6.2540: Report on Silicon Photonics at MIT.nano

Dimitar Dimitrov[1,2], Kenneth Chap[1], Sabrina Corsetti[1], Varun Hariprasad[1], and Xudong (Sheldon) Zheng[1]

[1]Experimentation
[2]Report

December, 2025

1 Introduction

Silicon photonics is the study of creating photonic circuits in silicon. A relatively new field, it is revolutionizing photonics by enabling an orders-of-magnitude decrease in size for many systems – including ones which would have been prohibitively large thus far. At the same time, it is easy to reduce to practice, because it leverages silicon, a material which has been studied more extensively than any other. In particular, because of the ubiquity of complementary metal–oxide–semiconductor (CMOS) technology in modern electronic systems, processes used to create 2-dimensional devices based on silicon are extremely mature, with new generations being developed perpetually. Silicon photonics benefits from these technologies in the forms of small feature sizes, fast development cycles and massive scaling. This has enabled a myriad of applications. Early adopters were optical communications systems and light detection and ranging (LiDAR), which provided commercial motivation for further development of silicon photonics. Today, applications include quantum computing, augmented reality, optogenetics and ultra-low-power computing.

Each of the authors was interested in exploring different applications from this list, but all shared a desire to leverage silicon photonics in the lab. This report serves as a semi-formal review of our progress towards developing a platform for silicon photonics experimentation at MIT.nano – what went well, what didn't go well, and what we would do next. Our hope is that it can be used as a starting point for further development by others in the future.

Section 2 provides motivation for the devices explored by briefly discussion photonic computing. Section 3 describes how we designed our devices. Section 4 describes how we fabricated samples with our devices. Section 5 describes how we tested our samples. Section 6 provides guidance on future work. Section 7 concludes this report.

2 Application Highlight: Computing

In lieu of surveying the diverse landscape of technologies leveraging silicon photonics, we focus on one particularly promising and important application: computing.

2.1 The Case for Photonic Computers

Over the past few years, a tremendous gap between software demand and hardware capability has emerged. The voraciousness with which new large language models are being developed and the breadth of the fields they are applied to has largely been met with incremental improvements to a technology originally designed for something entirely different (worse

yet, this technology has constrained the way in which we develop these algorithms). This gap has come at various costs unrelated to the development of the technology itself, not the least of which is huge energy waste impacting both the environment and the local communities which are becoming homes to new data centers seemingly every week. There is no reason to believe that compute demands will drastically decrease any time soon. So, it will be our computers that eventually improve to close this gap.

There are currently a half dozen candidates for how these new, ultra-low-power computers will work. Photonic computing is a compelling one, because light and silicon photonics already play a critical role in modern systems, as the interconnect between compute nodes. Huge strides have been made towards making this technology a reality [1, 2].

2.2 Principle of Operation

The computation of current machine learning algorithms consists predominantly of matrix multiplication (this is not true of certain neuromorphic approaches, but it is true of virtually all "AI" algorithms used commercially). While much of the energy required for current approaches is spent on data movement, which is primarily where photonic computing has an advantage, any alternative approach needs to be able to perform matrix multiplication efficiently.

The basis for this capability in photonic computing is derived from [3]. In this 1994 paper, the authors show that any discrete unitary operator can be realized using a network of Mach–Zehnder interferometers (MZIs; two 2x2 couplers in series). They prove this by showing that a single MZI can realize any 2x2 unitary operator. The two degrees of freedom required can be implemented in a few ways. In this paper, they are realized as a phase shift on one of the arms of the MZI and a phase shift on one of the outputs. Note that the topology presented in this paper (the so-called Reck topology) has since been improved upon [4]. Critically also, [5] shows that it is possible to construct network of MZIs that function as intended out of imperfect splitters.

2.3 Recent Work

More recently, this approach has been used for machine learning directly. [6] break down matrices in deep learning architectures using single value decomposition (which has to happen only once, and is thus inexpensive), and then uses MZI networks to compute the unitary sub-matrices, thereby demonstrating optical artificial neural networks inference. Newer works building on this concept include [7] and [8] [1].

Common to all of these approaches to photonic computing is the MZI building block. As we discuss in Section 3, the MZI itself is often composed of an even more ubiquitous building block in silicon photonics: the multi-mode interferometers (MMI). The application of photonic computing, along with many others, served as motivation to design and fabricate these building blocks in the lab.

3 Design

3.1 Wavelength Selection

We chose to work with a target wavelength of 1550 nm (in the infrared spectrum). Due to its low absorption in glass, this wavelength (along with 1310 nm) is extremely popular in the field of communications. Because of this, it is easy to find test equipment for this wavelength. Silicon is also transparent at 1500 nm, making it ideal for silicon photonics applications. This wavelength has thus become a de facto standard, and is a natural starting point for this work.

3.2 Substrate Selection

We picked silicon as the material out of which to construct devices as a primitive to the investigation (other common options are Si_3N_4 and InP). Photonic devices require cladding made of a material with a lower refractive index. A natural selection when working with silicon is silicon dioxide, which is not

[1] [6] shows that in addition to matrix multiplication, photonic circuits can be used for activation functions. Progress has been made in this direction as well, e.g. [9]. This is however not the focus of this report.

only a common insulator in CMOS, but forms naturally on silicon (in thin layers).

We chose to use a silicon-on-insulator wafer as our starting substrate. This type of wafer consists of three layers: silicon, silicon dioxide, and silicon again. The first silicon layer is the thickest layer, on the order of hundreds of microns, and provides structural stability. The "buried" silicon dioxide layer (the BOX) is generally a few microns thick. Finally, the top silicon layer, the "device" layer, is generally on the order of 100 nm. SOI wafers are used as a precursor to some metal–oxide–semiconductor field-effect transistor (MOSFET) fabrication processes and are common for micro-electromechanical systems (MEMS) applications. They have also become a popular choice for silicon photonics, because the device layer is in the correct thickness range and the BOX provides cladding on one side.

We reached out to a number of suppliers and found University Wafer (`https://universitywafer.com/`) to be the most affordable and convenient. At the time of writing, this supplier offers 6 options for ordering SOI wafers, in sizes ranging from 25 mm × 25 mm samples to 200 mm-diameter wafers and small-quantity unit prices ranging from $164.28 to $1,127.34. The two sample-sized options differ in the thickness of the device layer. The cheaper of the two has a 220 µm device layer, which is another de facto standard for silicon photonics applications. The other sample has a 70 nm device layer. A thinner device layer allows for larger features laterally, which in turn translates to more flexibility in the lab. The price difference is small, so this was the device we chose (University Wafer ID 4015). The total thickness of this sample is 725 µm and the BOX thickness is 2 µm.

Unfortunately, all samples offered are doped with Boron. An ideal sample for silicon photonics work would be intrinsically doped. Doping isn't specified, but the resistivity is reported as between 1 Ω cm and 20 Ω cm, which means that the doping is between 7×10^{14} cm^{-3} and 1.7×10^{16} cm^{-3} [10]. We were not able to find refraction index models for these doping levels, but extrapolation suggested that the refraction index would be similar to that of silicon [11].

3.3 Waveguides

Having selected our materials, we moved on to designing individual devices, starting with waveguides. The waveguides should only support one mode, but be as large as possible to minimize insertion loss. We used Ansys Lumerical to simulate the behavior of 1500 nm light in our material, and found that a 750 nm-wide waveguide was just small enough to support a single mode.

3.4 Couplers

Next, we designed couplers by simulating the overlap with a standard multi-mode optical fiber with 125 nm of cladding. We found that 250 nm was a good coupler width. We used a taper between the coupler and the waveguide of length 500 µm.

3.5 Antennas

We designed antennas as perturbed waveguides (equivalent to diffraction gratings). To aid in testing, we also designed larger antennas, which taper out of the waveguide before the perturbations begin.

3.6 1×2 MMI

Having designed the basic input and output systems for our chip, we moved up a level of complexity by designing MMIs. An MMI functions on the principle of interference, similarly to an interferometer, but acts as an $M \times N$ coupler, not an interferometer. It is an important building block in all the applications listed in Section 1. The principles of operation are described in detail in [12], which we used to design our MMIs.

The first MMI we designed was a 1×2 MMI ($M = 1$, $N = 2$). For this design, we took advantage of symmetric interference, as described in section V.B. in the paper. The requirement for this type of MMI is that the field excitation coefficients $c_\nu = 0$ for modes $\nu = 1, 3, 5, ...$ This is not a problem in our case, since our waveguides support only 1 mode ($c_\nu = 0$ for $\nu \neq 0$).

