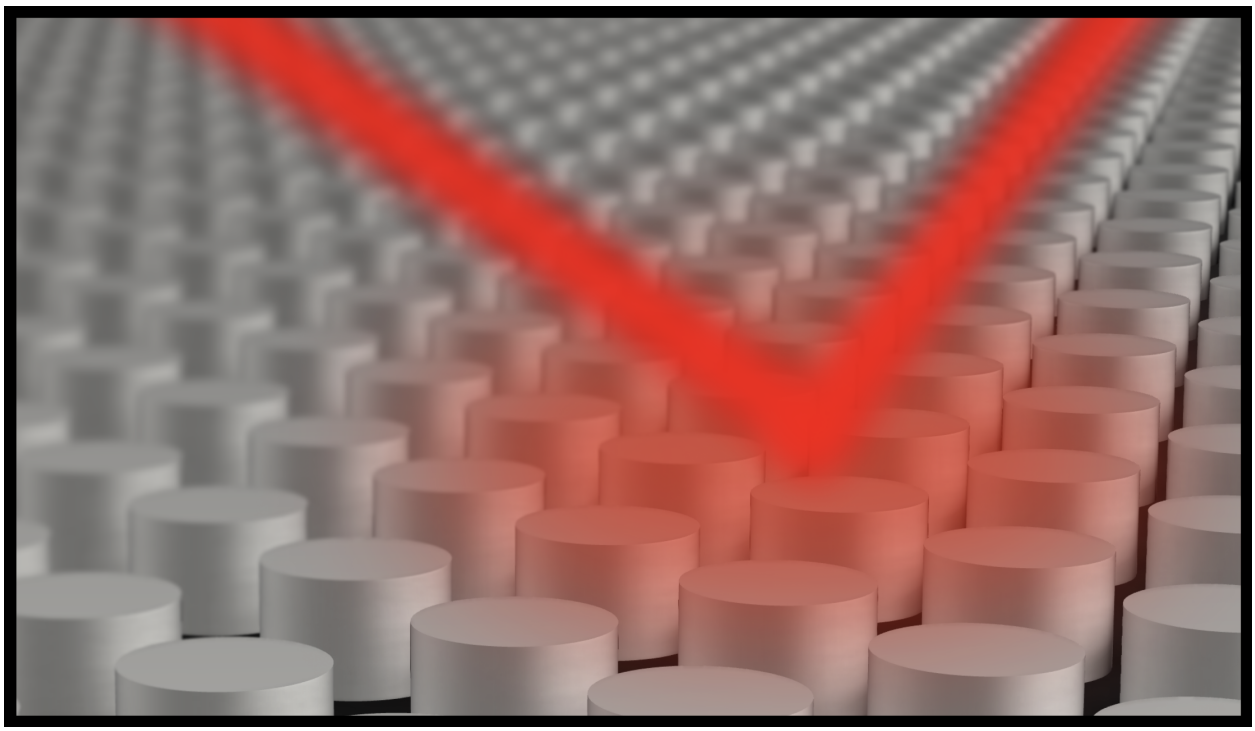

Fabrication optimization of high-reflectivity photonic crystals

Master's thesis in Nanotechnology

ALBERT HANSSON

DEPARTMENT OF MICROTECHNOLOGY AND NANOSCIENCE

CHALMERS UNIVERSITY OF TECHNOLOGY

Gothenburg, Sweden 2026

www.chalmers.se

Master's thesis 2026

Fabrication optimization of high-reflectivity photonic crystals

ALBERT HANSSON

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Microtechnology and Nanoscience
Division of Quantum Technology
Wieczorek Lab
Chalmers University of Technology
Gothenburg, Sweden 2026

Fabrication optimization of high-reflectivity photonic crystals

Supervisor: Alexander Jung
Examiner: Witlef Wieczorek

Master's Thesis 2026
Department of microtechnology and nanoscience
Division of Quantum Technology
Wieczorek Lab
Chalmers University of Technology
SE-412 96 Gothenburg

Cover: A laser beam reflected against a photonic crystal.

Abstract

With cavity optomechanics, one can cool mechanical resonators to their quantum ground state, generate squeezed light, and transduce signals between microwave and optical domains. However, to reach the strong coupling needed for this, a compact cavity with large mechanical zero-point motion is required. High-stress Si_3N_4 membranes are well suited for this due to their high mechanical quality factors, but a conventional dielectric mirror is too heavy to place on such a membrane without degrading its mechanical performance. Photonic crystals, patterned from a single sub-wavelength-thick layer, can act as light-weight, high-reflectivity mirrors through the interference of overlapping Fano resonances, offering a combination of high-quality mechanical resonators with high-reflectivity mirrors. In this work, a membrane-at-the-end fiber Fabry-Pérot cavity is built in which the far mirror consists of an array of amorphous silicon nano-pillars patterned directly onto a suspended Si_3N_4 membrane. The dry etching recipe used to fabricate the pillars was optimized through parameter sweeps of temperature, RF platen power, chamber pressure, and gas composition, and the resulting structures were characterized with scanning electron microscopy and atomic force microscopy to quantify sidewall and top-surface roughness. The optical performance of the fabricated photonic crystals was evaluated both through free-space reflectivity measurements and by placing the samples inside a fiber cavity to measure finesse. Samples fabricated from LPCVD-deposited silicon showed lower surface roughness and consistently outperformed those made from sputtered silicon, reaching finesse values on the order of a few thousand. Comparing the measured finesse to calculations of clipping, scattering, and absorption losses identifies absorption as the dominant loss channel, with an absorption-limited finesse close to the values actually measured. These results demonstrate that a photonic-crystal-patterned membrane can function as an end mirror in a fiber cavity, and indicate that further improvements in reflectivity are primarily a materials challenge rather than a fabrication-precision or geometry problem.

Keywords: Photonic crystal, cavity optomechanics, fiber Fabry-Pérot cavity, amorphous silicon, reactive ion etching.

Acknowledgements

I would like to thank Prof. Witlef Wieczorek for giving me the opportunity to undertake this project and to Alexander Wolfgang Martin Jung and Anastasiia Ciers, my supervisors, for their guidance and valuable feedback throughout this thesis work.

I would also like to direct a special thanks to Fabian, Rishab, Einar, Joram and Tom, whose warmth and good humor made the office a welcoming and enjoyable place to work.

Albert Hansson, Gothenburg, August 2026

Contents

1. Introduction

The field of cavity optomechanics explores the interaction between a mechanical resonator and an optical resonator through radiation pressure. In the prototypical cavity optomechanical system, the mirror that forms one end of the cavity is also a mechanical resonator. When the photons inside the optical resonator impinge on the mechanical resonator, they displace it through the radiation pressure force. When the length of the cavity changes, the resonance frequency, ω_{cav}, of the optical cavity changes. Therefore, $\omega_{\mathrm{cav}} = \omega_{\mathrm{cav}}(x)$, where x is the displacement of the mechanical resonator. Tracking that cavity shift reads out the motion of the mechanical resonator. This interaction can be used to cool the mechanical mode close to its ground state [1], generate squeezed light [2], enable force and displacement measurements at the standard quantum limit [3] or transduce signals between the microwave and optical domains [4].

The single photon optomechanical coupling rate g_0 is an important figure of merit for such a system and depends on how much the cavity frequency shifts per unit of mechanical displacement, and on the size of the mechanical zero-point motion. It is defined as,

$$g_0 = G x_{\mathrm{ZPF}}, \qquad x_{\mathrm{ZPF}} = \sqrt{\frac{\hbar}{2 m_{\mathrm{eff}} \Omega_m}}, \tag{1.1}$$

where x_{ZPF} is the zero-point fluctuation of the displacement, G is the cavity frequency shift per displacement defined as $G = \frac{d\omega_c}{dx}$, and m_{eff} and Ω_m are the effective mass and frequency of the mechanical resonator, respectively. Looking at equation (1.1), it is obvious that the lower the effective mass, the larger the zero-point fluctuation gets. Additionally, a shorter cavity turns a given displacement into a larger fractional length change. Taken together, the optomechanical coupling rate benefits from two things, a short cavity and a light mirror. In this work we will study a system that combines these two features.

For the mechanical resonator, thin films of high-stress, stoichiometric silicon nitride (Si_3N_4) are a common choice. The tensile stress dilutes mechanical dissipation, in an effect called dissipation dilution, and this increases the mechanical quality factor of modes in the kHz to MHz range to high values [5]. Confining these modes within a phononic bandgap, by modulating the stress or the mass density across the membrane, has pushed quality factors into the hundreds of millions [1, 6]. The resulting low thermal decoherence rate makes Si_3N_4 membranes useful both as motion sensors and as platforms for quantum experiments at room temperature.

The obvious way to form an optomechanical system would be to place a mirror in the center of such a membrane. However, a conventional high-reflectivity dielectric mirror is so heavy that it degrades the quality factor, and reduces the size of the zero-point fluctuation. One solution is to pattern the membrane with a photonic crystal, which consists of a two-dimensional array of unit cells with alternating high-index and low-index material. The photonic crystal supports a guided resonance that interferes with an incident plane wave and produces sharp, asymmetric Fano lineshapes in the reflectivity response [7, 8, 9, 10]. Tuned correctly, the Fano resonances reach near-unity reflectivity from a single patterned layer a fraction of a wavelength thick, making the membrane itself locally perform as a mirror. Reported reflectivities exceed 99.9 %, with cavity finesses in the tens of thousands, but only in macroscopic cavities several millimeters to centimeters long [11].

To achieve a short cavity, this thesis uses a fiber-based cavity, where light is guided to the membrane through an optical fiber with a concave, mirror-coated end facet. Although the fiber-based cavity solves the cavity length problem, it comes with other issues that need to be taken into consideration. Due to the necessary mode matching to the single mode fiber, and short cavity lengths, the cavity waist on the membrane is on the order of less than 10 micrometers. Such a tightly focused beam carries a wide spread of wavevectors, and the interference based photonic crystal mirror loses in reflectivity as the angle of incidence becomes more oblique. Due to these issues, fiber cavities have up to now only been used with membrane-in-the-middle configurations [12], where the membrane sits inside a cavity formed by two other, macroscopic mirrors, rather than forming a cavity mirror itself.

The goal of this project is to optimize the fabrication of the photonic crystal and then verify its performance in a membrane-at-the-end optomechanical fiber cavity setup. Instead of etching holes into the membrane, as done in previous implementations, the design here uses amorphous silicon pillars placed on top of the Si_3N_4 membrane. We aim for a reflectivity of more than 99.9 %, which corresponds to a finesse on the order of 10^4. To reach this goal, the focus in this work will lie on optimizing the quality of the photonic crystal fabrication process. This will be done by

- optimizing the architecture of the photonic crystal by investigating its design parameters,
- finding a dry etching process that produces smooth photonic crystal pillars with vertical sidewalls, and
- characterizing the pillars with scanning electron microscopy and atomic force microscopy.

Once the photonic crystal has been fabricated, it will be integrated with a fiber mirror to form a cavity. Since the reflectivity of the fiber mirror is a known quantity, we can use the cavity measurement to extract the reflectivity of the photonic crystal as a function of wavelength. The cavity measurement gives a much more sensitive picture of the mirror's quality than the reflectivity measurement alone, since even tiny imperfections show up clearly in how well the cavity holds onto light. Together, these measurements make it possible to evaluate whether the photonic crystal works, and what is limiting its performance.

The thesis is structured as follows. Chapter 2 introduces the relevant theory, covering photonic crystals and the Fano resonances that allow them to act as lightweight mirrors, followed by the fiber Fabry-Pérot cavities used to characterize them. Chapter 3 describes the microfabrication process, with focus on the dry etching process developed to pattern the amorphous silicon pillars, and the performance metrics used to evaluate the etch. The simulations used to design the photonic crystal and estimate its losses are presented in Chapter 4. Chapter 5 describes the experimental methods used to characterize the fabricated devices, both through direct reflectivity measurements and through the fiber cavity. These methods are then used in Chapter 6 to present the results, covering the microfabrication optimization as well as the measured reflectivity and finesse. Finally, Chapter 7 concludes the thesis with a discussion of the results and an outlook on possible directions for future work.

2. Theory

This chapter introduces photonic crystals and how they can be used as light-weight mirrors. Afterwards, Fabry-Pérot cavities and the parameters used to characterize them are introduced.

2.1 Photonic crystals

2.1.1 Basic properties

A photonic crystal is a structure in which the dielectric function $\varepsilon(\boldsymbol{r})$ is modulated periodically in space. This project will be limited to two-dimensional photonic crystals, where the periodicity is confined to the xy-plane. Figure 2.1a shows an example of such a structure with the corresponding unit cell in figure 2.1b. This example is an illustration of pillars with a large refractive index surrounded by a lower refractive index ($\varepsilon = 1$) material (e.g. air).

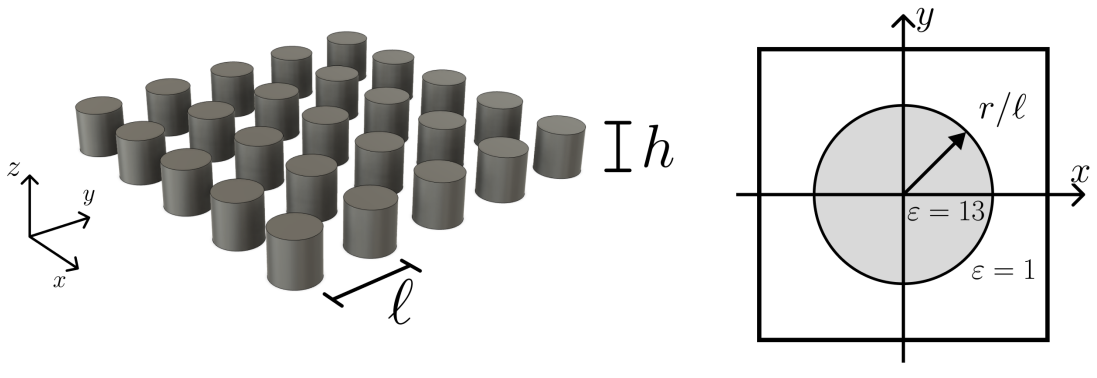

(a) **(b)**

Figure 2.1: (a) Illustration of a photonic crystal with pillars with a large refractive index. (b) The unit cell of the photonic crystal in (a). The white area has a lower dielectric constant while the shaded area has a larger one.

Because $\varepsilon(\boldsymbol{r})$ is periodic, Bloch's theorem restricts the available modes [9]. Every magnetic field mode $\boldsymbol{H_k}(\boldsymbol{r})$ can be written as a combination of a plane wave and a periodic function $\boldsymbol{u_k}(\boldsymbol{r})$ with the periodicity of the lattice: $\boldsymbol{H_k}(\boldsymbol{r}) = e^{i\boldsymbol{k}\cdot\boldsymbol{r}}\boldsymbol{u_k}(\boldsymbol{r})$, where $\boldsymbol{u_k}(\boldsymbol{r}+\boldsymbol{R}) = \boldsymbol{u_k}(\boldsymbol{r})$ for any lattice vector $\boldsymbol{R}$ and $\boldsymbol{k}$ is the wavevector [9]. All modes are labeled with a wave vector $\boldsymbol{k}$ defined in reciprocal space, and all physically distinct values of $\boldsymbol{k}$ are contained within the Brillouin zone.

In figure 2.2, the first Brillouin zone (FBZ) of a square lattice is depicted. Because of the symmetry of the lattice, many modes within the FBZ are equivalent. The smallest region from which the entire FBZ can be reconstructed is the irreducible Brillouin zone (IBZ), shown as the shaded triangle in figure 2.2 bounded by the high-symmetry points Γ, X and M. Because every other $\boldsymbol{k}$-point in the FBZ has a symmetry-equivalent point on or within this triangle, it is sufficient to generate the band structure by computing the dispersion relation along the boundary of the IBZ [9]. By doing so, following the positive orientation of the IBZ, $\Gamma \to X \to M \to \Gamma$, all band extrema of the crystal are captured.

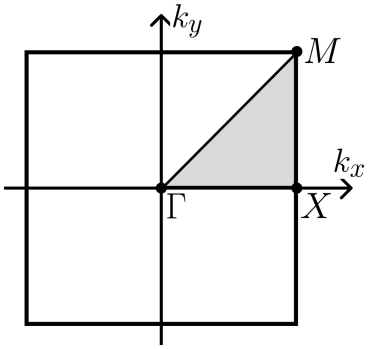

Figure 2.2: The Brillouin zone of a square-shaped unit cell with the irreducible Brillouin zone marked in gray.

A real photonic crystal, however, is not uniform in the z direction, but has a finite thickness. Yet, there are no Bloch modes in the z direction as there is no periodicity in $\varepsilon(z)$. Instead, confinement in the z direction is dependent on total internal reflection. This introduces the *light line*, $\omega = ck_{\parallel}$, which separates the modes into guided resonances (above the light line) and guided modes (below the light line). A guided resonance can couple to a plane wave ($k_x = k_y = 0$) at the Γ point [8].

The existence of guided resonances is the origin of the sharp features observed when light is incident on a photonic crystal from the outside (see figure 2.3) [8]. The

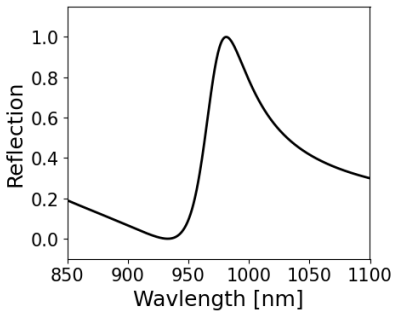

Figure 2.3: The reflectivity curve of a Fano resonance that appears due to the interference between a plane wave and a photonic crystal mode.

incoming plane wave has two ways of interacting with the photonic crystal. The first is the direct non-resonant way, where the incoming wave reflects from and transmits through the crystal. Here the photonic crystal is treated as a simple slab of dielectric material with an effective refractive index. The second way of interacting is the resonant way, in which part of the light in the plane wave couples into the guided resonance. The guided resonance couples back the light with a delay. A simple model to describe the lineshape for this interaction is called the Fano resonance,

$$t = t_d + f\frac{\gamma}{i(\omega - \omega_0) + \gamma}$$

$$r = r_d \pm f\frac{\gamma}{i(\omega - \omega_0) + \gamma},$$

where t_d and r_d are the direct transmission and reflection coefficients, ω_0 is the center frequency and γ is the width of the Lorentzian from the resonance and f is the complex amplitude of the resonant mode [8].

Since the directly reflected and time-delayed light radiate into the same channel, they interfere. Therefore, by sweeping across the resonance frequency of the leaky mode, the relative phase between the non-resonant and the resonant contributions goes through a full cycle. This produces constructive interference on one side of the resonance frequency and destructive resonance on the other. The result is the characteristic asymmetric Fano resonance lineshape seen in figure 2.3 [8, 7].

2.1.2 Overlapping Fano resonances

If we choose the right photonic crystal parameters, r, ℓ and h (see figure 2.1), we can engineer two Fano modes to lie close to each other and have opposite symmetry.

