | 60 sccm | 0 sccm | 8sccm |
| Chamber pressure | 50 mTorr | 10 mTorr | 9 mTorr |
| RF power | 50 W | 40 W | 40 W |
| ICP power | 0 W | 0 W | 400 W |
| Temperature | -20 C° | 40 C° | 20 C° |
| Helium back flow | 10 Torr | 10 Torr | 10 Torr |

Table B.0.1: Three of the ICP-RIE etching recipes used.

The cryogenic silicon etch is a two-step process and is given in table B.0.2. The chamber pressure of the cryogenic etching recipe has been changed in comparison to its original value. The original values were 2.5 mTorr for the first step and 7.5 mTorr for the second step. Unfortunately, these values were too low to obtain a stable plasma with the ICP-RIE of the NTNU cleanroom, so the pressures were raised to 10mTorr.

| Parameter | step 1 | step 2 |
|-------------------------|----------|----------|
| O₂ | 4 sccm | 13 sccm |
| SF₆ | 12 sccm | 60 sccm |
| Pressure | 10 mTorr | 10 mTorr |
| ICP Power | 600 W | 600 W |
| RIE Power | 15 W | 8 W |
| Temperature | -125 C° | -125 C° |
| Helium back flow | 10 Torr | 10 Torr |

Table B.0.2: The cryogenic ICP-RIE used for etching silicon.

| Parameter | Gas flow rate |
|-------------------------|---------------|
| Ar | 33.0 sccm |
| CHF₃ | 17 sccm |
| Pressure | 50 mTorr |
| RIE Power | 200 W |
| Temperature | 20 C° |
| Helium back flow | 0 Torr |

Table B.0.3: The RIE recipe used for etching oxide.

Appendix C

Tutorial on MPB photonic bands

All of the information needed in order to compile and run MPB can be found at http://ab-initio.mit.edu/wiki/index.php/MIT_Photonic_Bands. The installation guide for windows is found at <http://robert.hovden.googlepages.com/MPBonCygwin.pdf>.

A modest amount of linux knowledge is required to install and make MPB run properly, and from experience many errors will likely be encountered during the installation process. The MPB photonics package uses its own scripting language called CTL, which defines the structure of the crystal and what simulations to run on said structure. The computation a band gap for a photonic crystal composed of silicon rods is explained in this tutorial. The code usually starts with defining the model parameters.

```
(define-param r 0.2)
(define-param eps 3.48*3.48)
(define-param k-interp 4)
(define Si (make dielectric (epsilon eps)))
```

A relative radius r is first set. All sizes are with respect to the fundamental lattice constant a , so in this case the rod radius is set to $0.2a$. A dielectric constant eps is then defined. Furthermore, the amount of points the program will calculate for needs to be decided. This specifies the accuracy of the results and thereby also the smoothness of the graph. Naturally, this has an adverse effect on the computational time of the band diagram.

Finally, the material of which the rods are composed is defined. Once all of this is done, the geometry is created.

```
(set! geometry-lattice (make lattice (size 1 1 no-size)))
```

```
(set! geometry
  (list
    (make cylinder
      (material Si)
      (center 0 0) (radius r) (height infinity))))
```

The first line sets a 2d lattice. The second part of the code makes a cylinder, specifies its material, origin, radius, and height. In this case the height is set to infinite, as the current simulation is bound to two dimensions. The reciprocal points over which the bands are graphed are then set. These are defined in section 2.3, but should be familiar for anyone comfortable with solid state physics. [10].

```
(define Gamma (vector3 0 0 0))
(define X (vector3 0.5 0 0))
(define M (vector3 0.5 0.5 0))
(set! k-points (interpolate k-interp (list Gamma X M Gamma)))
```

The reciprocal points Gamma (Γ), X (χ), and M are defined. MPB then sweeps from Gamma to X to M and then back to Gamma with a set step size, in this case of 0.1. The geometry is now created but the program still needs to be told what to do with it. This is quickly defined:

```
(set-param! resolution 32)
(set-param! num-bands 8)
(run-te)
(run-tm)
```

The resolution parameter defines to what degree the cell is discretized, the num-bands parameter defines how many bands to calculate for, and finally run-te and run-tm runs the simulation to find te and tm bands, respectively.

All the above code is then assembled in a single .ctl file, and executed by typing "mpb file.ctl >& outputfile.out" in the linux command line. This compiles and finds the TE and TM bands and puts them in the file outputfile.out, along with considerable amounts of other data. The information of interest is then extracted and finally plotted in any plotting program. The results are given in figure C.0.1.