The MMI is a rectangle of width W_M and length L. The paper suggest choosing W_M such that at least

$N+1 = 3$ mode are supported. We opted to support 6 modes instead, and chose $W_M = 5\,\mu\text{m}$ using Lumerical for simulation. For MMIs taking advantage of symmetric interference, L is given by Equation (37) in the paper:

$$L = \frac{p}{N}\left(\frac{3L_\pi}{4}\right) \tag{1}$$

where p is any integer (we use $p = 1$ to minimize size) and L_π is the beat length of the two lowest-order modes, given by Equation (6) in the paper:

$$L_\pi = \frac{\pi}{\beta_0 - \beta_1} \tag{2}$$

Using Lumerical, we found the propagation constants β_ν of the first two modes to be:

$$\beta_0 = 7.7620 \times 10^6\,\text{m}^{-1} \tag{3}$$

$$\beta_1 = 7.6886 \times 10^6\,\text{m}^{-1} \tag{4}$$

(derived from their effective indices of refraction and the target wavelength $\lambda_0 = 1550\,\text{nm}$). Plugging in, that gives:

$$L = 16.0566\,\mu\text{m} \tag{5}$$

Next, we derived where the output signals will appear. For MMIs taking advantage of symmetric interference, they appear "symmetrically located along the y-axis with equal spacings W_e/N". W_e is given by Equation (4) in the paper:

$$W_{ev} \approx W_e = W_M + \left(\frac{\lambda_0}{\pi}\right)\left(\frac{n_c}{n_r}\right)^{2\sigma}\left(n_r^2 - n_c^2\right)^{-(1/2)} \tag{6}$$

$\sigma = 0$ for TE modes, which is how we will polarize our light ($\sigma = 1$ for TM modes). n_r and n_c are the effective refractive indices of the device ("ridge") and cladding, respectively. Effective modes are required because this analysis is for a 2-dimensional structure, while our actual device is 3-dimensional (it has a height of 70 nm). This is called the effective index method. Using Lumerical, we have:

$$n_r = 1.921 \tag{7}$$

$$n_c = 1.444 \tag{8}$$

Plugging in, we have:

$$W_e = 5.3894\,\mu\text{m} \tag{9}$$

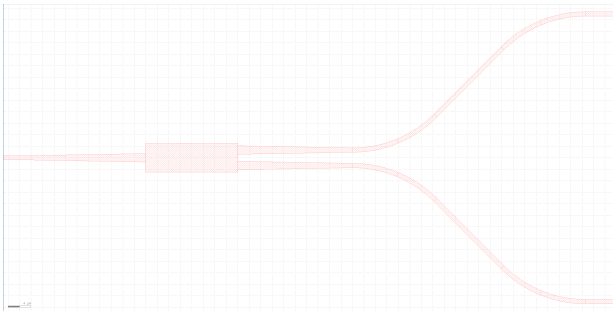

Figure 1: Layout of 1×2 MMI

meaning that the two outputs are centered at $\pm 1.3474\,\mu\text{m}$ relative to the center.

Next, we added light tapers to the inputs an outputs [13].

Finally, we spread the two outputs out to reduce cross-talk and make output coupling easier. We used a separation of 50 µm and a bend radius of 20 µm.

The result is shown in Figure 1.

3.7 2×2 MMI

2×2 MMIs cannot take advantage of symmetric interference, because $M \neq 1$. They can, however, take advantage of paired interference. The requirement for this type of MMI is that the field excitation coefficients $c_\nu = 0$ for modes $\nu = 2, 5, 8, \ldots$. As before, this is not a problem. Length for this type of MMI is not given by Equation (37) but by Equation (33):

$$L = \frac{p}{N}\left(L_\pi\right) \tag{10}$$

However, this imposes new requirements for the locations of input and output waveguides. So, we instead designed our 2×2 MMI as a generic $N \times N$ MMI, for which Equation (24) holds:

$$L = \frac{p}{N}\left(3L_\pi\right) \tag{11}$$

Note that this is simply 4 times larger than the length of our 1×2 MMI.

The result is shown in Figure 2.

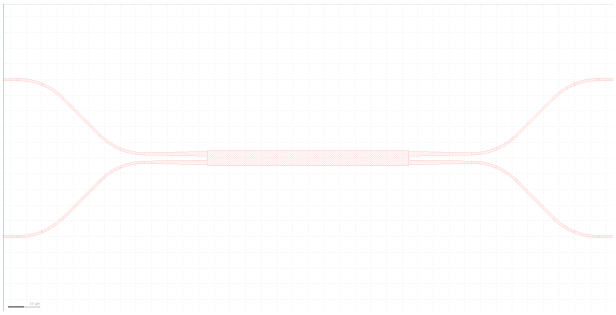

Figure 2: Layout of 2×2 MMI

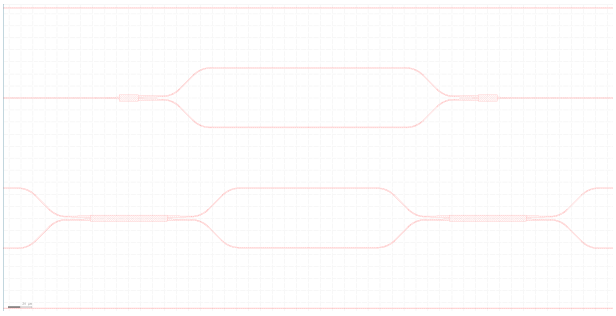

Figure 3: Layout of two MZIs

3.8 MZIs

MZIs are, unlike MMIs, interferometers, in that they produce an output as a response to the phase difference in two or more arms. In the simplest configuration, an MZI can be realized using two MMIs facing opposite directions. In our case, we implemented MZIs using both our 1×2 MMIs and our 2×2 MMIs (Figure 3).

Generally, MZIs include active phase shifting in at least one arm. Active controls were out of scope for this project, but we did implement static phase shifts. The spacing of the arms between the MMIs is 50 µm by default. Changing the spacing slightly changes the total length of the arm, and thus the phase. For a target phase-shift of ϕ, the required reduction in length is given by:

$$\Delta l = \frac{\phi \lambda_0}{2\pi n_{eff}} \tag{12}$$

where $n_{eff} = 1.716$ is the effective refractive index of the waveguide. Because the arms travel at an angle of $\frac{\pi}{4}$ while separating, the required change in separation is given by:

$$\Delta s = \frac{\Delta l}{2\left(\sqrt{2} - 1\right)} \tag{13}$$

3.9 Chip

Having designed individual devices, we generated a single chip demonstrating various combinations.

The chip size cannot be arbitrarily small, because it has to be handled physically. However, devices with coupled outputs have to span the entirety of one dimension of the chip, so at least one dimension (the "length") benefits from a smaller size. We set this value at 5 mm as a balance between both requirements. There is no upper limit to the other dimension (the "width"), but there is an overall limit to how many devices we can fabricate, because in our fabrication process, patterning time is proportional to final silicon area.

One of the risks in fabrication is rough edges. For devices with coupled outputs, both edges have to be smooth. But, for devices with antenna outputs, only the input edge has to be smooth. To improve

our chances of success, we implemented most[2] device twice: once from left to right, and once from right to left.

Another risk is disruptions along the waveguide, For devices with coupled outputs, the waveguide must be intact for the full 5 mm. But, for devices with antenna outputs, only the input waveguide and the start of the antenna have to be intact. To improve our chances of success, we placed the devices closer to the input side than the output side.

For the purposes of calibration and control, we added a waveguides with no device after every second device. We also created antennas of varying sizes.

Finally, we added test structures to validate fabrication of various feature sizes. The "solid" test structure is a long rectangle. The "perturbed" test structure is a long series of squares.

The final list of devices, in order, is given in Appendix A. All couplers are spaced at least 50 µm apart. The final chip is shown in Figure 4.

3.10 GDS Generation

All devices as well as the full chip were generated programmatically using Python 3 and the `gdspy` library. The script, input parameters, and generated Graphic Design System (GDS) files[3] are publicly available at `https://gitlab.cba.mit.edu/dimitar/kitchen-sink/-/tree/main/mmi_gds_generator`. (The ELS-HS50 uses a different format. The converted files are uploaded as well for convenience.)

4 Fabrication

We worked with three separate processes to fabricate our chip. Each is described below, followed by a brief discussion. Each process is also documented on FabuBlox, and a link is provided.

All samples came from a single 25 mm × 25 mm sample, acquired as described in Section 3.2. The composition of the sample was: 722 930 nm of Si, 2000 nm of SiO_2, 70 nm of Si. The following steps were performed to the full sample before any of the steps described below were performed:

1. Spin-coat the entire sample with PMMA
2. Dice the sample into various square and rectangular sub-samples
3. Rinse all sub-samples in acetone for 6 minutes
4. Plasma-clean all samples using the Tergeo Pro. Pre-clean for 2 minutes with `O2FastAsh`. Vent, insert, and pump down. Run `O2FastAsh` for 2 minutes. Vent and remove.

