AFRL-RY-WP-TR-2021-0143

VISIBLE INTEGRATED PHOTONICS ENHANCED REALITY (VIPER)

Michael R. Watts
Massachusetts Institute of Technology

SEPTEMBER 2021
Final Report

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.
See additional restrictions described on inside pages

© 2021 Massachusetts Institute of Technology

STINFO COPY

AIR FORCE RESEARCH LABORATORY
SENSORS DIRECTORATE
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7320
AIR FORCE MATERIEL COMMAND
UNITED STATES AIR FORCE

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the USAF 88th Air Base Wing (88 ABW) Public Affairs Office (PAO) and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (http://www.dtic.mil).

AFRL-RY-WP-TR-2021-0143 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

//Signature//

NICHOLAS G. USECHAK
Program Manager
Highly Integrated Microsystems Branch
Aerospace Components & Subsystems Division

//Signature//

STEPHEN L. HARY
Chief
Highly Integrated Microsystems Branch
Aerospace Components & Subsystems Division

//Signature//

LESTER C. LONG, Lt Col, USAF
Deputy
Aerospace Components & Subsystems Division
Sensors Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

*Disseminated copies will show "//Signature//" stamped or typed above the signature blocks.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS**.

1. REPORT DATE *(DD-MM-YY)*	**2. REPORT TYPE**	**3. DATES COVERED *(From - To)***
September 2021	Final	29 September 2017 – 30 November 2020

4. TITLE AND SUBTITLE
VISIBLE INTEGRATED PHOTONICS ENHANCED REALITY (VIPER)

5a. CONTRACT NUMBER
FA8650-17-1-7713

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
DARPA

6. AUTHOR(S)
Michael R. Watts

5d. PROJECT NUMBER
N/A

5e. TASK NUMBER
N/A

5f. WORK UNIT NUMBER
Y1M9

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Massachusetts Institute of Technology
77 Massachusetts Ave.
Cambridge, MA 02139-4301

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Air Force Research Laboratory
Sensors Directorate
Wright-Patterson Air Force Base, OH 45433-7320
Air Force Materiel Command
United States Air Force

Defense Advanced Research Projects Agency
(DARPA/MTO)
675 North Randolph Street
Arlington, VA 22203

10. SPONSORING/MONITORING AGENCY ACRONYM(S)
AFRL/RYDI

11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S)
AFRL-RY-WP-TR-2021-0143

12. DISTRIBUTION/AVAILABILITY STATEMENT
DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

13. SUPPLEMENTARY NOTES
© 2021 Massachusetts Institute of Technology. This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with The Under Secretary of Defense memorandum dated 24 May 2010 and AFRL/DSO policy clarification email dated 13 January 2020. This material is based on research sponsored by the Air Force Research Lab (AFRL) and the Defense Advanced Research Agency (DARPA) under agreement number FA8650-17-1-7713. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of AFRL and DARPA or the U.S. Government. Report contains color.

14. ABSTRACT
This report was developed under a DARPA-funded contract. This program seeks to develop the technologies necessary to demonstrate visible direct-view near-eye augmented reality (AR) display enabled by a grid of actively steerable integrated photonic phased array pixels for directly reconstructing the wave fronts of virtual scenes.

15. SUBJECT TERMS
OPA, integrated photonics, LiDAR, augmented reality, heads-up displays

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT:**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON** (Monitor)
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Nicholas Usechak
Unclassified	Unclassified	Unclassified	SAR	82	**19b. TELEPHONE NUMBER *(Include Area Code)***
					N/A

Table of Contents

Section **Page**

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

Section **Page**

List of Figures

Figure **Page**

Figure **Page**

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

Figure **Page**

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

Figure **Page**

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

1 INTRODUCTION

Augmented-reality head-mounted displays that display information directly in the user's field of view have many wide-reaching applications in defense, medicine, engineering, gaming, etc. However, current commercial head-mounted displays are bulky, heavy, and indiscreet. Moreover, these current displays are not capable of producing holographic images with full depth cues; this lack of depth information results in users experiencing eyestrain and headaches that limit long-term and wide-spread use of these displays (an effect known as the vergence-accommodation conflict).

In this report, recent advances in the development of Visible Integrated Photonics Enhanced Reality (VIPER), a novel integrated-photonics-based holographic display, will be reviewed. The VIPER display consists of a single discreet transparent chip that sits directly in front of the user's eye and projects visible-light 3D holograms that only the user can see. It presents a highly-discreet and fully-holographic solution for the next generation of augmented-reality displays.

First, a variety of novel integrated visible-light components, required for the VIPER display, will be demonstrated. The first integrated visible-light liquid-crystal-based phase [1,2] and amplitude [3,4] modulators (with device lengths an order of magnitude smaller than traditional inefficient thermo-optic visible-light modulators) will be shown. Additionally, the first actively-tunable visible-light integrated optical phased array [5] will be presented (prior visible-light integrated optical phased arrays have been limited to passive demonstrations).

Next, the VIPER display will be demonstrated. First, a novel transparent 300-mm-wafer foundry platform on glass for visible-light integrated photonics will be presented [6]. Second, a novel large-scale passive VIPER display that generates a holographic image of a wire-frame cube using 1024 optical-phased-array-based pixels passively encoded to emit light with the appropriate amplitudes and phases will be discussed [6,7]. Third, a novel active VIPER display consisting of cascaded compact active optical-phased-array-based pixels will be shown [6].

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

2 LIQUID-CRYSTAL THEORY, INTEGRATION, AND FABRICATION

2.1 Introduction to Visible-Light Applications and Challenges

Integrated photonics systems at visible wavelengths have many wide-reaching potential applications, including image-projection systems [7-10], underwater optical communications [11], *in vivo* biological imaging [12], and optogenetics [13-15].

Traditionally, in integrated photonics, silicon is used as the waveguiding material, since it is the foundation of mature complementary metal-oxide-semiconductor (CMOS) fabrication processes, it enables high index contrast for compact devices, and it is relatively easy to modulate; however, silicon absorbs at visible wavelengths. As such, visible-light integrated systems are generally based on silicon-nitride waveguides, since silicon nitride has a low absorption coefficient within the visible spectrum and is CMOS compatible. Although silicon nitride has the benefit of transparency in the visible wavelength range, it also has a few material drawbacks. Specifically, silicon nitride has a low thermo-optic coefficient and does not exhibit significant electro-optic properties, which makes integrated phase tuning at visible wavelengths a challenge. Consequently, typical current silicon-nitride phase modulators based on heaters are on the order of several millimeters long and require high power consumption [16]. Furthermore, a conventional integrated visible-light amplitude modulator would consist of these heaters integrated into a Mach–Zehnder interferometer; consequently, a typical amplitude modulator would also be on the order of several millimeters long, which is unreasonable for many systems. Therefore, compact and low-power integrated silicon-nitride phase and amplitude tuning at visible wavelengths is a challenge and an essential advancement for the field.

2.2 Introduction to Integrated Liquid-Crystal-Based Devices

As a solution to the challenge of compact integrated visible-light modulation, nematic liquid crystal, with strong birefringence in the visible spectrum, can be integrated and used for phase and amplitude modulation. These liquid-crystal-based devices can achieve modulation with a smaller physical footprint and lower power consumption than conventional heater-based devices.

Integrated liquid-crystal-based devices, including slot-waveguide phase shifters [17], rectangular-waveguide phase shifters [18], and ring resonators [19-24], have been explored in the past. However, these liquid-crystal-device demonstrations have been largely limited to silicon waveguides and operate within the infrared wavelength range. Hence, a compact liquid-crystal-based phase shifter based in silicon nitride and operating in the visible wavelength range would be a beneficial advancement.

Furthermore, liquid-crystal amplitude modulation has been achieved by directly integrating these phase shifters in a Mach-Zehnder-interferometer [2,17]. However, this type of configuration has a large form factor, since it requires a splitter at both the input and output of the device, phase shifters in each arm of the interferometer, and transitions connecting the liquid-crystal region to the splitters. Hence, a more compact configuration for liquid-crystal-based amplitude modulation would be beneficial for many applications.

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

2.3 Liquid-Crystal Background Theory and Operation

Liquid crystals have many unique properties [25-30], which make them interesting for integrated visible-light devices. One such property is that liquid crystals have an anisotropic index of refraction. This anisotropy is due to the fact that a liquid-crystal molecule has an ordinary and an extraordinary axis, as shown in Figure 1a, and there is a strong birefringence between these two axes in the visible wavelength range. Light propagating in a direction parallel to the liquid-crystal director (shown in Figure 1a) sees the ordinary index of refraction, n_o, while light propagating perpendicular to the liquid-crystal director sees the extraordinary index of refraction,n_e.

The birefringent property of liquid crystals allows for the manipulation of propagating light; however, the orientation of the liquid-crystal director with respect to the propagation direction of the light must be known in order to achieve the desired result. A second unique property that allows for the director orientation to be known, is that, although liquid crystal is indeed a liquid, the molecules can be aligned in a crystalline fashion [31]. In a nematic liquid crystal, the molecules align in one dimension, in the direction parallel to the directors, as shown in Figure 1b. Furthermore, in the presence of an alignment structure, such as a rubbed polyimide layer, the liquid-crystal molecules mechanically anchor to the alignment layer. Hence, the alignment of the liquid-crystal directors can be intentionally set to a desired direction.

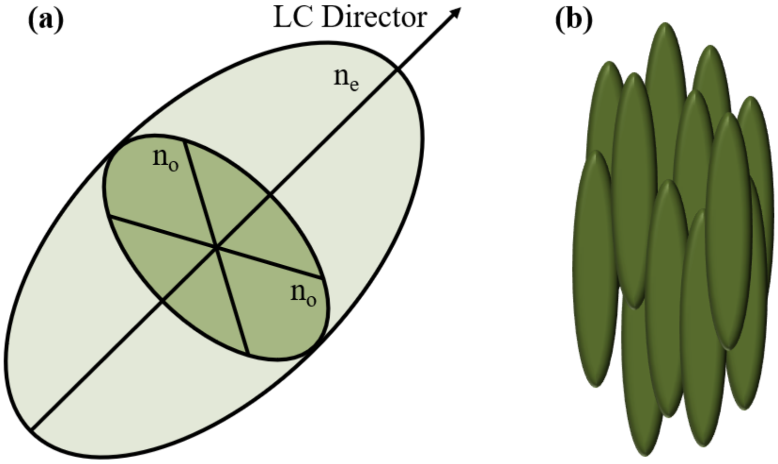

Figure 1: Ordinary and Extraordinary Axes of a Liquid-crystal Molecule (a) and (b) 1-dimensional Alignment of Nematic Liquid-crystal Molecules

Furthermore, not only can the liquid-crystal director be anchored in an initial direction via a mechanical alignment layer and the nematic property of the molecules themselves, but the molecules can be dynamically rotated. This rotation can be achieved via a third unique property: liquid-crystal molecules align to external electric fields. The molecules behave like electric dipoles, in the direction of the director, so the director aligns to the external electric field. Figure 2a shows the liquid-crystal molecules initially anchored via a mechanical alignment layer, when no external electric field is applied across the region. Figure 2b shows the liquid-crystal molecules beginning to align to an applied electric field, once the field strength is increased past a threshold to be great enough to overcome the mechanical anchoring strength. Figure 2c shows the liquid-crystal molecules fully aligned to an external electric field (note that the molecules close to the alignment layer remain more strongly anchored than molecules farther away). When the external electric field is removed, the molecules revert back to their initial anchored state. In this way, the liquid-crystal molecules can be dynamically rotated to take advantage of the birefringence and vary the index of refraction with respect to light propagating in a set direction.

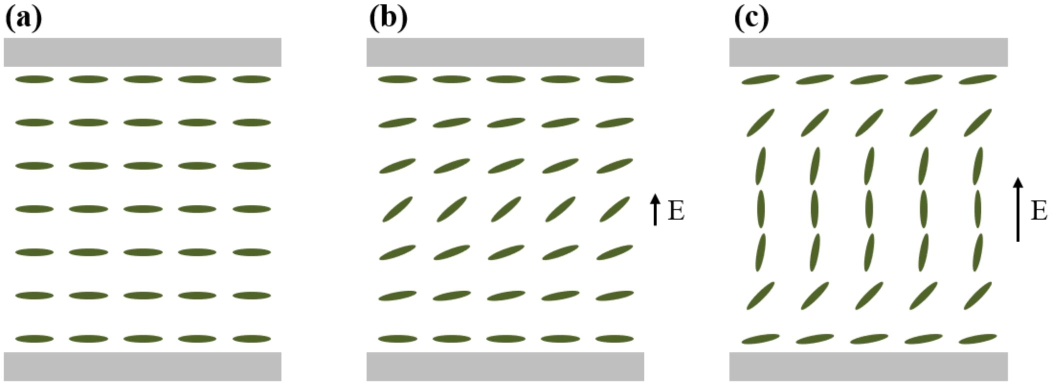

Figure 2: Liquid-crystal Molecule Alignment

No External Electric Field (a), (b) an External Electric Field Exceeding the Required Threshold to Overcome the Mechanical Anchoring Strength, and (c) an External Electric Field Strong Enough to Completely Rotate the Molecules

2.4 Liquid-Crystal Integration

The useful properties of liquid crystals can be leveraged in an integrated-photonics platform. In this platform, a silicon-nitride waveguide is recessed in a silicon-dioxide cladding, which confines and guides the light on the chip. As shown in Figure 3a, the fundamental mode in the waveguide is fairly well confined in the silicon nitride. Then, a trench is made in the silicon-dioxide cladding above the waveguide and the trench is filled with liquid crystal. Although the silicon-nitride waveguide is still recessed within a silicon-dioxide cladding, a liquid-crystal region can be deposited on top of the waveguide, as shown in Figure 3b-d. The mode in the waveguide is impacted by the liquid crystal and the mode is slightly pulled up into the liquid-crystal region. Furthermore, the mode in the waveguide varies based on the refractive index of the liquid-crystal region. The mode varies in two ways; the effective index of the mode changes and the mode shape gets pulled up farther into the liquid-crystal region as the effective index of the liquid crystal is increased, as shown in Figure 3b-d. The liquid crystal used here is 4-cyano-4'- pentylbiphenyl (5CB), which has an ordinary refractive index of 1.53 and an extraordinary refractive index of 1.7. However, because the mode becomes less confined as the liquid-crystal index increases, the loss also increases, so the refractive index is limited to 1.62 to minimize loss.