Figure C.0.1: The calculated band diagram of silicon rods in air.

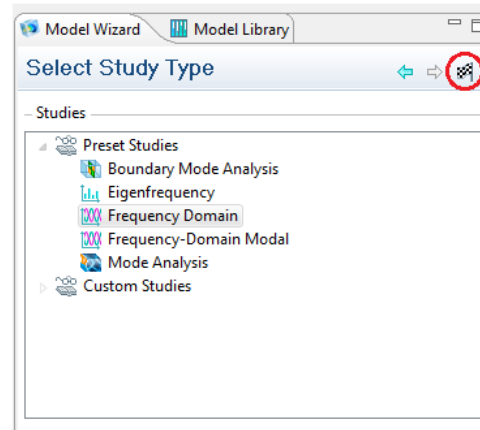
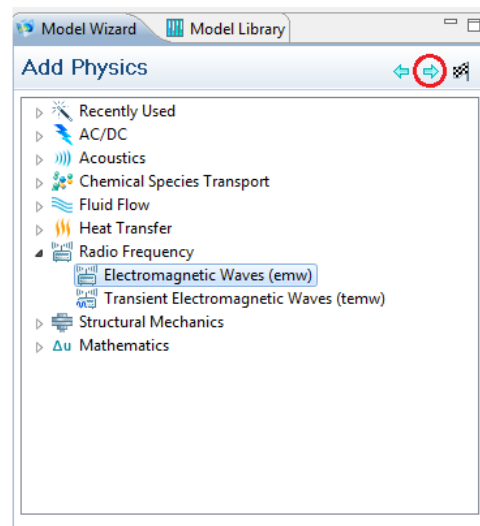
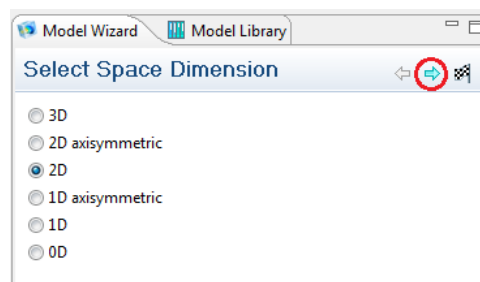
Appendix D

Simulating PhCs with COMSOL 4.2

This guide will assume that the user already has COMSOL 4.2 up and running. When COMSOL is started one must first decide the study environment. How many dimensions, what kind of study, and in which domain the study will be performed. The first screen you will notice is shown in the top part of figure D.0.1. Click "2D" then the arrow to move to the next step.

Select "Radio Frequency" then "Electromagnetic waves" as the middle part of figure D.0.1 to indicate that we will be working with EM waves. Click next.

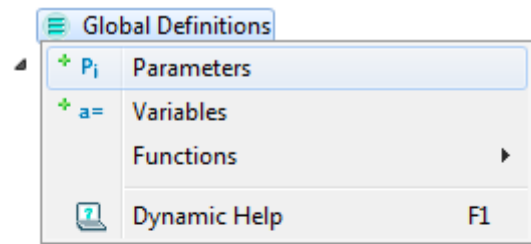
Finally, in the last step we tell the program what to study within the selected Physics. For this particular simulation we will choose "Frequency Domain". The frequency domain study type is the one generally used to study propagation. The eigenfrequency study type was also used to study field distributions, but we will focus on propaga-



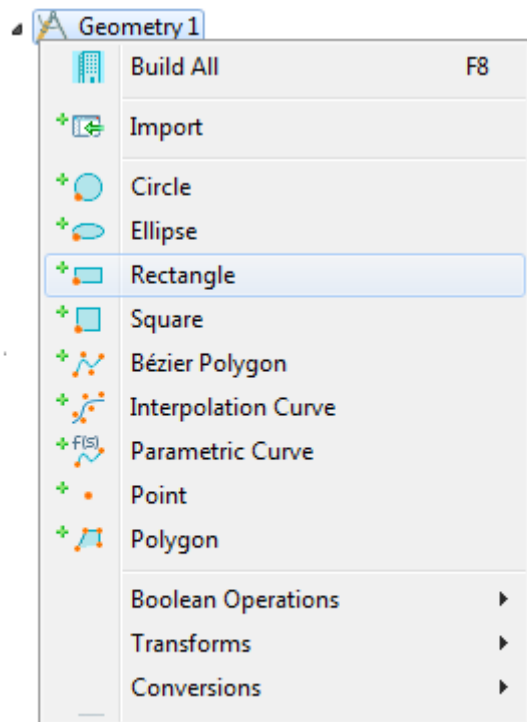
tion here. After selecting frequency domain, click the flag to indicate that you are done selecting your simulation parameters. You can both add physics and change the study type at a later time.