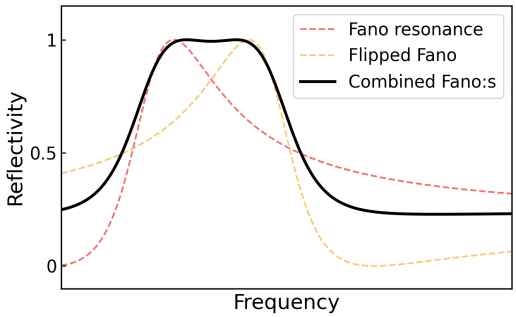

Figure 2.4: The combination of two Fano resonances forms a region of high reflectivity that is broader than a single Fano resonance.

Figure 2.4 gives a simplified illustration of such a situation, in which two Fano resonances overlap and create a broader region of high reflectivity compared to them both separated. In this way we can increase the wavelength region with large reflectivity significantly. This allows us to achieve large reflectivities for thin membranes.

2.2 Fiber Fabry-Pérot cavities

2.2.1 Standard Fabry-Pérot cavity

The simplest realization of an optical cavity is the ideal Fabry-Pérot cavity (FPC), which consists of two highly reflective planar mirrors facing each other separated by the *cavity length*, L_c (see figure 2.5a). An incoming field, E_{in}, resonantly interferes with itself inside the cavity and thereby builds up a mode amplitude, E. The criterion for resonance is to have a cavity round-trip phase, i.e. the phase of a photon that has traveled back and forth in the cavity, equal to a multiple of 2π. This is achieved when $L_c = m \cdot \frac{\lambda}{2}$, where λ is the wavelength of the incoming light and m is the mode number and a positive integer [13].

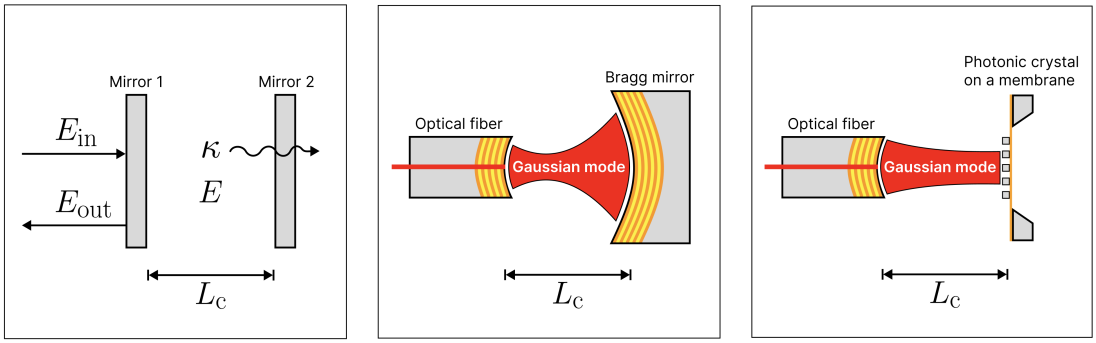

(a) **(b)** **(c)**

Figure 2.5: (a) A Fabry-Pérot cavity with ingoing and outgoing fields E_{in} and E_{out}, mode amplitude E, length L_c and energy decay rate κ. (b) A fiber Fabry-Pérot cavity where one of the mirrors is switched for a coated optical fiber facet. (c) A fiber Fabry-Pérot cavity where the Bragg mirror has been replaced with a wavelength dependent photonic crystal.

The round-trip time is calculated by

$$\tau_R = \frac{2L_{\text{c}}}{c}$$

where c denotes the speed of light in air [14]. Because a resonance occurs every time the round-trip phase changes by 2π, modes are equally spaced in frequency by the so called *free spectral range* [14],

$$\text{FSR} = \frac{1}{\tau_R} = \frac{c}{2L_{\text{c}}}.$$

In the following, ω_c denotes the cavity frequency and is taken as one of the n cavity modes.

An ideal Fabry-Pérot cavity would be able to trap a photon for an infinite amount of time. Yet, real cavities lose photons each round trip through mirror transmission, scattering and absorption. These losses are described by an energy decay rate, κ, such that the number of photons trapped inside the cavity decays as $e^{-\kappa t}$. With the decay rate, the average time a photon spends inside the cavity before it is lost, is

$$\tau = \frac{1}{\kappa}.$$

Another quantity that is used to characterize a cavity is the optical quality factor, Q, which is defined as

$$Q = \omega_c \tau,$$

and expresses how well the cavity confines energy [13, 14]. In this project, another figure of merit, the finesse $\mathcal{F}$, will be used to characterize the cavity. The finesse is defined as the ratio of the FSR to the linewidth, $\frac{\kappa}{2\pi}$, and can be rewritten to only depend on the reflectivity,

$$\mathcal{F} = \frac{\text{FSR}}{\kappa/2\pi} = \frac{\pi\sqrt[4]{R_1 R_2}}{1 - \sqrt{R_1 R_2}}, \tag{2.1}$$

where R_1 and R_2 are the reflectivities of the mirrors [13]. Physically, $\mathcal{F}$ represents the average number of round-trips a photon makes before it escapes the cavity. Because the finesse does not depend on the cavity length, it is natural to use it when comparing cavities of different geometries.

2.2.2 Fiber Fabry-Pérot cavity

We now replace both mirrors with distributed Bragg reflectors (DBR:s). However, one DBR is in this case special. It is an optical fiber with a concave indentation. Afterwards, it is covered with a dielectric coating to form a DBR as well. The resulting cavity is called fiber Fabry-Pérot cavity (FFPC) and is illustrated in figure 2.5b. Since we now deal with curved mirrors, the eigenmodes are Gaussian modes, with finite waists rather than the plane-wave mode captured by the ideal cavity [13].

This geometry is of interest because it allows the fiber to be directly coupled to the cavity. Additionally, the cavity volume is reduced, which is important to be able to involve a structure such as a photonic crystal on a membrane, into the cavity.

2.2.3 Fiber Fabry-Pérot cavity with a photonic crystal

The system studied in this thesis is an FFPC where the Bragg mirror has been replaced with a photonic crystal sitting on top of a suspended membrane. A schematic illustration of this system is found in figure 2.5c. Rather than relying on a relatively thick stack of dielectric layers, the photonic crystal achieves high reflectivity from a single patterned layer with a thickness that is only a fraction of a wavelength [14]. Its small weight is what makes it suitable for performing optomechanics, by placing it on a thin suspended membrane. Compared to the Bragg mirror which has a broad

frequency range over which its reflectivity is very close to unity, the photonic crystal, achieving its high reflectivity through Fano resonances, is more wavelength dependent and typically less reflective. Additionally, it introduces new loss mechanisms such as clipping loss, where if the waist of the mode living inside the cavity is larger than the photonic crystal, the parts of the mode waist outside the photonic crystal are lost. Further discussion of the losses tied to the fiber Fabry-Pérot cavity is found in section 4.4.

3. Microfabrication

In this chapter, the most important tools that have been used in the microfabrication of devices are described. Also, the entire microfabrication process of the devices is described in detail.

3.1 Inductively Coupled Plasma Reactive Ion Etching

Inductively Coupled Plasma Reactive Ion Etching (ICP-RIE) is a central technique in nanofabrication, as it enables precise fabrication of nanoscale features. The technique combines the chemical processes of reactive-gas chemistry with the directional energy transfer of physical ion bombardment to etch patterns into various materials. Its advantage lies in the ability to achieve controllable and anisotropic etch profiles by finding the right balance between physical and chemical etching [15].

In our optimized ICP-RIE process, hydrogen bromide (HBr) and chlorine gas (Cl_2) are used as *etchants*, oxygen gas (O_2) is used as *passivation* gas and the etched substrate is amorphous silicon (a-Si).

3.1.1 Working principle

An ICP-RIE system (see figure 3.1) combines two independently powered radio-frequency (RF) sources. One RF source is responsible for delivering energy into a plasma created above the sample, by sending an alternating current through coils surrounding the chamber. By the flux law, this alternating magnetic RF field gives rise to an electric field inside the chamber, accelerating the electrons in the gas volume [16]. The second RF source is connected to the substrate platen holding the sample. When this RF source is switched on, the ions and electrons are accelerated in the field. Because electrons are far more mobile than ions, they accumulate

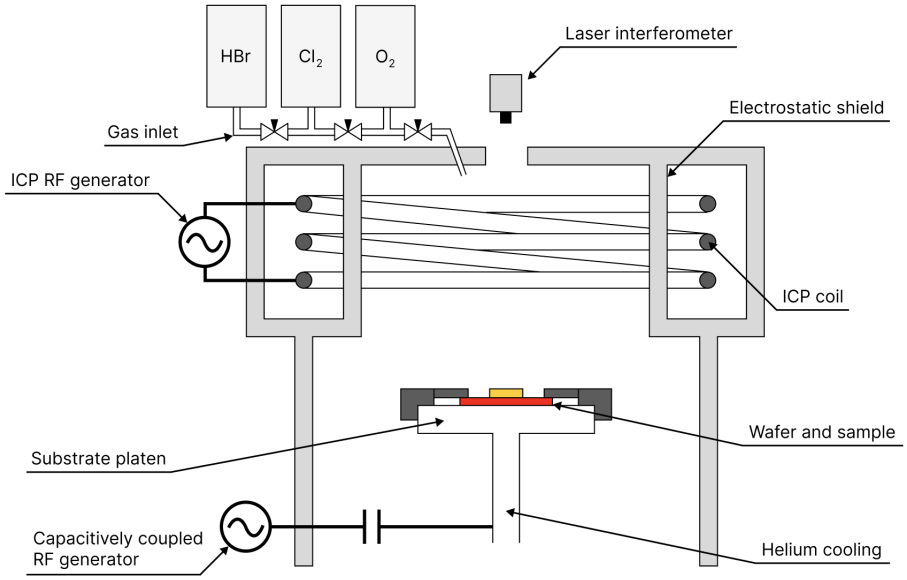

Figure 3.1: Schematic representation of the Oxford PlasmaPro 100 ICP/RIE used for optimizing the dry etching recipe in this project.

on the substrate during each RF half-cycle, which causes a negative DC bias on the platen. This negative bias then accelerates the positively charged ions from the bulk plasma to the surface of the substrate. To prevent the substrate from overheating under the bombardment of etchants, backside cooling of the Si carrier wafer is achieved by letting helium flow between the carrier and the platen. Helium is well suited for this purpose due to its high thermal conductivity. We apply Santovac oil between the Si carrier and the samples for proper thermal contact. Figure 3.1 shows a schematic drawing of the specific ICP-RIE system used in this project.

3.1.2 Physics and chemistry of material detachment

Material removal in ICP-RIE is governed by a sequence of consecutive surface processes. We distinguish these different steps, because the etch results depend on the speed of each process. The speed can in turn be influenced by choosing the right process parameter values.

Adsorption

The first step in material removal is the arrival and adsorption of reactive chemicals onto the substrate. Upon arrival, the radicals can either be reflected or adsorbed by the sample, depending on the type of radical as well as the temperature of the

sample. In the case of HBr/Cl_2 chemistry, the bromine and chlorine radicals stick to the silicon surface by chemisorption, meaning that they form chemical bonds with the silicon atoms [17]. Since there are in principle four sites for a halogen to bind to a silicon atom, mono-, di-, tri- and tetrahalide configurations are formed between the silicon and the bromine/chlorine radicals.

Surface reactions

Once the halogen radicals have adsorbed on the silicon surface, they will start breaking Si-Si bonds. The general reaction for silicon halide formation is

$$\mathrm{Si(s)} + 4\mathrm{X(g)} \rightarrow \mathrm{SiX_4(g)}, \text{ where X = Br or Cl.}$$

Although this reaction is thermodynamically favorable for both bromine and chlorine, a proportionate activation energy is needed to overcome the barrier where the Si-Si bonds break. Translated into thermal energy, the Si-Si bonds start breaking at around 600 °C, which is above the sidewall temperature of a normal ICP-RIE reactor. This means that there is insufficient thermal energy to overcome the activation barrier. Yet, the observed etch rates in ICP-RIE are orders of magnitude higher than what this thermal limit would allow. The necessary energy is supplied by the ion bombardment. The effect of ion bombardment is discussed further in section 3.1.3. After the reactions have taken place, the volatile byproducts must be removed from the surface.

Desorption and removal of byproducts

Once a volatile reaction byproduct has formed, it must desorb from the surface to make place for new reactions. Both $SiBr_4$ and $SiCl_4$ are volatile compounds, but their volatilities differ significantly. Volatility is commonly assessed by comparing boiling points, as a higher boiling point generally corresponds to a lower volatility. $SiCl_4$ has a boiling point of 57.6 °C compared to 154 °C for $SiBr_4$ [18]. Therefore, $SiBr_4$ remains on the surface for longer before desorbing. This leads to the formation of a thin byproduct layer that acts as an addition to the passivation layer.

3.1.3 Ion-assisted etching

In 1979, the chemist John Coburn and physicist Harold Winters discovered the benefits of letting both ions and reactive radicals act on a surface simultaneously (see figure 3.2b) [19]. They observed that the etch rate of combined physical and

chemical etching is considerably higher than the etch rate of the two separated. In figure 3.2a, the four most important material removal processes are illustrated and the ones encircled by a dashed line represent the ion-assisted processes whereas the others represent either physical or chemical etching.

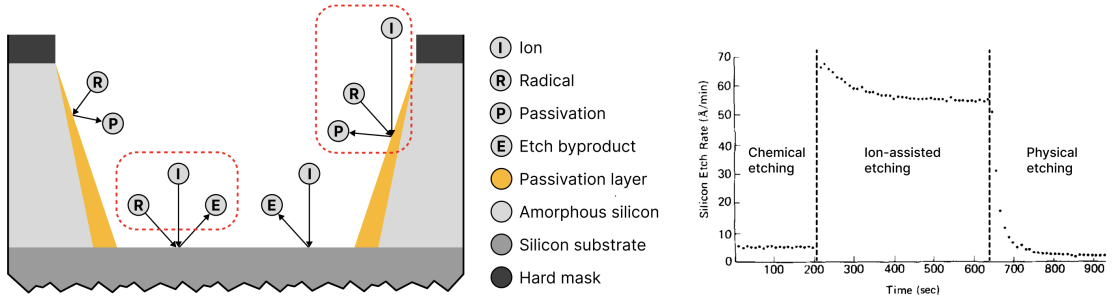

(a) **(b)**

Figure 3.2: (a) The two main etching mechanisms are radical desorption (illustrated on the left structure wall) and ion bombardment (illustrated to the right on the substrate floor). The two interactions encircled by red dashed lines are ion-assisted mechanism where ion bombardment enhances the chemical etch rate. (b) The classic Corburn and Winters experiment shows the synergetic effect of ion-assisted etching. The figure is adapted from Ref [19].

The explanation for the synergy of ion-assisted etching is the modification of the surface through ion bombardment. When an energetic ion strikes the surface, it creates a disordered zone at the very top of the material. This ion-damaged zone has a much higher density of dangling bonds, and is therefore more susceptible to halogen absorption [20]. The ion bombardment reduces the activation energy for the reactions driving the material removal by converting the inert crystalline surface to a chemically activated one. This means that the physical and chemical etching processes work in a synergetic way to overcome the sum of their separated efficiencies.

3.1.4 Process parameters

The large number of independent and tunable parameters during ICP-RIE etching is a blessing and a curse. Each parameter affects the plasma chemistry and surface processes in distinct ways, which results in a large process parameter space. This parameter space then has to be navigated to optimize for the sought-after features of the device, like etch uniformity, anisotropy or surface smoothness. In this section, the impact of the main parameters are discussed.

Gas composition

In HBr-Cl_2-O_2 based etching, chlorine is the primary etchant in terms of producing volatile byproducts. In the plasma, the chlorine and hydrogen bromine dissociate and produce Cl^-, Br^-, and H^+ radicals. The halides react with silicon to form silicon tetrahalides. However, since the bromine is significantly less reactive than the chlorine, the etch rate by Br^- radicals is comparatively low, which means that etching under HBr-dominated conditions is largely ion-driven, i.e. physical. This has two important consequences. First, it reduces the isotropic chemical etching element, contributing to improved profile anisotropy. Second, it dramatically increases the selectivity of silicon (the etched material) over silicon dioxide (the mask material) since the bromine radicals have low reactivity with SiO_2.

When creating an etching recipe, the relative flow rates between the different gases are more important than the total flow. The relative flow between HBr, Cl_2 and O_2 determines the chemical composition of the plasma and, consequently, the nature of all surface reactions. Increasing the fraction of Cl_2 raises the density of Cl^- radicals, which increases the etch rate but at the same time increases the isotropic chemical etching element. On the contrary, shifting the balance toward HBr reduces chemical etch rate and relies more on ion bombardment. This would theoretically generate a more anisotropic structure at the cost of an overall lower etch rate.

The oxygen fraction, or in the case of HBr/Cl_2-based etching more specifically the ratio between HBr and O_2, controls the thickness and stability of the passivation layer. Insufficient passivation will allow chemical etching of the sidewalls and thereby reduce the anisotropy, whereas excessive passivation will reduce the etch rate. Therefore, balance with respect to both the ion energy and the HBr flow is important, as the hydrogen radicals compete for surface space and the ion energy determines whether the passivation can be removed to maintain the etching process.

Chamber pressure

Once a constant gas flow is chosen, the pressure is controlled by the rate at which these gases, along with the detached material, are pumped out of the chamber. Higher pressure means that the overall gas-phase particle density is large, which generally raises the chemical etch rate since more particles impinge the substrate.

On the other hand, higher pressure also allows for more ion-to-ion collisions which broaden the angular distribution of the ion flux arriving at the substrate. As a direct consequence, the anisotropy reduces. Conversely, setting the pressure low will guide the ions straight down to the substrate and yield good directionality. The downside of a smaller pressure set point is that there will be a smaller supply of reactive species in the chamber, which may result in an etch rate too low to be useful.

ICP RF power

The ICP RF power controls the density of the plasma. Increasing the ICP power raises the flux of both ions and reactive species arriving at the substrate. This increases the etch rate, since both the supply of etchants and the ion flux surface activation are enhanced. The ICP power also controls the ratio of different radicals at the substrate surface since it controls the energy distribution of the electrons responsible for the plasma-maintaining dissociations. Nevertheless, this effect is relatively subtle in comparison to the gas flow control. The primary purpose of controlling the ICP is therefore to set the total throughput of the process without affecting the ion energy.

RF platen power

The RF platen power controls the self-bias voltage that builds up at the substrate, and thereby also the kinetic energy that ions gain when they accelerate toward the substrate surface. High platen power means higher ion energies and therefore also a higher etch rate, which comes with several consequences. Since the ions deposit more energy in the collision with the top layer of the substrate, the passivation layer is removed more aggressively. This removes the passivation layer even at an increased O_2-to-HBr ratio. At very high bias voltages, the process shifts toward a physical etching regime dominated by momentum transfer-based removal of material. In this regime, substrate damage becomes the most significant source of defects and the selectivity decreases. On the contrary, too low platen power will cause the passivation layer to build up faster than it can be removed, which will terminate the etching process. This means that there is a limit for the platen power under which the etching cannot continue.