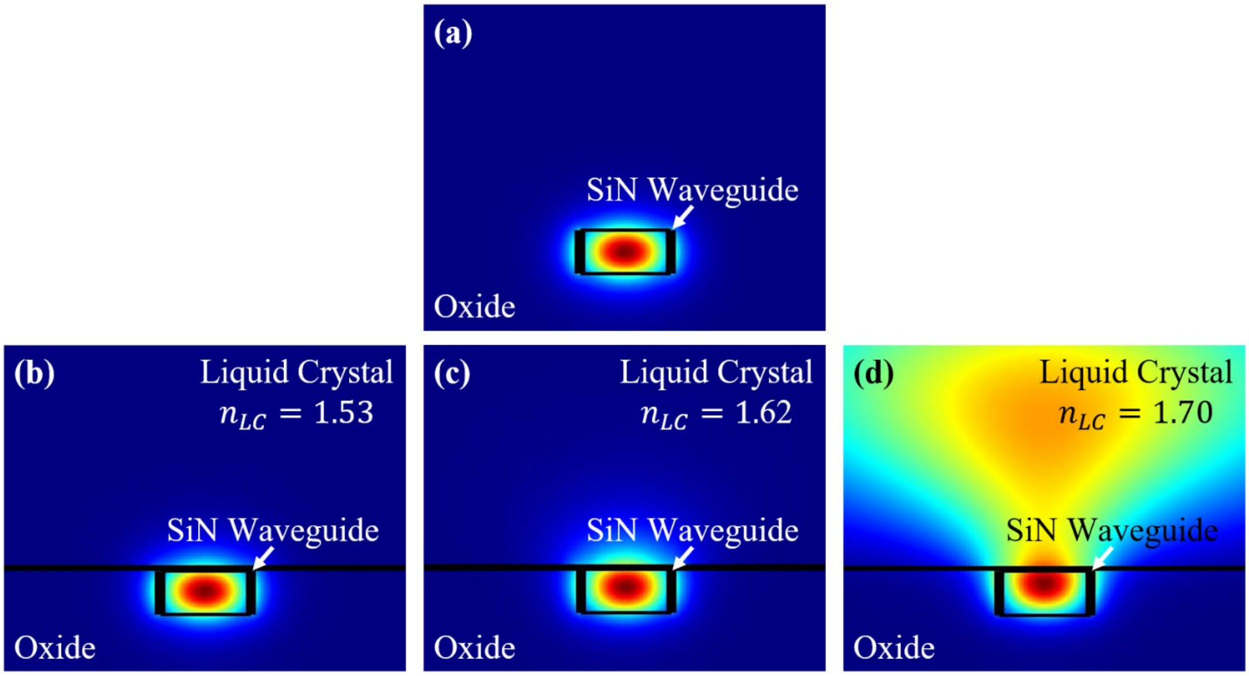

Figure 3: Fundamental Mode of a Silicon-nitride Waveguide Recessed within a Silicon-dioxide Cladding

With no liquid crystal (a) and (b) with a liquid-crystal region of refractive index n_{LC}*=1.53, (c)* n_{LC}*=1.62, and (d)* n_{LC}*=1.70*

As stated above, the liquid-crystal molecules align in the presence of an external electric field. In this integrated platform, an electric field is applied across the liquid-crystal region via metal electrodes on either side of the trench by applying a voltage to one electrode with respect to the other, as shown in Figure 4a. The electrodes run alongside the entire length of the liquid-crystal-filled trench, as shown in Figure 4b, to uniformly vary the refractive index throughout the entire region.

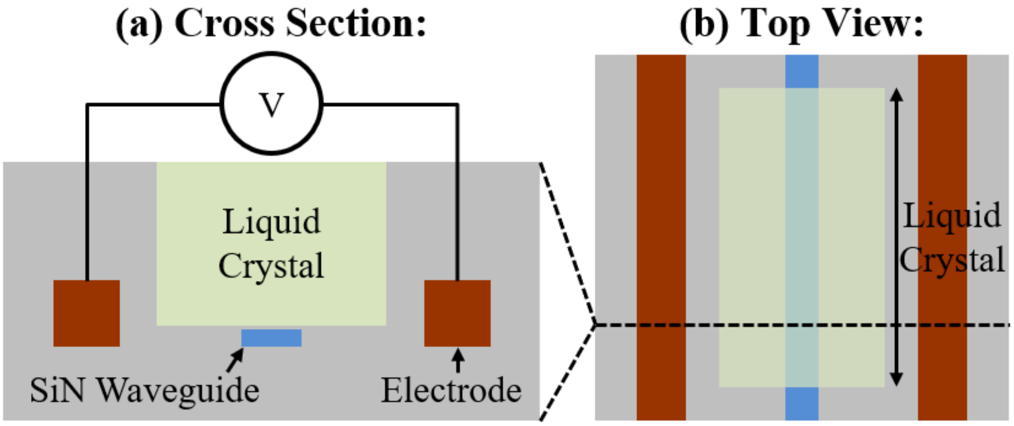

Figure 1:Cross Section (a) and (b) Top View of a Silicon-nitride Waveguide Below a Liquid-crystal-filled Trench with Electrodes on Either Side

2.5 Device Fabrication and Packaging

The devices are fabricated in a CMOS-compatible 300-mm wafer-scale process at the College of Nanoscale Science and Engineering (CNSE) at the State University of New York (SUNY) Polytechnic Institute in Albany, New York. A photograph of the fabricated wafer is shown in Figure 5.

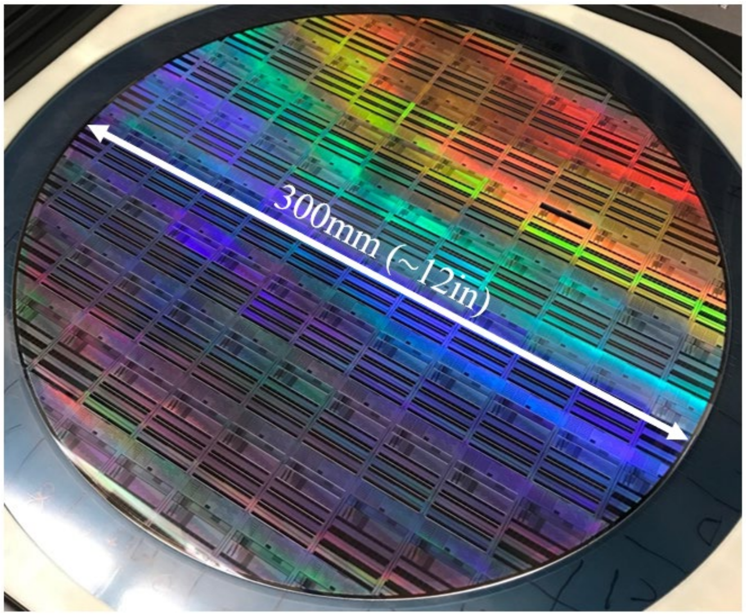

Figure 2: Photograph of a Photonic Wafer, Fabricated at the CNSE

Following this additional chip-scale fabrication and packaging is done in house at the Massachusetts Institute of Technology (MIT) fabrication facilities as back-end steps. An overview of the in-house fabrication and packaging process is shown in Figure 6. The initial cross section of the chip received from CNSE consists of a silicon-nitride (SiN) waveguide recessed within a silicon-dioxide cladding, with an empty trench above the waveguide (which is achieved via a second SiN etch stop layer), as well as electrodes on either side of the empty trench. A more detailed description of the in-house fabrication and packaging process is given below; however, a brief overview of the process consists of four steps, as shown in Figure 6. The first step is to perform a dry etch to bring the trench closer to the waveguide. The second step is to pattern an SU-8 photoresist spacer layer. The third step is to epoxy a glass chip with an alignment layer on top of the SU-8 spacer layer. The fourth step is to inject the liquid crystal (LC) into the formed cavity and seal the cavity with ultraviolet (UV)-cured epoxy. More details on all of these steps are provided in what follows below.

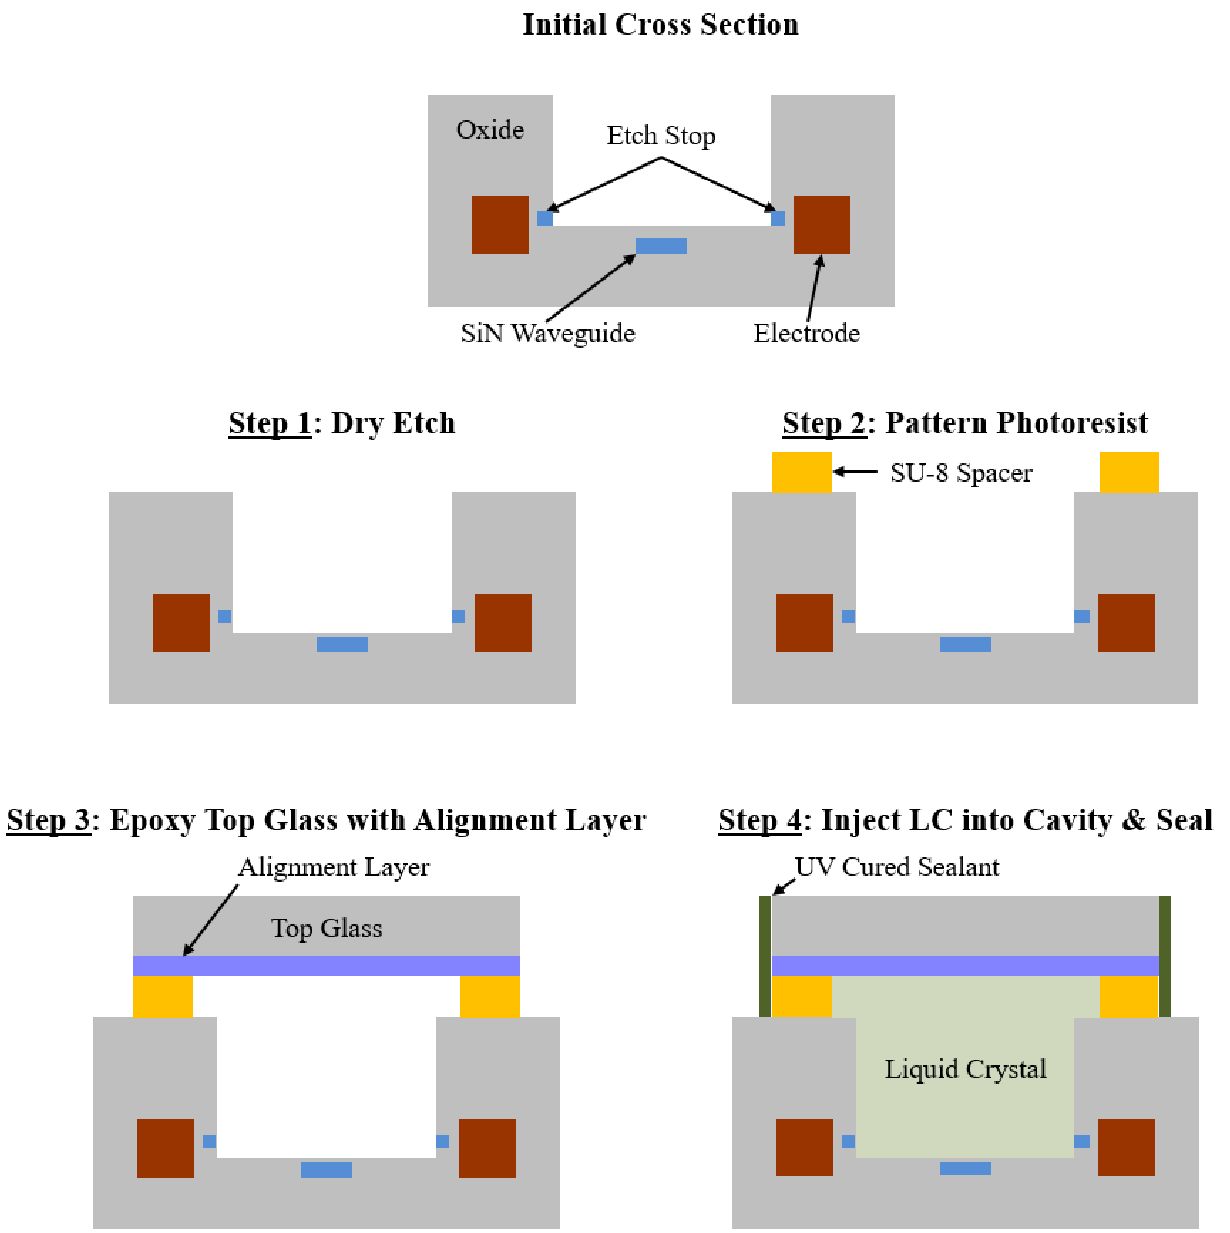

Figure 3: Cross-sectional Diagram of the Liquid-crystal Packaging Process
The initial cross section consists of the SiN waveguide, empty trench, and electrodes. Step 1: dry etch to bring trench closer to waveguide. Step 2: pattern SU-8 photoresist spacer layer. Step 3: epoxy top glass chip with alignment layer on top of spacer layer. Step 4: inject liquid crystal into cavity and seal with UV-cured epoxy