4.1 ma-N

`https://www.fabublox.com/process-editor/821e59b2-b1ca-491e-90f9-b34fbbd27867`

4.1.1 Steps

1. Start with a 15 mm × 5 mm sub-sample
2. Spin-coat ma-N 2403 (a negative photoresist) at 3000 rpm for 60 seconds. Bake at 90°C for 120 seconds.
3. Use the ELS-HS50 to pattern the GDS (generated and published as `main` as described in Section 3.10) onto the sample[4]. Perform a height mapping. Use a beam current of 10 nA and a target exposure dose of 1000 µC cm^{-2}. The total run time was 9 minutes.
4. Develop in AZ 300 MIF for 60 seconds. Rinse in DIW for 120 second. Rinse in DIW for 120 seconds again.

[2] We only implemented the largest antenna once, because it takes more time to pattern than the rest of the design combined. We also only implemented the phase-shifted MZIs once per phase shift.

[3] The files `main.*` were used for the ma-N and PMMA/Cr processes. The files `double.*` were used for the HSQ process. Programmatically, the only difference is the chip length. However, note that `double.gds` was modified manually during fabrication to reduce patterning time. `double.car` is generated from the modified version. See Section 4 for more details.

[4] Note that features in the GDS are extended by 500 µm on each side of the sample, so the total length to be patterned is 6 mm. Care should be taken to ensure that the original 5 mm design aligns well with the sample, but the extension ensures that features run all the way to each edge of the chip; the ELS-HS50 allows for this

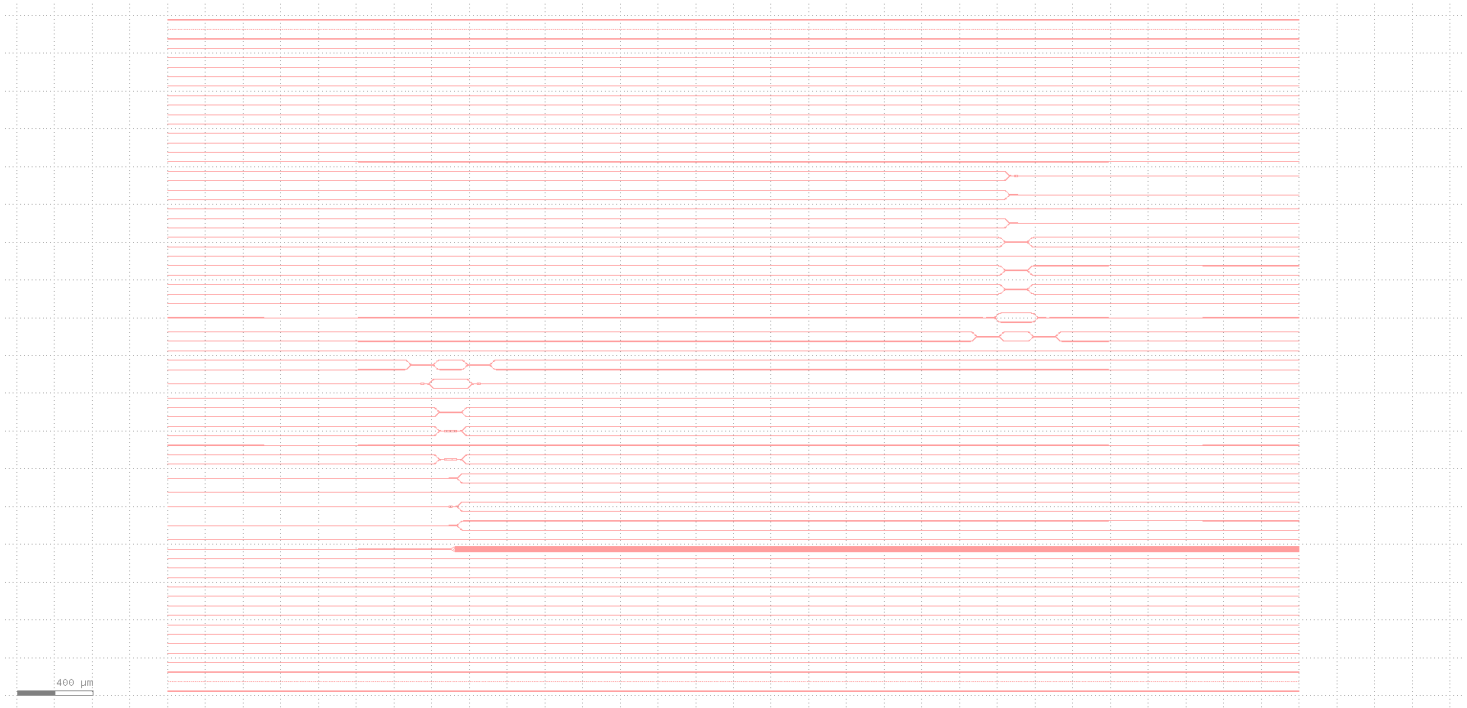

Figure 4: Final Chip Layout

5. Etch using the Corial 210RL. Start by cleaning using `SF6/O2 Clean`. Then, etch using `Si Etch 100W , 20mT` (an SF_6/Ar recipe). Stop Step 5 when `PM1 Laser Signal` transitions from sinusoidal to exponential (approximately 15 seconds).

6. Strip ma-N using the Tergeo Pro. Pre-clean for 2 minutes with `O2FastAsh`. Vent, insert, and pump down. Run `O2FastAsh` for 2 minutes. Vent and remove.

7. Deposit 2 µm of SiO_2 using PE-CVD.

4.1.2 Discussion

Note that a negative photoresist is required. The device area (which is not etched) is small compared to the overall area. The non-device area cannot be filled with silicon, because that would adversely affect the photonic properties of the devices. Using electron-beam lithography, the patterning time (and therefore the patterning cost) are proportional to the area patterned, so the area patterned must be the device area, which will have to be covered by the mask. It is also not practical to use other methods for patterning, because the required feature sizes are small (as explained in Section 3) and require a high resolution. So, the solution is to use electron-beam lithography with a negative photoresist.

Devices were manufactured successfully using this process. See, for example, Figure 5 and Figure 6. However, many of the couplers delaminated at the edges of the chip. See, for example, Figure 7.

We attempted the ma-N process twice. The above description is of the second, more successful attempt. The first time, we used two 5 mm×5 mm sub-samples. Neither was as uniform as the 10 mm × 5 mm sample during the second attempt, but one was more uniform than the other. We used an earlier version of the GDS, with a few differences, including no antennas larger than the waveguide width of 750 nm. We used 5 nA instead of 10 nA for the beam current, because the total GDS area was about twice as small. We used 1050 µC cm^{-2} instead of 1000 µC cm^{-2} for the exposure dose. Perhaps most importantly, we used IPA:MIBK 3:1 instead of AZ 300 MIF as the developer.

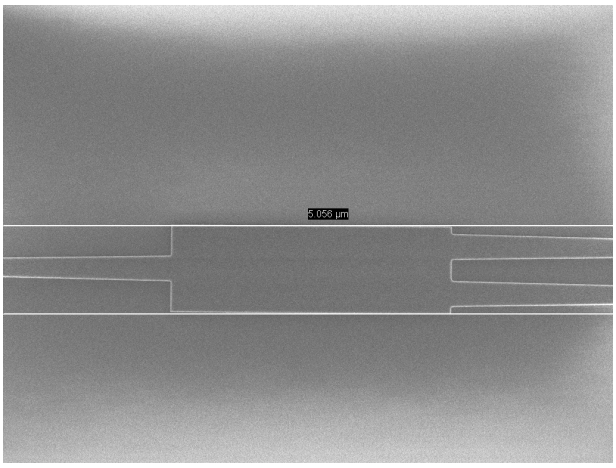

Figure 5: SEM Image of an MMI fabricated using the ma-N process, before SiO_2 deposition

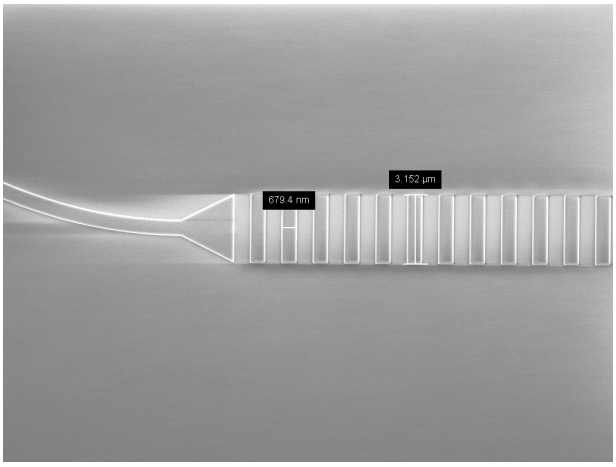

Figure 6: SEM Image of an antenna fabricated using the ma-N process, before SiO_2 deposition