Substrate temperature

The substrate temperature is controlled by the backside cooling of the platen on which the wafer sits and is an important control parameter to ensure balance between chemical and physical etching. Because the plasma transfers heat to the sample, it

is important to actively stabilize, and in most processes lower, the temperature so as to not overheat the substrate and thereby impair the etching process. With a higher temperature, a more chemically driven etching process is achieved, because thermally activated bond-breaking and material desorption is promoted. Additionally, the passivation layer will be less stable since a higher temperature will cause SiO_XBr_Y compounds to also desorb from the surface.

3.1.5 Performance metrics

Anisotropy in etching refers to the degree to which the material removal is directional. A perfectly anisotropic etch removes material only in the vertical direction, producing straight sidewalls, while a perfectly isotropic etch removes material equally in all directions resulting in a rounded, undercut profile. In practice, ICP-RIE achieves a high but not perfect anisotropy. However in comparison to wet etching, where isotropy results in undercutting, ICP-RIE typically generates slanted sidewalls.

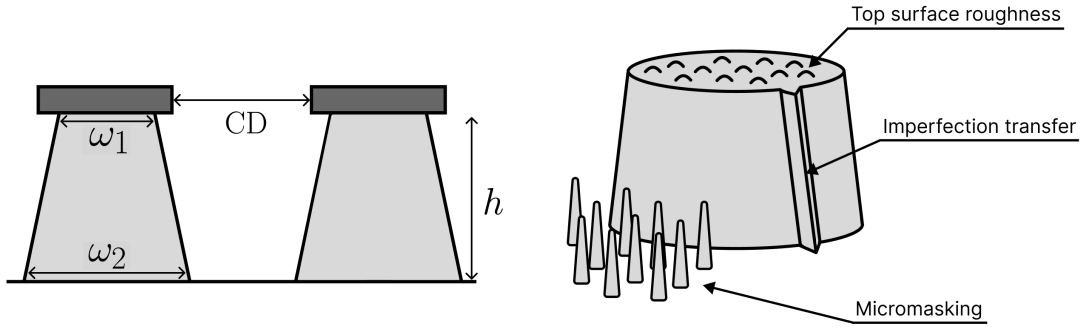

(a) **(b)**

Figure 3.3: (a) In a typical ICP-RIE etch, perfect anisotropy is unattainable due to imperfect directionality and unbalanced passivation, resulting in slanted sidewalls and undercutting. (b) The three main types of roughness occurring in this project are unwanted structures inherited from the hard mask etch, top surface features and micromasking.

Figure 3.3a shows a schematic image of the typical anisotropy achieved in ICP-RIE. The anisotropy, $\mathcal{A}$, is quantified as

$$\mathcal{A} = 1 - \left| \frac{\omega_1 + \omega_2}{2} - \mathrm{CD} \right| h^{-1} \tag{3.1}$$

where ω_1 and ω_2 are the widths of the feature at its narrowest and widest point, CD is the critical dimension (in this context referring to the smallest feature size of the

mask) and h is the height of the feature. When we optimize for anisotropy, we want to maximize equation (3.1). The maximum is 1, when $\frac{\omega_1+\omega_2}{2}$, an estimation of the average structure feature size, is equal to CD.

For photonic applications, alongside anisotropy, surface roughness is an important metric for optimization. A rough surface, even on the order of nanometers, scatters light, which can limit the reflectivity of a mirror [21]. The two most common sources of roughness in our case are *roughness transfer* and *micromasking*. Roughness transfer refers to the edge roughness present in the mask, which will be reproduced on the etched sidewall and cause unwanted lateral etch components.

In this fabrication process, a sputtered SiO_2 hard mask is used. This mask is itself patterned and etched before it is used, which means that any roughness, residues, or impurities might transfer directly onto the amorphous silicon etch beneath it, reproducing the mask edge roughness in the etched sidewall profile. Figure 3.4 shows an SEM image of the hard mask (with the e-beam resist on top) after patterning. Since no impurities or defects are visible in the image, the quality of the hard mask is assumed to be sufficiently high not to be a significant source of roughness transfer in the subsequent silicon etch.

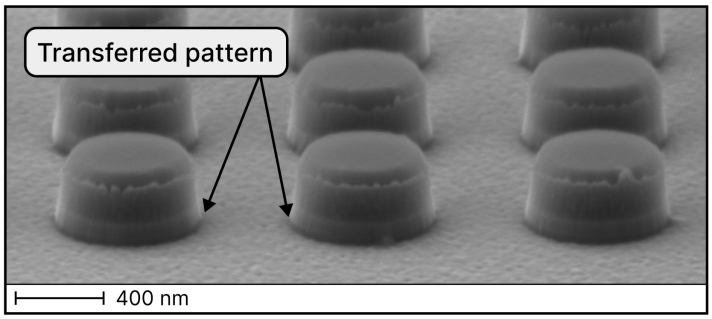

Figure 3.4: SEM image of the hard mask.

Micromasking is a mechanism in which involatile species locally inhibit the etching, resulting in narrow unetched pillars. In the particular system used in this project, micromasking will mostly appear as a result of redeposition of silicon dioxide or excessive passivation. Figure 3.3b depicts how roughness transfer shows up as ridges on the sidewalls and micromasking shows up as a “grass”-like structure. These three types of defects (isotropy, micromasking and roughness transfer) can be evaluated with SEM imaging and atomic force microscopy.

3.1.6 Laser end-point detection

Since the etching rates fluctuate, it is useful to monitor the etching progress in-situ. To that end, a laser-endpoint detection system is used. The chart in figure 3.5 shows a laser endpoint interferometry trace from a silicon etch performed during the project.

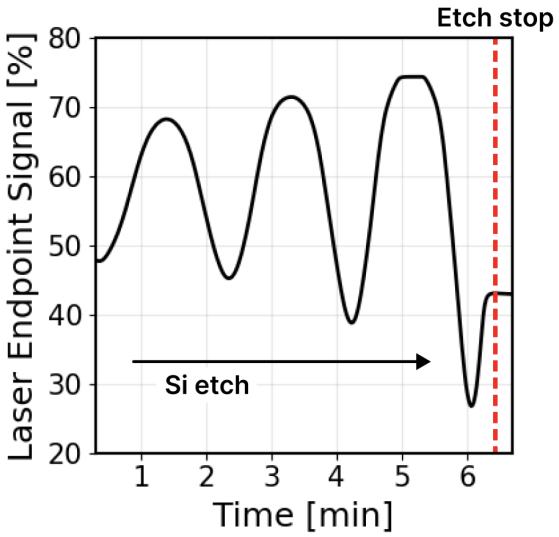

Figure 3.5: Laser endpoint signal during silicon etching. Interference fringes disappear at 6.2 - 6.6 min, indicating the etch endpoint.

As the plasma etches the silicon, the reflected laser beam interferes constructively and destructively with light reflected from the silicon substrate underneath, producing the sinusoidal oscillations seen from 0 to 6 min. Each full oscillation corresponds to etching through a fixed increment of amorphous silicon thickness determined by the laser wavelength and the refractive index of the amorphous silicon. We can see that the etch-stop layer has been reached when the endpoint signal saturates at around 6.4 min (see figure 3.5). At this point, the etch is manually interrupted.

3.2 Scanning Electron Microscopy

Scanning electron microscopy (SEM) is a surface characterization technique in which a tightly focused electron beam is rastered across a sample and collected by one or multiple detectors to form a two-dimensional image. In contrast to optical microscopy, where the resolution is limited by the wavelength of visible light (400-700 nm), SEM uses electrons whose de Broglie wavelength at regular *acceleration voltages* is on the

order of picometers. Diffraction is however not anymore what limits the resolution. Rather, it is the illumination system and beam-shaping mechanisms that determine the beam diameter at the sample [22]. For the Zeiss Sigma 360 used in this work, a resolution of 1 - 2 nm at 1 kV acceleration voltage is achieved [23]. This makes SEM imaging suitable for characterizing micro- and even nanostructured surfaces such as photonic crystals, where the feature dimensions are below the optical diffraction limit.

3.2.1 Operating principle

The primary beam is generated by an electron gun and accelerated through a potential difference. The acceleration voltage defines the kinetic energy of the incident beam. The beam is then guided through multiple sets of electromagnetic condenser lenses to achieve collimation and finally through an objective lens to focus it on the sample plane. Figure 3.6 shows a schematic representation of the SEM column from the source to the sample.

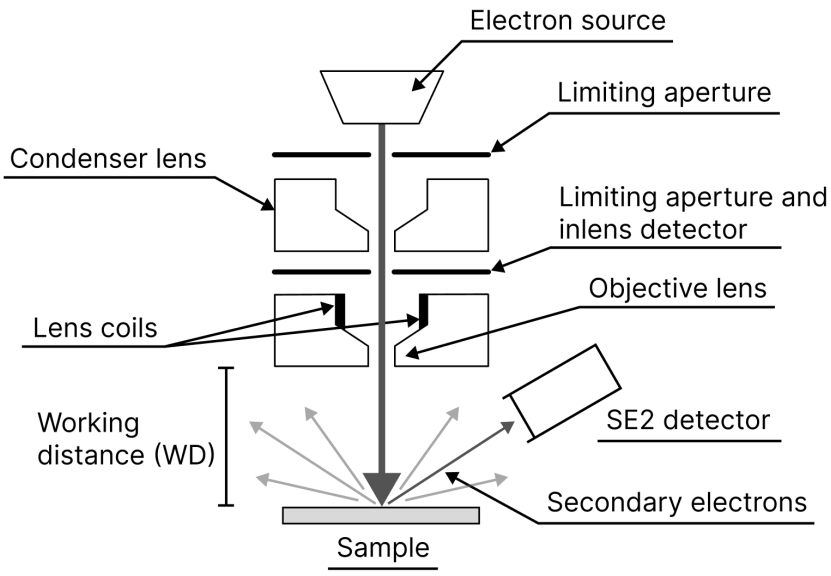

Figure 3.6: The Zeiss Sigma 360 SEM has been used to image samples. In this figure, only the inlens and the secondary electron detector (SE2) are illustrated as they are the only functions used during this project.

The collected signal from the sample consists of three main parts: the secondary electrons (SE), the backscattered electrons (BSE) and characteristic X-rays. The secondary electrons own their name to the fact that they are not part of the incident

beam, but are ejected via inelastic collisions in the valence and conduction bands of the materials in the sample. These electrons carry low energy (normally not more than 50 eV) and therefore only escape from a thin surface layer. Because of this property, the SE signal is a suitable choice for high-resolution surface and topographic imaging. Another advantage of the secondary electrons is their sensitivity to surface tilt. This effect can be explained geometrically by the fact that a tilted incidence will lead to a larger amount of electrons occupying a volume close to the sample surface. In this way, areas with a normal orientation with respect to the beam will appear darker and vice versa.

The backscattered electrons are electrons that originate from the illumination beam, that have undergone scattering to an extent where they again leave the sample. These electrons carry an energy close to the beam energy and penetrate the sample considerably deeper than the secondary electrons and carry information about the material composition. Lastly, characteristic X-rays are emitted when incident electrons create inner-shell vacancies in the sample atoms. The X-rays can originate from every depth of the interaction and therefore carry little topographic information. Therefore, in this project, the samples are imaged with the SE detector. The working distance defined in figure 3.6 is set between 7 - 8 mm to allow for the secondary electrons to reach the detector.

3.2.2 Influence of acceleration voltage

When the electrons enter the sample, they scatter repeatedly with the sample. The spatial distribution of this process, called the *interaction volume*, is responsible for determining both the information depth and the lateral resolution of the detector signal. For most relevant materials, the interaction volume takes a waterdrop-like shape (see figure 3.7a), widening with depth as the electrons spread laterally before eventually losing all of their energy to phonons or plasmons.

For materials such as silicon, with a relatively low atomic number, the properties of this interaction volume (and thereby the acceleration voltage) play a great role in determining the information in the resulting SEM image. The penetration depth is

$$R = \frac{0.028A}{Z^{8/9}\rho} E_0^{5/3},$$

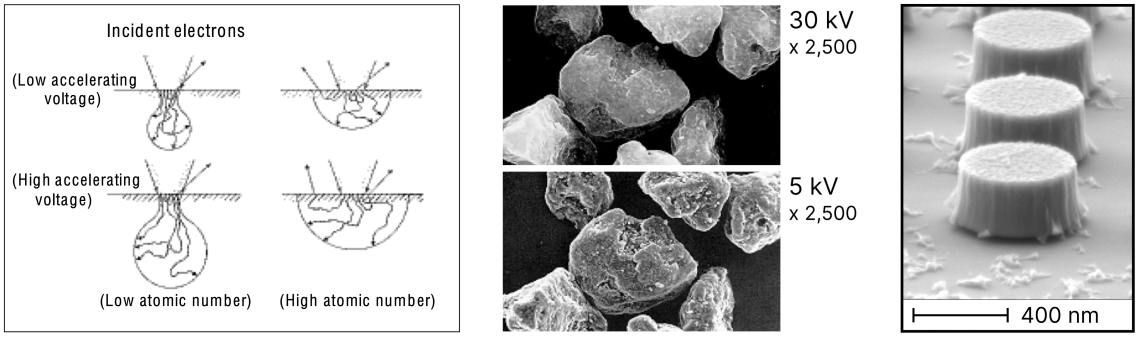

(a) **(b)** **(c)**

Figure 3.7: (a) The interaction volume of the incident electron beam takes a drop-like shape in low atomic number materials such as silicon. (b) High acceleration voltages capture information about material composition rather than surface topography. Low acceleration voltages do the opposite. The top image is an example of the former and the bottom one of the latter. (c) Image taken in the optimization process, with an acceleration voltage of 2 kV. The images in (a) and (b) are adapted from Ref. [24].

where R is the maximum penetration depth, A is the atomic weight, Z is the atomic number, ρ is the material density and E_0 is the acceleration voltage [25]. If the acceleration voltage is low, the penetration depth is smaller. This means that the reflected electrons carry more topographic information than at higher acceleration voltages. However, at low acceleration voltages the energy spread of the incident beam is greater, yielding a comparatively lower lateral resolution. This conflict between information depth and lateral resolution is an important trade-off to account for when taking SEM images. Figure 3.7b shows an example of the difference between high and low acceleration voltage. The bottom image has a lower acceleration voltage and displays much more detailed surface features than the top image. Figure 3.7c shows an image taken during the etch optimization in this project. It is taken with an acceleration voltage of 2 kV, allowing for examination of topography features on the order of 10 nm which is enough to determine the quality of the etching process.

3.2.3 Surface roughness imaging

Figure 3.8 includes a schematic of the SEM imaging of the samples along with an example of such an image. To capture both the sidewall and top surface roughness, the samples were tilted at an angle θ and observed solely with the SE2 (secondary electrons) detector. By choosing a tilt angle such that the pillars were mainly viewed from the side, sufficient information about the anisotropy could be achieved from

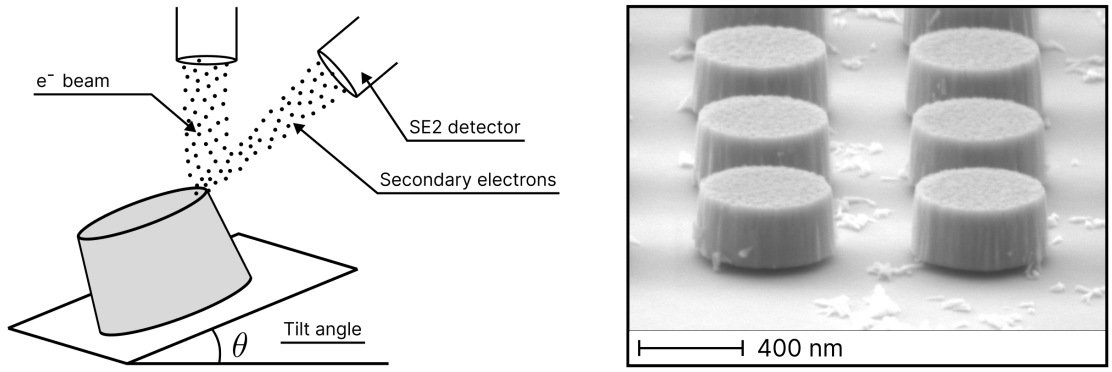

(a) **(b)**

Figure 3.8: (a) The secondary electron detector (SE2) was used to image the photonic crystal pillars. To capture both the sidewalls and the top surface, a tilt angle of 60° was introduced by using a tilted sample holder. (b) In this example micrograph, the top surface roughness as well as the sidewall imperfections are discernible.

these images along with a rough estimate of the roughness of pillar surfaces. Images like the example micrograph in figure 3.8b were taken of every sample to evaluate and compare their sidewall roughness and anisotropy with the goal of refining the etching recipe. The acceleration voltage is kept at 2 kV with a working distance of 7 - 8 mm to achieve enough topographic information as well and a sufficiently strong signal.

3.2.4 Surface roughness analysis

To achieve an accurate comparison of the sidewall roughness between the samples, a process in which the circularity of the pillars is quantified was set up. This process consists of three main steps: capturing an SEM image of the horizontal cross section of the pillars, find the RMS roughness by fitting an ellipse to the pillar sidewall and lastly average this roughness metric over multiple pillars on each sample. Figure 3.9a shows an example of a horizontal cross section of a pillar. These images were captured with the InLens detector on the SEM with a perpendicular orientation with respect to the pillars.

After the image was captured, points forming the border between the top of the pillar and its sidewall were detected by searching outward from the image center until a pixel saturation of 10 % was reached. 720 points of this sort were cast onto

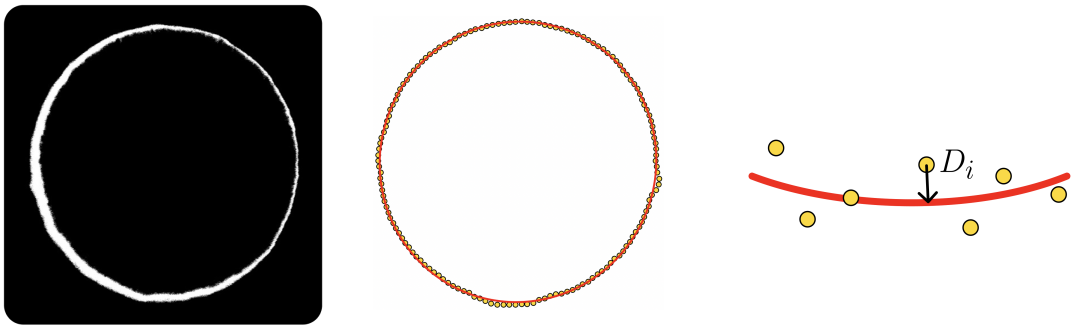

(a) **(b)** **(c)**

Figure 3.9: (a) The micrographs of the unwanted sidewall features were captured with the InLens detector at an acceleration voltage of 1.5 kV and a working distance of 2 - 3 mm. (b) From the SEM images, an ellipse was fit to detected points following the inner border between the top surface and the sidewall. From this fit an RMS roughness of the sidewall was determined. The pillar in this figure has a roughness of 8.3 pm. (c) In a close-up of (b) the detected points become distinguishable from each other. D_i is defined as the closest geometric distance from point i to the ellipse.

the image to ensure a resolution high enough to capture the lateral variations of the sidewall (see figure 3.9a). An ellipse was then fit to these points using the Fitzgibbon direct least-squares method [26], which minimizes the algebraic distance between the conic equation

$$Ax^2 + Bxy + Cy^2 + Dx + Ey + F = 0$$

and the detected points under the constraint $B^2 - 4AC < 0$, forcing an elliptic solution. This gives a rough elliptic fit, which can then be refined by minimizing the true geometric distance between the ellipse and the points using the least-squares method. Once the ellipse has an optimal shape, the RMS roughness Rq is calculated with the following formula:

$$\mathrm{Rq} = \left(\frac{1}{Nab} \sum_{i=1}^{N} D_i \right)^{1/2}, \tag{3.2}$$

where $N = 720$, a and b are the axis lengths of the ellipse and D_i is the shortest geometric distance between point i and the ellipse.