We now define the simulation geometry, and start off by adding some global constants, which makes it easier to change the geometry at a later time, should it be needed. Right click on "Global Definitions" on your toolbar on the lefthand side then click Parameters to add a set of parameters. Click on the newly added parameters tab and enter the parameters as shown in the figure. This makes it easier to go back and change them later on when you wish to change your simulation variables.

Once this is done, right click Geometry 1 and create a Rectangle. Set its base position as *Corner* and input your variable *width* for both width and height, since we are creating a square. Right click Geometry again and add a circle. Have its base be *Center* and input the location $a/2$ for both its y and x position. Right click *Geometry* and click *Build all* and you should see a square with a circle in the bottom left corner. If you do not, click the zoom extents button, (red circle of the figure D.0.2). Right click *Geometry* once again and navigate yourself down to *Transforms* and select *Array*. Click the one circle and click the + icon to add it to the list of geometries to copy. Under the size tab input "10" for x and "5" for y, and enter a for both x, and y displacement.



| Name | Expression |
|-------|------------|
| a | 0.5425[um] |
| width | 10*a |
| rad | 0.2*a |
| l | 1.55[um] |
| f | c_const/l |



Add another circle with radius rad and place it at x location $a/2$ and y location $6.5a$. Then finally make another array, add the circle you just created and type "10" for x size and "4" for y size and again a for both x and y displacement. Your figure should now look like the one below.

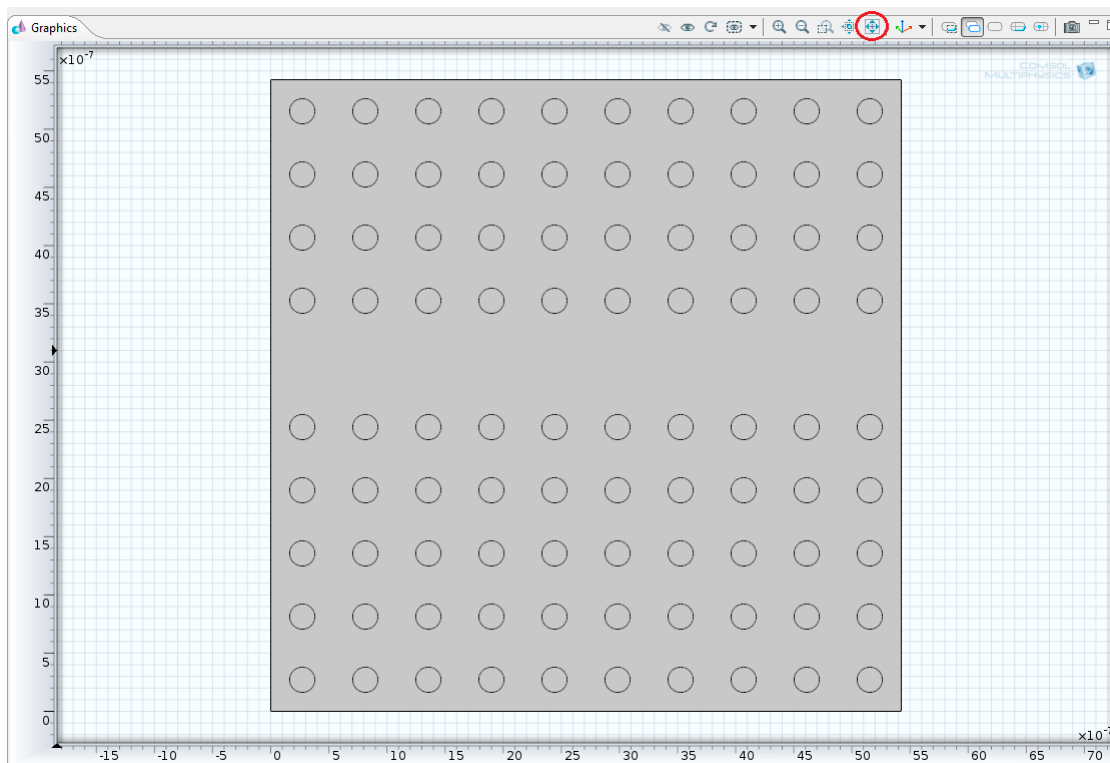


Figure D.0.2

Before continuing we should change one of the settings under Electromagnetic waves \rightarrow Wave Equation and change the *Electric Displacement Field* option from Relative permittivity to Refractive index, as depicted in figure D.0.3.