2.5.1 Reactive-Ion Etching

The first in-house fabrication step (step 1 in Figure 6) is to perform a reactive-ion etch, which is a type of dry etch [32,33]. This is done to reduce the thickness of the oxide layer in between the waveguide and the bottom of the trench. Initially, the trench is 60 nm above the SiN waveguide. The thinness of this oxide gap is limited in the wafer-scale fabrication by how close the two SiN layers (the waveguide layer and the etch stop layer) can be successfully patterned. Hence, this oxide gap is further reduced in-house on the chip scale via a reactive-ion etch. This is an important step, because lowering the bottom of the trench brings the liquid-crystal region closer to the top of the waveguide. This maximizes the interaction of the mode in the waveguide with the liquid crystal, allowing for more compact devices. (The significance of this heightened interaction is explained in greater detail in Sec. 3.1.2.)

The resulting thickness of the oxide gap is very important, because the device performance is significantly impacted by the gap thickness. Hence, it is imperative to precisely characterize the etch rate of the implemented process. This is achieved by etching four different chips with four incrementally increasing etch durations, then taking cross-section images to infer the resulting oxide gap thickness for each of the four chips, from which the etch rate can be calculated. However, due to typical wafer-scale fabrication variation, the initial oxide gap thickness can vary across the wafer. Hence, it is necessary to not only measure the resulting gap thickness after etching, but also the initial gap thickness before etching. Transmission electron microscopy (TEM) cross-section images are taken of the four chips in two locations (TEM images taken by CNSE). The initial oxide gap thickness is measured from a location on each chip that is not exposed to the reactive-ion etch, as highlighted in Figure 7. The initial oxide gap thickness varies across these four chips by > 10 nm, which demonstrates the importance of taking the initial gap thickness into account when calculating the etch rate.

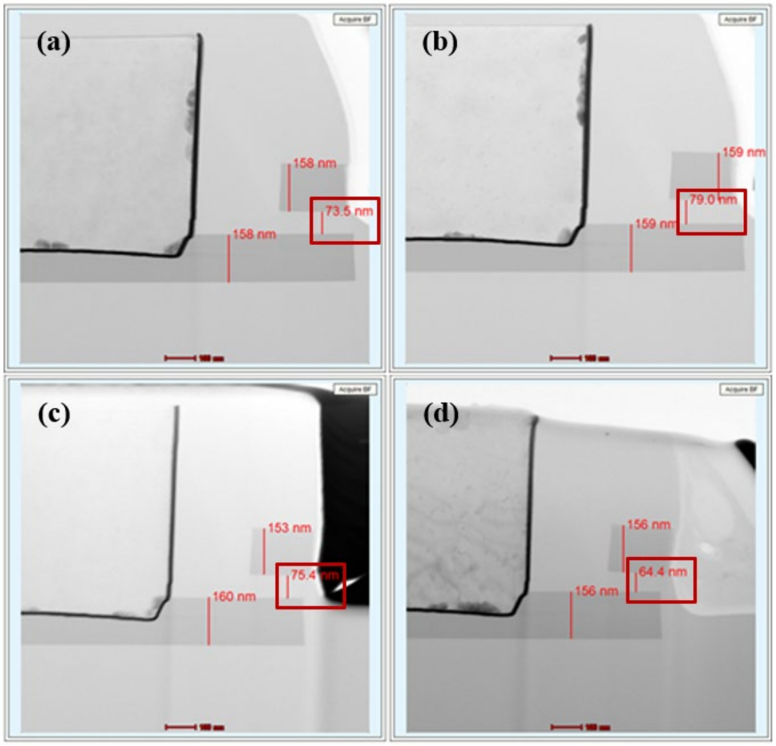

Figure 4: TEM Cross-section Images on a Section of the Chip that is Not Exposed to the Reactive-ion Etch, Highlighting Original Oxide Gap Thickness between the SiN Waveguide Layer and SiN Etch Stop Layer on Four Different Chips Prior to (a) Four Minutes, (b) Five Minutes, (c) Six Minutes, and (d) Seven Minutes of reactive-ion etching
TEM images taken by CNSE

Transmission electron microscopy cross-section images of the four chips are taken in a second location, at the site of a trench, to measure the resulting oxide thickness after etching, as shown in Figure 8. The chips are etched with a recipe using 3 standard cubic centimeters per minute (sccm) of O_2 and 30 sccm of CF_4 at 30 mT pressure with 100 W applied radio-frequency power. The four chips are etched for four, five, six, and seven minutes, respectively. The first two chips are only etched into the oxide gap layer, so they can be used to calculate the etch rate of oxide, while the last two chips are etched partially into the SiN waveguide, so they can be used to calculate the etch rate of SiN. The selectivity of this etch recipe is evident in Figure 8d, where the SiN waveguide is seen to be etched significantly more than the oxide cladding surrounding it. Taking into account the initial gap thicknesses of the four chips (from Figure 7), the etch rate of oxide is approximately 25 nm/min and the etch rate of SiN is approximately 54 nm/min. This etch rate characterization allows for precise chip processing.

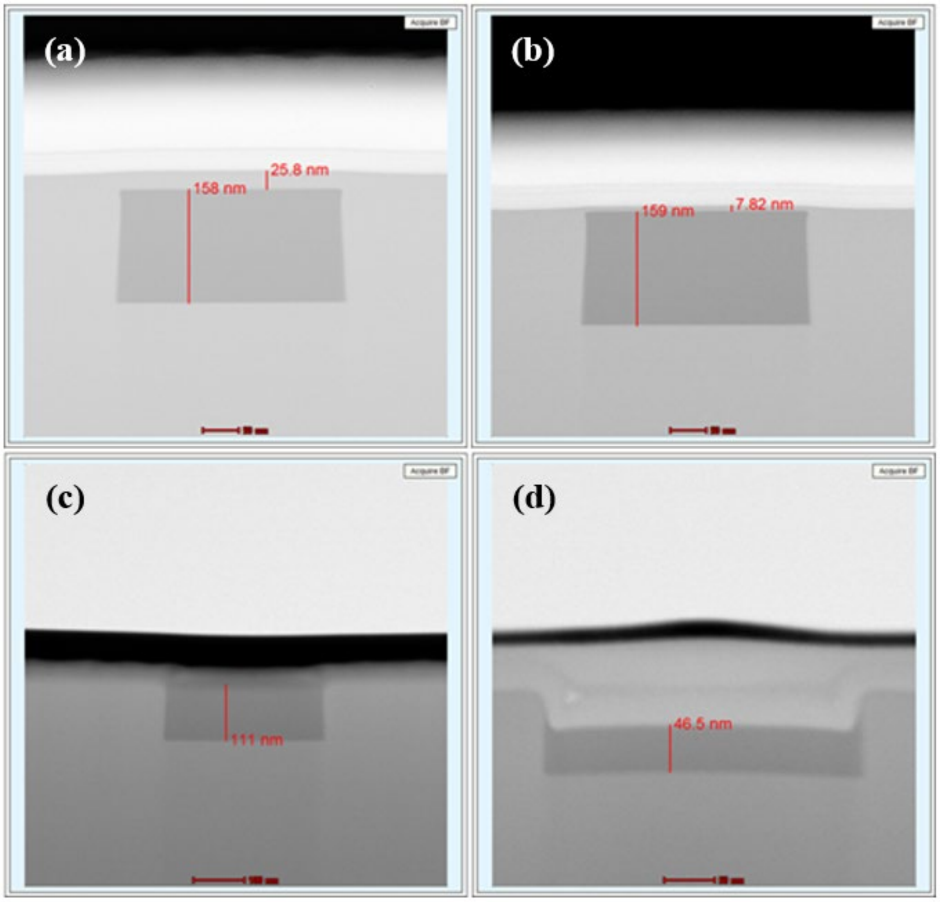

Figure 5: TEM Cross-section Images of SiN Waveguide within Oxide Cladding after (a) Four Minutes, (b) Five Minutes, (c) Six Minutes, and (d) Seven Minutes of Reactive-ion Etching

TEM images taken by CNSE

2.5.2 Waveguide Reveal

The reactive-ion etch allows for the oxide gap between the waveguide and the liquid-crystal trench to be minimized. However, a step which further increases the interaction between the waveguide mode and the liquid crystal is to reveal the waveguide to the trench, as shown in Figure 9. The significance of revealing the waveguide to the liquid-crystal trench is described in Sec. 3.1.4.

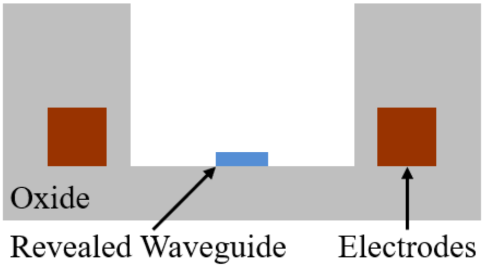

Figure 6: Diagram of the Initial Cross Section with the Waveguide Revealed to the Trench

Initially, this waveguide reveal process was intended to be performed in house at the MIT facilities via a wet-etch process using hydrofluoric acid. However, it was discovered that, since the wet etch is isotropic (meaning it etches at the same rate in every direction, rather than just straight down) [34,35], the waveguide is reshaped by the wet etch from its desired rectangle profile to a trapezoidal shape. Furthermore, the hydrofluoric-acid etch was discovered to not be selective enough to the silicon dioxide over the silicon nitride, which results in a reduction in the waveguide height. The approximate resulting waveguide shape and size reduction after the wet etch are shown in Figure 10.

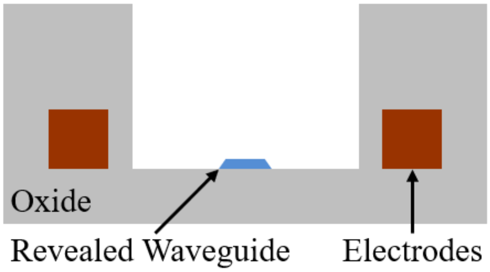

Figure 7: Diagram of the Cross Section of the Revealed Waveguide After the Hydrofluoric-acid Wet Etch, Showing the Undesired Trapezoidal Waveguide Shape and Reduced Dimensions

As a result of the in-house wet etch proving ineffective at revealing the waveguide, a waveguide reveal process was developed by CNSE. This process successfully reveals the waveguide to the trench, while maintaining the integrity of the original waveguide shape and dimensions. A scanning-electron-microscopy (SEM) in-line image of the waveguide revealed to the trench is shown in Figure 11 (SEM image taken by CNSE). Overall, there are two different CNSE fabrication processes. One process results in an embedded waveguide with a silicon-dioxide gap between the waveguide and the trench, which is then thinned down via an in-house reactive-ion etch at MIT, and a second process results in a fully revealed waveguide which does not need to go through the in-house etch.

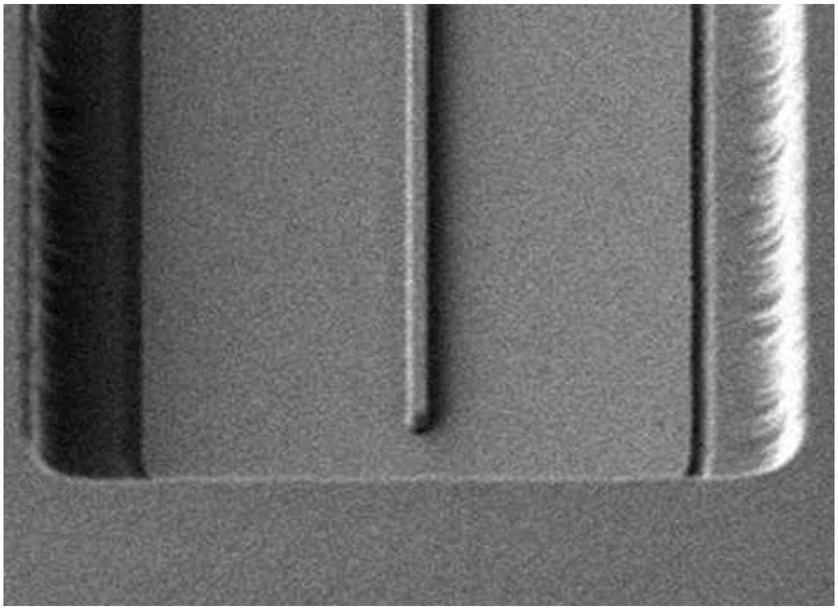

Figure 8: SEM Image of the Waveguide Revealed in the Trench
SEM image taken by CNSE

2.5.3 Photolithography

The second in-house fabrication step (step 2 in Figure 6) is to pattern an SU-8 photoresist layer on top of the photonic chip. This photoresist layer acts as a spacer to offset the top glass chip from the surface of the photonic chip, such that a cavity is formed into which the liquid crystal will be injected. The SU-8 is not patterned as a continuous strip around the chip; instead, there is one input gap at one end of the chip and two output gaps at the other end of the chip, as shown in Figure 12a.