3.3 Atomic Force Microscopy

Atomic force microscopy (AFM) is a standard surface characterization technique in microfabrication that can be used to measure roughness of deposited materials as well as features of patterned designs. In this project, it has been used for its ability to gauge the roughness of smooth surfaces with high precision.

3.3.1 Working principle

The core component of the AFM is its probe, which consists of a sharp tip mounted at the end of a cantilever. The cantilever acts as a mechanical spring that deflects upon interaction between the tip and the sample. This deflection is then measured by reflecting a laser against the backside of the cantilever. The reflection is then directed to a photodiode on which the trace of the reflection is recorded. This trace can then be converted into a topographic signal that represents the trace followed by the probe. Figure 3.10 shows a basic schematic of the AFM process.

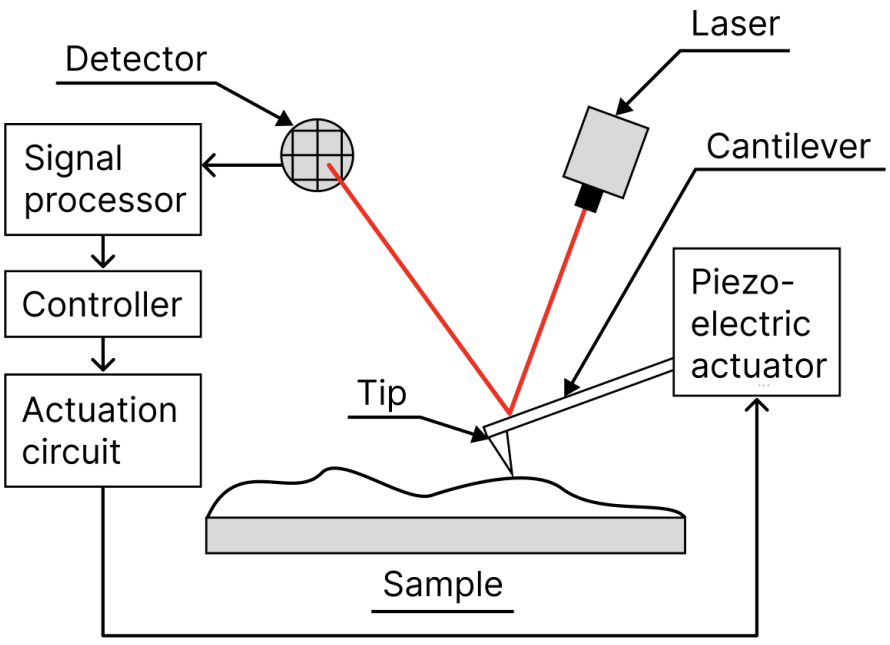

Figure 3.10: For imaging the top surface of different structures, the Dimension Icon Atomic Force Microscope was used. This schematic representation of its working principle shows how the surface variations are captured via the trace of a laser reflected off the cantilever holding the scanning tip.

In this work, the type of interaction that was used is called tapping mode. Tapping mode, also referred to as amplitude modulation AFM, is a standard approach for surface characterization in microfabrication. In this mode, the cantilever is driven to oscillate close to its fundamental resonance frequency. The *free oscillation amplitude* is set to a value that provides sufficient amplitude for the tip to periodically contact the surface on each oscillation cycle while still being able to fully retract between contacts. The advantage of tapping mode compared to contact mode is that lateral forces on the sample are almost entirely eliminated. This reduces the risk of sample destruction and slightly shifts the limiting resolution factor of the AFM from the tip geometry to the properties of the cantilever oscillation.

3.3.2 Roughness gauging and practical limitations

The vertical resolution is primarily determined by the noise floor of the deflection detection system. With the optical lever arrangement described above, the thermal noise of the cantilever itself sets a fundamental lower limit on the detectable deflection. In practice, a modern AFM can achieve a vertical noise floor on the order of tens of an Ångström, meaning that surface height variation below 1 nm can be resolved [27]. This highly precise vertical resolution is what makes an AFM advantageous compared to, for example, an SEM, when assessing surface roughness.

The lateral resolution of an AFM is limited by the geometry of the tip itself and in particular the radius of the tip apex. The image recorded by an AFM is not a direct map of the true surface, but rather a convolution between the true surface and the shape of the tip. Whenever the tip encounters a surface feature, like a step, a trench or a piece of dirt, the size of the tip prevents it from following the exact contour. This makes sharp edges appear rounded, narrow trenches appear wider and steep walls may be incompletely traced. The practical consequence is that smaller tip radii yield better lateral resolution. Modern AFM tips have radii of a few nanometers and can resolve lateral features on the order of $\sim$10 nm.

To quantify the surface roughness, a common quantity is Rq, which is also known as RMS roughness. In this case, however, D_i in equation (3.2) is defined as the difference between point i in the AFM image and the average surface height and $a = b = 1$ since the mean height is a flat surface. The Rq parameter now represents the standard deviation of the height values within the scanned region from the mean

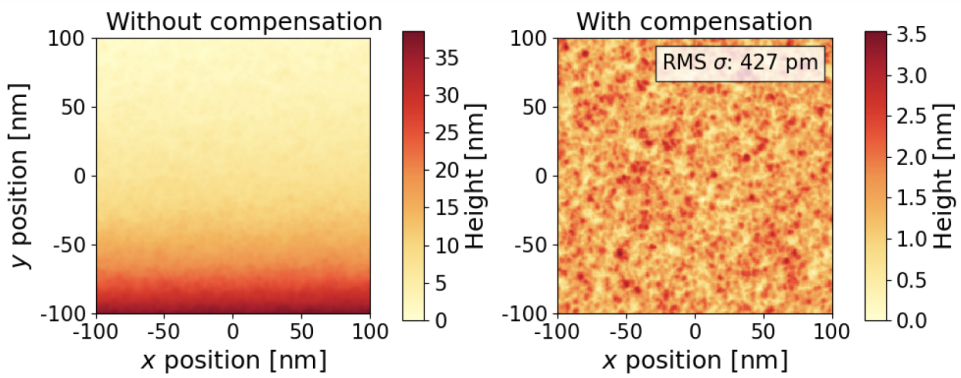

Figure 3.11: The left panel displays the raw data from the AFM, showing an unreasonably high roughness. The right panel shows the corrected data with a lower roughness.

height. The lower the Rq that is obtained, the smoother the surface is.

Before the Rq is extracted from the AFM image, the image first has to be corrected with respect to the tilt and bowing of the substrate surface that arise from a possibly tilted stage and a deflected sample. This is done by fitting a two-dimensional second degree polynomial to the data points in the AFM image and subtracting the value of the polynomial in every data point. Figure 3.11 shows an example of a raw AFM image (left panel) where tilt and/or bowing is clearly visible. The right panel in figure 3.11 shows the corrected image, where we evaluate the roughness to 427 pm.

3.4 Sample preparation

3.4.1 Fabrication

Figure 3.12 gives a step-by-step overview of the microfabrication process of the photonic crystal on a membrane. We start off with an intrinsic, high-resistivity, 4 inch Si⟨100⟩ wafer with a thickness of 525 µm. We clean the wafer in SC1 (10 min), HF (2 %, 20 s), and SC2 (10 min). Then we deposit between 20 and 35 nm of stoichiometric, high-stress Si_3N_4 with low pressure chemical vapour deposition (LPCVD) at 780 °C within 6 to 12 minutes. Then, a 6 nm thick layer of HfO_2 is deposited with atomic layer deposition (ALD) at 250 °C. This layer acts as an etch stop later during the reactive ion etching. After the HfO_2, the amorphous silicon is deposited. Two different deposition methods have been tried in this study, LPCVD and sputtered silicon.

The silicon is sputtered from a target for 485 s (1kW RF, 40 sccm Ar) and the LPCVD is made at 550 °C for 2 h 37 min. Both of these deposition methods result in a silicon layer with a thickness of about 300 nm. We also sputter a 50 nm thick layer of SiO_2 as a hard mask and end up with the material stack in figure 3.12.

The first step is to write gold markers on top of the material stack to assist alignment of the different lithography layers. After dehydrating, we coat the wafer with LOR3A (2.6 µm) and a layer of S1813 (1.5 µm). We write the markers with a maskless optical lithography system and develop with MF-CD26. This resist combination creates an undercut that allows smooth liftoff. We then evaporate a layer of Ti(3 nm)/Au(50 nm) and lift off everything except the markers in REM400.

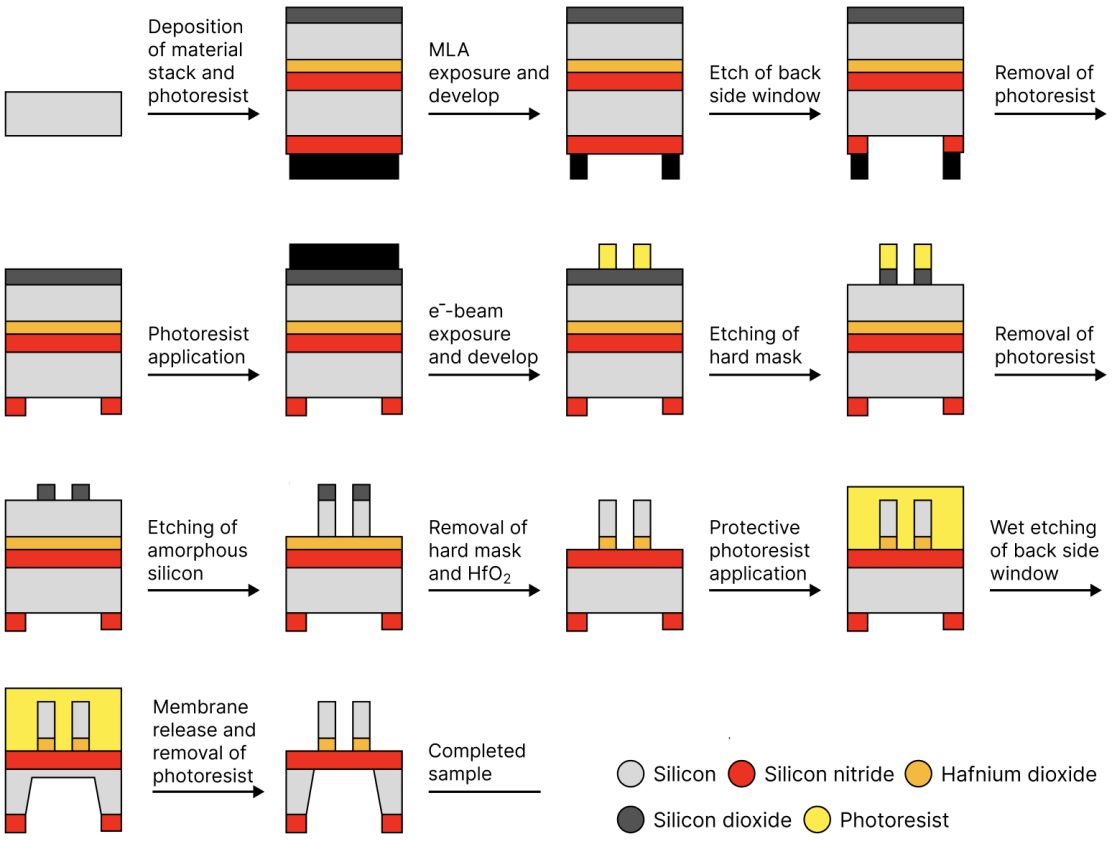

Figure 3.12: The fabrication of the samples was split in an e^--beam lithography part (for the photonic crystals) and a maskless photolithography part (for releasing the samples).

Then we coat the backside with a thin layer of HMDS for adhesion, and then 1.5 µm of S1813. We then write the backside windows with the same maskless optical lithography system and develop with MF-CD26. Then we etch the pattern into the

Si_3N_4 layer with CHF_3 and O_2. We then clean the sample to make it ready for the e^--beam lithography.

After applying Ti-Prime as adhesion promoter, the SiO_2 gets covered with a 300 nm thick layer of the negative photoresist ma-N which is then exposed with an e^--beam lithography machine (Raith EBPG 5200) and later developed in maD 525 to produce the photonic crystal pattern. Once this is done, the hard mask is etched with an Oxford PlasmaPro 100 ICP/RIE system which when finished allows the removal of the ma-N. Now is the time to etch the a-Si. This is done with an Oxford PlasmaPro 100 Cl_2 ICP/RIE system. The details of this etching step are an important result of this thesis and we discuss it in detail later. By removing the hard mask and the HfO_2 with hydrofluoric acid (HF), the photonic crystal side of the structure is completed.

The stack is then flipped upside down to make a window on the backside of the silicon wafer and thereby release the membrane. For this we first cover the photonic crystal side with photoresist to protect it. Then we mount the 4-inch wafer in a special PEEK holder that protects the front side during the KOH etch. We then etch through the majority of the Si substrate for about 8 h in 80 °C hot (40 %) KOH. We stop with roughly 25 µm Si left. Then we dice the chips and release the membranes on chip-by-chip basis. For the release, we protect the Si pillars from the KOH by coating them with ProTEK primer and ProTEK B3-25. We then release the chips in 55 °C hot KOH (40 %) and after neutralizing it in diluted HCl, we remove the ProTEK film. The resulting chip is now covered with Si_3N_4 on which the photonic crystal sits in the center of the backside window. Figure 3.13b shows an SEM image of a membrane with a photonic crystal produced in this manner.

3.4.2 Devices

The final samples made in this project are illustrated in figure 3.13. The released membrane is 300 × 300 µm and the photonic crystal measures 150 µm in diameter (see figure 3.13a).

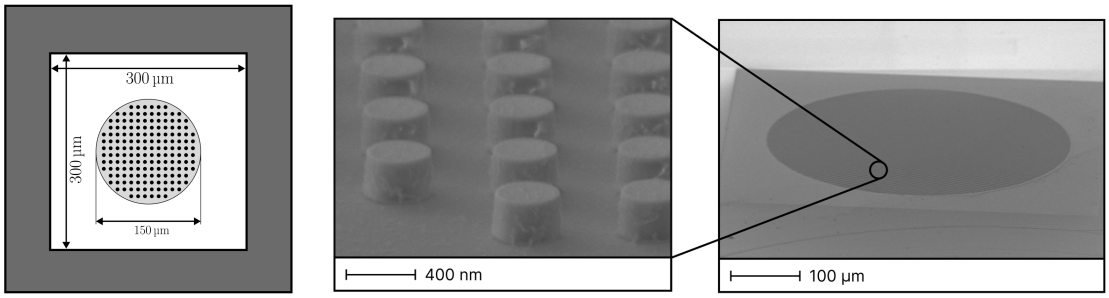

(a) **(b)**

Figure 3.13: (a) The photonic crystal patterns used to perform measurements are based on a square-shaped Si_3N_4 membrane (white) with a circular region of a square lattice pillar structure (light gray with black dots) suspended on a silicon substrate (dark gray). (b) The test device in 3D with a close-up of the pillar pattern.

In figure 3.13b, the design is represented in 3D with a close-up showing the placement of the etched pillars. The design parameters for the photonic crystal and the optimization of those are left to section 4.3.

4. Simulations of photonic crystals

The experimental results in this thesis are complemented with numerical simulations carried out using rigorous coupled-wave analysis (RCWA). RCWA is a semi-analytical frequency-domain technique for solving Maxwell's equations in structures that are periodic in one or more in-plane directions and formed into a stack of homogeneous layers along the propagation direction [28]. The photonic crystal slabs, which are homogeneous in the out-of-plane direction, and periodic in-plane, fall into this category. Therefore RCWA is a natural choice of simulation framework, and is the one used for all simulations shown in this thesis work.

The specific implementation of RCWA used in this thesis is GRCWA, which is a version based on the S-matrix algorithm [29]. It has served four main purposes throughout this project: simulation of the reflectivity under plane-wave illumination, simulation of response to a Gaussian beam through decomposition into plane-wave components, optimization of the photonic crystal pattern geometry and estimation of the losses in the Fabry-Pérot cavity.

4.1 Plane wave simulations

The most basic building block of every simulation in this chapter is the response of the photonic crystal to a single, monochromatic plane wave. Figure 4.1a illustrates the geometry where a plane wave with wave vector $\boldsymbol{k}$ is incident on the photonic crystal at an angle θ.

The method through which RCWA works is by expanding the permittivity of every layer in the structure together with the electromagnetic fields inside them as a Fourier series in the in-plane reciprocal lattice vectors. This expansion is then substituted into Maxwell's equations and the problem is thereby converted into an eigenvalue problem for a finite set of coupled plane-wave harmonics whose eigenmodes and

eigenvalues describe how each Fourier component propagates through the layer [30]. In practice, this Fourier expansion is reduced to a finite number of harmonics, and the accuracy of the simulation is determined by how many are included.

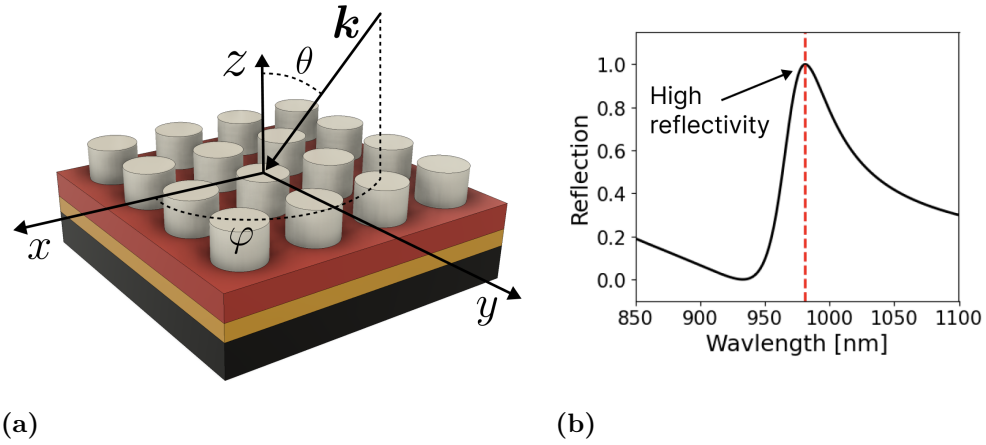

(a) **(b)**

Figure 4.1: (a) The typical use case for RCWA throughout this thesis is a plane wave impinging the surface of a slab of patterned or unpatterned material under an angle θ. (b) When the plane wave interacts with a photonic crystal, the reflectivity spectrum mimics a Fano resonance.