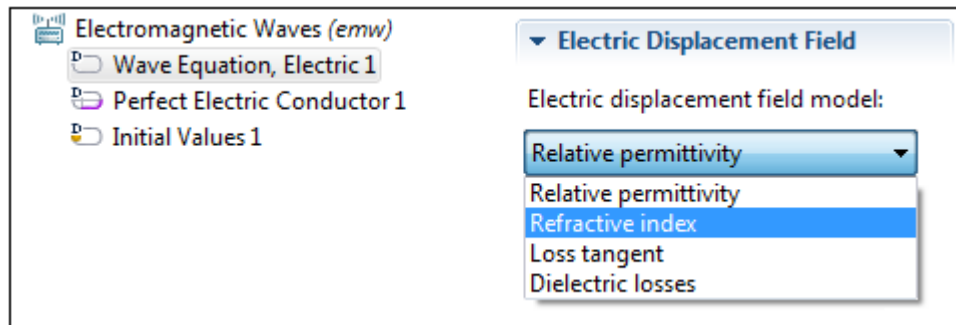


Figure D.0.3

Now that the geometry is done we need to assign materials, some boundary conditions as well as setting up our study. We begin by assigning materials. Right click the materials tab and click the +material button to add a material. The first material added to your model is automatically assigned to your entire geometry, and since we are creating silicon rods in air let us define the first material to be silicon. Select *Material 1* and click F2 in order to rename it to Silicon.

Once this is done, find the *Refractive index* property and change it to 3.48 for silicon. Add another material, right click the square and click the + button to add this object as your second material, which you can rename Air. Change its refractive index to 1.

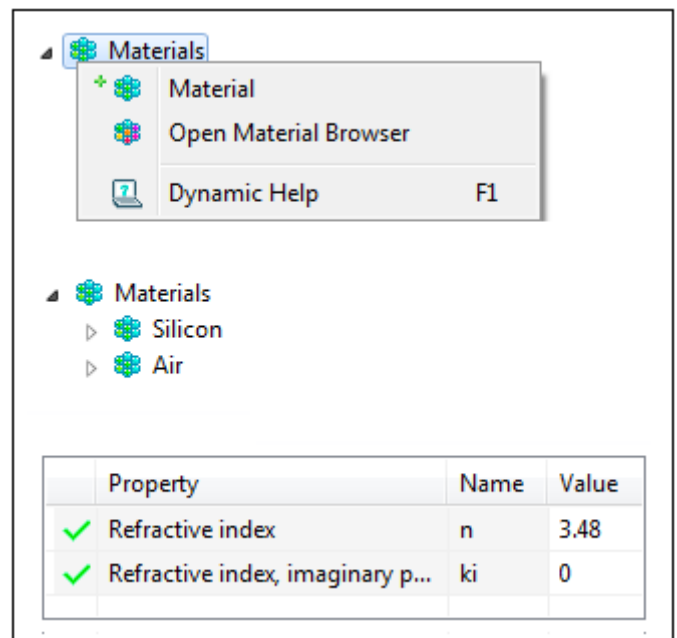


Figure D.0.4

We now impose the boundary conditions. Right click the **Electromagnetic Waves (emw)** tab and add two scattering boundary conditions, as shown in the figure. For your Scattering Boundary Condition 1, under the selection tab choose **All boundaries**, leave the rest of the settings unchanged. On your Scattering Boundary Condition 2 leave the Selection to **Manual**, click the left boundary of your model and click the + sign to add it. Since Scattering Boundary Condition 2 is defined after the Scattering Boundary Condition 1 your selections here will take precedence, and the left boundary will be removed from the Scattering Boundary Condition 1. In the list of boundaries for scattering boundary condition 2 only the boundary numbered 1 should be selected. Now define an incoming wave with $E_z = 1V/m$ from the menu given by selecting scattering boundary condition 2, so that an incoming plane TM wave will hit the photonic crystal from its leftmost boundary. The scattering boundary conditions make the boundary transparent for a scattered wave. It is a simple approach to simulate an incoming plane wave to the crystal and will suffice here.

