

Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich

Matteo Stefanini 22-937-890

Characterization of integrated photonic structures for trapped ion quantum computing

Semester Thesis

MSc Quantum Engineering Trapped Ion Quantum Information Group (TIQI) Institute for Quantum Electronics Swiss Federal Institute of Technology (ETH) Zurich

Supervision

Gillenhaal Beck Prof. Dr. Jonathan Home

March 2024

Abstract

Integrated photonic circuits provide a more robust and parallelizable way to address trapped ions for quantum computation. Light in planar-fabricated waveguides made from conventional silicon nitride (Si_3N_4) or alumina (Al_2O_3) is outcoupled to free-space using diffraction gratings.

In the first part of this work we build a beam profiling setup and analysis suite. These tools allow precise beam characterizations and comparisons of the fabricated diffraction gratings with the simulations used for their design.

In the second part, we investigate loss in straight waveguide sections of alumina using both a transmission measurement as well as a differential measurement with grating outcouplers. The preliminary measurements produced loss values in the range of 1.6 - 7.25 dB/cm.

Lastly, we point out ways to improve the setup to yield better and more reproducible results, with little additional effort.

Contents

A	bstra	ct	i
1	Intr 1.1 1.2	oduction Motivation	1 . 1 . 1
2	Exp 2.1 2.2	erimental Setup Outcoupler Characterization	3 . 3 . 5
3	Ana 3.1 3.2	lysis Outcoupler Beam Profiling 3.1.1 Image Rotation 3.1.2 Visualization 3.1.3 Emission Angle 3.1.4 Beam Waist 3.1.5 Other Examples Waveguide Loss Measurement 3.2.1 Waveguide Loss Using Grating Outcouplers 3.2.2 Waveguide Loss Using Transmission	8 . 8 . 10 . 11 . 12 . 13 . 16 . 16
4	Cor	clusion and Outlook	23 23

Chapter 1

Introduction

1.1 Motivation

Utilizing integrated optics for addressing trapped ions has many benefits compared to free-space optics. It provides a stable optical path to the ion, which is robust towards vibrations. Moreover, large free-space optical setups become a challenging task as the number of ions increases, and are difficult to reproduce. Integrated waveguides have a small footprint and can leverage industrial CMOS processes. Thus, together with the surface-electrode QCCD architecture for shuttling ions, integrated photonics are a promising approach towards scaling trapped ions as a platform for quantum information processing [1], [2], [3]. In addition to allowing high-fidelity operations in a parallelizable manner, grating outcouplers allow the emission of complex spatial field profiles with large field gradients [4] or, capitalizing on the stable optical path, the creation of passively phase-stable standing waves [5]. These capabilities augment the toolbox for manipulating trapped ions and enable novel interaction schemes.

While much work has been done for telecom wavelengths using silicon nitride (Si_3N_4) , alumina (Al_2O_3) has only recently surfaced as a low-loss waveguide material for UV-wavelengths [6]. The required small feature size for diffraction gratings due to the short wavelengths poses additional challenges for the fabrication. It is therefore instrumental to characterize the light emission from these grating structures, as well as the losses in the specific waveguide architecture of our chips. Such capabilities can also be employed to validate recent simulations of waveguide bends and integrated power splitters [7].

1.2 Grating Outcouplers

Figure 1.1 illustrates the device cross section of a trap chip with a backward-emitting diffraction grating. To allow the emitted beam to pass, there are holes in the ground plane and trap electrodes. A conductive but transparent layer of indium tin oxide (ITO) on the surface helps to prevent charging of exposed dielectric surfaces.

The periodic pattern for a grating is etched into the waveguide which is widened with a taper. This periodic variation of refractive index causes diffraction of the light out of the chip plane. While asymmetric etchings leading to directional emission are possible [8], our designs emit symmetrically towards the top of the chip and towards the substrate.

The design of the grating is split into two parts: the "longitudinal" design which restricts the behavior to the xz-plane in Fig. 1.2, and the "transverse" design along y.

For the longitudinal direction, the light emission is modified using the emission angle θ and the grating strength α , and their respective spatial evolution from x_1 to x_2 . α is defined such that the guided field amplitude decays as $e^{-\alpha x}$ when propagating through the centerline of the grating



Figure 1.1: Schematic cross section of a chip design with a backward-emitting diffraction grating, indicating the different layers and their thickness. On the left, the intensity profiles of the fundamental quasi-TE mode in the waveguide are shown for both silicon nitride and alumina. The width of the waveguides for single-mode operation ranges from 400-600 nm. Taken from [4].

along x. θ and α depend on the etching period Λ and the duty cycle DC, which describes the ratio of etch length to period. Thus, varying Λ and DC as a function of x tailors the longitudinal emission.

For the transverse direction, the obtained 2D grating structure is extruded along specific grating line curvatures. These curvatures are determined from holographic interference between waveguide field and desired emitted field, which ensures their phase matching in the waveguide plane. To accurately describe the phase evolution required for the holographic interference, the in-grating effective index n_{eff} as a function of Λ and DC needs to be determined. This ultimately allows transverse focusing, as well as adding additional phase profiles to create higher order modes. The detailed design procedure is described in [4].



Figure 1.2: Schematic of grating and coordinate system. The cyan and red arrows indicate the longitudinal and transverse beam axes, respectively. Taken from [4].

Chapter 2

Experimental Setup

2.1 Outcoupler Characterization



Figure 2.1: Photo of the setup used for beam profiling. The main components are a 50x microscope with a high-NA objective (NA = 0.95) attached to a vertical motorized stage, an aluminum slab to hold the chip, and a 6-axis stage to align a fiber with waveguides on the edge of the chip. The 3-axis translation stage on the left-hand side was not used in the measurements.