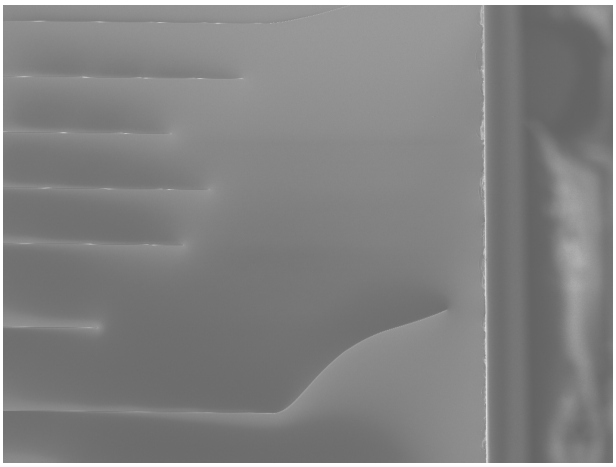

Figure 7: SEM Image of some of the couplers fabricated using the ma-N process, before SiO_2 deposition

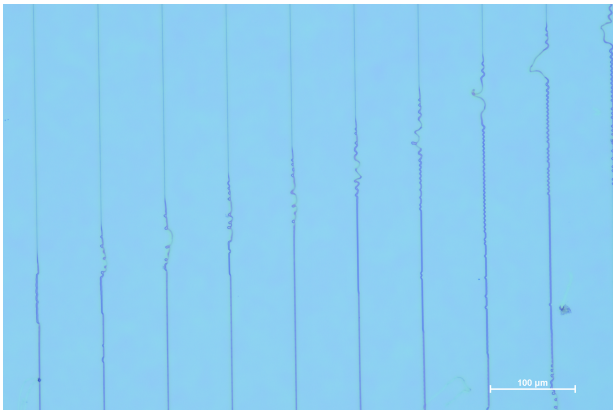

Figure 8: SEM Image of some of the waveguides fabricated using an earlier ma-N process

The results for the sample that visually looked better before patterning also appeared better on the SEM. However, neither sample looked promising. See, for example, a few waveguides in Figure 8. We chose not to test these samples, and instead focused on developing the ma-N process further. We also chose to develop two other processes in parallel in improve chances of success.

4.2 PMMA / Cr

`https://www.fabublox.com/process-editor/90cb603d-88ac-4162-8148-51048e9d37d3`

4.2.1 Steps

1. Start with a 15 mm × 5 mm sub-sample
2. Spin-coat 950k PMMA (a positive photoresist) at 3000 rpm for 60 seconds. Bake at 150°C for 240 seconds.
3. Use the ELS-HS50 to pattern the GDS (generated and published as `main` as described in Section 3.10) onto the sample[5]. Perform a height mapping. Use a beam current of 10 nA and a target exposure dose of 1000 µC cm^{-2}. The total run time was 9 minutes.
4. Develop in IPA:MIBK 3:1 for 60 seconds. Spritz with IPA. Dry with N_2.
5. Deposit chromium using the FC-2000. Pump down. Set to 400 Å thickness (deposition rate is nominally 1 Å s^{-1}).
6. Lift off chromium. Use N-Methyl-Pyrrolidone. Process in 69°C water bath for 300 seconds. Rinse in acetone. Rinse in acetone again. Rinse in IPA. Dry with N_2.
7. Etch using the Corial 210RL. Start by cleaning using `SF6/O2 Clean`. Then, etch using `Si Etch 100 W, 20mT` (an SF_6/Ar recipe). Stop Step 5 when PM1 Laser Signal transitions from sinusoidal to exponential (approximately 15 seconds).

[5]Note that features in the GDS are extended by 500 µm on each side of the sample; see footnote in Section 4.1.1

Figure 9: SEM Image of an MMI fabricated using the PMMA/Cr process, before SiO_2 deposition

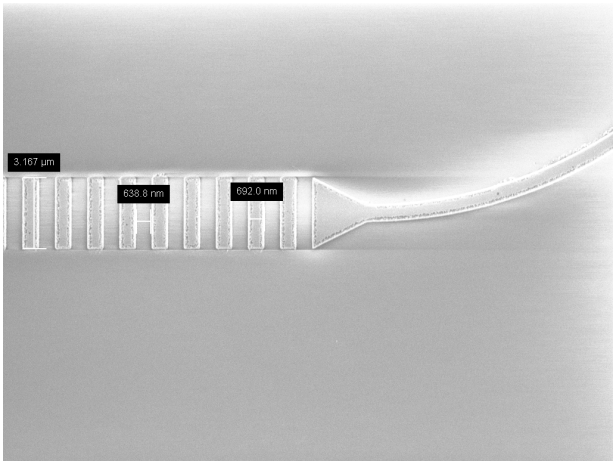

Figure 10: SEM Image of an antenna fabricated using the PMMA/Cr process, before SiO_2 deposition

8. Strip chromium. Use Transene Chromium Etchant 1020. (Make sure to use proper fume hood and proper PPE.) Etch for 30-35 seconds at 25°C (etch rate is nominally $1.4\,\mathrm{nm\,s^{-1}}$).
9. Deposit 2 µm of SiO_2 using PE-CVD.

4.2.2 Discussion

This is a hard-mask process. Note that because the photoresist is not the final mask, a positive photoresist is required.

Devices were manufactured successfully using this process. See, for example, Figure 9 and Figure 10. The couplers at the edges of the chip had a higher rate of success than those fabricated using the maN process. However, there were deposits of leftover chromium, which can absorb infrared light. See, for example, Figure 11. Also, the large antenna was not fully etched.

We initially spin-coated a $5\,\mathrm{mm} \times 5\,\mathrm{mm}$ sub-sample, but the results were not satisfactory. Larger pieces lead to a more even layer of photoresist.

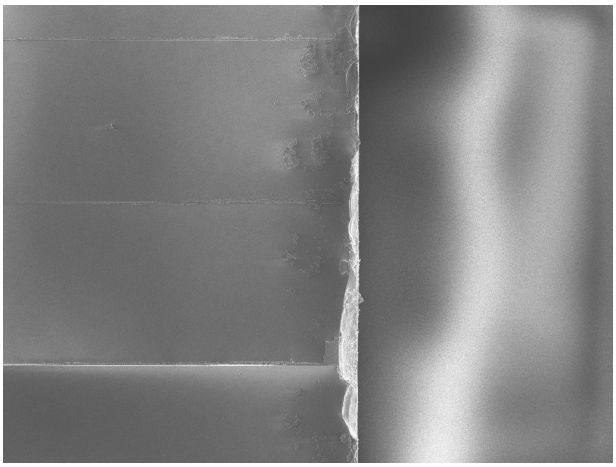

Figure 11: SEM Image of some of the couplers fabricated using the PMMA/Cr process, before SiO_2 deposition

4.3 HSQ

`https://www.fabublox.com/process-editor/56d3a658-9220-4e8e-8382-553dd1fbf350`

4.3.1 Steps

1. Start with a 15 mm × 10 mm sub-sample
2. Spin-coat 2% HSQ at 4000 rpm for 60 seconds. Do not bake. Proceed to next step quickly.
3. Use the ELS-HS50 to pattern the GDS (generated and published as `double` as described in Section 3.10) onto the sample[6][7]. Perform a height mapping. Use a beam current of 2 nA and a target exposure dose of 6000 µC cm^{-2}. The total run time was 18 minutes.
4. Develop
5. Etch using the `SI_ETCHC` Cl_2-based recipe on the Samco RIE-230iP.
6. Deposit 2 µm of SiO_2 using PE-CVD.

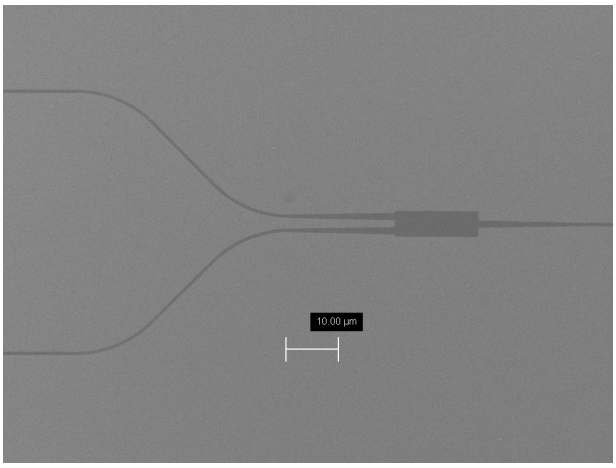

Figure 12: SEM Image of an MMI fabricated using the HSQ process, before etching

4.3.2 Discussion

Using Hydrogen silsesquioxane (HSQ) for the mask was a very promising option [14, 15]. HSQ is slightly more difficult to work with (it is stored cold, and degrades quickly between coating and patterning). However, patterned HSQ becomes SiO_2, so it does not need to be removed after etching for this application!