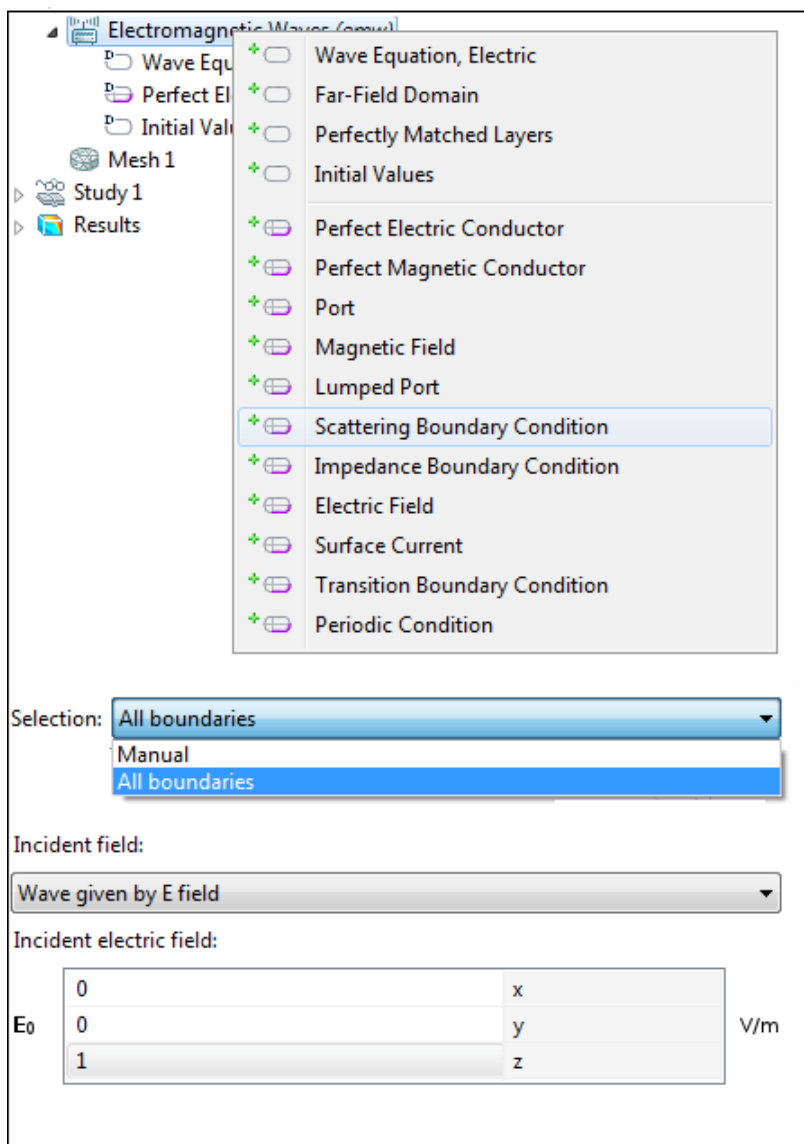


Figure D.0.5

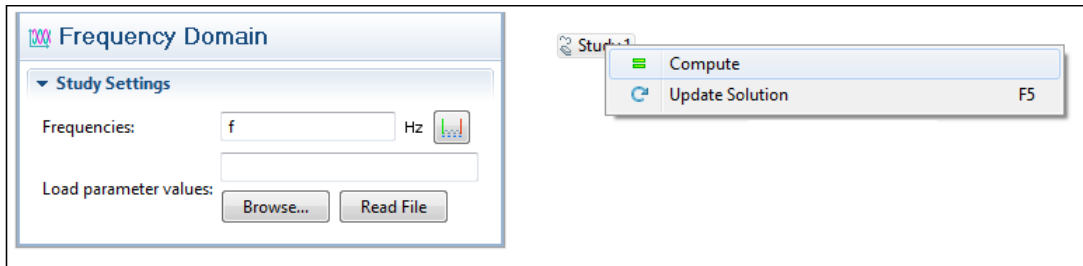


Figure D.0.6

Finally we need to define the frequency study parameters. Open the study node, click "Frequency Domain" and input the frequency f as was defined in our parameter list at the beginning of this tutorial. Finally right click the study node and click "Compute". The simulation may take a few minutes or more depending on your hardware. Once it is done, your graphical results should look like figure D.0.7. If not, well, try again.




Figure D.0.7

This concludes the tutorial on creating the geometry of the simplest photonic crystal and simulating it. After the simulation is complete one typically needs to collect results by looking at various other properties such as the Poyntings vectors, the displacement and magnetic field as well as gathering results such as field intensities at various locations of the model and comparing them with each other. Usually you may end up doing a lot of testing before your crystal will even confine the light within the defect, so careful selections of rod size, wavelength and defect width also need to be performed. Another important thing to note is that this is the idealized case of a two-dimensional structure. Any real photonic crystal will naturally be in three dimensions, something that not only severely changes its properties but makes our simulation space and result analysis a lot more complicated.