To characterize the output of the fabricated grating outcouplers, the setup shown in Fig. 2.1 is used. The chip with grating outcouplers can be placed on an aluminum slab. Above the chip, there is a microscope with a 50x magnification objective¹ attached to a motorized stage which moves the objective vertically, perpendicular to the chip's surface. Together with the shallow focal depth created by its high numerical aperture (NA = 0.95), this allows to image thin slices parallel to the chip's surface at different heights (see Fig. 2.2). The output of the microscope is imaged by

 $^{^1 \}mathrm{Olympus}$ MPLANAPON $50 \mathrm{x}/0.95$ NA

a CCD camera². The camera delivers pictures with a resolution of 1392x1040 pixels and a single 8-bit intensity channel (no colors), providing 256 discrete intensity levels. After this project, it was discovered that the camera can also be toggled to 16-bit, distinguishing 65,536 intensity values. This will significantly improve the results obtainable with the analyses described in Sec. 3 as many are limited by the dynamic range and the resulting discretization artifacts at low beam intensities.

A 6-axis stage³ is used to couple laser light from a fiber into different waveguides on the chip. For the tests performed in this work, we use a 729 nm laser diode for silicon nitride waveguides and 532, 423, and 375 nm laser diodes for alumina waveguides.

A typical measurement of the beam profile is conducted by first using the microscope with a white light source to align the fiber with the waveguide of interest at the edge of the chip. Once scattered light can be seen from the waveguide, indicating successful incoupling, the microscope is moved along the waveguide from the edge of the chip to the respective grating outcoupler. Now, the height of the microscope is manually adjusted such that the features of the grating are in focus. We define this height as $z = 0 \,\mu$ m. To reduce the influence of vibrations during the imaging, the white light source is removed from the microscope. An automated script interfaces with the camera and the motorized stage: Typically, a picture is taken for every increment of $1 \,\mu$ m, up to a height of $z = 100 \,\mu$ m. Example pictures at different heights are shown in Fig. 2.3. This stack of 100 images contains the full 3D information for the intensity distribution of the beam emitted by the grating, similar to an MRI scan that records 3D information as a set of 2D slices through the object of interest.



Figure 2.2: Schematic of taking z-slice images with a high numerical aperture microscope by vertically scanning different heights above the chip. Image taken from [3].

To calibrate the camera in the horizontal plane, i.e, devise a conversion from pixels to μ m, a chip similar to the design shown in Fig. 2.4 is examined under the microscope. The given vertical separation between the waveguides on the left can be compared to the number of pixels that cover this distance in the camera image. This analysis (done using "ImageJ" [9]) leads to the conversion that 1μ m corresponds to 8.974 ± 0.032 pixels.

²Lumenera INFINITY3S-1UR

 $^{^3\}mathrm{Thorlabs}$ NanoMax600



Figure 2.3: Camera images through the microscope at different heights above the chip. At $z = 0 \,\mu$ m, one can see the structure of the grating. For larger heights, the beam is focused to a narrower spot. For $z = 30 \,\mu$ m and $z = 70 \,\mu$ m the second order emission of this (forward emitting) grating is visible in the bottom of the frame. The diagonal fringes are artifacts from the beam splitter in the imaging system that is needed to inject white light into the imaging path. At the time of writing, the setup has been modified such that the beam splitter can be removed prior to each measurement, thus removing the diagonal fringes.



Figure 2.4: Chip design for loss measurements using grating outcouplers. The camera can be calibrated with the given vertical separation of the straight waveguides on the left $(254 \,\mu\text{m})$.

2.2 Waveguide Loss Measurement Using Transmission

The setup described in Sec. 2.1 can be used for one method to measure the waveguide loss. A second method involves placing a chip with the waveguide design shown in Fig. 2.5 into the setup shown in Fig. 2.6. The design features a set of structures with different lengths of straight waveguide sections, also known as paperclip structures, guiding light from one end of the chip to the other. By relating the power transmitted through those different waveguides to their length, the waveguide loss per unit length can be extracted. Note that the number and radius of the bends does not change from one structure to another. Therefore, the relative transmitted powers yield information about the loss in straight waveguide, decoupled from the loss induced by the bends. The setup in Fig. 2.6 allows to approach the chip with fibers from both sides: On the right, a fiber carrying the input light is placed on a 6-axis stage. At can be adjusted to maximize the incoupling efficiency into the chip. On the left side, a fiber on a 3-axis translation stage is used to collect the light coming out of the on-chip waveguides. A pitch/yaw stage was present but not used due to its coarse resolution. Additionally, the chip holder can be translated perpendicular to the fibers. A digital microscope⁴ attached to a translation stage and suspended above the chip facilitates the initial alignment.

⁴Dino-Lite Digital Microscope



Figure 2.5: Chip design for measuring the waveguide loss via transmission with in- and outcoupling. Such waveguides with identical bends and variable straight waveguide length are often called "paperclip" structures.

A typical loss measurement consists of the following steps: First, the microscope is placed to view the right-hand facet of the chip. Using the live image (see Fig. 2.7 (a)), the position of the fiber is adjusted such that it is aligned with the waveguide at the chip edge. Successful incoupling is discerned by the waveguide lighting up, as seen in Fig. 2.7 (b). Now, the microscope is moved to the left facet, where the light exits the waveguide. The fiber, which is connected to a power meter, is visually aligned with the waveguide on the chip until the power meter indicates the presence of light. Both the in- and outcoupling fibers can now be fine-tuned to maximize the power shown on the power meter. Smaller distances between fibers and chip facets lead to better in- and outcoupling. However, one needs to be careful not to bump the chip, as this ruins the alignment. After recording the maximum achieved value, the chip is moved perpendicular to the fiber until the next, longer waveguide is reached and the power meter shows a signal. The fine-tuning is repeated to optimize the in- and outcoupling.



Figure 2.6: Setup to measure waveguide losses by coupling light into the chip on one side and collecting the output on the other side.

This procedure is repeated for all five waveguide sections. Since it is difficult to maintain a constant coupling efficiency, the measurements are repeated several times for each waveguide, moving the chip all the way up and down several times. To further combat this issue, the chip is polished on



Figure 2.7: Images from the digital microscope. (a) depicts the incoupling fiber and the chip facet, whereas (b) shows a waveguide with successfully incoupled light.

both sides before the measurement. The chip is clamped into an aluminum holder and manually polished with four polishing films with decreasing coarseness. During this procedure, which should be conducted in a clean room, one needs to be careful not to introduce a curvature to the facet, i.e., keeping the chip edge parallel to the polishing films at all times.