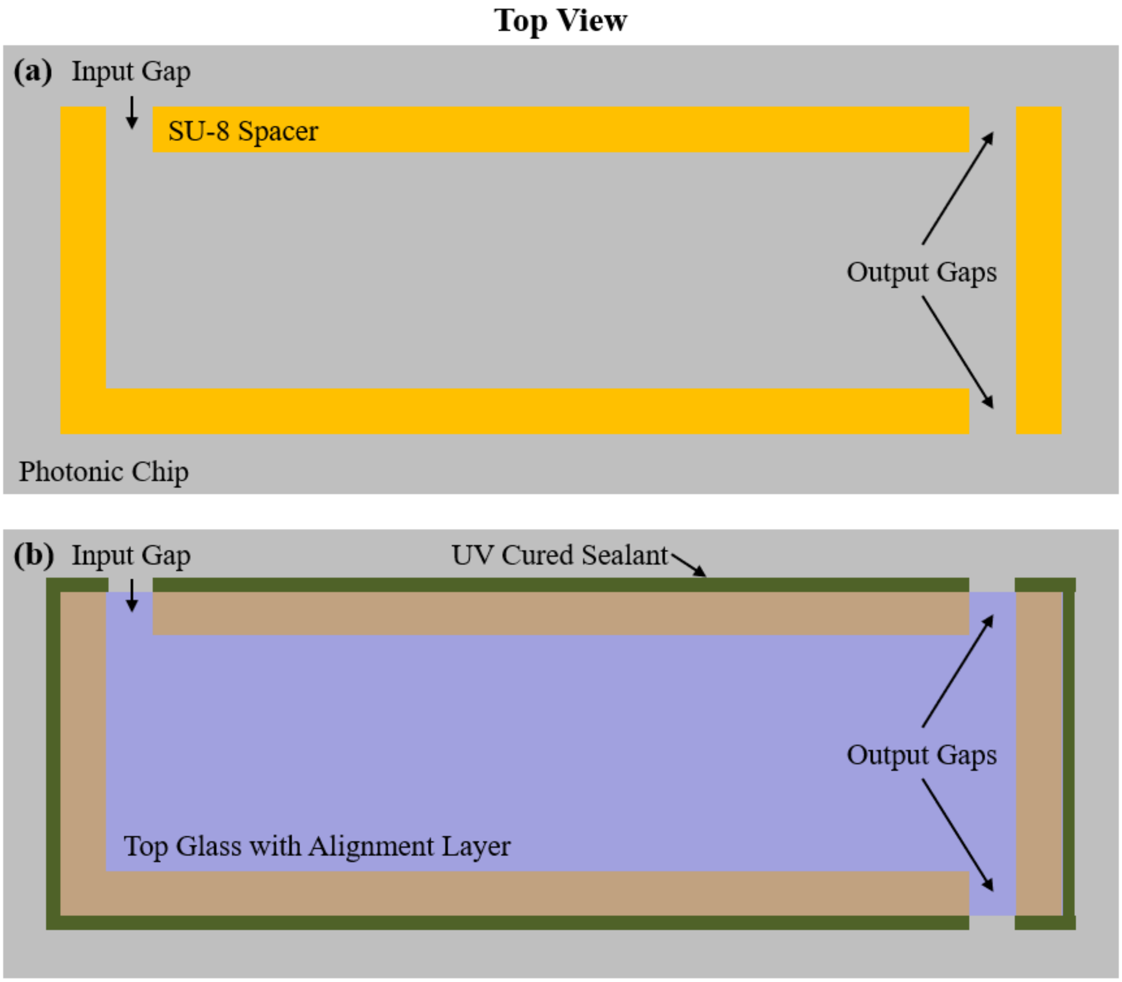

Figure 9: Top View of Packaging Process

Showing input and output gaps in the SU-8 spacer layer to allow for liquid-crystal injection and uniform filling (a) and (b) showing glass alignment chip placement and UV-cured epoxy sealant

The photolithography process used to pattern the photoresist spacer layer varies based on the type of photoresist used, the substrate onto which it is patterned, the type of mask used, and the specific fabrication tools used during the process. Hence, the proper photolithography recipe had to be developed to ensure a precise, uniform, and repeatable process. SU-8 photoresist is used because it is permanent and the layer thickness can be controlled and varied fairly easily. The photolithography process for SU-8 photoresist consists of coating, soft bake, UV exposure, post-exposure bake, and development.

The coating speed and duration determines the thickness of the photoresist layer. The desired photoresist thickness is 600 nm, so SU-8 2000.5 photoresist is coated at 2000 rpm for 30 sec, with a ramp up of 10 sec from an initial speed of 500 rpm. These coating specifications were taken from the SU-8 2000 data sheet. To achieve consistent results, the soft bake duration was significantly increased from the recommended duration, to 4 min at 100°C, via a hot plate with a silicon wafer underneath the chip to ensure uniform heat distribution.

Since SU-8 is a negative photoresist, the UV exposure step causes molecular cross linking and, hence, the areas that are exposed will remain in tact after the chip is developed. The exposure duration depends heavily on the type of mask used and the output energy of the tool being used. The exposure duration had to be varied until the proper conditions were found. As shown in Figure 13a, if the chip is under exposed, the resulting photoresist layer will not be of uniform thickness and the photoresist will be undesirably removed from some locations, resulting in improper photoresist patterning and structures. As shown in Figure 13c, if the chip is over exposed, the resist will not be removed properly in regions where it should be removed completely, which may also result in the photoresist structures having non-perpendicular walls or edges. In contrast, a properly exposed chip will result in high-resolution structures with nearly-perpendicular walls and uniform thickness, as shown in Figure 13b. Similar to the soft bake, the post-exposure-bake duration was increased as well to 4 min at 100°C, to ensure consistency.

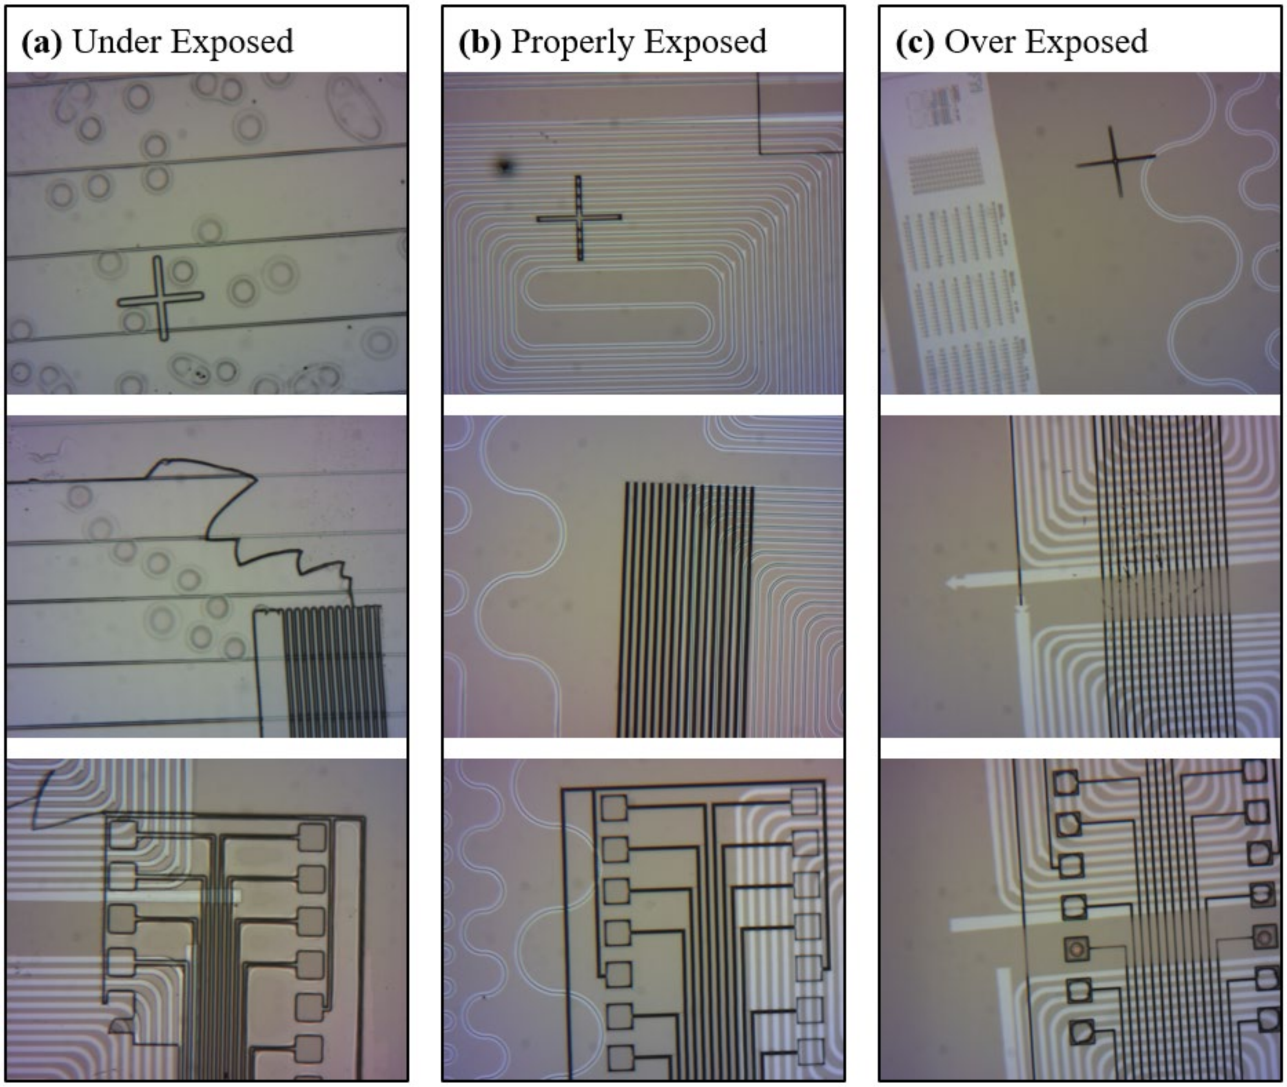

Figure 10: Resulting Resist Pattern Due to Chip Being Under Exposed (a) and (b) Properly Exposed, and (c) Over Exposed

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

During the development step, the photoresist that was not cross linked during exposure will be dissolved by the solvent. If the chip is overdeveloped, even the photoresist that was cross linked during the exposure step will begin to dissolve, as shown in Figure 14. This causes the resulting photoresist layer to become damaged and the thickness becomes very inconsistent. For development using the solvent propylene glycol monomethyl ether acetate, 1 min was found to achieve consistent results.

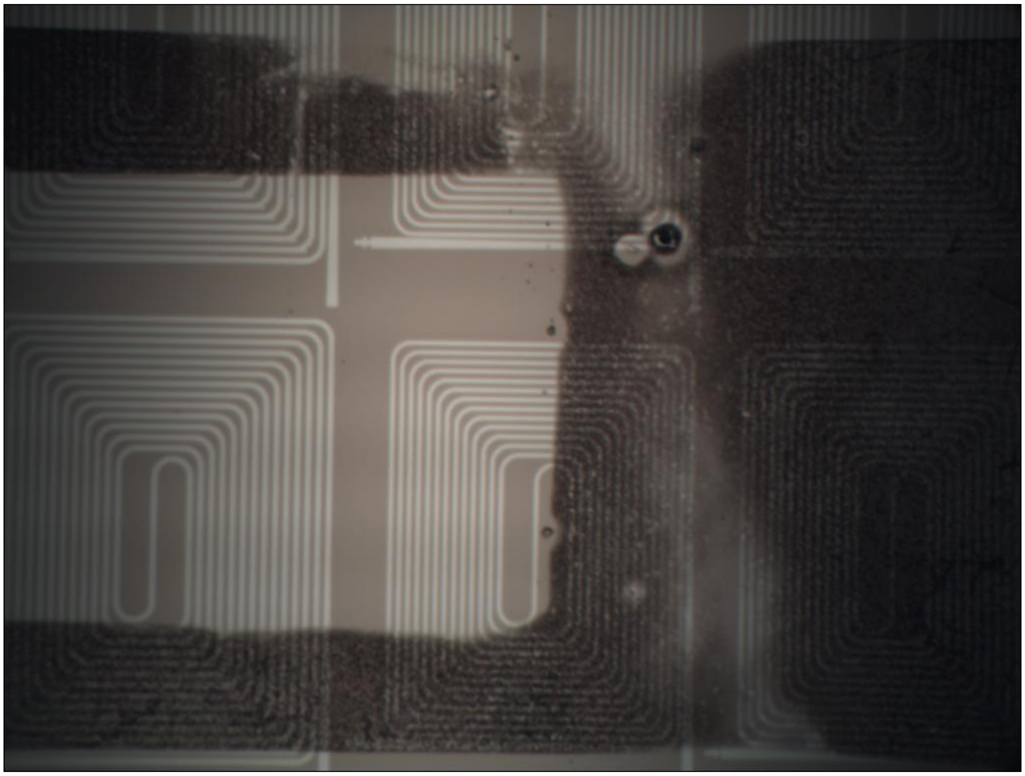

Figure 11: Resulting Photoresist Pattern Due to Over-development

The above photolithography recipe targeted a final photoresist thickness of 600 nm. The thickness of the resulting photoresist was measured using a Dektak stylus metrology measurement. The thickness across a 400-µm-wide photoresist structure was found to be approximately 620 nm, as shown in Figure 15.