Appendix E

Sample parameters

| Parameter | Sample 1 | Sample 2 | Sample 3 | Sample 4 | Sample 5 | Sample 6 |
|---------------------------------|----------|-----------|-----------|-----------|-----------|-----------|
| PMMA | 70 nm | 70 nm | 70 nm | 70 nm | 70 nm | 70 nm |
| SiO ₂ layer | 30 nm | 48 nm | 48 nm | 48 nm | 30 nm | 30 nm |
| Chromium layer | 30 nm | 30 nm | 30 nm | 30 nm | 30 nm | 30 nm |
| SiO ₂ stopping layer | 20 nm | 45 ± 3 nm | 45 ± 3 nm | 45 ± 3 nm | 45 ± 3 nm | 45 ± 3 nm |
| HF etch | 23s | 60s | 50s | 40s | 40s | 50s |
| Chromium etch | 4min | 5min | 5min | 5min | 5min | 5min |
| Cryo etch step 1 | 2min | 2min | 2min | 2min | 2min | 2min |
| Cryo etch step 2 | 2min | 2min | 2min | 2min | 2min | 2min |

Table E.0.1: The samples used when researching the opening of the chromium mask with a wet oxide etch.

| Parameter | Sample 7 | Sample 8 | Sample 9 | Sample 10 | Sample 11 |
|---------------------------------|-----------|-----------|-----------|-----------|------------------|
| PMMA | 70 nm | 70 nm | 100 nm | 100 nm | 100 nm |
| SiO ₂ layer | 30 nm | 30 nm | 20 nm | 20 nm | 20 nm |
| Chromium layer | 30 nm | 30 nm | 25 nm | 25 nm | 25 nm |
| SiO ₂ stopping layer | 45 ± 3 nm | 45 ± 3 nm | 12 ± 3 nm | 12 ± 3 nm | 12 ± 3 nm |
| Dry oxide etch | 45s | 30s | 45s | 44s | 44 +30s stopping |
| Chromium etch | 6min | 6min | 5min | 4min | 5min |
| Cryo etch step 1 | 2min | 2min | 2min | 2min | 2min |
| Cryo etch step 2 | 2min | 2min | 2min | 2min | 2min |

Table E.0.2: The samples used when researching the opening of the chromium mask with a dry oxide etch.

| Parameter | Sample 12 | Sample 13 | Sample 14 | Sample 15 | Sample 16 | Sample 17 |
|---------------|-----------|-----------|-----------|-----------|-----------|-----------|
| STEP 1 | | | | | | |
| Time | 4min | - | 2 min | 2min | 2min | 2min |
| O_2 | 4 sccm | - | 4 sccm | 4 sccm | 4 sccm | 4 sccm |
| SF_6 | 12 sccm | - | 12 sccm | 12 sccm | 12 sccm | 12 sccm |
| Pressure | 10 mTorr | - | 10 mTorr | 10 mTorr | 10 mTorr | 10 mTorr |
| ICP Power | 600 W | - | 600 W | 600 W | 600 W | 600 W |
| RIE Power | 15W | - | 15W | 15W | 15W | 15W |
| Temperature | -125 C° | - | -125 C° | -125 C° | -115 C° | -110 C° |
| STEP 2 | | | | | | |
| Time | - | 4min | 2 min | 2min | 2min | 2min |
| O_2 | - | 13 sccm | 6.5 sccm | 4 sccm | 4 sccm | 4 sccm |
| SF_6 | - | 60 sccm | 60 sccm | 60 sccm | 60 sccm | 60 sccm |
| Pressure | - | 10 mTorr | 10 mTorr | 10 mTorr | 10 mTorr | 10 mTorr |
| ICP Power | - | 600 W | 600 W | 600 W | 600 W | 600 W |
| RIE Power | - | 8 W | 8 W | 8 W | 8 W | 8 W |
| Temperature | - | -125 C° | -125 C° | -125 C° | -115 C° | -110 C° |

Table E.0.3: Process parameters for samples used when researching the cryogenic etching recipe.