While the setup in Fig. 2.1 existed already and only needed minor modifications, the setup in Fig. 2.6 was built during the project.

Chapter 3

Analysis

3.1 Outcoupler Beam Profiling

As described in Sec. 1, outcoupling beams of tailored shapes from integrated waveguides is a relatively recent development. Particularly for the short wavelengths used in trapped ions, the required fine grating structures pose difficulties when fabricating these devices. To characterize the beam that the ion will interact with, we need to extract beam parameters and the 3D intensity profile from the camera images, see Sec. 2.1. These can also be used to compare the fabricated designs to the simulations (obtained using a software like Ansys Lumerical [10]) and inform the design process for future gratings.

In the following, the analysis steps are visualized using measurements of a grating designed to emit a loosely backward-focusing beam of 732 nm light at an angle of 32° . Starting from the grating at height $z = 0 \,\mu$ m, 100 images in height increments of $1 \,\mu$ m were taken. A 729 nm laser diode was used for this measurement. In an experiment, this grating will be used to deliver both 729 nm and 732 nm light as both wavelengths are relevant for Ca ions.

3.1.1 Image Rotation

Since the emission is not perpendicular to the chip surface, the beams position appears to follow a linear trajectory in the z-slices imaged by the camera for increasing heights above the grating. This trajectory is aligned with the centerline of the grating. To facilitate the analysis of the pictures, they are rotated such that the trajectory is horizontal. In addition, the are cropped/extended so that the image stacks from all measurements have the same format.



Figure 3.1: Pictures taken by the camera at three selected heights above the grating outcoupler, without any post-processing.

Figure 3.1 shows the original pictures taken by the camera for three different heights. One

can see that the beam position moves from the top and slight right towards the bottom center of the image. To get the precise angle of this trajectory, the position of each pixel in an image is weighted with its intensity value. The weighted average of all pixels yields the average position of the intensity. This process is repeated for all 100 images in the stack and the average positions are plotted in Fig. 3.2, where the darker markers indicate lower heights. The slope of the linear fit gives the rotation angle required to make the trajectory horizontal, in this example 82.7° .



Figure 3.2: Scatter plot of the location of the average intensity for all 100 pictures. The higher above the chip, the lighter the symbol color. The slope of the linear fit yields the angle with respect to the horizontal.

The rotated pixel grid need to be mapped to a new grid with horizontal/vertical pixel rows and columns. The Python Imaging Library (PIL) which is used for the rotations offers different methods for this resampling, including nearest neighbor and interpolation techniques. By visually assessing the different methods, it was determined that the bilinear method (linear interpolation in a 2x2 environment) minimizes resampling artifacts in the rotated image. The rotated and resized images are shown in Fig. 3.3. The coordinate system indicated on the right matches the grating frame of Fig. 1.2, with the z-axis pointing out of the image plane.



Figure 3.3: Same pictures as in Fig. 3.1, but rotated and cropped such that the beam propagates horizontally backwards.

3.1.2 Visualization

The stack of images contains the 3D information of the beam's intensity profile. Different visualizations are useful to compare this information to the simulations.

First, it is useful to look at slices parallel to the chip surface, together with line plots through the pixel with maximum intensity along the x- and y-direction, see Fig. 3.4. For the slice at $z = 7 \,\mu$ m, one can see the influence of the grating structure along the x-direction. For the slice at $z = 55 \,\mu$ m on the other hand, which is close to the focus of the emitted beam, the distribution appears roughly Gaussian along both axes.



Figure 3.4: Heatmaps of intensity distributions at $z = 7 \,\mu$ m and $z = 55 \,\mu$ m. On the right, line plots through the maximum pixel along the x- and y-direction show the profile in more detail. Circular artifacts in the measurements, e.g., for $(x, y) \approx (3, -1) \mu$ m in the $z = 7 \,\mu$ m slice, appear due to dust in the imaging system.

Secondly, the side view in the xz-plane is useful. It provides a more natural image of the beam compared to the z-slices. There are two ways to approach this: One can either integrate the 3D intensity array along the y-axis, or take a slice of it through the y-coordinate of the maximum beam intensity (see Fig. 3.5 (a) and (b), respectively). The first is more natural since it corresponds to what we would see in an actual side view if the beam was made visible. The latter highlights the focusing nature of the beam along the y-direction. Additionally, it is necessary for characterizing the optical vortex beam shapes shown in Sec. 3.1.5, since the tubular profile is lost when integrating the intensity along y.



Figure 3.5: Side view of the beam in the xz-plane. (a) shows the intensity integrated along the y-axis, whereas (b) shows a slice through the intensity profile perpendicular to y.

The third view, the xz-plane, is useful to assess the focusing behavior along the y-axis (see Fig. 3.6). Since the beam is emitted at an angle in the xz-plane, the sliced version perpendicular to the x-axis only displays one part of the beam.

For many measurements, particularly around the focus, the pictures are overexposed. This means that the distribution is clipped around the location of peak intensity, where many pixels assume the maximum value of 255 (considering the 8-bit mode of the camera). In these cases, I define the position of the maximum intensity as the average position of all pixels with intensity 255. Fig. 3.7 shows such a case. While this is undesirable, for tightly focusing beams the limited dynamic range of the camera makes it impossible to both not clip the intensity around the focus and also keep details at low intensity closer to the grating at the same time. Hence the exposure time should be carefully chosen for the use case of the measurement: If one is more interested in the beam shape close to the grating, a long exposure time should be chosen to avoid discretization artifacts at low intensities. On the other hand, if the region of interest is the focus of the beam, an exposure time short enough to not clip the intensity is needed. These requirements can likely be relaxed in the future when measurements are taken with the camera's intensity channel toggled to 16-bit.

3.1.3 Emission Angle

An important beam parameter is the emission angle θ with respect to the normal on the chip surface (see Fig. 1.2). Making sure that the fabricated devices produce the expected angle is essential for hitting the ion, especially when more beams need to coincide at the ion's position.