We were planning to use 4% HSQ for this experiment, but only 2% was available. We used a larger chip to ensure a more even film. Between that, and the recommended electron beam settings, the patterning time was prohibitively large. So, we removed half of the structures and shortened the antennas. To develop, we used a solution of 1% NaOH by weight and 4% NaCl by weight in DIW. We developed for 120 second, then performed two 120-second DIW baths.

After this, we intended to etch the silicon as described in Step 5. However, we found out that once a sample has been processed using Na^+, **it**

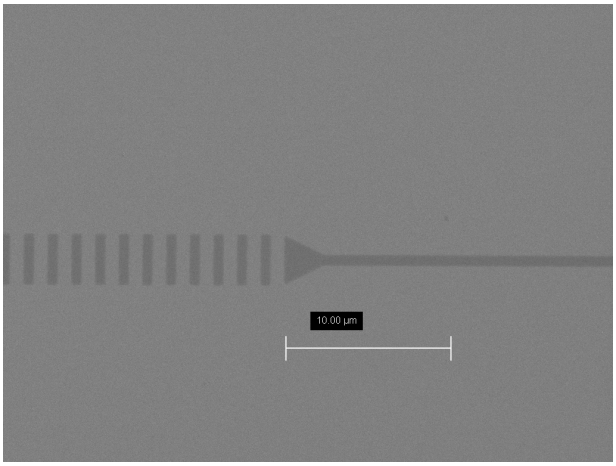

Figure 13: SEM Image of an antenna fabricated using the HSQ process, before etching

[6]Note that during fabrication, the output file was reduced dramatically from the programmed values to account for the slower etch rate and longer chip

[7]Note that features in the GDS are extended by 500 µm on each side of the sample; see footnote in Section 4.1.1

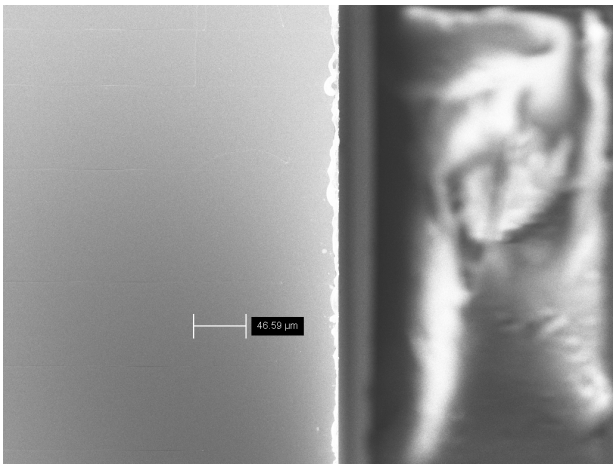

Figure 14: SEM Image of some of the couplers fabricated using the HSQ process, before etching

is no longer compatible with most tools at MIT.nano, including the Samco RIE-230iP. In the future, a TMAH-based process can be used.

We imaged the sample using the SEM at this stage. Optically, the features were difficult to spot, but under the SEM they were clear. This process achieved a very high resolution. See, for example, Figure 12 and Figure 13 (c.f. Figure 5, Figure 6, Figure 9, Figure 10). The edge couplers also seemed intact (Figure 14).

The HSQ process is promising for the future, but we were not able to evaluate it further.

5 Testing

We tested the sample fabricated using the ma-N process and the sample fabricated using the PMMA / Cr process. We did not attempt to test the sample fabricated using the HSQ process, because we were not able to etch that sample. We tested the chips at a wavelength of 1500 nm. We generated the input using a Keysight (Agilent) 81600B inside a 8164B chassis. For coupled output testing, we used a 81635A module in the same chassis. For antenna output testing, we used an IR camera through a microscope. We interfaced with the chip using stripped multi-mode fiber

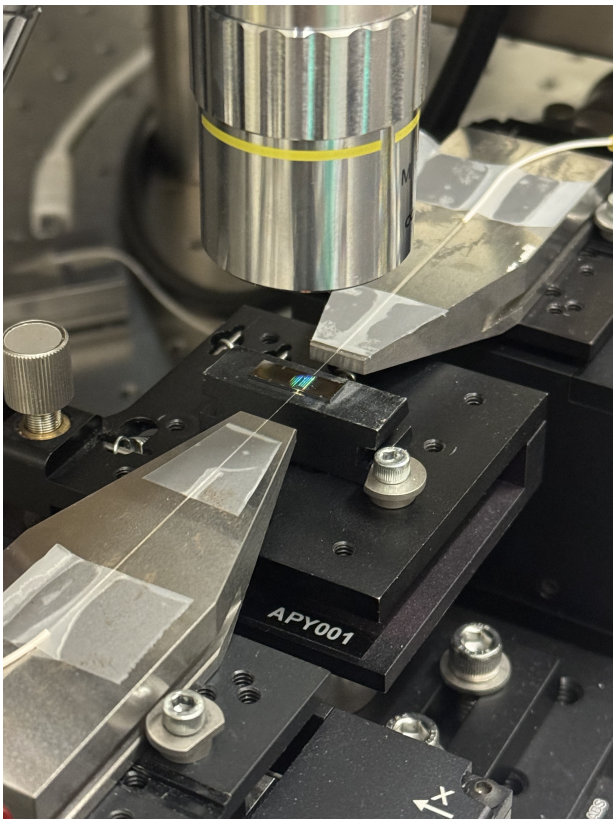

Figure 15: A fabricated sample under test

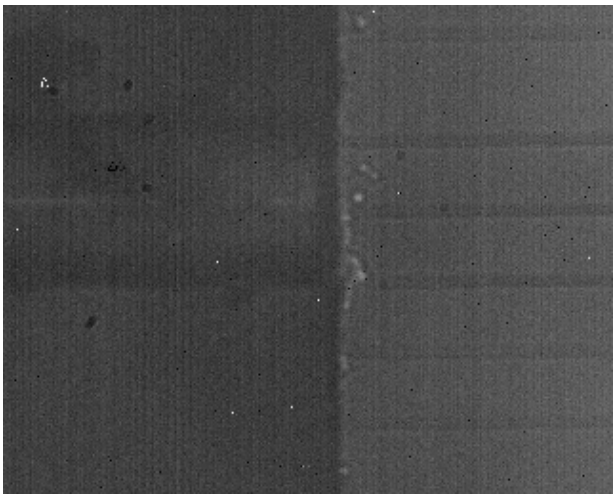

Figure 16: Fiber coupled to large left-to-right antenna on PMMA / Cr sample, input off

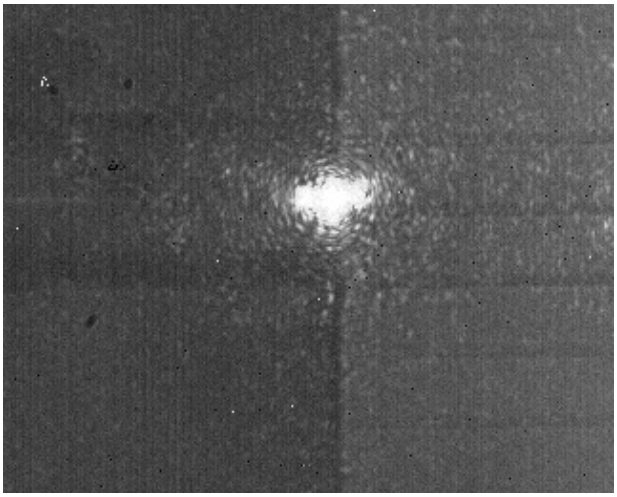

Figure 17: Fiber coupled to large left-to-right antenna on PMMA / Cr sample, source on

cables (Figure 15). We set the source power at the instrument to 0dBm, and used three manual fiber polarization controllers to tune the polarization at the input of the chip.

5.1 PMMA / Cr

We began testing with the large, 25 µm left-to-right antenna (see Section 3.9 and Appendix A) on the PMMA / Cr chip, because it gives the strongest response. Figure 16 shows the input fiber coupled to the antenna's input coupler, with the input disabled at the source. Figure 17 shows it with the input enabled. Figure 18 shows the antenna with the input disabled at the source. Figure 19 shows it with the input enabled. It is clear that the antenna illuminates. Cross-talk can also be observed with the neighboring 3.2 µm antenna.