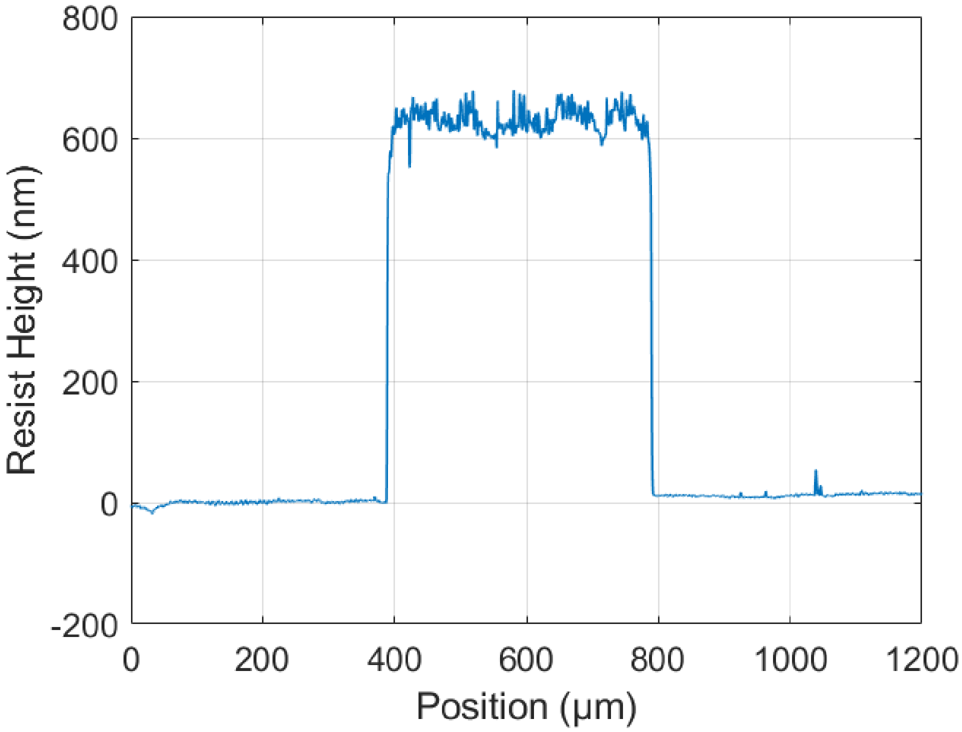

Figure 12: Dektak Stylus Metrology Measurement of Resulting SU-8 Resist Thickness

2.5.4 Liquid-Crystal Packaging

The final set of in-house fabrication steps (steps 3 and 4 in Figure 6) is to epoxy a glass chip on top of the SU-8 spacer layer, inject the liquid crystal into the formed cavity, and seal the cavity. A glass chip, which has an organic polyimide alignment layer on the bottom of it is placed on top of the SU-8 photoresist spacer layer, as shown in Figure 16a. The glass chip is obtained from Instec and is diced in-house at MIT to the appropriate size. The polyimide alignment layer mechanically anchors the liquid-crystal molecules when no external electric field is applied across the liquid-crystal region. Once the glass alignment chip is properly placed and aligned on top of the SU-8 photoresist spacer layer, the two chips are carefully transferred to a clamp, as shown in Figure 16b. The clamp applies light pressure and ensures that all edges of the two chips are uniformly touching. Then, Norland optical adhesive, a UV-cured epoxy, is used to epoxy the alignment chip to the photonic chip, as shown in Figure 12b. The input and output gaps in the SU-8 spacer layer remain open with no epoxy covering them, to allow for the liquid crystal to be injected into the cavity.

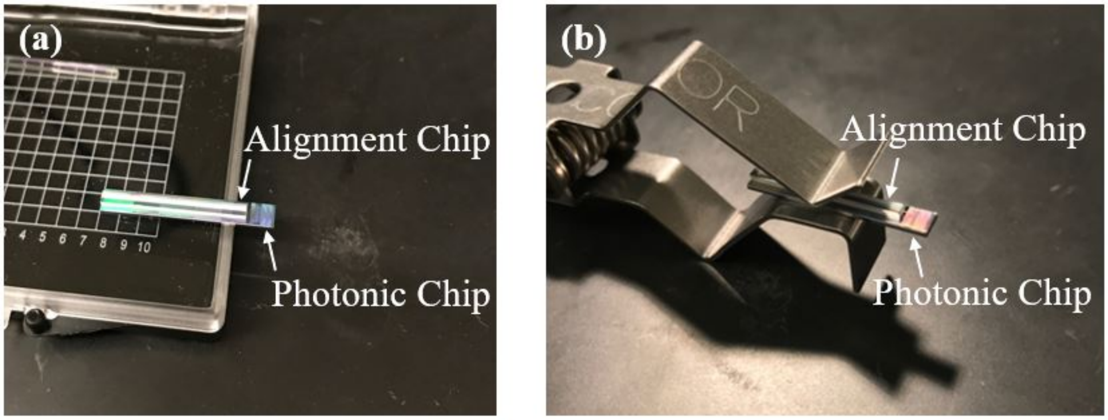

Figure 13: Photograph of a Glass Chip
Containing alignment layer on top of a photonic chip (a) during chip alignment process and (b) transferred to a clamp to provide light pressure during the epoxy process

At this point, a cavity has been formed between the photonic chip and the glass alignment chip. 5CB liquid crystal is then injected into this formed cavity through the input gap on one end of the chip. The chip is placed in a vacuum during this process to ensure to air bubbles remain in the cavity or trenches. The liquid crystal flows into the cavity via capillary action and, when put into the vacuum, the air gets removed from the cavity via the two output gaps on the other end of the chip. Once the liquid crystal fills up the entire cavity, the input and output gaps are sealed with the UV-cured epoxy.

2.6 Conclusion

In this chapter, liquid-crystal theory, including birefringence, nematic molecular alignment, anchoring via a mechanical alignment layer, and molecular rotation via an external electric field were introduced. As discussed, these properties can be leveraged in an integrated-photonics platform via integrating a liquid-crystal-filled trench above a silicon-nitride waveguide to take advantage of the birefringence and by placing electrodes on either side of this trench to actively tune the refractive index of the liquid-crystal region. Furthermore, the wafer-scale device fabrication at CNSE and the in-house chip-scale back-end fabrication and liquid-crystal packaging were described.

3 LIQUID-CRYSTAL PHASE MODULATOR

In this chapter, an integrated visible-light liquid-crystal phase modulator will be introduced. As discussed in Sec. 2.1, visible-light phase modulation is difficult to achieve, because silicon nitride has a low thermo-optic coefficient and insignificant electro-optic properties. Conventional heater-based silicon-nitride phase modulators are very long and require a high amount of power to operate [16], which is impractical for many applications.

The developed phase-modulator device leverages the birefringence of liquid crystal to vary the effective refractive index of the waveguide mode to induce a phase shift. Multiple design parameters, which affect the performance of the device, will be discussed. Finally, the experimental setup used to characterize the device and measured results will be shown.

3.1 Liquid-Crystal Phase-Shifter Theory and Design

3.1.1 Device Structure and Operation

The birefringent property of liquid crystal, introduced in Sec. 2.3, can be leveraged to design an integrated visible-light phase modulator. The final cross section of the device, after the fabrication and packaging process outlined in Sec. 2.5, is shown in Figure 1. The device consists of a SiN waveguide recessed within a silicon-dioxide cladding, a liquid-crystal-filled trench above the waveguide, and metal electrodes on either side of the trench. The operation of the liquid-crystal phase modulator is based on the fact that as the liquid-crystal refractive index increases and the waveguide mode gets pulled up into the liquid-crystal region, as shown in Figure 3b-d, the effective refractive index of the waveguide mode changes accordingly.

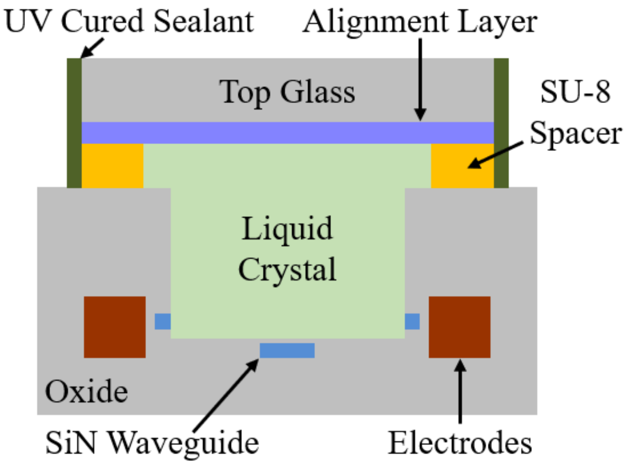

Figure 17: Final Liquid-crystal phase-modulator Device Cross Section

At the operating wavelength of 632.8 nm, the silicon nitride used in the developed platform has a refractive index of 1.995, silicon dioxide has a refractive index of 1.46, and 5CB liquid crystal has a refractive index varying from 1.53 to 1.70. However, this entire liquid-crystal refractive index range is not used, because, as mentioned in Sec. 2.4, the farther the waveguide mode is pulled up into the liquid-crystal region, the greater the loss. The loss as a function of liquid-crystal refractive index is shown in Figure 18 and, consequently, the liquid-crystal refractive index is limited to 1.62 to maintain low loss. Note that this simulated loss is only taking into account the dissipation into the electrodes. It does not take into account material loss in the waveguide, cladding, and liquid crystal or surface roughness of the waveguide.

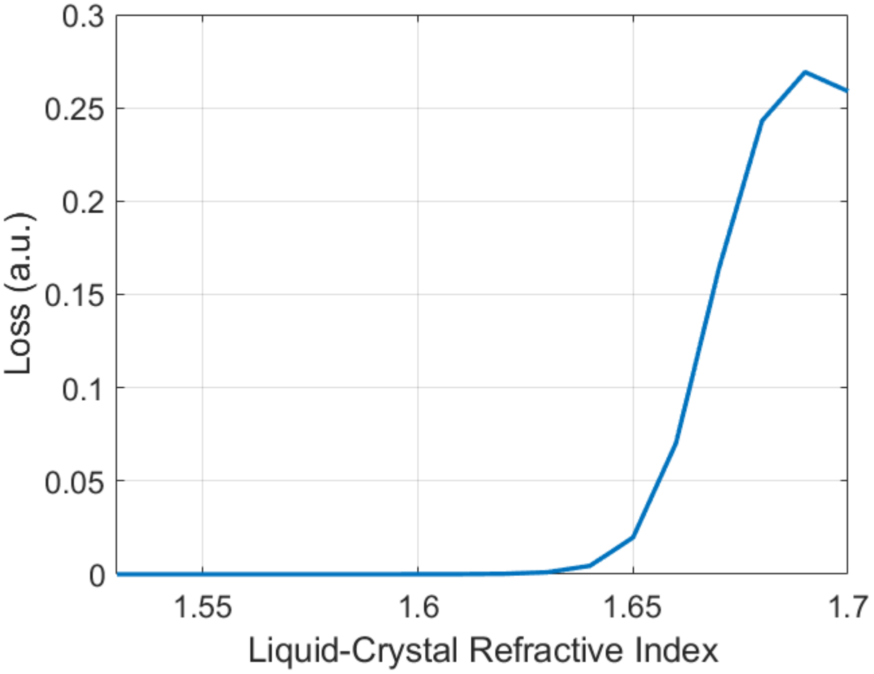

Figure 18: Simulated Loss Versus Liquid-crystal Refractive Index, Showing the Significance of Limiting the Implemented Liquid-crystal Refractive-index Range to a Maximum Value of 1.62

Now, with the liquid-crystal refractive-index range set to minimize loss, the effective refractive index of the fundamental mode in the waveguide can be simulated as the liquid-crystal index is varied, as shown in Figure 19a. As the liquid-crystal index increases and the mode is pulled up farther into the liquid-crystal region, the effective index increases to account for more of the mode propagating through the liquid crystal, which has a higher refractive index than the surrounding silicon-dioxide cladding.