To extract the emission angle, we consider the z-slices taken by the camera. We want to track by how much the intensity moves along the x-axis when increasing the height. After translating the pixel position to μ m using the calibration data from Sec. 2.1, the slope of a fitted line to the height vs x-position plot yields the angle.

The naive way to track this shift in x-position is to only consider the pixel with maximum intensity in each slice (see. Fig. 3.8 (a)). However, this is subject to noise and, for backwards emitting gratings, underestimates the angle because most of the emission happens at the beginning of the grating.



Figure 3.6: Side view of the beam in the xz-plane. (a) shows the intensity integrated along the y-axis, whereas (b) shows a slice through the intensity profile perpendicular to y.

A second option is to fit a 1D Gaussian to each slice, after integrating over the y-direction to get a 1D-array of intensity values. This roughly corresponds to fitting Gaussians to the blue line plots in Fig. 3.4, however, using values integrated along y rather than sliced. The peak positions of the Gaussians deliver a smoother trajectory than using the max method, see Fig. 3.8 (b). However, close to the grating, a single Gaussian does not fit the intensity distribution well (see Fig. 3.4 for $z = 7 \mu m$). Again, this leads to underestimating the emission angle in the fit.

The best results with the lowest fit errors are obtained when tracking the weighted average position of the intensity for each slice, as described in Sec. 3.1.1. The corresponding data and fit is shown in Fig. 3.8 (c), with the emission angle resulting in $\theta_{\text{meas}} = 31.475 \pm 0.023^{\circ}$. The uncertainty given is solely from the fit of the slope. Additional sources of error are introduced by possible tilts within the imaging system relative to the chip and the near-field effects close to the chip surface. The result is slightly below the intended emission angle of $\theta_{\text{sim}} = 32^{\circ}$ in this particular example. However, the direction and scale of the observed deviation can be explained by the fact that we are using a laser diode with a center frequency 2-3 nm smaller than the 732 nm the grating was designed for.

3.1.4 Beam Waist

Another useful metric is the evolution of the beam waist. We differentiate between the waist in the xz-plane (cyan arrow in Fig. 1.2, denoted with x') and the waist along y (red arrow in Fig. 1.2). The first is governed by the period and duty cycle of the grating and the second depends on the curved extrusion shape of the 2D grating pattern (see Sec. 1).

To extract the beam waists, we again consider each z-slice independently. The goal is to fit 1D Gaussians along both axes of one image. In analogy to Sec. 3.1.2, we perform the analysis both for the 1D slices of the image through the pixel with maximum intensity, i.e., fitting Gaussians to both line plots displayed in Fig. 3.4 for each height, as well as for the images contracted (integrated) along one of the axes. The waist is defined as twice the distance from the peak position to the point where the intensity has fallen to $1/e^2$ of its peak value.

The results are shown in Fig. 3.9. The plot omits heights below $z = 32 \,\mu\text{m}$, because the Gaussian fit cannot describe the intensity distribution along x for these images (as described in Sec. 3.1.3).



Figure 3.7: Same plot as in 3.4, but with measurements of a different grating. The line plots are indicate the clipped intensity profile.



Figure 3.8: Plots of height vs x-position of the intensity. The latter is extracted via (a) the position of the pixel with maximum intensity value, (b) the peak position of a 1D Gaussian fit to the z-slice integrated along the y-axis, and (c) the weighted average of the intensity. The subtitle states the emission angle (relative to the normal) and its error determined by the slope of the linear fit.

Note that the waists shown are in the beam frame with axes perpendicular to the propagation direction, i.e., the waists along x have been corrected by taking the emission angle into account. One effect of this correction can be seen when comparing Fig. 3.9 to Fig. 3.4. At $z = 55 \,\mu\text{m}$ in Fig. 3.4, which shows a slice perpendicular to the z-axis, the two waists appear to be similar. However, after moving to the beam frame, the longitudinal waist (blue curves in Fig. 3.9) is in fact smaller than the transverse one (red) at $z = 55 \,\mu\text{m}$.

Fig. 3.9 reveals that the beam has a tighter focus along the longitudinal $(w_{x',\min} = 2.27 \,\mu\text{m})$ compared to the transverse direction $(w_{y,\min} = 2.82 \,\mu\text{m})$. Additionally, the two focuses do not meet at the same height $(z_x = 50 \,\mu\text{m} \text{ vs } z_y = 57 \,\mu\text{m})$. These findings match the simulated values for this grating.

3.1.5 Other Examples

The tools described in Sec. 3.1.2 can also be used to visualize beams that do not result in a single, focusing emission.

Figure 3.10 shows the emission of a forward-emitting grating designed for UV-light. The goal of this particular grating design is to emit two separate beams at different angles from a single grating. This can be used to photo-ionize ${}^{40}\text{Ca}^+$ ions in a common two-step process using one 423 nm laser to excite the S_0 to P_1 dipole transition and one laser below 389 nm (in our case 375 nm) to lift the excitation to the continuum [11]. With this grating, the two required wavelengths can



Figure 3.9: Beam waists along both axes vs height above grating. Waists are extracted from 1D Gaussian fits to the of the z-slice images both integrated (solid lines) and sliced through the pixel with maximum intensity (dashed lines). The waists are given in the beam frame. The subtitle shows the minimum values of the waists, as well as the height at which they occur.

be delivered via the same waveguide. This is achieved by designing the two-fold emission of the grating such that one of the beams emitted at 423 nm overlaps with the other beam emitted at 375 nm at the height of the ion. The measurements shown in 3.10 confirm this expected behavior and yield information about the exact emission angles of the two beams for the two different wavelengths.



Figure 3.10: Integrated side view pictures analogous to Fig. 3.5 (a), but for a forward-emitting grating intended for photo-ionization. The measurements have been conducted with (a) 375 nm light and (b) 423 nm light.

Another interesting example is the emission of Laguerre-Gaussian (LG) modes. These cylindrical modes, also called "optical vortex" modes, feature a phase that is spiraling about the beam axis.