Because we designed the chip for TE, we knew that TE would give the strongest response. We thus tuned the fiber polarization controllers at this stage to optimize the response of the large antenna. We kept this setting for the remainder of testing.

Figure 20 shows the response when coupled to the adjacent 3.2 µm antenna. Figure 21 shows the response when coupled to the nearby 750 nm antenna.

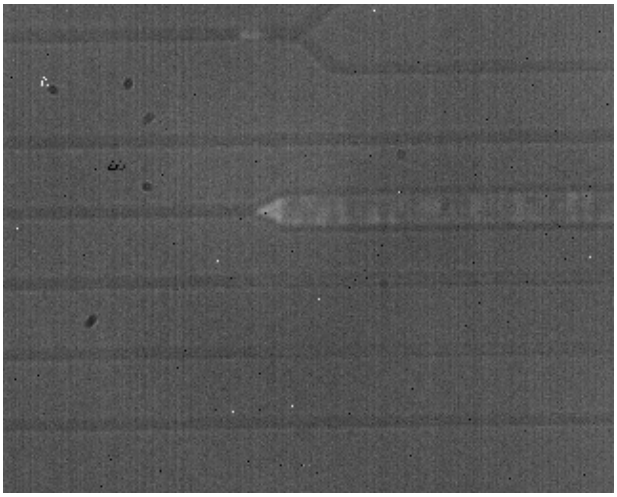

Figure 18: Large left-to-right antenna on PMMA / Cr sample, input off

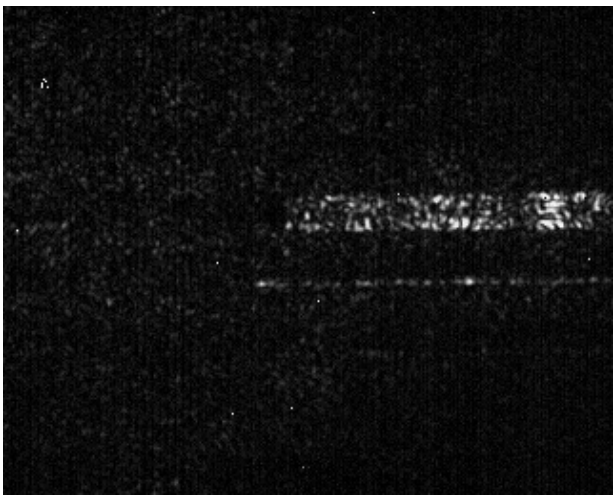

Figure 19: Large left-to-right antenna on PMMA / Cr sample, source on

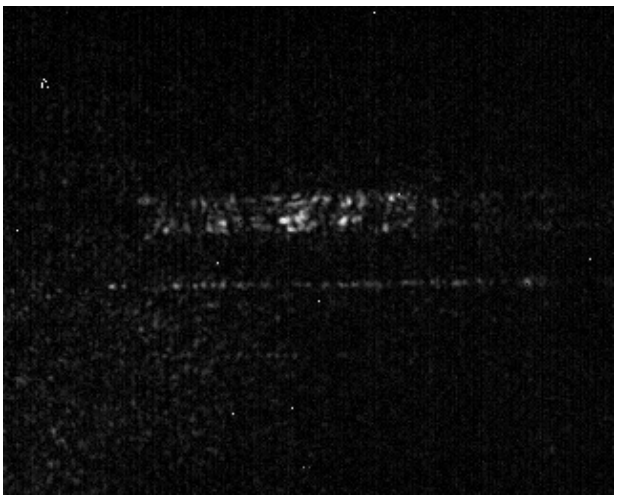

Figure 21: Small left-to-right antenna on the PMMA / Cr sample, source on

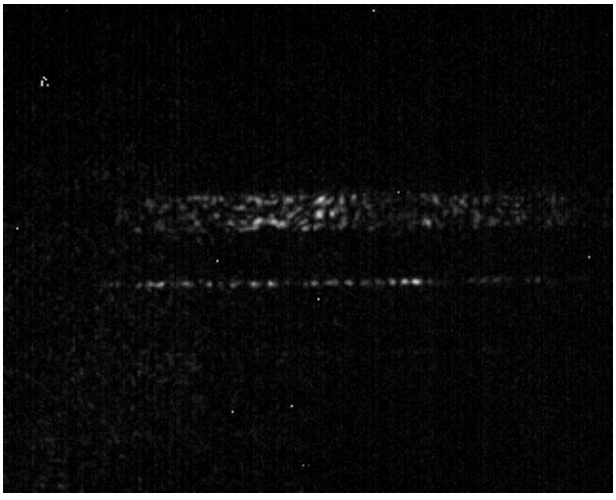

Figure 20: Medium left-to-right antenna on the PMMA / Cr sample, source on

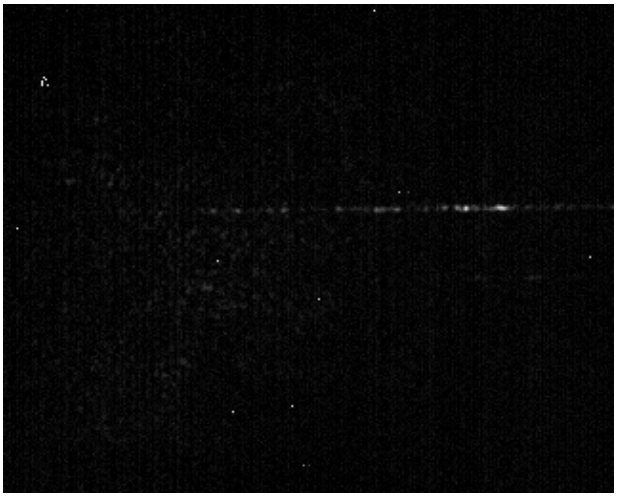

Figure 22: Medium right-to-left antenna on the PMMA / Cr sample, source on

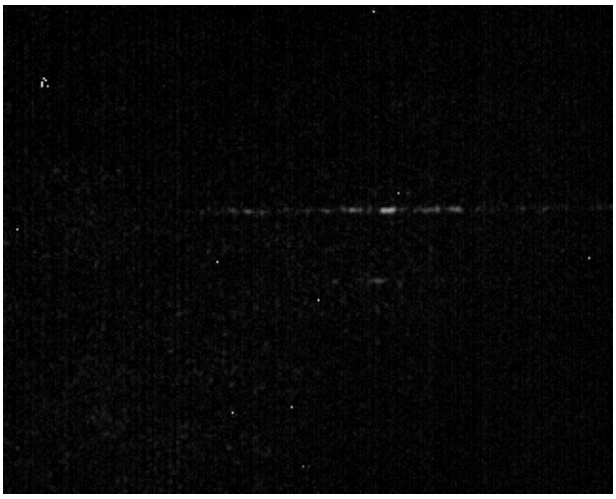

Figure 23: Small right-to-left antenna on the PMMA / Cr sample, source on

Note that the bright lines are the large and medium antennas illuminated by cross-talk, not the small antenna. The small antenna is visible faintly further down (c.f. Figure 20). Figure 22 shows the response when coupled to the 3.2 µm antenna on the other side of the chip. Figure 23 shows the response when coupled to the nearby 750 nm antenna. Note again that the brighter line is the medium antenna coupling in, but that the small antenna is faintly visible in the latter image. Note also that there is no large antenna on that side (to reduce patterning time). It is evident that coupling occurs with the large and medium antennas, and that they emit light. It seems that the small antennas are perhaps coupling as well, but results are less conclusive.

Figure 24 shows the output of one of the 1x2 MMIs with medium antennas. It is clear that the signal couples and emits, and splitting seems (qualitatively) even. Figure 25 shows the output of one of the 2x2 MMIs with medium antennas, with the source coupled to input 0. Figure 26 shows the output of the same MMI with the source coupled to input 1. This suggests that both the output and the input of the MMI are evenly splitting the signal.