As the effective index changes, the propagation constant changes, so the light propagates through the waveguide with a different speed, depending on the liquid-crystal refractive index. This variable propagation speed means that, for a length of liquid-crystal region, there will be a phase shift as the index of the region changes. A longer liquid-crystal region enables more phase shift, because there is more length for this change in propagation constant to have an impact on the phase at the end of the liquid-crystal region. The resulting phase shift depends on this liquid-crystal-region length, as well as the maximum change in effective index, which can be found from Figure 3.3a. Hence, the equation for phase shift is

$$\Delta\phi = \frac{2\pi L \Delta n_{eff}}{\lambda_o}$$

where L is the length of the liquid-crystal region, Δn_{eff} is the change in effective refractive index, and λ_o is the free-space operating wavelength. The phase shift versus length can be seen in Figure 3.3b.

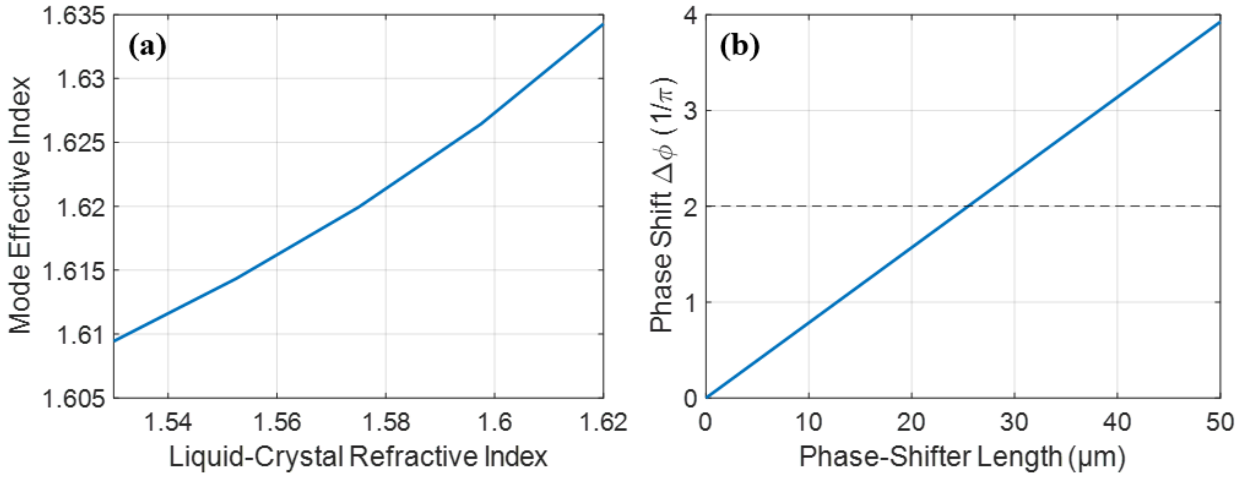

Figure 19: Change in Fundamental-waveguide-mode Effective Index of Refraction as the Liquid-crystal Refractive Index is Tuned (a) and (b) Resulting Phase Shift Due to a Varying Length of Liquid-crystal Region

3.1.2 Impact of Gap Separating Waveguide and Liquid-Crystal Region

There are many variables that affect how the waveguide mode effective index is impacted by the liquid-crystal region and hence, the amount of phase shift that is possible for a given shifter length. One such variable that impacts the device performance is the thickness of the layer of silicon dioxide in between the waveguide and the bottom of the liquid-crystal-filled trench, as mentioned in Sec. 2.5. The devices are fabricated with a 60-nm-thick gap layer between the top of the waveguide and the bottom of the liquid-crystal trench, due to fabrication limitations. This silicon-dioxide gap needs to be minimized in order to maximize the interaction between the waveguide mode and the liquid crystal. If the gap thickness is left at its original height of 60 nm, only a small portion of the mode is in the liquid-crystal region, as shown in Figure 20a. However, if the gap thickness is reduced via a reactive-ion etch described in Sec. 2.5.1, the mode in the waveguide reaches farther into the liquid-crystal region, as shown in Fgure 20b, which allows the effective index of the waveguide mode to be more significantly impacted by the liquid crystal.

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

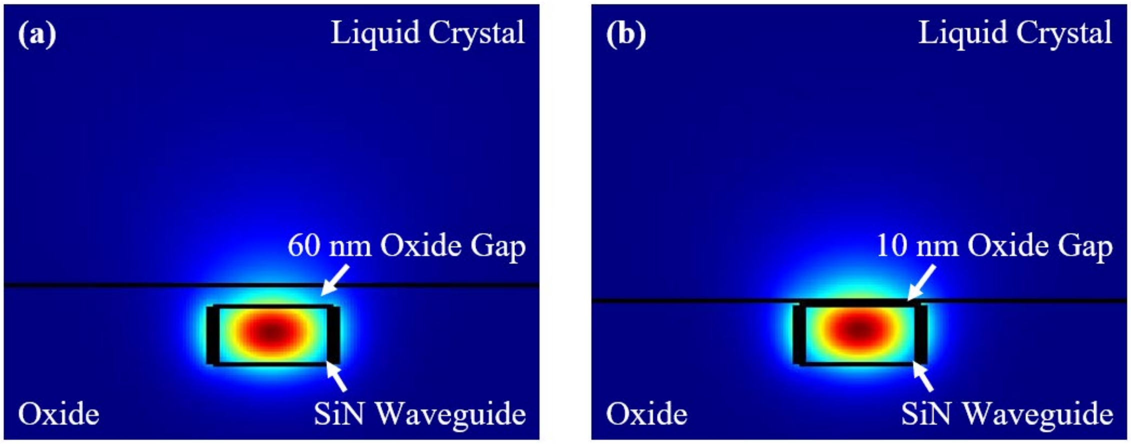

Figure 20: Fundamental Mode with (a) Initial 60-nm-thick Gap Layer and (b) Gap Thickness Etched Down to 10 nm

Note that the minimum gap thickness is limited to 10 nm, rather than 0 nm, even though completely removing the gap would result in even better performance. This is because, as discussed in Sec. 2.5.1, the reactive-ion etch used to minimize the gap is quite selective to silicon nitride (meaning it etches silicon nitride at a faster rate than silicon dioxide), so, if the target is to completely remove the silicon-dioxide gap, it is very possible that the top of the waveguide would be undesirably etched down as well. This is a viable possibility, because although the etch rate was calibrated, the initial gap thickness varies chip to chip, across the wafer, due to fabrication tolerances, also discussed in Sec. 2.5.1. Consequently, it would be very difficult to consistently completely remove the silicon-dioxide gap layer. Hence, the minimum target gap thickness is limited to 10 nm.

If the silicon-dioxide gap thickness is reduced and the waveguide mode is pulled up farther into the liquid-crystal region, the effective index of the waveguide mode will be affected more drastically by a change in the liquid-crystal refractive index. This means that, for a thinner gap, there is a greater Δn_{eff}, as shown in Figure 21a, which results in a greater phase shift for a set liquid-crystal shifter length, as shown in Figure 21b. This is critical, because the necessary shifter length to get 2π phase shift can be reduced. For example, for a device with an original gap thickness of 60 nm, a 38-µm-long shifter length is necessary to achieve 2π phase shift, whereas, for a device with an etched gap thickness of 10 nm, a 25.5-µm-long shifter length is needed to achieve 2π phase shift. Overall, reducing the silicon-dioxide gap thickness allows for a more compact device.

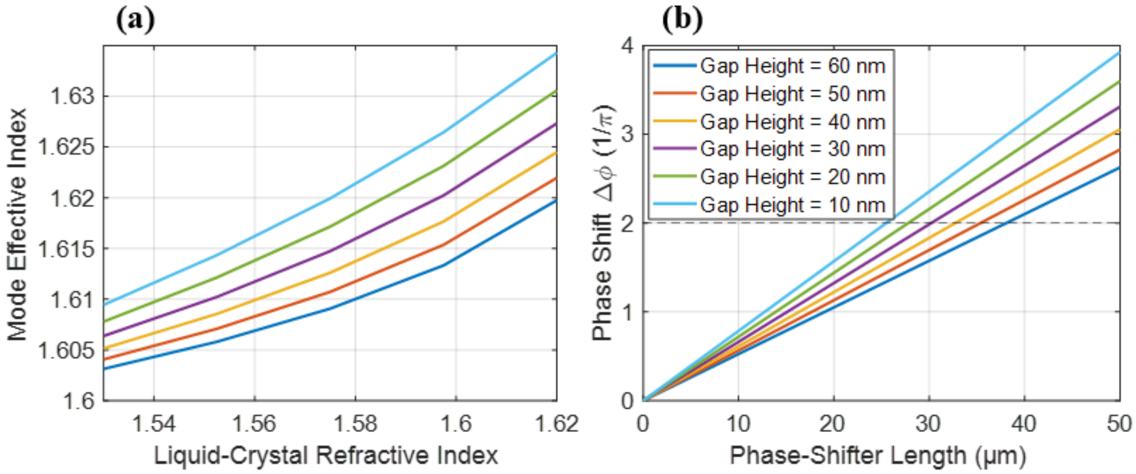

Figure 21: Waveguide Mode Effective Refractive Index as a Function of Liquid-crystal Index (a) and (b) Phase Shift as a Function of Shifter Length for Various Silicon-dioxide Gap Thicknesses

3.1.3 Impact of Waveguide Width

Another aspect which affects how much the waveguide mode is pulled up into the liquid-crystal region is the width of the waveguide. In general, for a simple waveguide recessed within a cladding, the mode is more confined in a wide waveguide and less confined in a narrow waveguide. This phenomenon can be leveraged in this device to further increase the interaction of the waveguide mode with the liquid crystal. Because the mode is less confined for a narrow waveguide, the mode gets pulled up farther into the liquid-crystal region, as shown in Figure 22. This results in a larger Δn_{eff} for a narrower waveguide, as shown in Figure 3.7a, and hence more phase shift for a given liquid-crystal length, as shown inFigure 23. For example, a 420-nm-wide waveguide requires a 34-μm-long liquid-crystal region to achieve 2π phase shift, while a 320-nm-wide waveguide requires a 25.5-μm -long liquid-crystal region to achieve 2π phase shift.

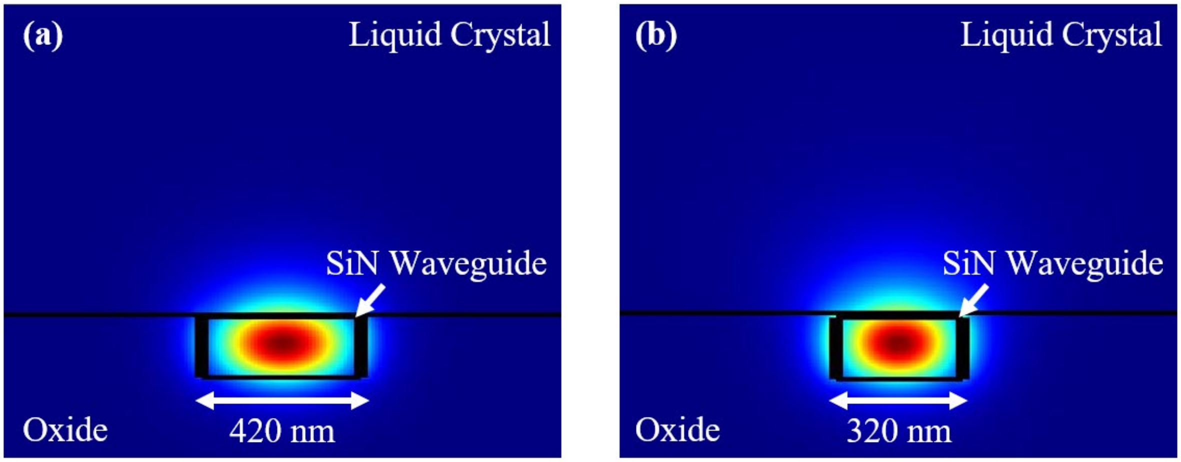

Figure 22: Fundamental Mode for a (a) 420-nm-wide Waveguide and (b) 320-nm-wide Waveguide

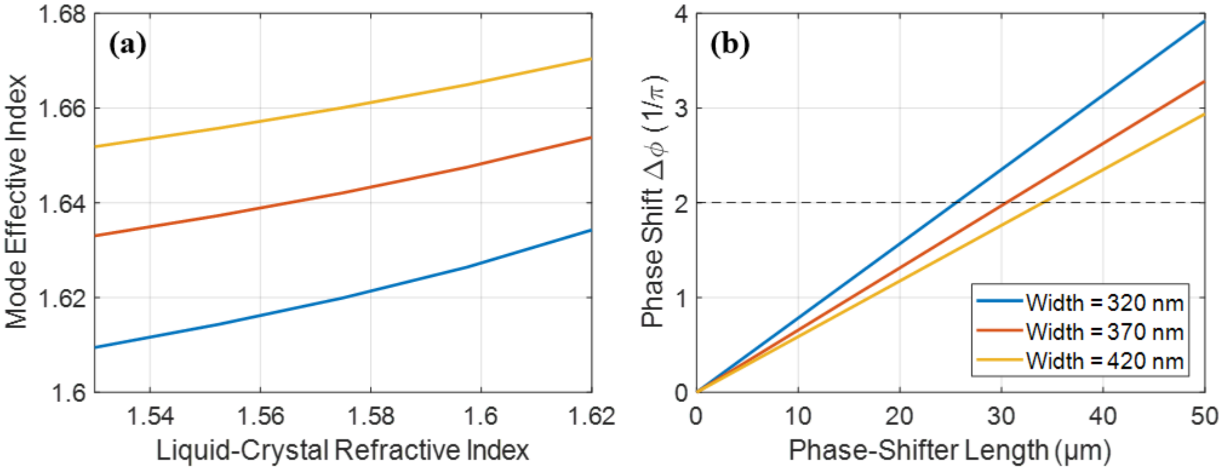

Figure 23: Waveguide Mode Effective Refractive Index as a Function of Liquid-crystal Index (a) and (b) Phase Shift as a Function of Shifter Length for Various Waveguide Widths

Although a narrower waveguide allows for more efficient phase-shift capabilities, it also results in more loss. The less-confined mode of the narrow waveguide is what allows for more interaction with the liquid crystal; however, it is also the cause of increased loss. Nevertheless, as shown in Figure 24, the increase in loss is not very substantial. Note that the simulated loss only depends on absorption into the electrodes. It does not take into account material loss in the waveguide, cladding, and liquid crystal or waveguide roughness. Overall, there is a trade off between the amount of phase shift and loss that must be balanced, and the optimal solution likely depends on the application.