APPENDIX E. SAMPLE PARAMETERS

| Parameter | Sample 18 | Sample 19 | Sample 20 | Sample 21 | Sample 22 | Sample 23 | Sample 24 |
|---------------|-----------|-----------------|-----------|-----------|-----------|-----------|-----------|
| STEP 1 | | | | | | | |
| Time | 2min | - | - | 15s | 15s | - | - |
| O_2 | 4 sccm | - | - | 4 sccm | 4 sccm | - | - |
| SF_6 | 12 sccm | - | - | 12 sccm | 12 sccm | - | - |
| Pressure | 10 mTorr | - | - | 10 mTorr | 10 mTorr | - | - |
| ICP Power | 600 W | - | - | 600 W | 600 W | - | - |
| RIE Power | 15 W | - | - | 15 W | 15 W | - | - |
| Temperature | -120 C° | -115 C° | -115 C° | -115 C° | -115 C° | - | - |
| STEP 2 | | | | | | | |
| Time | 2min | 2+2+2min | 1min | 3min 30s | 1min | 1min | 30s |
| O_2 | 13 sccm | 20-15-10 sccm | 13 sccm | 13 sccm | 13 sccm | 4.3 sccm | 13 sccm |
| SF_6 | 60 sccm | 60-60-60 sccm | 60 sccm | 60 sccm | 60 sccm | 20 sccm | 60 sccm |
| Pressure | 10 mTorr | 10-10-10 mTorr | 10 mTorr | 10 mTorr | 10 mTorr | 10 mTorr | 10 mTorr |
| ICP Power | 600 W | 600-600-600 W | 600 W | 600 W | 600 W | 600 W | 600 W |
| RIE Power | 8 W | 8-8-8 W | 8 W | 8 W | 8 W | 8 W | 20 W |
| Temperature | -120 C° | -120-120-120 C° | -115 C° | -115 C° | -115 C° | -115 C° | -115 C° |

Table E.0.4: Process parameters for samples used when researching the cryogenic etching recipe.

Bibliography

- [1] MPB.http://ab-initio.mit.edu/wiki/index.php/MIT_Photonic_Bands
- [2] R. Stoffer, H.J.W.M. Hoekstra, R.M. De Ridder, E. Van Groesen, F.P.H. Van Becku, *Optical and Quantum Electronics*, **32**, pp. 947-961 (2000).
- [3] Marko Lončar, Theodor Doll, Jelena Vučković, Axel Scherer, *Journal of Light-wave Technology*, **18**, 10 (2000).
- [4] Steven G. Johnson, Shanhui Fan, Pierre R. Villeneuve, J.D. Joannopoulos. *Phys. Rev.* **B.60**, 8, (1999).
- [5] *Fundamentals of Photonics*, 2nd Edition by B.E.A. Saleh, M.C. Teich, (2007).
- [6] Wonjoo Suh, M.F. Yanik, Olav Solgaard, Shanhui Fan, *Appl. Phys. Lett.* **82**, 13, (2003).
- [7] Virginie Lousse, Wonjoo Suh, Onur Kilic, Sora Kim, Olav Solgaard, Shanhui Fan, *Optics Express*, **12**, 1575-1582, (2004).
- [8] Jie Gao, James F. McMillan, Ming-Chung Wu, Jiangjun Zheng, Solomon Assefa et al. *Appl. Phys. Lett.* **96**, 051123 (2010).
- [9] Tatsuya Nakamura, Takashi Asano, Kazunobu Kojima, Takanori Kojima, Susumu Noda, *Phys. Rev.* **B 84**, 245309 (2011).
- [10] *Introduction to Solid State Physics*, 8th Edition by Charles Kittel.
- [11] *Photonic Crystals Molding the Flow of Light*, 2nd Edition by John D. Joannopoulos, Steven G. Johnson, Joshua N. Winn, Robert D. Meade
- [12] Richard Martin, Ahmed Sharkawy, Eric Kelmelis. *Photonic Crystals reduce the size of optical sensors*. SPIE newsroom DOI: 10.1117/2.1200610.0413 (2006)
- [13] D. Biallo, A. D’Orazio, M. De Sario, V. Marrocco, V. Petruzzelli, *IEEE: ICTON*, Mo. D2.1, 2006, pp.44-48.