Thus, the interference leads to an intensity null along the axis. This intensity profile, e.g., in a configuration where two such beams cross, could be used to optically trap ions [12]. Fig. 3.11 shows the sideview of two such beams, focusing (a) and non-focusing (b). Note that it is necessary to use the slice in the xz-plane, as in Fig. 3.5 (b). The picture integrated along y loses the information about the tubular intensity profile.



Figure 3.11: Sliced side view pictures analogous to Fig. 3.5 (b), but for gratings designed to emit an (a) focusing and (b) non-focusing Laguerre-Gaussian mode. (a) can be compared to simulation results of a similar grating in Fig. 12 (a) of [4]. The staircase structures in the intensity profile in (b) are imaging artifacts from the beam splitter in the camera objective, as described in the caption of Fig. 2.3.

3.2 Waveguide Loss Measurement

The propagation loss in a photonic waveguide is influenced by two contributions: Absorption in the material and sidewall scattering [2]. The first is intrinsic to the materials and can only be reduced by limiting the field's presence in lossy material, particularly the Si substrate since both waveguide materials as well as the silica cladding exhibit negligible loss for our wavelengths. In our case, one can be use wider waveguides with better mode confinement. Sidewall roughness, on the other hand, can be improved with better lithography techniques in fabrication. We compare two methods for extracting loss in a straight waveguide.

3.2.1 Waveguide Loss Using Grating Outcouplers

The first method involves grating outcouplers in a chip design shown in Fig. 2.4. Each input waveguide coming from the left is split into two paths using a 50:50 integrated beam splitter: One path goes directly to a forward emitting grating, while the other path involves four bends, as well as a longer section of straight waveguide before being outcoupled by an identical grating. This structure is repeated four times, but with varying length of the straight waveguide section. In the following, each of the four sections is referred to as "waveguide", with waveguide 1 being the one with the shortest straight length (topmost in Fig. 2.4). Grating 1 and 2 describe the grating without bends (short) and with bends (long), respectively.

Data Acquisition

This setup allows for a differential measurement: Light from a 532 nm diode laser is coupled into a fiber and, from there, into one of the on-chip waveguides. Then, the camera is moved to grating 1 of that waveguide, imaging the emission for several different heights above the grating. We do the same with the camera moved to grating 2. Ideally, the sum of the intensity values of all pixels in an image is proportional to the power of the imaged beam. The ratios of these two powers, measured for the different waveguide lengths, provide a way to decouple the straight waveguide loss from the other losses present in the system, see Eq. (3.2).

Challenges

While straightforward in theory, the method poses some challenges in practice:

• Noise floor: As seen in Fig. 3.12 (b), the intensity does not reach the value 0 (out of 255), even when far away from the peak of the beam. Thus, integrating over a part of the image without beam intensity will still yield a non-zero value. To remove this noise floor, the background is examined in a corner of the image where there should not be any intensity from the grating emission (see green box in Fig. 3.12 (a)). In this box, the average intensity value of a pixel is calculated. This value ranges from 2.5 to 3.5 for different images, and, as a first post-processing step, is subtracted from each pixel in the image. However, this is not sufficient to make image areas far away from the beam integrate to zero. In Fig. 3.12 (b), one can also see a fluctuating noise component on top of a noise floor. Additionally, Fig. 3.12 (a) implies that the noise floor is not constant over the whole image, as the background appears brighter on the left-hand side of the image. Both effects can be combated by considering that the intensity can only take integer values, so the noise usually manifests itself in discrete jumps among only three intensity levels around the mean noise floor. Comparing different excerpts of the transverse line plot as in Fig. 3.13 shows that towards the right-hand side of the image (Fig. 3.13 (b)), the noise integrates to a value close to zero, as expected. However, when moving further to the left (Fig. 3.13 (b)), the rate of intensities above the mean increases, such that the integral over an equal width of the line plot increases by a factor of ten and is no longer negligible.

The simple remedy employed here is to set all pixels with one of those three distinct intensity values to exactly zero. This enforces that integrals over areas without intensity are zero,

without losing much information. In the future, when the camera is switched to an operation mode with 16-bit intensity values, it might be necessary to devise a more involved method to remove the noise floor. One such method could involve fitting a plane or a higher-order function to the noise floor and subsequently subtracting the fit.



Figure 3.12: (a) shows the image produced by the camera at a height of $z = 52 \,\mu\text{m}$ above a grating, before any post-processing. To better see the effects of noise with low intensity, the colorbar is scaled logarithmically. (b) illustrates the intensity along the row and column of the pixel with maximum intensity in the image.



Figure 3.13: Excerpts from the transverse line plot in Fig. 3.12 (b), but after subtracting the average background intensity value from each pixel. (a) shows a 200 pixel wide excerpt from the left hand side of the image and (b) an equally wide excerpt from the right hand side, closer to where the average background intensity was extracted. The subtile states the sum of all values in the respective excerpt, quantifying the increase in noise floor towards the left side of the frame.

• Second order emission: The forward-emitting gratings used in this design also feature a second order emission which is visible in Fig. 3.12 (a) below the main intensity peak, and also in Fig. 3.12 (b) by the side lobe of the longitudinal line plot at around 950 pixels. As we typically want to track only the emission into the first order, we cannot integrate the intensity over the whole image. Instead, the integration is performed only within a region of interest (ROI) around the peak intensity, shown as a black rectangle in Fig. 3.12 (a). Since the emission moves for pictures taken at different heights, the ROI is redefined for every image, tracking the movement of the peak intensity. Additionally, we consider only

pictures taken above a certain height that provides sufficient spatial separation between first and second order emission in the z-slice.

• Dynamic range: While it is clear that, to extract a proxy for beam power from the imaged emission, the intensity values must not saturate (see clipping in Fig. 3.7), one also needs to be careful about faint emission on the other hand. With the settings used in these measurements, the camera only provides 8-bit intensity values. Thus, when the maximum value drops below $\sim 30/255$, the resulting discretization of the emission shape starts to obscure the distribution and therefore influences the integrated intensity. Additionally, the zeroing of the noise described above cuts off a significant portion of the tail in this case.