Once the eigenmodes of the whole system are known, the layers are stitched together using a scattering matrix (S-matrix) formalism, which relates the amplitudes of the outgoing waves to the incoming waves [31]. The S-matrix of the full structure is built recursively by combining the S-matrices of adjacent layers one at a time, enforcing continuity of the tangential field components at each interface. From the S-matrix, the complex reflection amplitudes r_m, where m is the mode number, are calculated. In this case of a photonic crystal, the periodic structure of the material slab has a lattice parameter smaller than the wavelength of the incoming wave. This means that every diffracted order except $m = 0$ has $|k_{\parallel,m}| > k$ and is therefore evanescent, which means light is only coupled back into the zeroth order. The reflectivity is thereby calculated as $R = |r_0|^2$. Figure 4.1b shows a simulation of the reflectivity response from a plane wave incident on a photonic crystal. The lineshape mimics the characteristic Fano resonance depicted in figure 2.3.

4.2 Gaussian beam reconstruction

The plane-wave simulation holds true for an infinitely extended plane wave, but the actual beam used in experiments is a focused Gaussian beam of finite waist size ω_0 that carries a finite spread of propagation angles rather than a single well-defined θ. To predict the reflectivity of a Gaussian beam, the plane-wave results must be combined according to the angular content of the beam. The angular representation of the Gaussian beam is

$$\boldsymbol{E}_R(\boldsymbol{r}_\parallel, z) = \iint \frac{d^2k_\parallel^2}{(2\pi)^2} \left(\boldsymbol{E}_s(k_\parallel)\boldsymbol{n}(k_\parallel) + \boldsymbol{E}_p(k_\parallel)\boldsymbol{t}(k_\parallel)\right) e^{-i\boldsymbol{k}\cdot\boldsymbol{r}},$$

where $\boldsymbol{E}_R$ is the electric field, $\boldsymbol{r}_\parallel$ and z are the in-plane and out-of-plane position coordinates, $k_\parallel$ is the magnitude of the in-plane wavevector, $\boldsymbol{E}_s$ and $\boldsymbol{E}_p$ are the s- and p-polarized electric fields and $\boldsymbol{n}$ and $\boldsymbol{t}$ are the unit vectors along the s and p directions [32]. The connection to the angle of incidence θ is through $k_\parallel = k_0 \sin\theta$. Figure 4.2a shows the reflectivity dependence on the angle of incidence. From this graph it becomes evident that at oblique incidence, reflectivity is lost. In this case, a range four times as large as the divergence angle ($4 \cdot \theta_0 = 4 \cdot \frac{\lambda}{\pi\omega_0} = 0.17$ rad, where $\lambda = 1315$ nm and $\omega_0 = 10\,\mu$m) is plotted.

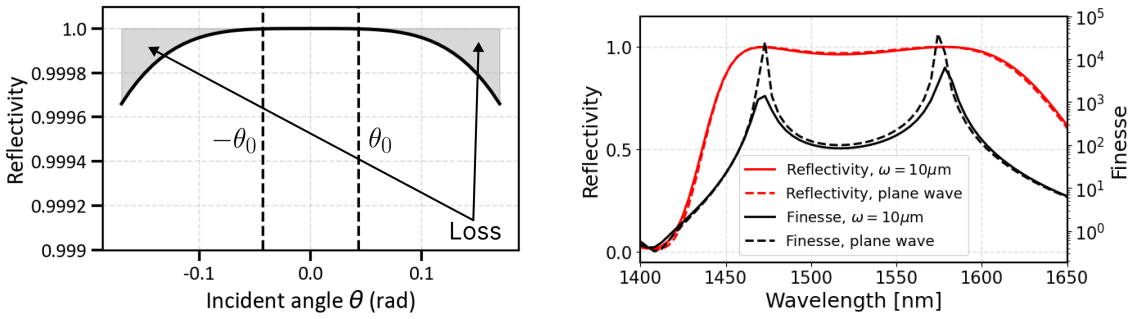

(a) **(b)**

Figure 4.2: (a) Tilting the incidence angle of a plane wave incident on a photonic crystal. The reflectivity drops for larger incidence angles. (b) The angular loss from the Gaussian beam is limiting the finesse of the cavity.

Each of the plane-wave components reflects independently and contributes with complex reflection coefficients $r_s(\theta, \lambda)$ and $r_p(\theta, \lambda)$, where λ is the wavelength of the incident light, that are computed as described above. After some calculations (that are described in ref. [33]), the reflectivity R at a specific wavelength is computed

with the formula

$$R(\alpha) = \frac{\iint d\theta d\varphi \sin\theta \cos\theta \left|E\right|^2 \left((r_s \cos\left(\alpha+\varphi\right))^2 + (r_p \sin\left(\alpha+\varphi\right))^2\right)}{\iint d\theta d\varphi \sin\theta \cos\theta \left|E\right|^2}, \text{ where}$$
$$E(\theta) = \sqrt{2\pi\omega_0} e^{-k^2 \sin^2\theta \frac{\omega_0^2}{4}},$$

and α is the polarization in the xy-plane of the incoming field according to the description $\boldsymbol{E} = E\sin\left(\alpha\right)\boldsymbol{x} + E\cos\left(\alpha\right)\boldsymbol{y}$ [33]. Figure 4.2b shows the response of a Gaussian beam incident on a photonic crystal. Similar to what is described in section 2.1.2, a reflectivity plateau appears between 1500 nm and 1600 nm. Assuming a cavity with two identical mirrors, the corresponding finesse calculated with equation (2.1) is plotted in figure 4.2b and reaches a maximum of $\sim 10^4$ at $\lambda = 1540$ nm for a Gaussian beam with a waist of 10 µm as compared to the plane wave reaching a finesse on the order of $\sim 10^5$.

4.3 Optimization

The design of the photonic crystal is defined by three parameters illustrated in figure 4.3a, the lattice parameter, ℓ, the radius of the pillars, r, and the thickness of the pillars, h. The lattice constant has, in this study, been set to a fixed value of 870 nm to keep the structure in the sub-wavelength regime while still in a region where its resonance frequency is in the telecom wavelength range (∼1550 nm). This leaves r and h as the two free parameters over which the reflectivity is optimized.

To perform this optimization, these two parameters have been individually swept over a wide range of values. In figure 4.3b, an example of such a sweep is depicted with h (as a ratio where it is compared to the lattice parameter) on the vertical axis and the wavelength on the horizontal axis. The r/ℓ ratio is kept constant at $r/\ell = 0.25$. In this sweep, the reflectivity of a plane wave is computed for every combination of thickness and wavelength. From figure 4.3b, we can see that at small thicknesses, the map is dominated by a narrow resonance feature. As the thickness increases, this feature broadens and eventually merges into a large, diamond-shaped region of high reflectivity for $\lambda = [1520, 2500]$ nm and $h = [0.25, 0.8] \cdot \ell$, which arises from the pair of overlapping Fano resonances discussed in section 2.1.2. In this region, we are interested in finding a thickness for which the plateau of high reflectivity is both

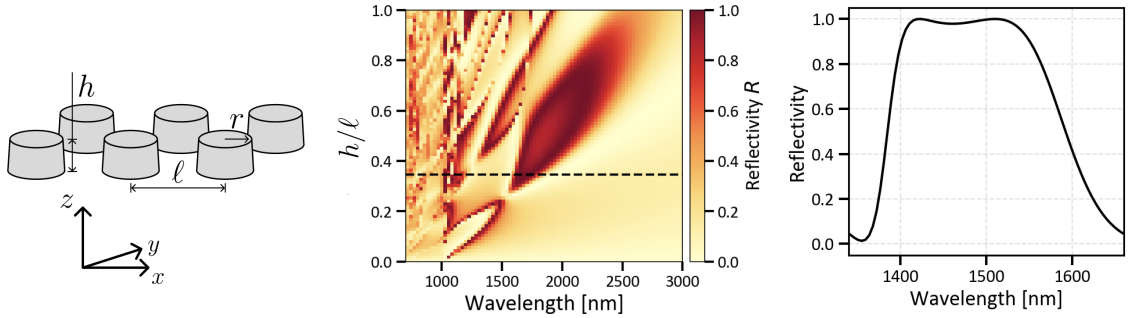

(a) (b) (c)

Figure 4.3: (a) The design parameters for a pillar-based photonic crystal are the thickness, h, the radius, r, and the lattice parameter, ℓ. (b) Contour plot of how the reflectivity-wavelength dependence is affected by the thickness of the pillars. (c) We plot a cut-through of the contour plot and present the reflectivity at a thickness of 0.345ℓ. A plateau forms with two maxima in reflectivity.

broad and consistently near 1. Figure 4.3c shows the reflectivity as a function of wavelength at $h = 0.345\ell$, where the reflectivity is far above 0.9 for a region between $\lambda = 1400$ nm and $\lambda = 1540$ nm. A similar optimization process was carried out for the radius of the pillars, and the final value was set to $r = 0.30\ell$.

4.4 Losses

The RCWA calculations described so far output the reflectivity of an idealized photonic crystal mirror. A fabricated mirror has additional loss mechanisms that are not captured by this ideal calculation, and this section describes how each of them can be estimated individually. All of these loss types share the same basic effect: on each round-trip, they remove a fraction of the intensity inside the cavity, which shows up as a loss in reflectivity. The reflectivity can be expressed as $R = e^{-\mathcal{L}}$, where $\mathcal{L}$ is the loss. Since $\mathcal{L}$ is assumed to be $\ll 1$, this expression can be approximated by $R \approx 1 - \mathcal{L}$. The loss term $\mathcal{L}$ can now be computed for every loss type to make an estimation of the reflectivity limit, which by equation (2.1) estimates the limit in finesse caused by the specific loss type.

Clipping loss

The photonic crystal pattern extends over a finite diameter D, while the mode incident on the crystal has a waist ω_m, defined by the cavity length and fiber facet geometry (see figure 4.4a). The result of this is that any power in the Gaussian

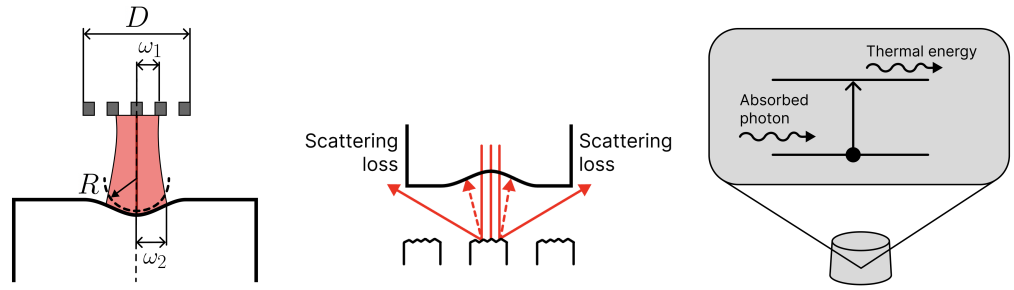

Figure 4.4: (a) Clipping loss comes from the Gaussian mode having a larger spread than the size of the photonic crystal. (b) Due to the sub-wavelength surface roughness of the structure, a fraction of the intensity is lost as incoherent scattering. (c) A fraction of the incoming photons gets absorbed through excitation of electrons depending on the material's extinction coefficient. The energy of the excited electrons then partially dissipates as heat and adds to the loss.

mode that falls outside the patterned region is lost due to an effect referred to as clipping loss. Clipping loss happens on both sides of the cavity and for a Gaussian beam of waist ω_m truncated by a circular mirror of diameter D, the clipping loss term $\mathcal{L}_{\text{clipping}}$ is computed by

$$\mathcal{L}_{\text{clipping}} = e^{-2(D/2)^2/\omega_m^2}, \text{ where}$$
$$\omega_{m,1} = \sqrt{\frac{L\lambda}{\pi}\sqrt{\frac{R}{L}-1}},$$
$$\omega_{m,2} = \sqrt{\frac{2L\lambda}{\pi}\sqrt{\frac{R}{L\left(1-\frac{L}{R}\right)}}}$$

and L is the cavity length, R is the radius of curvature defined as in figure 4.4a, λ is the wavelength of the light and $\omega_{m,1}$ and $\omega_{m,2}$ are the beam waists at the photonic crystal and at the fiber mirror, respectively [34]. An analogous clipping loss occurs on the fiber side of the cavity, since the parabolic fiber mirror (see figure 4.4a) only extends over a finite area.

Scattering loss

Residual roughness of the photonic crystal surfaces will scatter a fraction of the incident light into random directions rather than back into the cavity mode (see

figure 4.4b). For a surface with an RMS roughness σ, smaller than the wavelength of the incoming light, the scattering loss term, $\mathcal{L}_{\text{scattering}}$, is given by

$$\mathcal{L}_{\text{scattering}} = 1 - e^{-\left(\frac{4\pi\sigma}{\lambda}\right)^2}, \tag{4.1}$$

where again, λ is the wavelength of the light [35].

Absorption loss

The remaining loss channel, optical absorption, can be estimated with RCWA simulations. The way this is done is by allowing the refractive index of the patterned layer to be complex, $n = n_r + in'$, where the *extinction coefficient* n' quantifies the absorption (see figure 4.4c). The microscopic origin of absorption can be manifold, and includes light interaction with gap states, defects etc. With the addition of the extinction coefficient, RCWA simulations are performed to reproduce absorption dips recorded in reflectivity/transmission measurements of the photonic crystal.

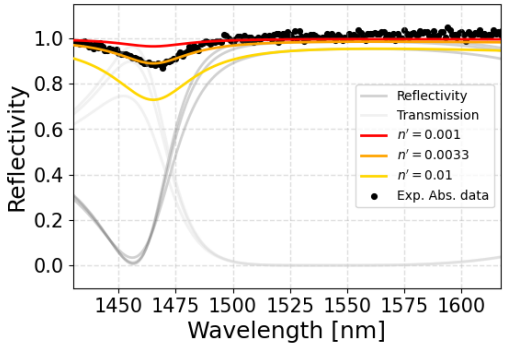

Figure 4.5: The absorption loss is estimated by fitting the extinction coefficient to absorption features in reflectivity measurements (see section 5.1).

Figure 4.5 shows an example of this, where three different values of n' have been simulated to find an accurate fit to the reflectivity trace. Once this has been done, the limiting finesse is computed from the maximum reflectivity of the fit to the specific sample for which the absorption was estimated. Contrary to the clipping and scattering losses, the absorption loss needs an actual sample to be quantified, since the extinction coefficient is not intrinsic to the material, but rather depends on the deposition of the material and the post-processing of it. However, since absorption loss is determined from a material property, it is assumed that samples fabricated from the same wafer will have equal absorption. This makes estimation of absorption loss necessary only upon changes in the material preparation process.

5. Experimental methods

5.1 Reflectivity measurement

To validate the optical performance of the fabricated photonic crystals, reflectivity measurements were performed across the wavelength range of interest. These measurements served two main purposes: first, to confirm that the measured reflectivity spectrum agreed with the simulated results, verifying that the fabrication process produced the intended optical structure; and second, to ensure that the reflectivity remained sufficiently high within a region of interesting wavelengths for the measurements with the fiber cavity setup presented in section 2.2.1.

5.1.1 Setup

The reflectivity measurements were carried out using the setup illustrated in figure 5.1. The laser light is created with a certain wavelength (between 1420 nm and 1620 nm), and linear polarization at the laser. We then attenuate it with a variable optical attenuator (VOA) to bring the power to an acceptable level for our photodiodes and transimpedance amplifiers. Afterwards, it is guided with an optical fiber to a collimator (COL), after which the light propagates in free space as a beam with a waist of 1.2 mm. This light then goes through a polarizing beam splitter (PBS1) to prevent polarization drifts. Before each measurement sweep, we optimize the polarization with the paddle wheel polarization controller (PWPC) to send the maximal power into the setup. The light is then split with a non-polarizing beam splitter (50:50) into two parts (BS1 in figure 5.1), one traveling straight, which goes into an isolator, and one that is reflected towards the sample. The part going into the isolator (constructed from a polarizer (POL) and a retarder (RET45)) is measured with a photodiode (PDreference) and acts as a reference for the incoming beam intensity. Like this we can normalize the power fluctuations in the laser in the end. On the sample side of the setup, the beam comes in and reflects against a mirror

mounted on a flippable mount. This separates the visible light imaging system from the infrared laser.

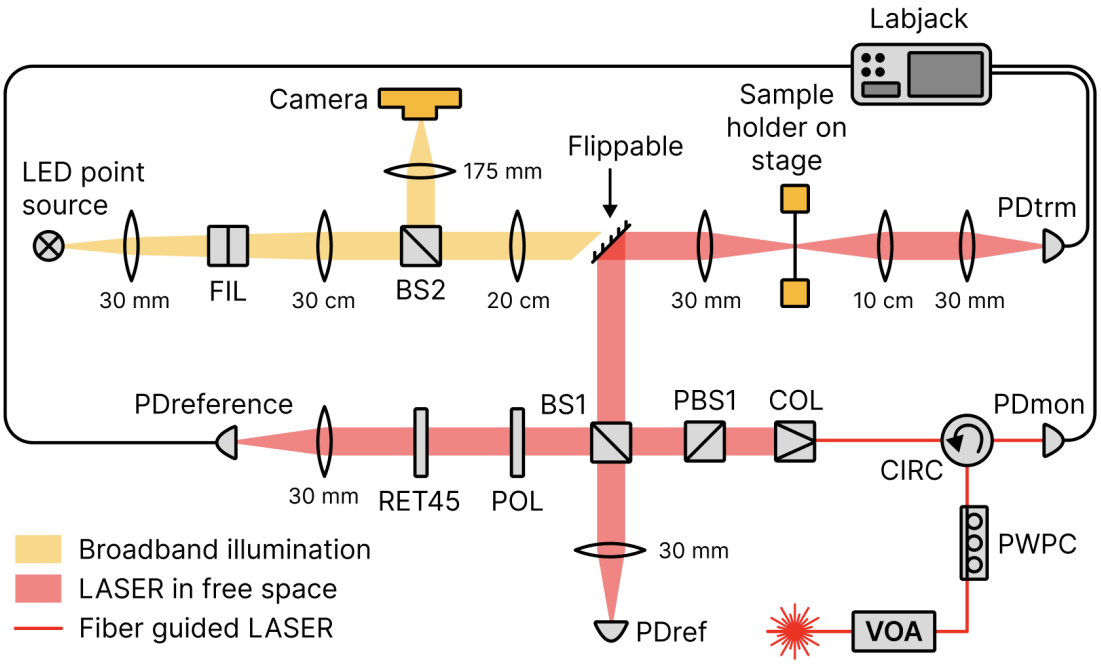

Figure 5.1: Illustration of the reflectivity measurement setup.