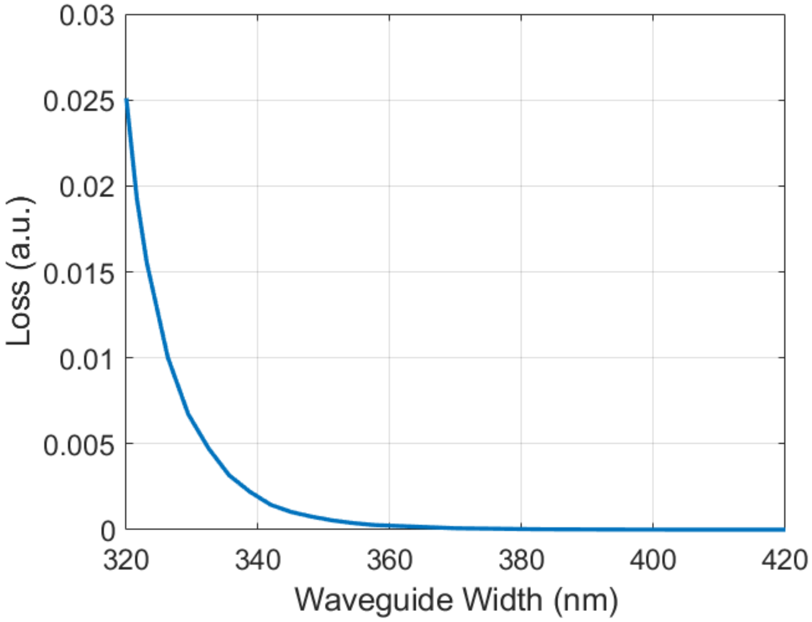

Figure 24: Loss as a Function of Waveguide Width

3.1.4 Impact of Waveguide Reveal

Another way to further improve the performance of the phase modulator is to reveal the waveguide to the liquid-crystal region. The liquid-crystal-filled trench can be etched down all the way to the bottom of the waveguide, to reveal the top and sides of the waveguide to the liquid crystal, as shown in Figure 25. Now, the waveguide is exposed to the liquid crystal on three sides, rather than just on top. This means that more of the waveguide mode will interact with the liquid-crystal region, as shown in Figure 26.

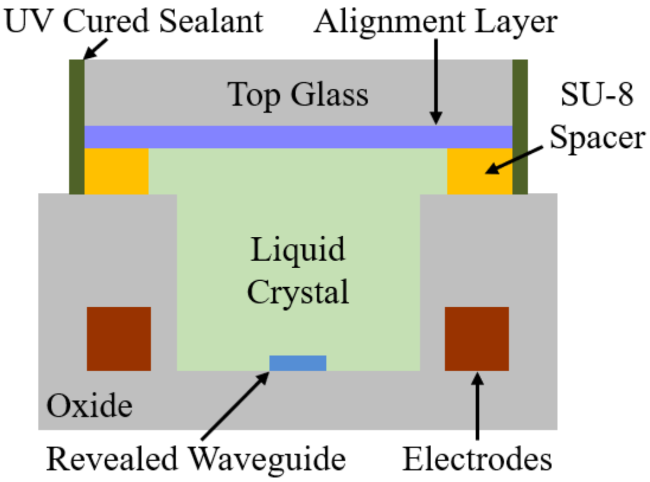

Figure 25: Liquid-crystal Phase-modulator Device Cross Section with Waveguide Revealed to the Liquid-crystal Region

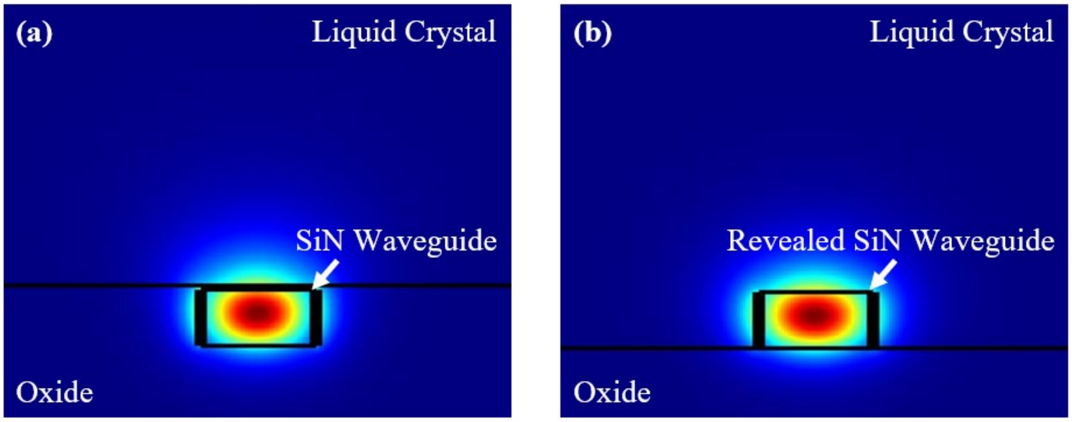

Figure 26: Fundamental Mode of the (a) Original Embedded Waveguide and (b) Revealed Waveguide

Similar to the parameters discussed above, when the waveguide mode is situated more within the liquid-crystal region, the effective refractive index of the mode will be impacted more by the variation in liquid-crystal index. As shown in Figure 27a, the effective refractive index of the waveguide mode is increased, because the liquid crystal has a higher refractive index than the silicon-dioxide cladding; however, more importantly, the range of the effective index, Δn_{eff}, increases as well, which can be seen through the greater slope of the waveguide mode effective index versus liquid-crystal index for the revealed waveguide. This greater range in the waveguide mode effective index results in greater possible phase shift for a set shifter length as shown in Figure 27b. For example, the original embedded 320-nm-wide waveguide requires a 25.5-µm-long liquid-crystal region to achieve 2π phase shift, while the revealed 320-nm-wide waveguide requires only a 16.6-µm-long liquid-crystal region to achieve 2π phase shift. Hence, the revealed waveguide allows for a more compact phase shifter. Additionally, with the revealed waveguide, the same trend is seen where a narrower waveguide can achieve more phase shift, as shown in Figure 27b.

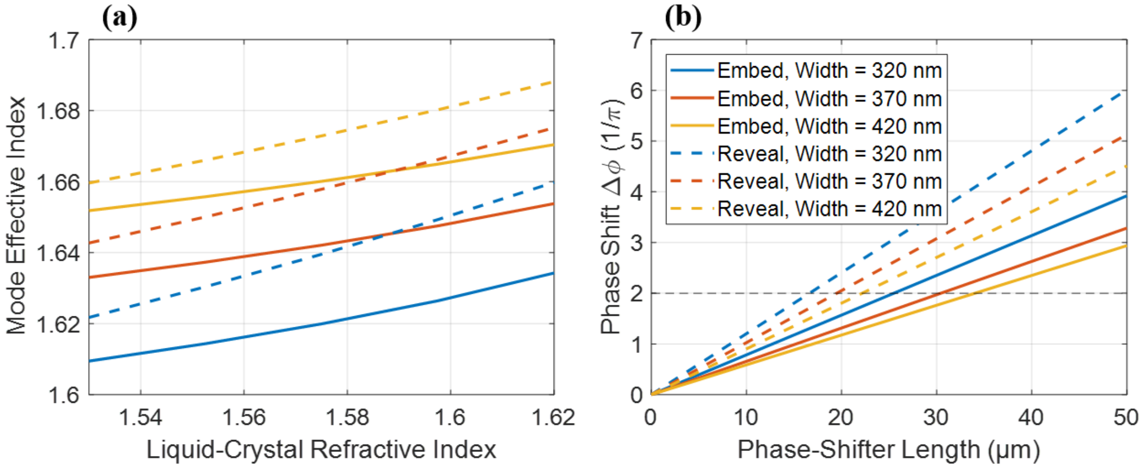

Figure 27: Comparison of (a) Waveguide Mode Effective Refractive Index as a Function of Liquid-crystal Index and (b) Phase Shift as a Function of Shifter Length, for a Waveguide Embedded within the Silicon-dioxide Cladding and Revealed to the Liquid-crystal Region

Similar to the original embedded waveguide loss that increases for a narrower waveguide, the revealed waveguide loss also increases for a narrower waveguide, as shown in Figure 28.

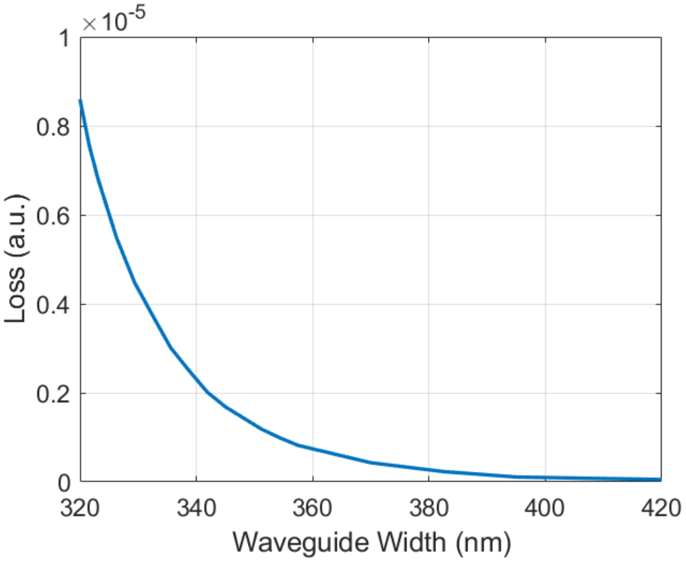

Figure 28: Loss as a Function of Waveguide Width for the Revealed Waveguide

3.2 Liquid-Crystal Phase-Shifter Experimental Setup and Results

The liquid-crystal phase modulator is fabricated and packaged via the process outlined in Sec. 2.5. It is difficult to directly experimentally measure phase shift, so the phase modulators are fabricated in an integrated Mach–Zehnder interferometer (MZI) to experimentally characterize their performance, as shown in Figure 3.13. A 632.8-nm Helium-Neon (He-Ne) laser is coupled onto the chip from a fiber via an input edge coupler [36] (a vertical coupler could also be used for input coupling [37]). A 1x2 multimode interference (MMI) splitter is used to transition from a single waveguide to the two separate arms of the MZI [38]. A liquid-crystal phase modulator is placed in each arm of the MZI. An adiabatic escalator is used to transition from the lower silicon-nitride waveguide (into which the light is coupled from the off-chip fiber) to the upper silicon-nitride waveguide (which interacts with the liquid-crystal region) and a second adiabatic escalator is used to transition after the phase modulator from the upper waveguide to the lower waveguide. The layer transition is achieved by slowly tapering down the width of the lower waveguide while simultaneously tapering up the width of the upper waveguide. This escalator layer transition is used because there is significant loss across the interface of a single-layer waveguide transitioning directly into a liquid-crystal region. This loss for a waveguide simply transitioning into and out of a liquid-crystal region of varying refractive index is simulated in Figure 30. This harsh interface is seen to have significant loss and reflection, with less than 55% of the light successfully transmitting into and out of the liquid-crystal region. The escalator layer transitions allow the light to gradually enter and exit the liquid-crystal region and hence limits the loss and reflections due to the change in refractive index. Finally, a 2x1 MMI coupler is used to combine the two arms of the MZI back into a single waveguide and a second edge coupler is used to couple the light from the chip to an output fiber.