Therefore, the exposure of the camera needs to be set for each waveguide in Fig. 2.4 such that the emission of grating 1 (short and no bends) does not overexpose the camera while at the same time maintaining sufficient intensities for grating 2 (long and with bends). Unfortunately, such a setting could only be achieved for waveguides one and two in Fig. 2.4.

• Sequential imaging: With the current chip design, the spatial separation of the two gratings for each waveguide prevents them from fitting into the frame of the camera simultaneously. We therefore have to move the camera from one grating to the other and do two measurements sequentially. This makes the power ratio susceptible to drifts in laser power and other changes in the system that occur when moving the microscope to the new location. For future designs, placing the gratings closer together would allow simultaneous measurements of the two gratings. The two emissions in one frame can be separated with different ROIs.



Figure 3.14: Same plot as in Fig. 3.12, but after post-processing for a grating that receives less power, and for a lower height. The line plot in (b) reveals that the discretization for this low peak power obscures the true shape of the emission.

Results

The difference in straight waveguide length before grating 1 and grating 2 for each of the four waveguide structures in Fig. 2.4 is given in Tab. 3.1.

	WG 1	WG 2	WG 3	WG 4
Straight section length [cm]	0.48	0.68	0.88	1.08

Table 3.1: Differences in straight waveguide length before grating 1 and grating 2 of the four waveguide structures in Fig. 2.4, from top to bottom.

The power emitted by grating 1 is given by

$$P_{\text{out},1} = \alpha_{\text{incoup}} \cdot \frac{1}{2} \cdot \alpha_{\text{grat}} \cdot P_{\text{in}},$$

where α_{incoup} contains the incoupling loss and the loss in the first straight section before grating 1, which stays constant for all waveguides. α_{grat} describes the loss when outcoupling the light from the on-chip waveguide into free-space via the grating. The factor 1/2 enters because of the 50:50 beam splitter, which we assume to be ideal for the purposes of this measurement [13].

On the other hand, the power emitted by grating 2 adds two sources of loss, the bend losses α_{bend} and the losses in the extra section of straight waveguide. The latter depends on the straight waveguide length l via an exponential law, characterized by the decay constant α_{wg} . Thus, the output power from grating 2 is given by

$$P_{\text{out},2}(l) = \alpha_{\text{incoup}} \cdot \frac{1}{2} \cdot \alpha_{\text{bend}} \cdot \exp(-\alpha_{\text{wg}} \cdot l) \cdot \alpha_{\text{grat}} \cdot P_{\text{in}}.$$
(3.1)

Assuming that the input power stays constant during the imaging of grating 1 and 2, the fraction of the two power values cancels all the terms except bend and waveguide loss:

$$\frac{P_{\text{out},2}(l)}{P_{\text{out},1}} = \alpha_{\text{bend}} \cdot \exp(-\alpha_{\text{wg}} \cdot l).$$

When this power ratio is expressed in dB, it takes a simple linear form:

$$\frac{P_{\text{out},2}(l)}{P_{\text{out},1}} \text{ [dB]} = 10 \log_{10} \left(\frac{P_{\text{out},2}(l)}{P_{\text{out},1}} \right) = 10 \log_{10}(\alpha_{\text{bend}}) - \frac{\alpha_{\text{wg}}}{\ln 10} \cdot l$$
$$= -\alpha_{\text{wg,dB/cm}} \cdot l \text{ [cm]} - \alpha_{\text{bend,dB}}.$$
(3.2)

Therefore, if we fit a line to the power ratios obtained from each of the waveguides and expressed in dB, we can extract the straight waveguide loss $\alpha_{wg,dB/cm}$ in dB/cm, as well as the total bend loss $\alpha_{bend,dB}$ in dB. The differential measurement decouples these values of interest from the other losses present in the system.

The fact that, due to the limited dynamic range of the camera in its 8-bit mode, only two out of the four available waveguide structures provide usable results, severely limits the analysis. Measurements have been taken on three different days, with different parameters. The results are summarized in Tab. 3.2.

1. In the first measurement, ten images of each grating emission were taken, in height increments of $10 \,\mu\text{m}$ from $z = 0 \,\mu\text{m}$ to $z = 100 \,\mu\text{m}$. As long as no part of the beam leaves the image, the total power should not change with respect to the height at which the picture was taken. The ten different measurements only provide statistics for the power measurement, i.e., its mean and standard deviation. Because there were only ten images per grating in this run, the ROI was set to cover the whole image, including the second order emission. As long as this is done in the same way for all gratings and the emission never leaves the image, the power ratios should not change compared to only using the first order emission.

Fig. 3.15 (a) shows the fit of Eq. (3.2) to the data. The error bars for each data point are obtained by propagating the standard deviations of the two individual power measurements through the conversion formula to dB. They help to evaluate the quality of the measurement but do not influence the fit in this case, because the two data points can be fitted by a line perfectly.

In this measurement, the laser was attenuated with a neutral density (ND) filter with an optical density (OD) of 3.0, and the exposure times were 2 ms for waveguide 1 and 5 ms for waveguide 2. This setting was too low as almost all images were underexposed, with max intensities below 30/255, leading to strong discretization artifact (compare to Fig. 3.14).

- 2. The second measurement aimed to resolve this issue by using a much larger exposure time of 50 ms for waveguide 1 and 80 ms for waveguide 2, however, introducing a larger laser attenuation with an ND filter with an OD of 4.0. Additionally, we now gathered the usual sweep of 100 images at increments of 1 μ m. This yields more statistics and allows to disregard all overexposed and underexposed images. I defined an image as underexposed if its max pixel intensity is below 35/255. Further, I disregard all images where the beam (partially) leaves the frame, and all images where first and second order emission are too close together and cannot be separated with a ROI. After these steps, there are 30-40 usable images out of the original 100 left for the analysis. The results are plotted in Fig. 3.15 (b).
- 3. A third measurement was performed with an ND filter with an OD of 1.0, and an exposure time of 0.5 ms for both waveguides 1 and 2. This measurement includes many overexposed images. All images with one of the issues described above where removed. The result of this measurement run is shown in Fig. 3.15 (c).