In the sample path, a lens focuses the beam so that the focal point aligns with the sample. The sample is mounted on an x-y-z-stage, which allows us to bring the sample into focus. We can ensure that the sample is in focus by maximizing the power that we measure in PDmon. Furthermore, the sample is mounted on a tilt-holder, that allows us to optimize for the angle of the incident light. When the incidence is straight and in focus the power at PDmon is maximal. In a reflectivity measurement, the light is reflected back onto the focusing lens (which then collimates the light), reflects against the mirror and ends up on the photodiode on the bottom of the figure. Conversely, in a transmission measurement, the light travels through the sample holder and is collimated by a lens similar to the focusing one. It is then focused by another lens onto a photodiode. These two photodiodes (PDref and PDtrm) are responsible for measuring the reflected and the transmitted intensity, respectively. The current from the photodiodes is amplified with a custom built transimpedance amplifier box, and the resulting voltages are sent to a Labjack T7 for digitization.

The part of the setup that is drawn with a yellow beam is used for positioning

the sample holder so that we can see the sample. It starts with a point source of broad-band visible light that goes through a set of collimating lenses and a filter (FIL) that removes undesired long-wavelength radiation, which will otherwise interfere with the laser light. This light is then split into two equal parts (BS2 in figure 5.1) whereafter the light is focused onto the sample and once returned, is reflected by the same beam splitter, into the camera.

5.1.2 Data processing

The raw output from the photodiodes in figure 5.1 is voltage signals that, after amplification, can be recorded by the Labjack. Before they can be converted to physically meaningful reflectivity values between 0 and 1, they must be calibrated against known reference measurements. This calibration was performed using two reference measurements presented in the upper row of figure 5.2.

First, a mirror with reflectivity close to unity (Thorlabs BB1-E04) was placed in the sample holder, and both the reflected and transmitted signals were recorded. As expected from this measurement, the upper left panel of figure 5.2 shows a transmission signal close to 0 for the whole wavelength range, while the reflectivity signal represents the voltage corresponding to a reflectivity of nearly 1. Second, the sample holder was measured empty, which means the transmission photodiode should measure the full beam intensity. These two measurements now provide references for the maximum and minimum reflectivity that can be measured with respect to the high-reflectivity DBR.

Using these two reference measurements, the raw voltage signals obtained from the actual sample were converted into reflectivity and transmission values. An example of this is presented in figure 5.2. The left panels show the voltage signals recorded when a photonic crystal occupies the sample holder and when the sample holder is empty. The right panels show a raw measurement of the voltage signal (top panel) and the converted data with reflectivity and transmission values (bottom panel).

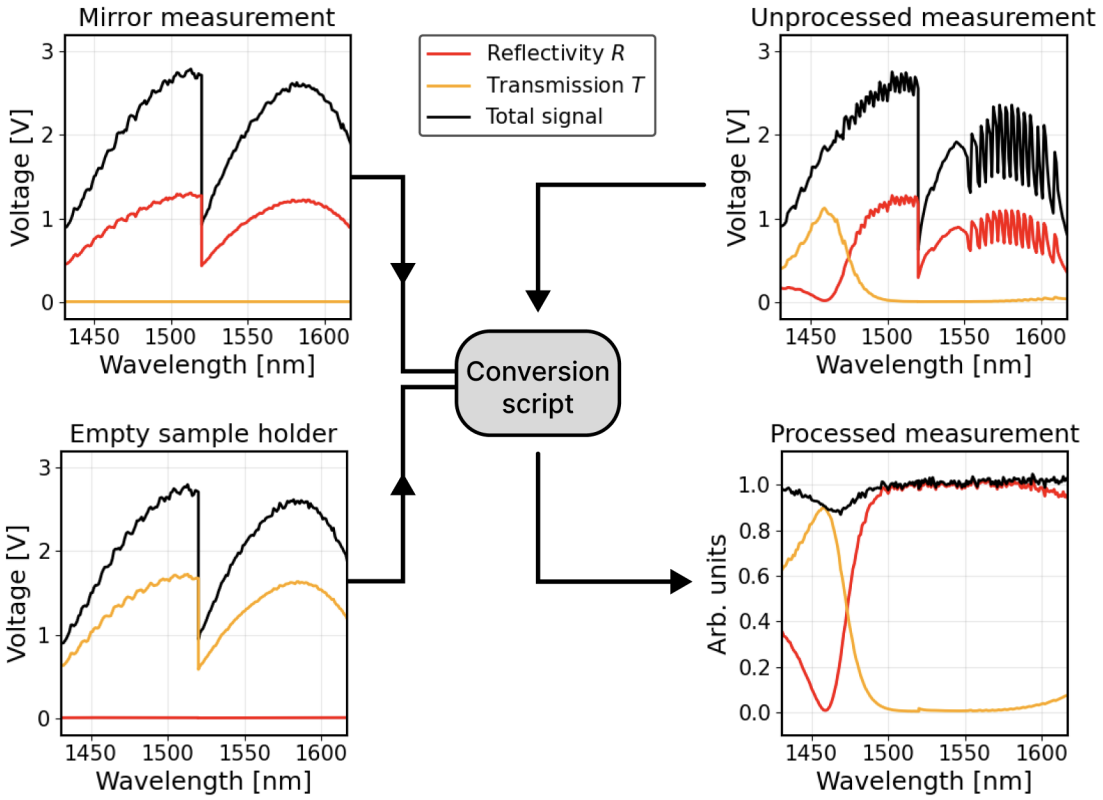

Figure 5.2: To obtain accurate reflectivity measurements, the setup records the trace of a highly reflective DBR as well as an empty sample holder. These two measurements are then used to convert the raw voltage signals from the photodiodes into physically meaningful reflectivity and transmission values. The black line corresponds to the reference signal, the red to the reflectivity and the orange to the transmission in all four panels. The abrupt dip in the panels containing unprocessed data is due to a laser switch, since two lasers are used to cover the whole wavelength region.

5.2 Fiber cavity measurement

Following the reflectivity characterization, which could confirm the qualitative performance of a sample, photonic crystals showing the intended characteristics were placed in the cavity setup (see figure 2.5c) for finesse measurements. These measurements are important for two reasons. First, the final goal of this study is to identify samples that perform well in a cavity setup, and thus a measurement where the crystals are actually measured within such a setup is necessary. Second, the finesse measurement provides a much more accurate determination of the reflectivity than the reflectivity measurement itself, by the use of equation (2.1), since the finesse is highly sensitive to small changes in reflectivity and thereby also to optical losses.

5.2.1 Setup

To carry out the finesse measurements, the setup depicted in figure 5.3 was used. In this setup, the laser first goes through a VOA and then through a paddle wheel polarization controller (PWPC). After the polarization controller, the laser goes into a phase modulating electro-optic modulator (EOM). The EOM is driven by a signal generator that adds a time-dependent phase to the light traveling through. This introduces two sidebands to the main tone guided through the fiber. This addition of sidebands will be discussed in the following section. After the EOM, the light is guided through a 90/10 fiber-optic splitter from which 10 % of the light goes through another paddle wheel polarization controller before it is sent to the cavity. At the cavity, the laser exits through a copper-coated, single-mode optical fiber with a concave mirror shot with a CO_2 laser into its facet. The reflected signal from the cavity goes back into the fiber splitter and the intensity is detected by a photodiode on the other side (PD1).

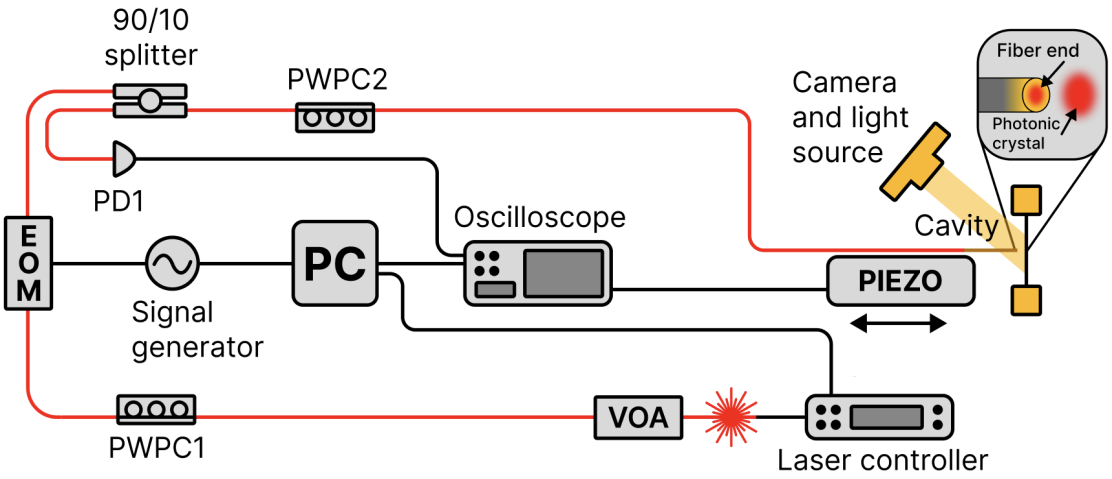

Figure 5.3: Illustration of the cavity setup used to measure the finesse. The setup uses a piezo crystal to modulate the length of the cavity in order to capture resonances. To calibrate the measurement, an EOM adds two sidebands to the incoming fiber mode.

The fiber end of the cavity is mounted on a shear plate actuator, modulating the cavity length by controlling a piezo crystal with a sinusoidal wave coming from the signal generator on the oscilloscope. On the other end of the cavity sits the sample, on a tiltable sample holder attached to a micrometer stage. The tiltable mount allows for alignment of the far mirror to reach optimal coupling of the optical mode in the fiber with the main optical mode in the cavity, while the micrometer stage is

used to align the fiber with respect to the samples on the chip.

5.2.2 Fit and rescaling of data

To ensure resonance inside the cavity is reached, the piezo modulating the cavity length is being sent a sinusoidal voltage signal, resulting in a back-and-forth motion of the fiber. When resonance occurs, the photodetector records a pattern consisting of three dips, a central dip corresponding to the main resonance, and two minor dips referred to as sidebands coming from the EOM (see figure 5.4). Because this trace is measured in time, the horizontal axis has to be converted into frequency. To achieve an accurate conversion, we assume linearity between the change in resonance frequency due to the length modulation and the length modulation itself, which is in turn approximately linear in time when the modulation is in the middle of two adjacent extrema [36]. Since the frequencies of the two sidebands are known, the time trace in the recording can be translated into frequency.

The shape of the trace depicted in figure 5.4 can be described by

$$\begin{aligned}
\frac{P_{\text{out}}}{P_{\text{in}}} &= \eta_r - D_{\text{main}} - D_{\text{SB1}} - D_{\text{SB2}}, \text{ where} \\
D_{\text{main}} &= \eta_L \left(\frac{1}{1+\nu^2} - \mathcal{S} \frac{\nu}{1+\nu^2} \right) \\
D_{\text{SB1}} &= \mathcal{S}_{\text{SB}} \eta_L \left(\frac{1}{1+\nu_{\text{SB1}}^2} - \mathcal{S} \frac{\nu_{\text{SB1}}}{1+\nu_{\text{SB1}}^2} \right) \\
D_{\text{SB2}} &= \mathcal{S}_{\text{SB}} \eta_L \left(\frac{1}{1+\nu_{\text{SB2}}^2} - \mathcal{S} \frac{\nu_{\text{SB2}}}{1+\nu_{\text{SB2}}^2} \right)
\end{aligned} \tag{5.1}$$

and ν, ν_{SB1} and ν_{SB2} are the frequencies of the main dip and the left and right sideband, respectively, $\mathcal{S}_{\text{SB}}$ is a scale factor compensating for the lower intensity of the sidebands compared to the main resonance and η_L, η_r and $\mathcal{S}$ are scaling factors depending on the spatial overlap between the cavity mode and the guided mode of the fiber and between the mode directly reflected by the incoupling mirror and the guided mode of the fiber [36]. To make the sidebands visible and not interfere with the main resonance dip, the frequency of the signal going into the EOM was tuned and the polarization of the light optimized with the paddle wheels (PWPC1 and PWPC2 in figure 5.3).

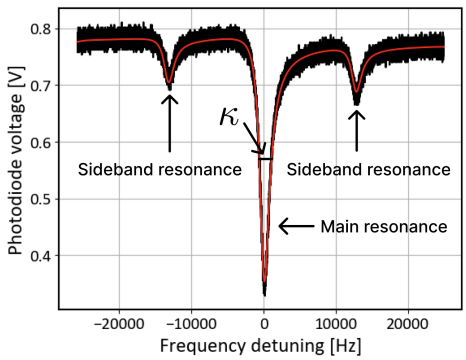

Figure 5.4: To calculate the finesse at a certain wavelength, a trace of the main resonance alongside its sidebands is captured and equation (5.1) is fit to the trace to extract the parameters needed to compute the finesse.

To retrieve a wavelength-dependent measurement of the cavity finesse, a laser frequency sweep was performed with this system. For every wavelength, a trace like the one depicted in figure 5.4 is recorded. Equation (5.1) is fit to the trace, from which κ is extracted as the FWHM of the main resonance dip. The FSR is determined by recording the amplitude of the main resonance dip while tuning the laser frequency between two adjacent and equally strong resonances, whereafter the finesse is calculated through equation (2.1). To calculate the reflectivity spectrum of the sample, equation (2.1) is used again. This time, either of R_1 and R_2 is taken as the photonic crystal reflectivity, and the other one as the fiber mirror reflectivity. The fiber mirror reflectivity is calculated by performing the same type of experiment with a near perfect far mirror.

6. Results

This chapter presents the results of the project, split into microfabrication and optical characterization. First, the etch recipe optimization is described, along with the resulting anisotropy and roughness of the fabricated pillars. Then, the reflectivity of the photonic crystals is measured and compared to simulations, followed by finesse measurements from the fiber cavity. Finally, the different loss mechanisms are compared to figure out which one limits the cavity performance the most.

6.1 Microfabrication optimization

6.1.1 Etch optimization

The optimization strategy for the etch recipe began with a baseline process known to produce reasonable results. The parameters for this recipe can be found in table 6.1. From this starting point, we performed parameter sweeps across five key process variables: temperature, RF platen power, pressure, chlorine gas content, and passivation gas content. Each of the parameters were tested separately while keeping the rest of the parameters unchanged. In this way, the etching effect of every parameter was isolated and could be evaluated individually.

Based on the trends observed in these sweeps, three candidate recipes were selected, each combining parameter values expected to improve on the baseline performance. These recipes were then used to fabricate new samples, which were subsequently characterized to evaluate the resulting quality. One of these improved recipes was then chosen for fabrication of the photonic crystals to be further characterized by assembling a cavity.

Temperature

In figure 6.1a, the temperature is kept at 25 °C, similar to in the baseline recipe. This seems to cause two main problems, re-deposition of Si/SiO_2 and slightly tapered sidewalls. These are both a result of a too low process temperature. Since a lower temperature decreases the surface reaction rate, the desorption of by-products is not efficient enough to remove detached material. This could explain the presence of the grass-like structures or islands surrounding the pillars (see figure 6.1a). In figure 6.1b, no trace of re-deposited material is present. This means that the reaction kinetics now allow for efficient desorption and removal of etched material. Additionally, the higher chemically driven etching process has affected the sidewall verticality positively, meaning that the tendency of tapered walls is removed.

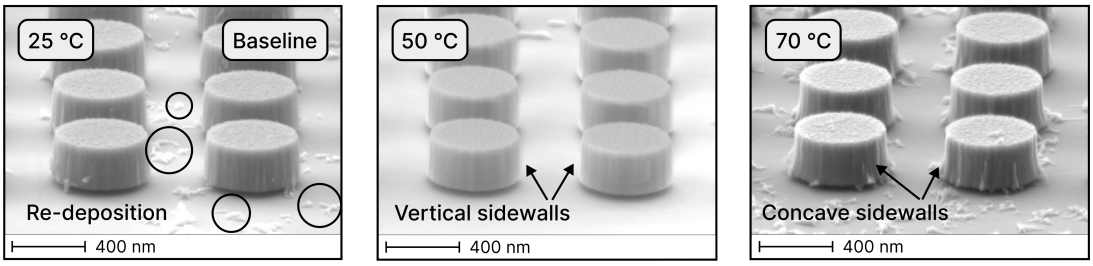

(a) **(b)** **(c)**

Figure 6.1: SEM images in (a), (b) and (c) show that the temperature sweep reveals a trade-off between re-deposition and reduced anisotropy. A higher temperature than in the baseline recipe (a) seems to be favorable.

At even larger temperatures, this is no longer the case. As can be seen in figure 6.1c, the walls start taking on a slightly concave shape. This is a natural result of an etch with higher temperature, since reduced anisotropy is a side effect of chemically driven etching. Furthermore, the re-deposition reappears, which might be a result of increased etch rate at which material removal once again becomes limiting. From this sweep, we see that a slight increase of the process temperature has a positive effect on the resulting structure.

RF platen power

As can be read in section 3.1.4, the RF platen power controls the kinetic energy of the ions and is therefore closely related to the etch rate. In general, the reasoning is that the higher the platen power, the higher the etch rate and the rougher the surfaces.

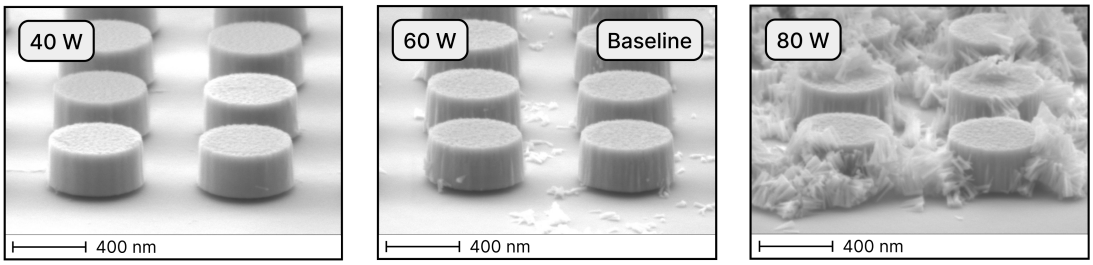

(a) **(b)** **(c)**

Figure 6.2: The SEM images in (a), (b) and (c) show that a lower platen power leaves less re-deposition of etched material.

In the baseline recipe (see figure 6.2b), the platen power is kept at 60 W. With this platen power choice, a vertical sidewall structured with visible ridges is achieved. It is worth noting that it is unclear whether the ridges are a result of the a-Si etch recipe or of the roughness present in the hard mask that has been transferred onto the amorphous silicon. Regardless, these ridges appear to be reduced when the platen power is decreased to 40 W (see figure 6.2a). This suggests that the amorphous silicon etch recipe is able to remove sidewall roughness in the hard mask, or that the lower platen power simply does not introduce roughness. In any of these cases, the lower platen power, despite producing a slightly tapered sidewall, seems to improve smoothness. This is also confirmed when looking at the 80 W case (see figure 6.2c), where the ridges are as pronounced as in the 60 W case. Furthermore, this recipe has led to heavy re-deposition of material. This could again be explained by the higher etch rate obtained by higher platen power, rendering the removal of the detached material insufficient. This leads us to onward reduce the platen power in our recipe.

Chamber pressure

Finding the right pressure is, as mentioned in section 3.1.4, about setting it low enough to achieve high anisotropy, but high enough to avoid substrate damage due to excessive bombardment of the surface. In this sweep, it becomes clear that the latter is more of a problem than the former.