Finally, we attempted coupling to the output of

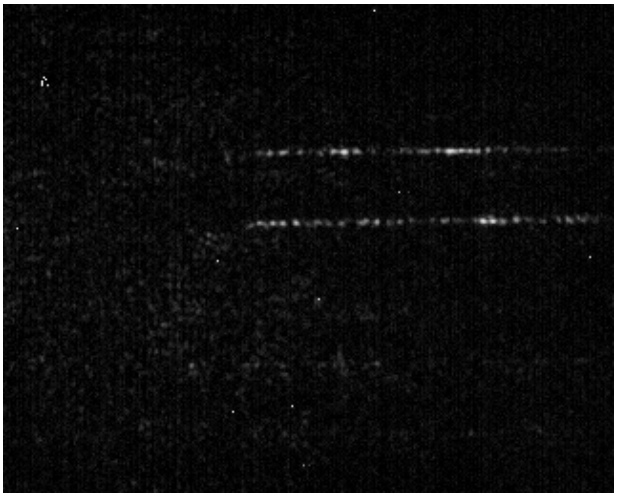

Figure 24: A 1x2 MMI on the PMMA / Cr sample, source on

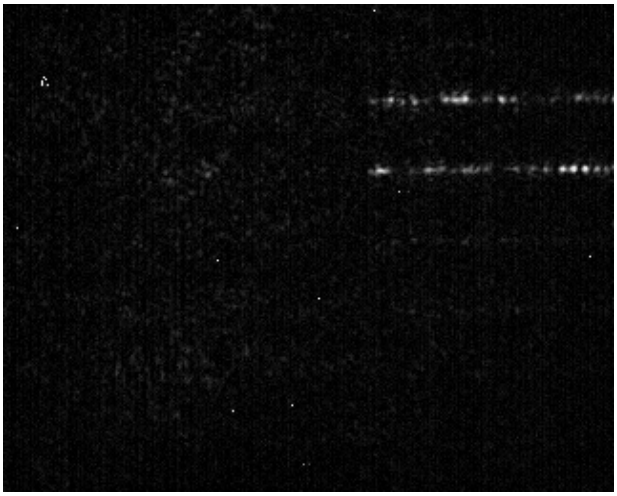

Figure 25: A 2x2 MMI on the PMMA / Cr sample, source on, input 0

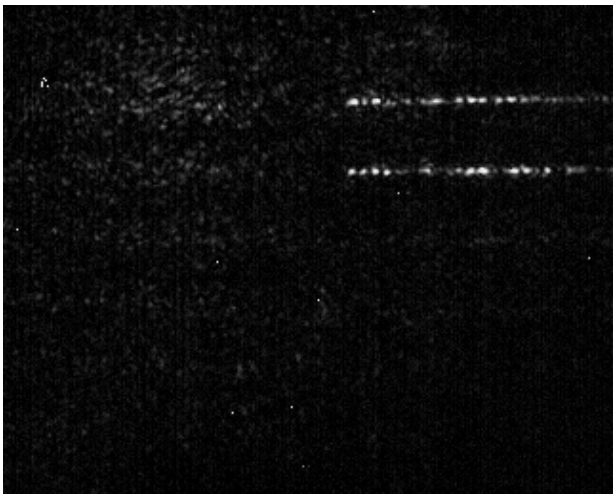

Figure 26: A 2x2 MMI on the PMMA / Cr sample, source on, input 1

the chip. We tested a few waveguides, because they would have the lowest insertion loss. The attenuation measured was approximately 50dB. One potential explanation is that the edges of the chip were not polished. Another is leftover Cr absorbing light.

5.2 ma-N

We attempted similar tests on the ma-N chip. However, many of the input couplers were delaminated (as shown previously in Figure 7). So, only a few tests were successful. Notably, one of the 1x2 MMIs with 3.2 µm antennas worked well. The MMI is shown in Figure 27 with the source off and in Figure 28 with the source on.

6 Next Steps

6.1 HSQ Process Development

The HSQ process showed great promise, and the authors believe it may be a viable process in the future. However, a non-Na^+ development method needs to be developed and tested. TMAH may be a good starting point. If such a method is found, the process

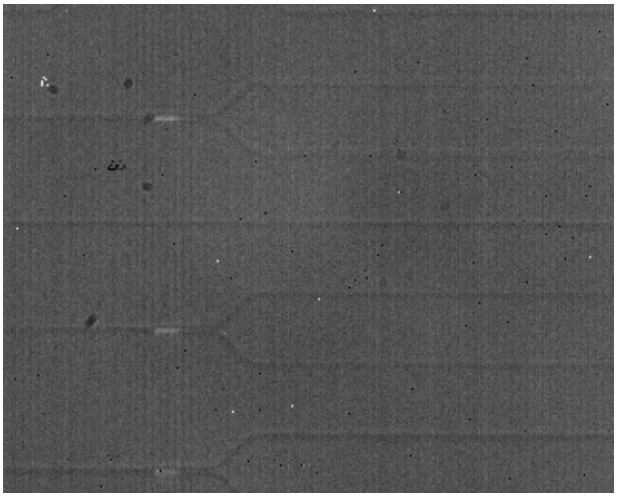

Figure 27: A 1x2 MMI on the ma-N sample, source off

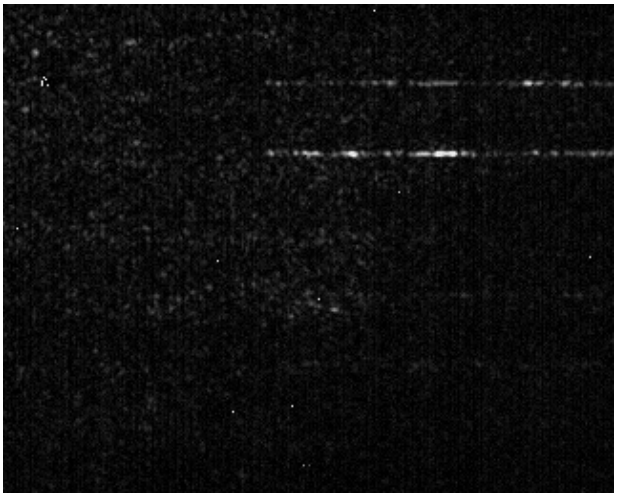

Figure 28: A 1x2 MMI on the ma-N sample, source on

may allow for finer feature resolution. Furthermore, because it consists of less steps, it may enable faster prototyping. If this becomes a viable path, attempting to speed up patterning will be beneficial as well. Also, it may be beneficial to use 4% HSQ instead of 2% HSQ.

6.2 Chromium Etching Improvements

We had the most success with the PMMA / Cr process. However, we suspect not all chromium was fully etched away, which may have affected our results. We suggest the chromium etching step be characterized more carefully in the future.

Note in particular that the 25 µm antenna on this sample wasn't fully etched. This may or may not have been an issue with developing.

6.3 ma-N Development Improvements

The ma-N photoresist may be capable of the small features required for silicon photonics. However, the reason for delamination during development needs to be determined. We already observed improvements when changing the development chemical. Perhaps further improvements can be made.

6.4 Polishing

One step we omitted for each process was polishing the edges of the chip. Polishing may reduce losses at the edges of the chip and improve coupling. Methods for polishing should be explored further.

6.5 Better Output Coupling Testing

Results from output coupling testing were poor. Our hope is that full chromium removal and polishing would help. To make testing more direct, we suggest fabricating MMIs which have an antenna on one output arm and a coupler on the other. Since we have verified that MMIs work reliably, activity on the antenna arm would provide confidence that a signal is reaching the output coupler, whether or not it makes it off the chip.

6.6 MZIs with Antennas

Similarly, it may be beneficial to add MZIs with antenna outputs. Because of the high attenuation when coupling to the output, we were unfortunately not able to test any of the MZIs. MZIs with antenna outputs would help verify their operation even if the attenuation problems are not resolved.

7 Conclusion

Silicon photonics enables a slew of new applications and improves greatly on existing ones, such as computing. Many of these applications rely only on a small number of building blocks. In this project, we demonstrated successful design, fabrication and testing of some of these building blocks. We also identified and documented challenges along the way. It is the authors' hope that the processes developed can be recreated and improved upon by future generations, so that silicon photonics at MIT.nano becomes more accessible to those interested in perusing it.

8 Acknowledgments

This work was performed as classwork for MIT 6.2540, with much support from its teaching team: instructors Prof. Farnaz Niroui and Prof. Akintunde Akinwande and teaching assistants Sheldon Zheng and Mateusz Czajka. Sheldon in particular facilitated all work at MIT.nano and provided invaluable guidance with process development. Though Mateusz was not assigned to this team, he helped greatly on numerous occasions. Help was also received from members of Dr. Niroui's group – Hohyeon Kim and Shelly Ben-David in particular. Finally, Eric Lim and the rest of the MIT.nano staff were indispensable in this project.