	Measurement 1	Measurement 2	Measurement 3
Straight Waveguide Loss $\alpha_{wg,dB/cm}$ [dB/cm]	2.33	4.12	7.25
Total Bend Loss $\alpha_{\text{bend,dB}}$ [dB]	4.18	2.22	0.13

Table 3.2: Summary of waveguide and bend losses of 532 nm light in alumina, extracted from power measurements using images of the grating emission.



Figure 3.15: Data obtained from the three loss measurements that were performed. The fit parameters are shown in the subtitle. The slope and intercept resemble $\alpha_{wg,dB/cm}$ and $\alpha_{bend,dB}$, respectively, see Eq. (3.2).

Discussion

The results are inconsistent in between different measurements and vary across a 5 dB-range. Unfortunately, the existence of only two data points makes it hard to discern measurement fluctuations from a general trend. Measurement 3 in particular appears to be faulty since the extrapolation to an offset length of 0 cm, i.e., the y-intercept, should reveal the bend losses. It would suppose very low bend losses of only 0.13 dB. The waveguide losses implied by measurement 1 and 2, on the other hand, are comparable but larger than values in literature: [6] found a loss of roughly $1.5 \,dB/cm$ for 542 nm light, and [14] even suggests a loss below $0.8 \,dB/cm$ for wavelengths larger than 458 nm.

More measurements need to be taken to achieve reproducible results. Recent modifications of the setup will improve future measurements:

• The beam splitter in the imaging system, which leads to the diagonal fringes in all images, can now be removed prior to the measurement. This will lead to a smoother shape of the intensity distribution that is easier to work with.

- Recently it was discovered that the camera supports 16-bit intensity values instead of only 8-bit, which will lift the resolution in intensity from 255 to 65536 discrete values per pixel, dramatically improving the dynamic range. This should allow to gather data from more than only two waveguides, which is the main limitation of the measurement described here. In addition, it is likely that the discretization that is currently present in the measurements of grating 2 underestimates the real power. This leads to higher loss values and could therefore be an explanation for the discrepancy compared to the literature.
- Initially, we had assumed that the TM-mode is very lossy, such that there is only an insignificant fraction of TM which reaches grating 1 in each of the waveguides. The loss measurement would then be restricted to only TE. However, recent insights regarding the waveguide input taper and dicing depth have revealed that TM can in fact be present in the waveguides. This could explain the difference in loss values between measurements: Any movement or strain in the (non-polarization-maintaining) input fiber can change the output polarization which, since TE and TM generally have different losses, would change the measured waveguide loss. We now added polarization control, allowing to measure TE and TM loss independently.

3.2.2 Waveguide Loss Using Transmission

The second method for determining the waveguide loss uses the paperclip structures shown in Fig. 2.5. There are two main differences compared to the chip used for the first method: Firstly, there is no beam splitter involved, so one directly measures the output power without converting it to a power ratio first. Secondly, the beam is not outcoupled via a grating towards the top of the chip, but uses regular tapered waveguides which exit on the side of the chip. This light is then collected with a fiber which is connected to a power meter.

Since we now need two fibers instead of only one, we place the chip into the setup in Fig. 2.6. The power received on the other side of the chip takes a similar form to Eq. (3.1):

$$P_{\text{out}}(l) = \alpha_{\text{incoup}} \cdot \alpha_{\text{bend}} \cdot \exp(-\alpha_{\text{wg}} \cdot l) \cdot \alpha_{\text{outcoup}} \cdot P_{\text{in}},$$

where α_{incoup} and $\alpha_{outcoup}$ are the losses for in- and outcoupling of the light. The main downside of this method compared to the one described above is the lack of a differential measurement, i.e., there is no second measurement to cancel out the coupling losses. For this reason, the transmitted power in dBm (now in absolute units because we do not use ratios) is given by

$$P_{\text{out}}(l)[\text{dBm}] = 10 \log_{10} \left(\frac{P_{\text{out}}(l)}{1 \text{ mW}} \right)$$
$$= -\alpha_{\text{wg,dB/cm}} \cdot l \text{ [cm]} + \alpha_{\text{incoup,dBm}} + \alpha_{\text{bend,dBm}} + \alpha_{\text{outcoup,dBm}} + P_{\text{in,dBm}}.$$
(3.3)

Compared to Eq. 3.2 this formula includes many additional terms that are not, unlike α_{bend} , constant by design. Thus, to extract only the straight waveguide loss $\alpha_{\text{wg,dB/cm}}$ from the slope of the linear fit, one needs to make sure that the in- and outcoupling efficiency, as well as the input power stay constant for all measured waveguides.

Data Acquisition

The incoupling fiber features manual 6-axis control, while the outcoupling fiber is mounted to a 3axis translation stage, with an additional stage to adjust pitch and yaw. However, the latter proved too coarse to be useful for maximizing outcoupling efficiency. A typical measurement procedure is described in Sec. 2.2.

The main challenge is to ensure a constant in- and outcoupling for all five transmission measurements. This coupling is adjusted manually, attempting to optimize the light received on the power meter. This is difficult to do reliably: Firstly, there are nine parameters to manually tune, six and three axes for the in- and outcoupling fiber, respectively. Secondly, the coupling only gets better the closer the fiber is to the chip, so the best coupling is achieved when chip and fiber touch. However, in the current setup the chip will move when bumped with the fiber, which negates all prior alignment efforts. Considering that the digital microscope used to monitor the process does not offer a sufficient resolution to examine the chip-fiber distance, it is hard to keep this distance constant at a good coupling between the measurements, without bumping the chip during the alignment.

Results and Discussion

The length of the straight waveguide section for the five waveguide structures in Fig. 2.5 are given in Tab. 3.3. The measured powers vs straight waveguide lengths, together with an exponential fit, are shown in Fig. 3.16. Since the transmitted powers were difficult to reproduce after subsequent out- and incoupling into the same waveguide, the measurements were repeated two or three times for some waveguides. The slope of the fit suggests a straight waveguide loss of

 $\alpha_{\rm wg,dB/cm} = 1.6 \pm 1.0 \, \rm dB/cm.$

While the fit parameter does lie close to the value of $1.5 \,\mathrm{dB/cm}$ for a similar wavelength stated in [6], the large standard deviation of $1 \,\mathrm{dB/cm}$ due the inconsistent measurements, and therefore poor fit quality, prevents any conclusive interpretation.