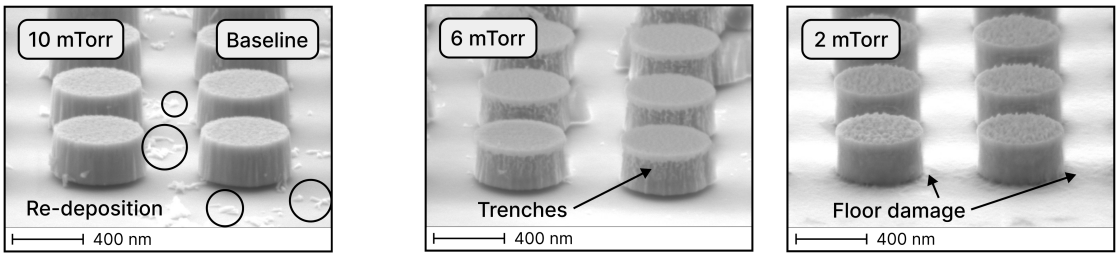

(a) **(b)** **(c)**

Figure 6.3: SEM images of the samples from the pressure sweep show that 10 mTorr (baseline) is not a too high pressure and that even higher pressures might improve the surface roughness.

In figure 6.3b, it is already evident that too low pressure damages the structure, where small irregular trenches appear on the surface. At 6 mTorr, they are not as pronounced, while at 2 mTorr, signs of damage to the bottom layer of the device appear in the images. This suggests that 10 mTorr is not limiting in terms of anisotropy and that even higher pressures may be worth investigating.

Gas composition

The gas composition balances physical and chemical etching effects. By varying the fraction of chlorine to bromine, the etching gas composition has been investigated.

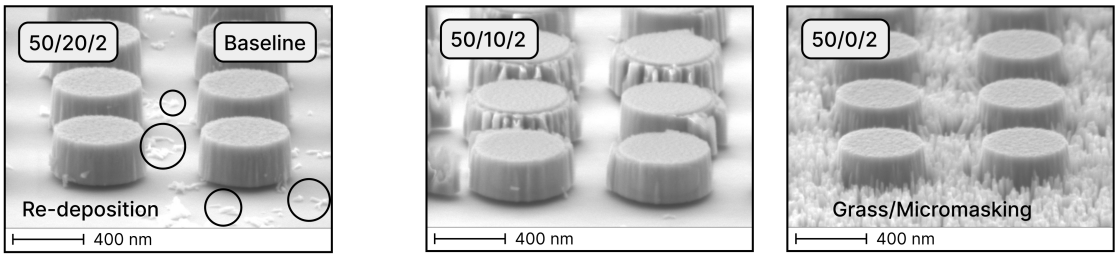

(a) **(b)** **(c)**

Figure 6.4: The gas composition sweep reveals that lowering the chlorine content from the flow rate in (a), the baseline recipe, either severely roughens the sidewalls or gives rise to micromasking. The composition is given as $HBr/Cl_2/O_2$ in units of sccm.

Figure 6.4 shows images of a decreasing chlorine fraction. As can be seen in figures 6.4b and 6.4c, the gas ratio does not affect the anisotropy and roughness of the pillars, but is rather responsible for the uniformity across the sample. In the case of 10 sccm Cl_2, a harsh structure surrounds many of the pillars. While the rough structures surrounding the pillars are gone in the 0 sccm chlorine recipe, this composition leads

to heavy micromasking effects. From a theoretical point of view, the formation of this "grass"-like structure can be explained. In contrast to the volatile compound-forming chlorine, the bromine promotes passivation. When the chlorine is removed, there will be more passivation and the removal of detached material slows down. The extreme case of this (see figure 6.4c) is when the etching is fully physical and microscopic islands of passivation remain due to lack of reactive species.

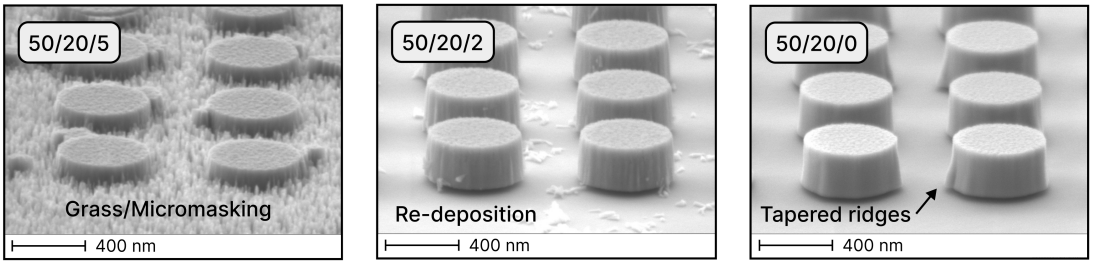

(a) **(b)** **(c)**

Figure 6.5: SEM images in (a), (b) and (c) of varying passivation gas flow show that micromasking occurs when the flow is increased from the baseline (b) and that tapered ridges appear when the passivation gas is removed.

The same logic can be applied to the flow rate of passivation gas. Enough oxygen is needed to passivate the surface and smooth out local bombardment fluctuations, while too much oxygen will cause a level of passivation that cannot be overcome by the etching gases, resulting in micromasking.

In figure 6.5, all of the above mentioned cases can be seen. In 6.5a, there is a comparatively high flow of oxygen (5 sccm) which shows heavy micromasking, whereas in 6.5c the passivation gas is removed, which has caused tapered ridges to form on the sidewalls.

Evaluation

Based on the findings from the parameter sweeps described above, three candidate recipes (referred to as recipe 1, 2 and 3, respectively) were selected, each combining parameter values expected to yield improved etch performance. In recipe 1, the pressure was set to 11 mTorr, exceeding the sweep range from above. This was done since the pressure sweep indicated that a higher pressure yielded a smoother sidewall surface. For all three recipes, the RF platen power was decreased and the temperature increased, in accordance with the results of the sweeps. The gas composition was held

constant. The full set of parameters for these three recipes is summarized in table 6.1.

Table 6.1: Summary of the etching parameters used in the baseline and the most optimal recipes. The baseline recipe is marked with an asterisk (*) and the recipe chosen for device fabrication is highlighted in gray.

	ICP power	RF power	Pressure	Flow of HBr/Cl2/O2	Temperature
*	360 W	60 W	10 mTorr	50/20/2 sccm	25 °C
1	360 W	40 W	11 mTorr	50/20/2 sccm	30 °C
2	360 W	40 W	10 mTorr	50/20/2 sccm	30 °C
3	360 W	45 W	10 mTorr	50/20/2 sccm	50 °C

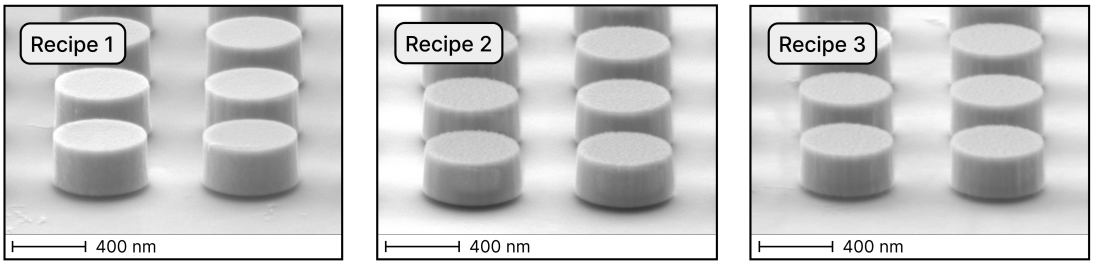

(a) **(b)** **(c)**

Figure 6.6: SEM images of pillars etched with the different parameters in table 6.1. From this visual inspection, recipe 1 yields the best results and is therefore used for the photonic crystal fabrication.

Figure 6.6 confirms that all three recipes led to improved etch quality compared to the baseline. Among these, Recipe 1 produced the smoothest sidewalls/surfaces, and was therefore selected as the process used to fabricate the photonic crystal samples.

6.1.2 Top surface roughness assessment

To verify that no significant roughness is introduced during the fabrication process, the topography of the sample surface was characterized by AFM after each deposition step. Figure 6.7 shows 200×200 nm^2 scans of the bare Si$\langle 100 \rangle$ substrate and after subsequent deposition of Si_3N_4 and HfO_2. The corresponding RMS roughness σ, computed with the method described in section 3.3.2, is given in each panel.

The bare silicon substrate shows a baseline roughness of $\sigma = 225$ pm, consistent with a polished wafer with a thin native oxide. After the Si_3N_4 deposition, the roughness

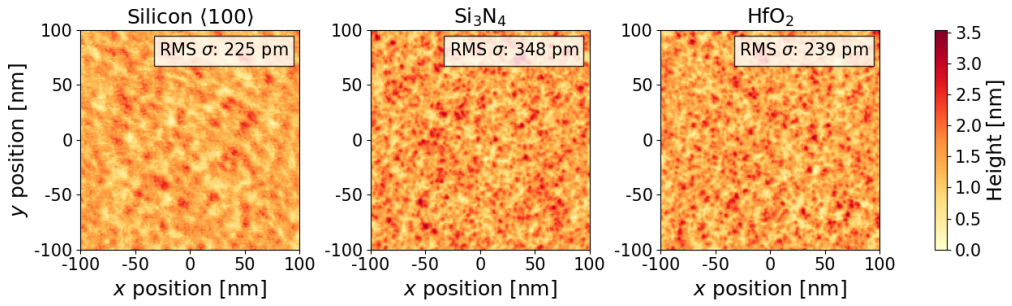

Figure 6.7: AFM height maps from every layer deposited prior to the amorphous silicon after the post-processing.

increases to $\sigma = 348$ pm. The Si_3N_4 layer is deposited by LPCVD at 780 °C. This high temperature generally promotes high surface mobility to smoothen the surface. However, since Si_3N_4 nucleates the surface at discrete, randomly distributed sites before growing laterally into a continuous film, grain boundaries and height mismatches contribute to roughness. Following HfO_2 deposition, the roughness decreases again to $\sigma = 239$ pm, close to the roughness of the starting substrate. This is expected, since the HfO_2 is deposited using atomic layer deposition (ALD). ALD is a layer-by-layer process in which each cycle coats the existing surface, smoothing over the underlying Si_3N_4 topography rather than adding further roughness.

Taken together, all three roughness values remain within a narrow sub-nanometer range. This indicates that none of the individual deposition steps introduced large defects and that the layer stack as a whole is of sufficient quality to proceed with pillar fabrication.

Figure 6.8 compares the roughness of pillar tops fabricated from sputtered silicon (σ = 690 pm) and from LPCVD-deposited silicon (σ = 330 pm). The result is that the LPCVD-deposited pillars are smoother, by more than a factor of two.

This difference is consistent with the growth mechanisms of the two deposition methods. Sputtering is a physical, non-equilibrium process in which atoms are ejected from a target and arrive at the substrate with high kinetic energy and a broad range of incidence angles. At the same time, the surface diffusion is limited in comparison to the high-temperature LPCVD settings. Under these conditions, adsorbed atoms tend to remain close to where they land rather than migrating to more energetically

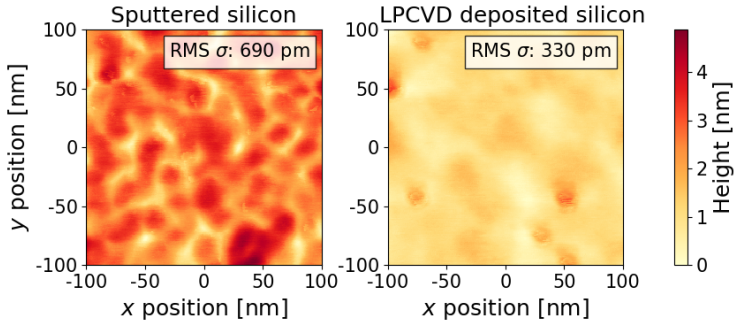

Figure 6.8: AFM height maps of the top of sputtered and LPCVD-deposited pillars.

favorable sites. This leads to columnar growth, that tends to increase with film thickness.

LPCVD on the other hand, is a chemically driven process carried out at high substrate temperature. The high temperature allows adsorbed atoms to diffuse laterally on the substrate to low-energy sites. The low pressure also gives a long mean free path in the gas surrounding the substrate, favoring uniform transport to the surface. Both these effects suppress the aggressive and geometry-dependent growth that contributes to the roughness seen in sputtering.

For both deposition methods, the measured roughness was reproducible. Pillars sampled from different locations on the same photonic crystal showed roughnesses within a 15 pm window for both deposition methods. This suggests that the pillar-top roughness is affected primarily by the growth kinetics rather than the patterning process, and supports treating the roughness values above as representative of each deposition method and not as artifacts coming from particular wafer areas.

6.1.3 Sidewall roughness assessment

In addition to the top-surface roughness discussed above, the sidewall roughness of the fabricated pillars was characterized, since this roughness is generated by the etching step itself rather than by any of the depositions. This characterization was carried out using the ellipse-fitting method described in section 3.2.4 applied to the three etching recipes that produced the smoothest-looking pillars by visual inspection in SEM. Figure 6.9 shows three pillars for each of these recipes, with the resulting

sidewall roughness Rq for every captured pillar.

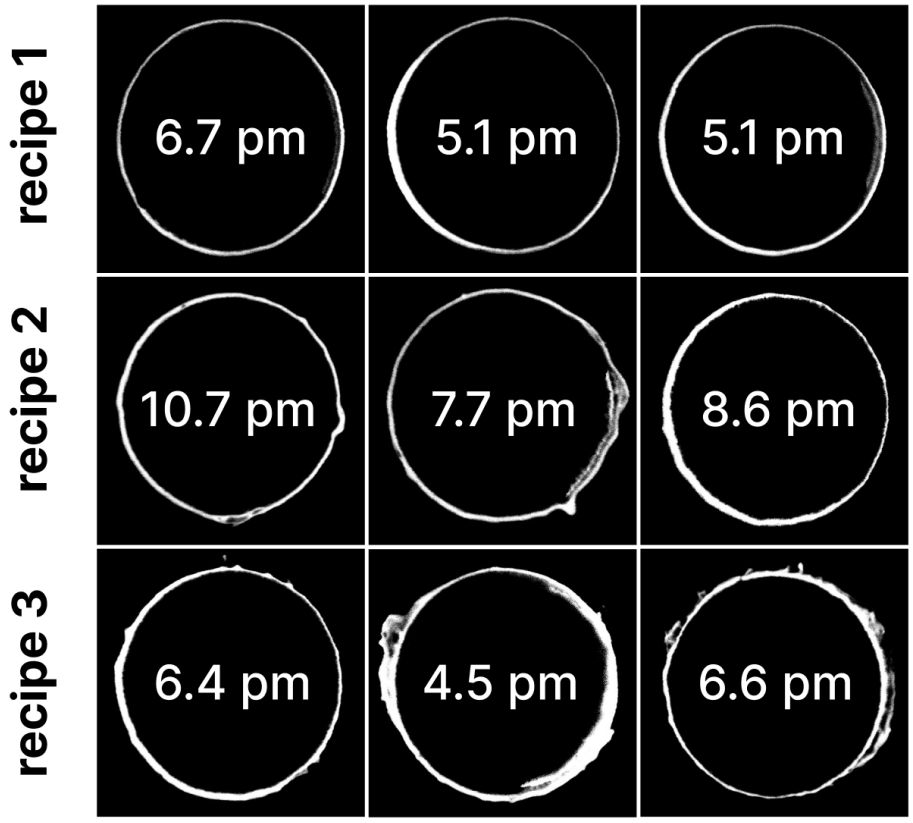

Figure 6.9: The sidewall roughness was captured by taking SEM images of the pillar tops and evaluating the smoothness of the edge line.

Recipe 1 gave an average sidewall roughness of Rq = 5.6 pm, recipe 2 gave Rq = 9.0 pm, and recipe 3 gave Rq = 5.8 pm. Recipes 1 and 3 are essentially indistinguishable within the precision of this measurement, while recipe 2 results in a rougher sidewall. Of these three, recipe 1 was selected as the recipe used for pillar fabrication going forward, based on its combination of low sidewall roughness and slightly greater anisotropy (see figure 6.6).

All three sidewall roughness values are two orders of magnitude smaller than the roughness measured on the pillar tops. This suggests that the sidewall roughness arising from the etch step is not a significant contributor to optical loss compared to the roughness already present on the top surface prior to etching, regardless of which of the three recipes is used. It should be noted, however, that this comparison is only qualitative, and that a quantitative loss estimate is not made for the sidewall

roughness in the same way as for the top-surface roughness. Equation (4.1) assumes light incident perpendicular to the rough surface, which is a reasonable approximation for the horizontal top surface of the pillar but does not apply to the vertical sidewalls.

6.2 Measurements of optical reflectivity

6.2.1 Nano-pillar design parameters

Since the thickness and lattice parameter for the photonic crystal design have already been fixed, the remaining free parameter available for optimization was the pillar radius r. We therefore performed a sweep over wavelength and radius, calculating the reflectivity for every combination. The resulting contour map (see figure 6.10a) reveals a clear transition region separating the high-reflectivity and low-reflectivity regimes, with the boundary shifting to longer wavelengths as r increases.

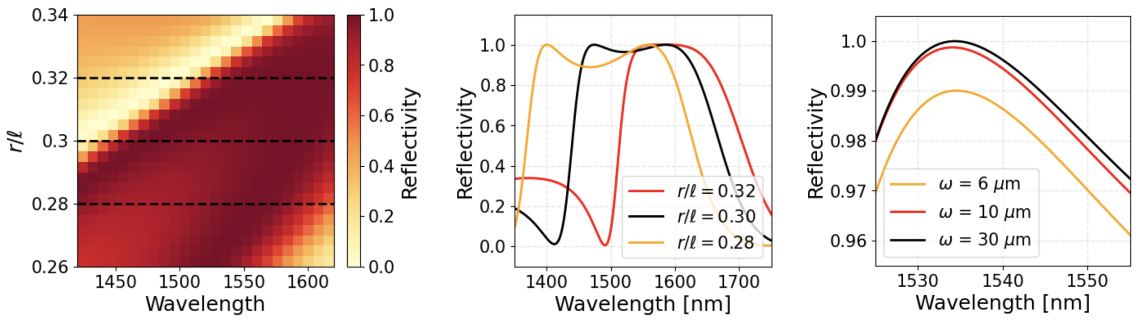

(a) **(b)** **(c)**

Figure 6.10: (a) Contour plot of the photonic crystal reflectivity. (b) Three alternatives for the pillar radius were extracted and out of these alternatives, $r = 0.3\ell$ performed best. (c) The $r = 0.3\ell$ design was then tested for incidence of Gaussian beams with three different beam waists.

To examine the dependence of reflectivity on wavelength, line cuts at three different radii were extracted ($r = 0.28$, $r = 0.3$ and $r = 0.32$) and plotted with respect to a larger wavelength region (see figure 6.10b). All three curves exhibit a high-reflectivity region, but its position and characteristics vary with r. Increasing the radius shifts the band toward longer wavelengths while narrowing it and decreasing the radius seems to do the opposite and additionally add a pronounced dip in the middle of the band.