- [14] Lina Shi, Pierre Pottier, Yves-Alain Peter, Maksim Skorobogatiy. Optics Express, **16**, 22, 17962 (2008).
- [15] Tushar Prasad, Daniel M. Mittleman, Vicki L. Colvin. Optical materials, **29**, Issue 1, pp. 56-59 (2006).
- [16] Sriram Venkataraman, Janusz Murakowski, Thomas N. Adam, James Kolodzey, Dennis W. Prather. Journal of Micro/Nanolithography, Mems and Moems 02(04), pp. 248-254.
- [17] M. Mulot, S. Anand, R. Ferrini, B. Wild, R. Houdré, J. Moosburger, A. Forchel. J. Vac. Sci. Technol. **B 22**(2), (2004).
- [18] Ying Liu, Shou Liu, Xiangsu Zhang, Applied Optics, **75**, 3, (2006).
- [19] COMSOL. <http://www.COMSOL.com> (2011).
- [20] Harke et al. IEEE Electronic letters vol 41 no 25 (2005).
- [21] Shunkar Kumar Selvaraja, PhD thesis at Ugent, (2012).
- [22] Sandra Rao, Francesco G. Della Corte, Caterina Summonte. JEOS RP 5, 10039s (2010).
- [23] Silicon Photonics II, 1st Edition by Lockwood, David J. Pavesi, Lorenzo (2011).
- [24] T Dobrev, D T Pham, S S Dimov. Manufacturing engineering centre, cardiff universtiy, Cardiff, CF24 3AA.
- [25] Takahiro Asukai, Makoto Inamoto, Takeo Maruyama, Koichi Iiyama, Keisuke Ohdaira, Hideki Matsumura. 2010 7th IEE international conference, 269-271 (2010).
- [26] Pentti Karioja, Dennis Howe, Applied Optics, Vol. 35, Issue 3, pp. 404-416 (1996)
- [27] Taillaert et al, JQE 38(7), p.949 (2002)
- [28] <http://ntnu.norfab.no/WebForms/Equipment/EquipmentView.aspx?toolId=33>
- [29] <http://ntnu.norfab.no/WebForms/Equipment/EquipmentView.aspx?toolId=32>
- [30] <http://ntnu.norfab.no/WebForms/Equipment/EquipmentView.aspx?toolId=26>
- [31] <http://ntnu.norfab.no/WebForms/Equipment/EquipmentView.aspx?toolId=6>
- [32] <http://ntnu.norfab.no/WebForms/Equipment/EquipmentView.aspx?toolId=23>

- [33] <http://ntnu.norfab.no/WebForms/Equipment/EquipmentView.aspx?toolId=37>
- [34] <http://ntnu.norfab.no/WebForms/Equipment/EquipmentView.aspx?toolId=39>
- [35] <http://ntnu.norfab.no/WebForms/Equipment/EquipmentView.aspx?toolId=10>
- [36] <http://ntnu.norfab.no/WebForms/Equipment/EquipmentView.aspx?toolId=9>
- [37] <http://ntnu.norfab.no/WebForms/Equipment/EquipmentView.aspx?toolId=11>
- [38] <http://ntnu.norfab.no/WebForms/Equipment/EquipmentView.aspx?toolId=42>
- [39] <http://ntnu.norfab.no/WebForms/Equipment/EquipmentView.aspx?toolId=38>
- [40] A.P Milenin, C. Jamois, T. Geppert, U. Gosele, R.B. Wehrspohn. *Microelectronic Engineering* 81 (2005) 15-21.
- [41] Chuan C. Cheng, Alex Scherer, Rong-chung Tyan, Yeshayahu Fai nman, George Witzgall, Eli yablonovitch. *American Vacuum Society* [S0734-211x(97)12006-6] (1997).
- [42] Antti Sayanatjoki, PhD dissertation Helsinki university of technology (2008).
- [43] Paul E. Barclay, Kartik Srinivasan, Matthew Borselli, Oskar Painter. *OPTICS LETTERS* vol. 29, No. 7 (2004).
- [44] L. O. Faolain, X. Yuan, D. McIntyre, S. Thoms, H. Chong, R.M. De La Rie, T.F. Krauss. *Electronics Letters*, vol. 42, No. 25 (2006).
- [45] Maria Makarova, Jelena Vuckovic, Hiroyuki Sanda, Yoshio Nishi. Department of Electrical Engineering, Stanford University, Stanford, CA 94305-4088.
- [46] Soeren Heinze, S(T)EM Introduction Course. <http://www.ntnu.no/nanolab>.
- [47] Bangqui Wu, Ajay Kumar, Sharma Pamarthy. *J. Appl. Phys.* 108, 051101 (2010).
- [48] T.W. Ang, G.T. Reed, A. Vonsovici, A.G.R. Evans, P.R. Routley, T. Blackburn, M.R. Josey, Department of Electronics and Computer Science, University of Southampton, Highfield, SO17 1BJ, UK.
- [49] Jens Høevik, "Characterization of photonic crystals using frequency-domain simulation methods." NTNU (2011).