Improvements to the setup are being implemented at the time of writing. Most notably, an imminent modification will allow the chip to be held in place by vacuum suction. This way, the chip can be bumped with the fiber without moving it, making the coupling easier to optimize and removing a frustrating issue of the current measurement procedure. Additionally, the polarization control described in Sec. 3.2.1 will also improve the measurement, as some of the fluctuations might stem from differing TE/TM fractions of the incoupled light in between measurements.

Table 3.3: Differences in straight waveguide length for the paperclip structures in Fig. 2.5, from top to bottom.



Figure 3.16: Data from the transmission measurement. The fitted slope, together with its standard deviation, is shown in the subtitle. Ideally, it reflects the waveguide loss $\alpha_{wg,dB/cm}$, see Eq. (3.3).

Chapter 4 Conclusion and Outlook

This semester project consisted of two tasks. First, I implemented tools to characterize beams emitted by grating outcouplers that were fabricated in both silicon nitride and alumina, using images taken with a custom beam profiling setup. We used wavelength of 729, 532, 423, and 375 nm. The aim is to extract important beam parameters from the raw image data, such as the emission angle and the waist size, as well as different views that facilitate the comparison with simulation data. The measured emission angles and beam waist evolutions are in good agreement with the ones that went into the design of the gratings. Minor deviations can be explained by the different laser wavelengths we use in the experiment. The purpose of different side views is illustrated by measurements of less conventional beams: two Laguerre-Gaussian beams (focusing and non-focusing) and a design with two emissions at different angles from a single grating, which is useful for the photo-ionization of calcium ions.

Secondly, I characterized the loss in straight waveguide sections of alumina with two methods. One features a differential measurement using grating outcouplers, where camera images serve as a power meter. To validate the results, the other method uses absolute transmission measurements without grating outcouplers. For this, I built a new setup that allows for in- and outcoupling on two sides of the chip using fibers and alignment stages. The loss measurements range from 1.6 to 7.25 dB/cm.

This discrepancy and the large uncertainties present in the second method guide the improvements that are made to the setups at the time of writing. The main issues are the 8-bit intensity values of the camera leading to discretization artifacts, and the lack of polarization control for the input light. Both are straightforward to fix, by toggling the camera to save 16-bit pictures and adding a fiber polarization controller to the incoupling fiber. These main improvements are accompanied by removing the beam splitter from the camera imaging path in the first setup (removing unwanted fringes in the images), and adding a suction holder for the chip in the transmission setup. Using the same measurement techniques and analysis procedures described in this work, these improvements should lead to much more consistent results. With this enhanced precision, one could possibly use the grating outcoupler method to compare losses between gratings with and without ITO layer. Additionally, simulations and recent designs for waveguide bends and integrated power splitters can be validated.

Bibliography

- K.K. Mehta, C. Bruzewicz, and R. McConnell et al. Integrated optical addressing of an ion qubit. *Nature Nanotech*, 11:1066–1070, 2016.
- [2] K.K. Mehta, C. Zhang, and M. Malinowski et al. Integrated optical multi-ion quantum logic. *Nature*, 586:533–537, 2020.
- [3] Niffenegger R.J., Stuart J., and Sorace-Agaskar C. et al. Integrated multi-wavelength control of an ion qubit. *Nature*, 586:538–542, 2010.
- [4] Gillenhaal J. Beck, Jonathan P. Home, and Karan K. Mehta. Grating design methodology for tailored free-space beam-forming, 2023.
- [5] Carmelo Mordini, Alfredo Ricci Vasquez, Yuto Motohashi, Mose Muller, Maciej Malinowski, Chi Zhang, Karan K. Mehta, Daniel Kienzler, and Jonathan P. Home. Multi-zone trapped-ion qubit control in an integrated photonics qccd device, 2024.
- [6] Mustafa M. Aslan, Nathan A. Webster, Courtney L. Byard, Marcelo B. Pereira, Colin M. Hayes, Rodrigo S. Wiederkehr, and Sergio B. Mendes. Low-loss optical waveguides for the near ultra-violet and visible spectral regions with al2o3 thin films from atomic layer deposition. *Thin Solid Films*, 518(17):4935–4940, 2010.
- [7] Sophie Cavallini (Semester thesis at ETH Zurich). Photonics building blocks in alumina, 2024.
- [8] Mark T. Wade, Fabio Pavanello, Rajesh Kumar, Cale M. Gentry, Amir Atabaki, Rajeev Ram, Vladimir Stojanovic, and Milos A. Popovic. 75In 2015 IEEE Optical Interconnects Conference (OI), pages 46–47, 2015.
- [9] ImageJ analysis software, URL: https://imagej.net/ij/index.html.
- [10] Ansys Lumerical software, URL: https://www.ansys.com/products/optics.
- [11] H. Shao et al. Laser ablation and two-step photo-ionization for the generation of 40ca+. J. Phys. Commun., 2, 2018.
- [12] Yijie Shen, Xuejiao Wang, and Xiaocong Yuan et al. Optical vortices 30 years on: Oam manipulation from topological charge to multiple singularities. *Light: Science & Applications*, 8, 2019.
- [13] Joonhyuk Kwon, William J. Setzer, Michael Gehl, Nicholas Karl, Jay Van Der Wall, Ryan Law, Daniel Stick, and Hayden J. McGuinness. Multi-site integrated optical addressing of trapped ions, 2023.
- [14] Cheryl Sorace-Agaskar, Suraj Bramhavar, Dave Kharas, William Loh, Paul W. Juodawlkis, John Chiaverini, and Jeremy M. Sage. Multi-layer integrated photonics from the ultraviolet to the infrared. In Amos Danielli, Benjamin L. Miller, and Sharon M. Weiss, editors, Frontiers in Biological Detection: From Nanosensors to Systems X, volume 10510, page 105100D. International Society for Optics and Photonics, SPIE, 2018.