A Appendix: Device List

- Solid Test Structure, 3.2 µm
- Perturbed Test Structure, 3.2 µm
- Solid Test Structure, 1.6 µm
- Perturbed Test Structure, 1.6 µm
- Solid Test Structure, 0.8 µm
- Perturbed Test Structure, 0.8 µm
- Solid Test Structure, 0.4 µm
- Perturbed Test Structure, 0.4 µm
- Solid Test Structure, 0.2 µm
- Perturbed Test Structure, 0.2 µm
- Solid Test Structure, 0.1 µm
- Perturbed Test Structure, 0.1 µm
- Coupler, 250 nm – Waveguide, 750 nm – Coupler, 250 nm
- Antenna, 750 nm – Waveguide, 750 nm – Coupler, 250 nm
- Antenna, 3.2 µm – Waveguide, 750 nm – Coupler, 250 nm
- Coupler, 250 nm – Waveguide, 750 nm – Coupler, 250 nm
- 2x Couplers, 250 nm – 2x1 MMI –Coupler, 250 nm
- 2x Antennas, 750 nm – 2x1 MMI
- Coupler, 250 nm – Waveguide, 750 nm – Coupler, 250 nm
- 2x Antennas, 3.2 µm – 2x1 MMI
- 2x Couplers, 250 nm – 2x2 MMI –2x Couplers, 250 nm
- Coupler, 250 nm – Waveguide, 750 nm – Coupler, 250 nm
- 2x Antennas, 750 nm – 2x2 MMI
- 2x Antennas, 3.2 µm – 2x2 MMI
- Coupler, 250 nm – Waveguide, 750 nm – Coupler, 250 nm
- Coupler, 250 nm – 1x1 MZI, pi delta – Coupler, 250 nm
- 2x Couplers, 250 nm – 2x2 MZI, pi delta – 2x Couplers, 250 nm
- Coupler, 250 nm – Waveguide, 750 nm – Coupler, 250 nm
- 2x Couplers, 250 nm – 2x2 MZI –2x Couplers, 250 nm
- Coupler, 250 nm – 1x1 MZI –Coupler, 250 nm
- Coupler, 250 nm – Waveguide, 750 nm – Coupler, 250 nm
- 2x Couplers, 250 nm – 2x2 MMI –2x Antennas, 3.2 µm
- 2x Couplers, 250 nm – 2x2 MMI –2x Antennas, 750 nm
- Coupler, 250 nm – Waveguide, 750 nm – Coupler, 250 nm
- 2x Couplers, 250 nm – 2x2 MMI –2x Couplers, 250 nm
- Coupler, 250 nm – 1x2 MMI – 2x Antennas, 3.2 µm
- Coupler, 250 nm – Waveguide, 750 nm – Coupler, 250 nm
- Coupler, 250 nm – 1x2 MMI –2x Antennas, 750 nm
- Coupler, 250 nm – 1x2 MMI –2x Couplers, 250 nm
- Coupler, 250 nm – Waveguide, 750 nm – Coupler, 250 nm
- Coupler, 250 nm – Waveguide, 750 nm – Antenna, 25 µm
- Coupler, 250 nm – Waveguide, 750 nm – Antenna, 3.2 µm
- Coupler, 250 nm – Waveguide, 750 nm – Antenna, 750 nm
- Perturbed Test Structure, 0.1 µm

- Solid Test Structure, 0.1 µm
- Perturbed Test Structure, 0.2 µm
- Solid Test Structure, 0.2 µm
- Perturbed Test Structure, 0.4 µm
- Solid Test Structure, 0.4 µm
- Perturbed Test Structure, 0.8 µm
- Solid Test Structure, 0.8 µm
- Perturbed Test Structure, 1.6 µm
- Solid Test Structure, 1.6 µm
- Perturbed Test Structure, 3.2 µm
- Solid Test Structure, 3.2 µm

References

[1] Nicholas C. Harris et al. "Linear programmable nanophotonic processors". In: *OSA Publishing* (2018). Accepted: 2021-10-27T20:11:02Z Publisher: The Optical Society. URL: https://dspace.mit.edu/handle/1721.1/135163 (visited on 10/31/2025).

[2] Wim Bogaerts et al. "Programmable photonic circuits". In: *Nature* 586.7828 (Oct. 8, 2020), pp. 207–216. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/s41586-020-2764-0. URL: https://www.nature.com/articles/s41586-020-2764-0 (visited on 10/31/2025).

[3] Michael Reck et al. "Experimental realization of any discrete unitary operator". In: *Physical Review Letters* 73.1 (July 4, 1994). Publisher: American Physical Society, pp. 58–61. DOI: 10.1103/PhysRevLett.73.58. URL: https://link.aps.org/doi/10.1103/PhysRevLett.73.58 (visited on 10/31/2025).

[4] William R. Clements et al. "Optimal design for universal multiport interferometers". In: *Optica* 3.12 (Dec. 20, 2016), p. 1460. ISSN: 2334-2536. DOI: 10.1364/OPTICA.3.001460. URL: https://opg.optica.org/abstract.cfm?URI=optica-3-12-1460 (visited on 10/31/2025).

[5] David A. B. Miller. "Perfect optics with imperfect components". In: *Optica* 2.8 (Aug. 20, 2015). Publisher: Optica Publishing Group, pp. 747–750. ISSN: 2334-2536. DOI: 10.1364/OPTICA.2.000747. URL: https://opg.optica.org/optica/abstract.cfm?uri=optica-2-8-747 (visited on 11/05/2025).

[6] Yichen Shen et al. "Deep learning with coherent nanophotonic circuits". In: *Nature Photonics* 11.7 (July 2017), pp. 441–446. ISSN: 1749-4885, 1749-4893. DOI: 10.1038/nphoton.2017.93. URL: https://www.nature.com/articles/nphoton.2017.93 (visited on 10/31/2025).

[7] Hailong Zhou et al. "Photonic matrix multiplication lights up photonic accelerator and beyond". In: *Light: Science & Applications* 11.1 (Feb. 3, 2022). Publisher: Nature Publishing Group, p. 30. ISSN: 2047-7538. DOI: 10.1038/s41377-022-00717-8. URL: https://www.nature.com/articles/s41377-022-00717-8 (visited on 11/06/2025).

[8] Bo Wu et al. "Real-Valued Optical Matrix Computing with Simplified MZI Mesh". In: *Intelligent Computing* 2 (Jan. 2023), p. 0047. ISSN: 2771-5892. DOI: 10.34133/icomputing.0047. URL: https://spj.science.org/doi/10.34133/icomputing.0047 (visited on 10/31/2025).

[9] Yang Shi et al. "Nonlinear germanium-silicon photodiode for activation and monitoring in photonic neuromorphic networks". In: *Nature Communications* 13.1 (Oct. 13, 2022). Publisher: Nature Publishing Group, p. 6048. ISSN: 2041-1723. DOI: 10.1038/s41467-022-33877-7. URL: https://www.nature.com/articles/s41467-022-33877-7 (visited on 11/06/2025).

[10] G. Masetti, M. Severi, and S. Solmi. "Modeling of carrier mobility against carrier concentration in arsenic-, phosphorus-, and boron-doped silicon". In: *IEEE Transactions on Electron Devices* 30.7 (July 1983), pp. 764–769. ISSN: 1557-9646. DOI: 10.1109/T-ED.1983.21207. URL: https://ieeexplore.ieee.org/document/1483108 (visited on 12/06/2025).

[11] Y. Chen. "Development of mid-infrared surface plasmon resonance-based sensors with highly-doped silicon for biomedical and chemical applications". In: *Optics Express* 17 (Feb. 17, 2009), pp. 3130–3140. DOI: 10.1364/OE.17.003130.

[12] L.B. Soldano and E.C.M. Pennings. "Optical multi-mode interference devices based on self-imaging: principles and applications". In: *Journal of Lightwave Technology* 13.4 (Apr. 1995), pp. 615–627. ISSN: 1558-2213. DOI: 10.1109/50.372474. URL: https://ieeexplore.ieee.org/document/372474/ (visited on 10/30/2025).

[13] B. C. Buchler et al. “Precision spectral manipulation of optical pulses using a coherent photon echo memory”. In: *Optics Letters* 35.7 (Mar. 2010), p. 1091. ISSN: 1539-4794. DOI: 10.1364/ol.35.001091. URL: http://dx.doi.org/10.1364/OL.35.001091.

[14] Maziar P. Nezhad et al. “Etch-free low loss silicon waveguides using hydrogen silsesquioxane oxidation masks”. In: *Optics Express* 19.20 (Sept. 26, 2011), p. 18827. ISSN: 1094-4087. DOI: 10.1364/OE.19.018827. URL: https://opg.optica.org/oe/abstract.cfm?uri=oe-19-20-18827 (visited on 12/10/2025).

[15] Michael G. Wood et al. “Optimization of electron beam patterned hydrogen silsesquioxane mask edge roughness for low-loss silicon waveguides”. In: *Journal of Nanophotonics* 8.1 (Jan. 3, 2014), p. 083098. ISSN: 1934-2608. DOI: 10.1117/1.JNP.8.083098. URL: http://nanophotonics.spiedigitallibrary.org/article.aspx?doi=10.1117/1.JNP.8.083098 (visited on 12/10/2025).