Having identified $r = 0.3\ell$ as the radius yielding the highest reflectivity, the effect of beam waist size was investigated by repeating the $r = 0.3\ell$ calculation for three incoming beam waists, $\omega = 6, 10$ and 30 µm (see figure 6.10c). The result shows that tighter focusing (smaller waist) reduces the reflectivity peak, whereas the $\omega = 30$ µm case approaches the plane-wave limit. Quantitatively, this simulation shows that with these design parameters, the samples are expected to be robust to realistic beam waists.

6.2.2 Cavity based measurements

6.2.2.1 Finesse

To characterize the cavity performance experimentally, we performed two sets of finesse measurements. In the first, the finesse as a function of the FSR was determined for a sample made from sputtered Si using $r = 0.325\ell$, $\ell = 870$ nm and $h = 300$ nm as parameters. The radius here is larger than expected from the simulations. The reason for why larger radii worked better outside the simulations is not further investigated, but became apparent during the characterization. As shown in figure 6.11a, the finesse increases with increasing FSR. This is consistent with the expectation, as a larger FSR means a shorter cavity, which in turn means that the beam waist impinging the photonic crystal is smaller and so is the size of the beam returning to the fiber end. In figure 6.11b, the beam waists on the photonic crystal and on the fiber facet mirror are plotted according to formulas adapted from ref. [34], and confirm that for cavity lengths below 140 µm, the waist grows with cavity length. At very small cavity lengths, the finesse drops abruptly, which may be caused by thermal effects. When the waist becomes too small the intensity inside the cavity increases and at some point the heating of the membrane causes losses large enough to collapse the finesse.

In the second measurement, we compared the finesse of samples fabricated with sputtered versus LPCVD-deposited Si (see figure 6.12). The LPCVD sample consistently outperforms the sputtered one. Part of this difference can be attributed to the lower surface roughness of LPCVD films, which reduces scattering loss. However, the scattering loss is not close to being the limiting loss channel in either case (see section 6.2.2.2), indicating that additional loss mechanisms, plausibly related to absorption,

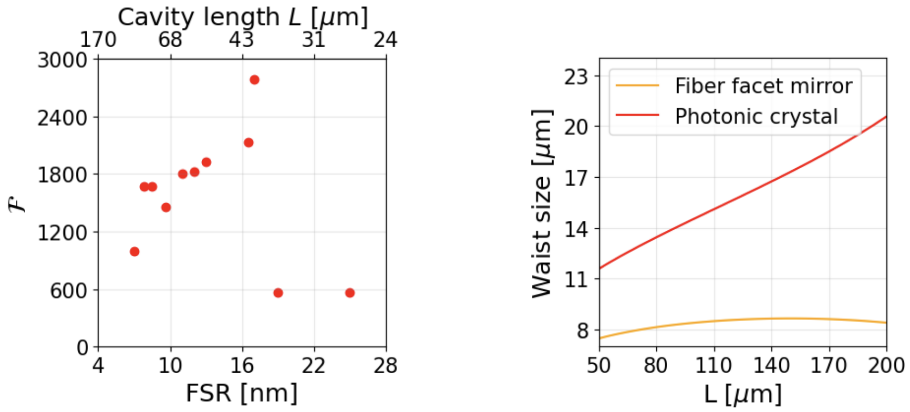

(a) **(b)**

Figure 6.11: (a) A measurement of finesse $\mathcal{F}$ in dependence of FSR and cavity length. (b) Theoretical graph of how the beam waists depend on the cavity length.

are also at play.

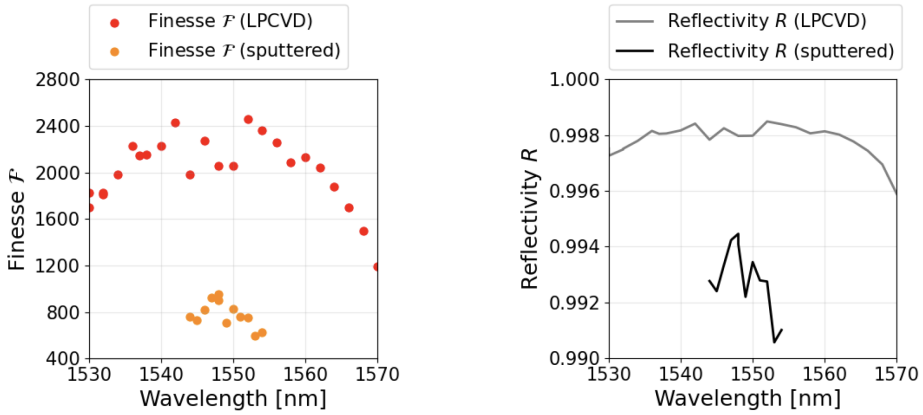

(a) **(b)**

Figure 6.12: (a) Wavelength sweep of the finesse for a sputtered and an LPCVD-deposited sample. The LPCVD-deposited sample yields a higher finesse compared to the sputtered one.

6.2.2.2 Losses

To assess the loss mechanisms of the cavity, calculations and simulations were performed in accordance with the methods presented in section 4.4. The clipping loss was evaluated by calculating the maximal achievable finesse assuming clipping loss as the only loss mechanism, as a function of cavity length, for two different photonic crystal mirror sizes (see figure 6.13a). The larger of the two corresponds

to the design used in this work, while the smaller is intended for the optomechanics version of the cavity. As expected for clipping loss, the finesse increases as the cavity length is decreased, supported by that a shorter cavity reduces beam spreading.

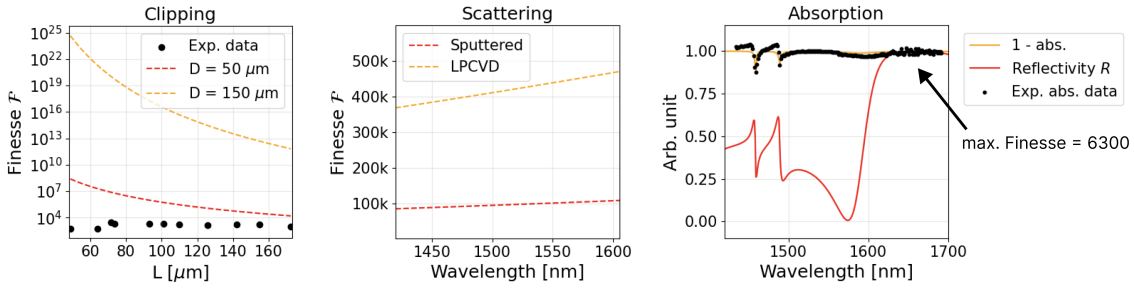

(a) (b) (c)

Figure 6.13: The loss mechanisms were computed separately and after comparison, absorption loss seems to be the most limiting loss channel.

The scattering loss was evaluated as a function of wavelength for both sputtered and LPCVD Si (see figure 6.13b). The finesse limited by scattering loss increases slightly with wavelength, consistent with reduced sensitivity to surface roughness at longer wavelengths. As expected, the rougher sputtered films show consistently higher scattering loss and correspondingly lower achievable finesse.

Taken together, these simulations indicate that neither clipping nor scattering loss is the limiting loss mechanism in the current devices since they limit the finesse approximately as much as the angular effect from the Gaussian beam. The absorption loss, however, seems to be the loss mechanism that limits the finesse the most. The rightmost panel in figure 6.13c shows a fit to experimental data from a reflectivity measurement according to the method described in section 4.4. When this fit is converted into finesse values, the limit is approximately 6000, which is on the same order of magnitude as the finesse measured in the actual samples. Given that the absorption-limited finesse lies so much closer to the measured values than any of the other computed loss channels, absorption is the most likely candidate for the mechanism currently limiting the cavity finesse.

7. Conclusion

This thesis improved a pillar-based amorphous silicon photonic crystal on a high-stress Si_3N_4 membrane. This photonic crystal acts as a mirror for out-of-plane light. We use it to build a short cavity together with a fiber mirror.

In the fabrication, we optimized the ICP-RIE etching recipe for the patterning of the pillars. We sweep the temperature, RF platen power, chamber pressure, and gas composition. The optimized recipe produced pillars with vertical sidewalls and a sidewall roughness of around 5-6 pm. AFM measurements showed that this sidewall roughness is two orders of magnitude smaller than the roughness of the pillar top surface, so it is not expected to contribute meaningfully to optical loss. The choice of deposition method for the amorphous silicon mattered more. LPCVD-deposited silicon gave a top surface roughness less than half that of sputtered silicon, and samples made from it consistently reached a higher finesse in the fiber cavity.

The measured finesse rose with increasing free spectral range consistent with a smaller mode waist on the photonic crystal for shorter cavities. The highest finesse values measured were 2500, below the target of 10^4 set out at the start of the project. Comparing the measured finesse against calculations of clipping loss, scattering loss, and absorption loss points to absorption as the dominant loss channel, with an absorption-limited finesse close to the values actually measured. Clipping and scattering loss, by contrast, would allow a finesse far beyond what was observed. This suggests that improving the finesse further is primarily a materials problem rather than a geometry or fabrication-precision problem, though the exact microscopic origin of the absorption was not identified in this work.

Overall, this project has shown that a pillar-based photonic crystal can be integrated onto a suspended Si_3N_4 membrane and used as an end mirror in a fiber cavity. The reflectivity achieved so far falls short of what is needed for the target application in

quantum cavity optomechanics, but the loss analysis gives a clear direction for the next steps rather than leaving the shortfall unexplained.

Future work should start by tracking down the source of the absorption loss. Post-deposition annealing is a natural first thing to try, since it is known to reduce absorption in amorphous silicon [37]. Once the reflectivity has been pushed closer to the 99.9 % target, the next natural step is to move from purely optical characterization toward an actual optomechanical measurement, where the coupling between the cavity light field and the membrane's mechanical motion is probed directly rather than assumed from the mirror's optical performance alone.

Bibliography

[1] Guanhao Huang et al. "Room-temperature quantum optomechanics using an ultralow noise cavity". In: *Nature* 626.7999 (2024), pp. 512–516. DOI: 10.1038/s41586-023-06997-3. URL: https://www.nature.com/articles/s41586-023-06997-3.

[2] S. Mancini and P. Tombesi. "Quantum noise reduction by radiation pressure". In: *Physical Review A* 49.5 (1994), pp. 4055–4065. DOI: 10.1103/PhysRevA.49.4055. URL: https://journals.aps.org/pra/abstract/10.1103/PhysRevA.49.4055.

[3] D. Mason et al. "Continuous force and displacement measurement below the standard quantum limit". In: *Nature Physics* 15 (2019), pp. 745–749. DOI: 10.1038/s41567-019-0533-5. URL: https://www.nature.com/articles/s41567-019-0533-5.

[4] B. M. Brubaker et al. "Optomechanical Ground-State Cooling in a Continuous and Efficient Electro-Optic Transducer". In: *Physical Review X* 12 (2022), p. 021062. DOI: 10.1103/PhysRevX.12.021062. URL: https://link.aps.org/doi/10.1103/PhysRevX.12.021062.

[5] Nils Johan Engelsen, Alberto Beccari, and Tobias Jan Kippenberg. "Ultrahigh-quality-factor micro- and nanomechanical resonators using dissipation dilution". In: *Nature Nanotechnology* 19.6 (2024), pp. 725–737. DOI: 10.1038/s41565-023-01597-8.

[6] Dennis Høj, Ulrich Busk Hoff, and Ulrik Lund Andersen. "Ultracoherent Nanomechanical Resonators Based on Density Phononic Crystal Engineering". In: *Physical Review X* 14.1 (2024), p. 011039. DOI: 10.1103/PhysRevX.14.011039.

[7] Ugo Fano. “Effects of Configuration Interaction on Intensities and Phase Shifts”. In: *Physical Review* 124.6 (1961), pp. 1866–1878. DOI: 10.1103/PhysRev.124.1866.

[8] Shanhui Fan and John D. Joannopoulos. “Analysis of guided resonances in photonic crystal slabs”. In: *Physical Review B* 65 (2002), p. 235112. DOI: 10.1103/PhysRevB.65.235112.

[9] John D. Joannopoulos et al. *Photonic Crystals: Molding the Flow of Light*. 2nd. Princeton, NJ: Princeton University Press, 2008.

[10] Kazuaki Sakoda. *Optical Properties of Photonic Crystals*. 2nd. Berlin: Springer, 2005.

[11] Georg Enzian et al. “Phononically shielded photonic-crystal mirror membranes for cavity quantum optomechanics”. In: *Optics Express* 31.8 (2023). arXiv:2212.12148, p. 13040. DOI: 10.1364/OE.484369.

[12] Sampo A. Saarinen et al. “Laser cooling a membrane-in-the-middle system close to the quantum ground state from room temperature”. In: *Optica* 10.3 (2023). arXiv:2206.11169, p. 364. DOI: 10.1364/OPTICA.468590.

[13] Bahaa E. A. Saleh and Malvin Carl Teich. *Fundamentals of Photonics*. 3rd. Hoboken, NJ: Wiley, 2019.

[14] Markus Aspelmeyer, Tobias J. Kippenberg, and Florian Marquardt. “Cavity optomechanics”. In: *Reviews of Modern Physics* 86.4 (2014), pp. 1391–1452. DOI: 10.1103/RevModPhys.86.1391.

[15] Michael Huff. “Recent Advances in Reactive Ion Etching and Applications of High-Aspect-Ratio Microfabrication”. In: *Micromachines* 12.8 (2021), p. 991. DOI: 10.3390/mi12080991. URL: https://www.mdpi.com/2072-666X/12/8/991.

[16] David J. Griffiths. *Introduction to Electrodynamics*. 3rd. Upper Saddle River, NJ: Prentice Hall, 1999.

[17] Vincent M. Donnelly and Avinoam Kornblit. “Plasma etching: Yesterday, today, and tomorrow”. In: *Journal of Vacuum Science & Technology A* 31.5 (2013), p. 050825. DOI: 10.1116/1.4819316.

[18] WebElements. *The periodic table of the elements*. https://www.webelements.com. Accessed: 2026-07-06.

[19] J. W. Coburn and Harold F. Winters. “Ion- and electron-assisted gas-surface chemistry — An important effect in plasma etching”. In: *Journal of Applied Physics* 50 (1979), pp. 3189–3196. DOI: 10.1063/1.326355.

[20] J. W. Coburn and Harold F. Winters. “Plasma-assisted etching: Ion-assisted surface chemistry”. In: *Applied Surface Science* 22–23 (1985), pp. 63–71. DOI: 10.1016/0378-5963(85)90037-6.

[21] Samantha Roberts et al. “Measurements and Modeling of Atomic-Scale Sidewall Roughness and Losses in Integrated Photonic Devices”. In: *Advanced Optical Materials* 10 (2022), p. 2102073. DOI: 10.1002/adom.202102073.

[22] Laurence Marks. “What Are the Resolution Limits in Electron Microscopes?” In: *Physics* 6 (2013), p. 82. DOI: 10.1103/Physics.6.82. URL: https://physics.aps.org/articles/v6/82.

[23] ZEISS. *ZEISS GeminiSEM*. https://www.zeiss.com/microscopy/us/products/sem-fib-sem/sem/geminisem-family.html#accordionItem-721420587. Accessed: 2026-07-06.

[24] JEOL. *A Guide To Scanning Microscope Observation*. https://www.jeolusa.com/RESOURCES/Electron-Optics/Documents-Downloads/a-guide-to-scanning-microscope-observation. Accessed: 2026-07-06. 2020.

[25] K. Kanaya and S. Okayama. “Penetration and energy-loss theory of electrons in solid targets”. In: *Journal of Physics D: Applied Physics* 5.1 (1972), pp. 43–58. DOI: 10.1088/0022-3727/5/1/308.

[26] Andrew W. Fitzgibbon, Maurizio Pilu, and Robert B. Fisher. “Direct Least Squares Fitting of Ellipses”. In: *Proceedings of the 13th International Conference on Pattern Recognition (ICPR)*. 1996.

[27] AFM Workshop. *Atomic Force Microscopy*. https://www.afmworkshop.com/learning-center/atomic-force-microscopy. Accessed: 2026-07-06.

[28] Victor Liu and Shanhui Fan. “S4: A free electromagnetic solver for layered periodic structures”. In: *Computer Physics Communications* 183.10 (2012), pp. 2233–2244. DOI: 10.1016/j.cpc.2012.04.026.

[29] Wei Jin et al. “Inverse Design of Lightweight Broadband Reflector for Relativistic Lightsail Propulsion”. In: *ACS Photonics* 7.9 (2020), pp. 2350–2355. DOI: 10.1021/acsphotonics.0c00768.

[30] M. G. Moharam et al. “Formulation for stable and efficient implementation of the rigorous coupled-wave analysis of binary gratings”. In: *Journal of the Optical Society of America A* 12.5 (1995), pp. 1068–1076. DOI: 10.1364/JOSAA.12.001068.

[31] D. M. Whittaker and I. S. Culshaw. “Scattering-matrix treatment of patterned multilayer photonic structures”. In: *Physical Review B* 60.4 (1999), pp. 2610–2618. DOI: 10.1103/PhysRevB.60.2610.

[32] Lukas Novotny and Bert Hecht. *Principles of Nano-Optics*. 2nd ed. Cambridge: Cambridge University Press, 2012. ISBN: 9780511794193. DOI: 10.1017/CBO9780511794193. URL: https://www.cambridge.org/core/books/principles-of-nanooptics/EAA3E0D04179A7252088E27483C4ACC1.

[33] Sushanth Kini Manjeshwar et al. “Suspended photonic crystal membranes in AlGaAs heterostructures for integrated multi-element optomechanics”. In: *Applied Physics Letters* 116.26 (2020), p. 264001. DOI: 10.1063/5.0012667.

[34] David Hunger et al. “A fiber Fabry–Perot cavity with high finesse”. In: *New Journal of Physics* 12.6 (2010), p. 065038. DOI: 10.1088/1367-2630/12/6/065038.

[35] John C. Stover. *Optical Scattering: Measurement and Analysis*. 3rd. Bellingham, WA: SPIE Press, 2012. ISBN: 9781628418408. DOI: 10.1117/3.975276.

[36] Johanna Örgård. “A vacuum integrated fibre Fabry-Pérot cavity setup for optical and optomechanical characterisation”. Department of Microtechnology and Nanoscience, Quantum Technology Laboratory. Master’s thesis. Gothenburg, Sweden: Chalmers University of Technology, 2025.

[37] Gravitational Wave International Committee (GWIC) 3G Committee, David Reitze, and Michele Punturo. *3G R&D: R&D for the Next Generation of Ground-based Gravitational-wave Detectors*. Tech. rep. Final release, April 2021. Gravitational Wave International Committee (GWIC), Apr. 2021. arXiv: 2111.06991 [gr-qc]. URL: https://arxiv.org/abs/2111.06991.

DEPARTMENT OF SOME SUBJECT OR TECHNOLOGY

CHALMERS UNIVERSITY OF TECHNOLOGY

Gothenburg, Sweden

CHALMERS
UNIVERSITY OF TECHNOLOGY