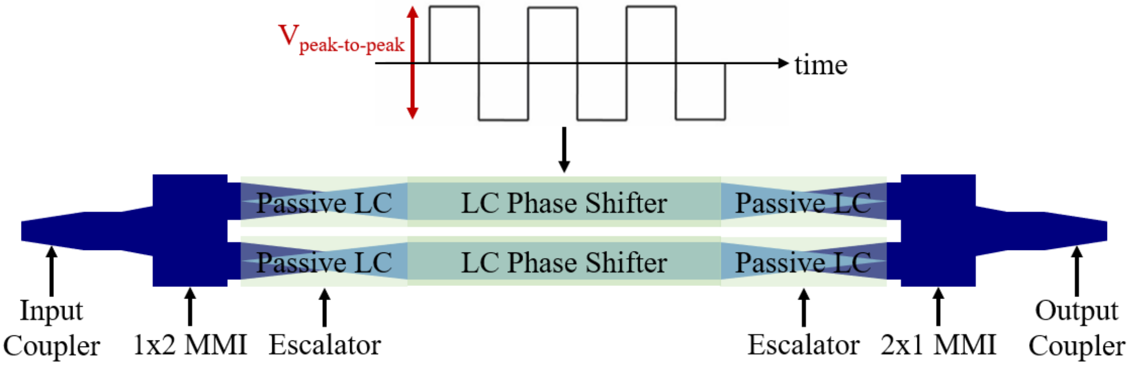

Figure 29: Diagram of the Integrated Mach-Zehnder-interferometer Test Structure and Modulation Scheme

Not to scale

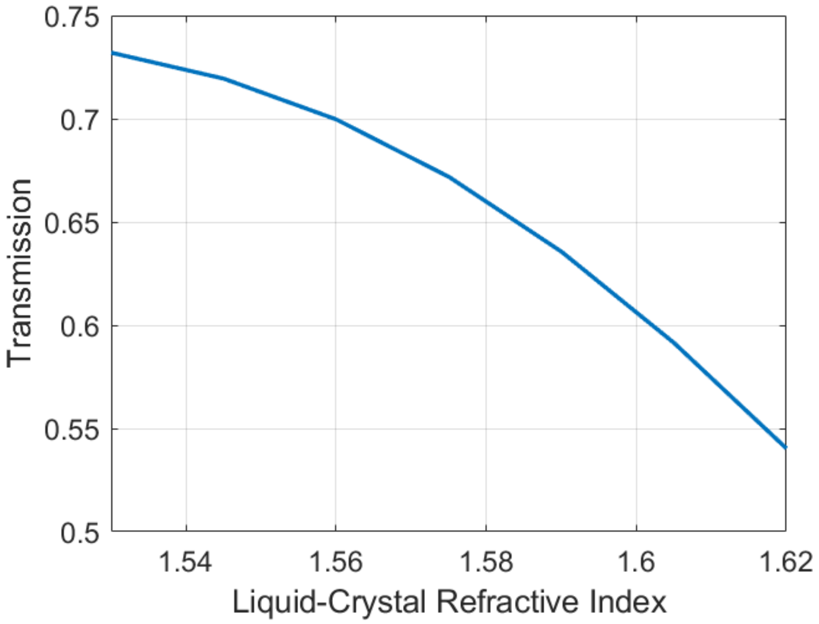

Figure 30: Simulated Transmission into and out of a Liquid-crystal Region as a Function of Liquid-crystal Refractive Index

A diagram of the experimental setup is shown in Figure 31 and a photograph of the experimental setup in shown in Figure 32. First, a He-Ne laser is coupled from free space to an optical fiber via bulk-optics mirrors and lenses. Then, the fiber is routed through a polarization controller to ensure that the light coupled onto the chip is transverse-electric polarization. Next, the light is coupled from the input fiber onto the packaged chip and goes through the MZI device described above. Then, the light is coupled from the chip to an output fiber, which is connected to an optical power meter. Simultaneously, a function generator is connected to an electronics breadboard, which then connects to the packaged chip via probes that touch down on the top surface of the chip to metal pads that route to the integrated electrodes of the phase modulators. A photograph of the packaged chip, with the input and output fibers coupling light onto and from the chip and the probe connecting to the integrated electrodes, is shown in Figure 33.

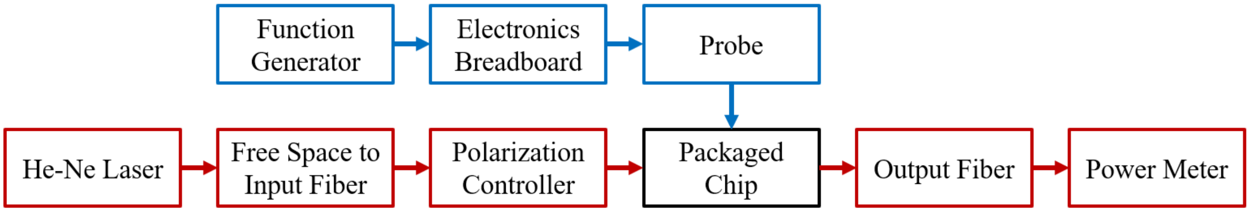

Figure 31: Diagram Showing the Experimental Setup

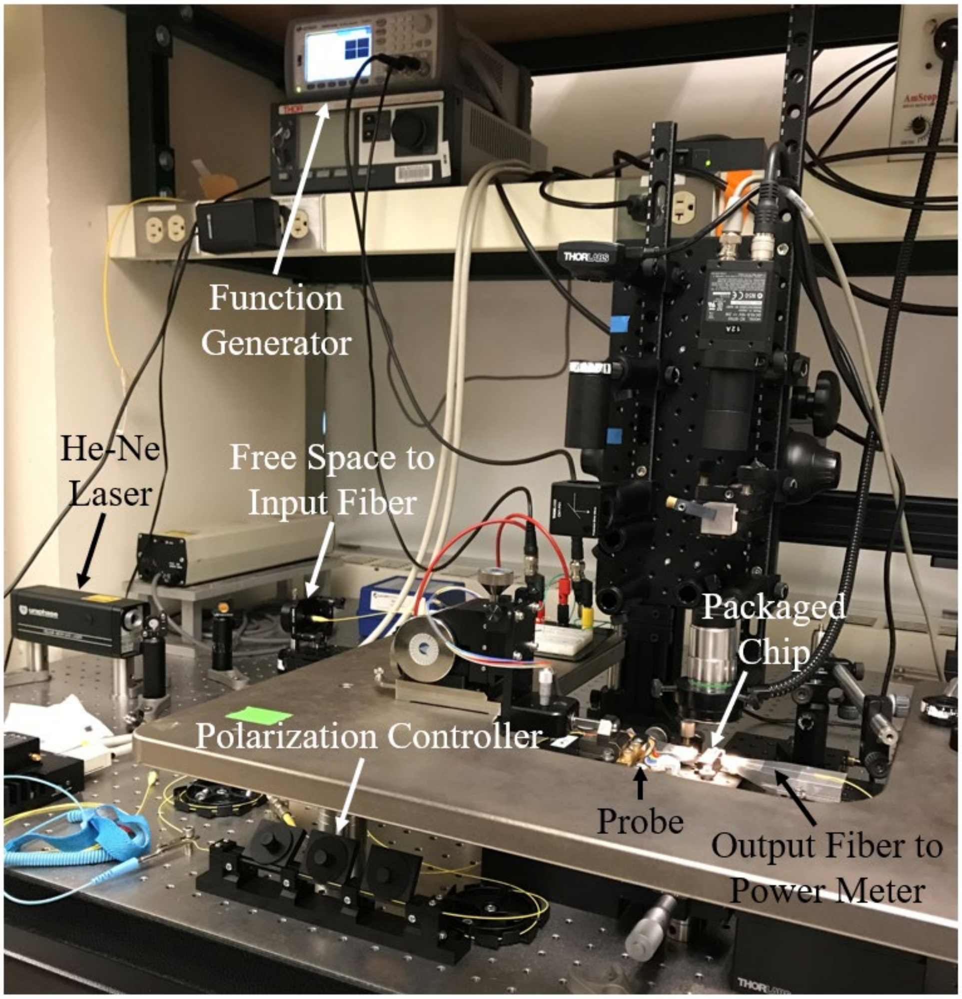

Figure 32: Photograph of the Experimental Setup

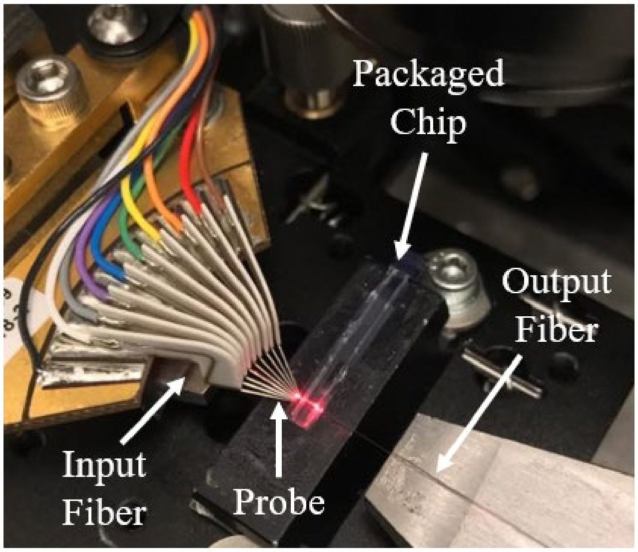

Figure 33: Photograph of the Packaged Chip on the Experimental Setup

The function generator connected to the electronic probes is used to apply a 10-kHz square wave across the electrodes of the phase modulator in the top arm of the MZI, as shown in Figure 29. The square wave has no offset voltage (the average is 0V) to ensure that charge does not accumulate in the liquid-crystal region and cause damage. No signal is applied across the phase modulator in the other arm of the MZI; the liquid-crystal molecules in this modulator are maintained in a constant orientation via the mechanical alignment layer, discussed in Sec. 2.5.4. The peak-to-peak voltage of the 10-kHz square wave is varied to vary the electric field across the liquid-crystal region of the top phase modulator. As the voltage is increased, the electric-field strength is increased, and the refractive index of the liquid-crystal region increases, which induces a phase shift in the top arm of the MZI relative to the bottom arm. If the two arms of an MZI are in phase, the power at the output of the MZI is maximized, whereas, if the two arms of the MZI are completely out of phase, the power at the output is minimized. Hence, the amount of phase shift can be measured via measuring the change in output power of the MZI device.

Experimental results of the power at the output of the MZI device, as the peak-to-peak voltage is varied, is shown in Figure 34. Since each dip in output power occurs when the two arms of the MZI are out of phase, there is 2π phase shift between each dip. Hence, this 500-µm-long device achieves 36π phase shift, within 6 V_{pp}. Consequently, a 2π phase shifter would need to be only 28-µm long.

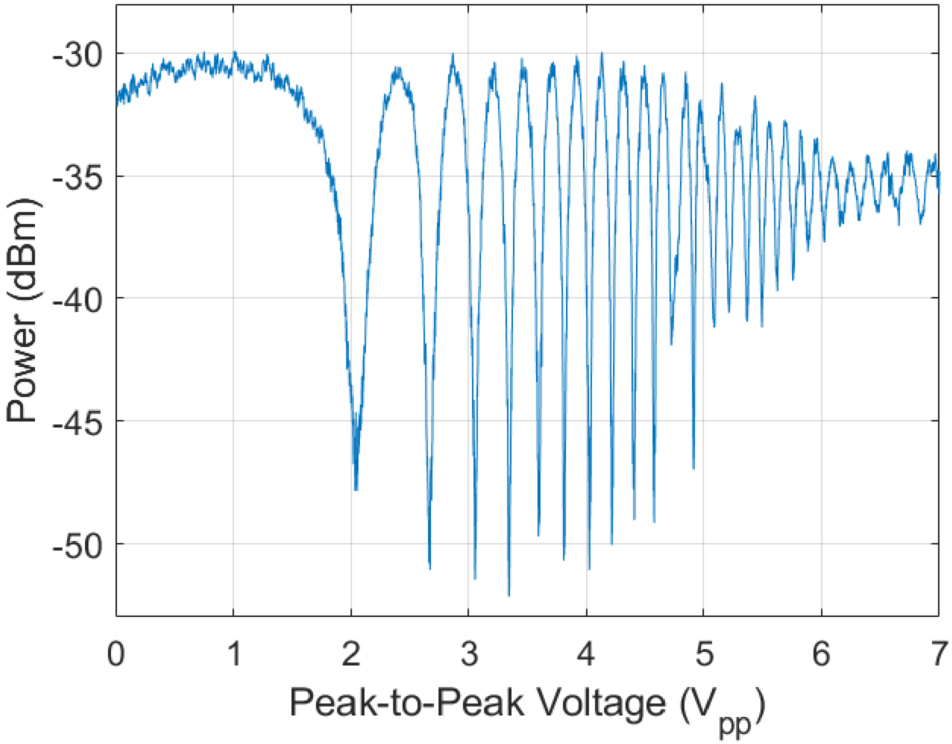

Figure 34: Experimentally measured power at the output of the MZI with phase shifters integrated into both arms with a 10 kHz square wave of varying peak-to-peak voltage applied across one arm of the MZI

The phase shift versus peak-to-peak voltage, inferred from the MZI power measurement shown in Figure 34, is shown in Figure 35. The phase shift occurs nonlinearly with respect to the applied voltage. The relationship is nonlinear because the effective impedance of the waveguide mode varies slightly nonlinearly with respect to the liquid-crystal refractive index. Furthermore, the liquid-crystal refractive index also varies nonlinearly with applied voltage. Figure 34and Figure 35 demonstrate the threshold phenomenon described in Sec. 2.3. Recall that the electric field applied across the liquid-crystal region has to be strong enough to overcome the mechanical anchoring strength of the alignment layer.

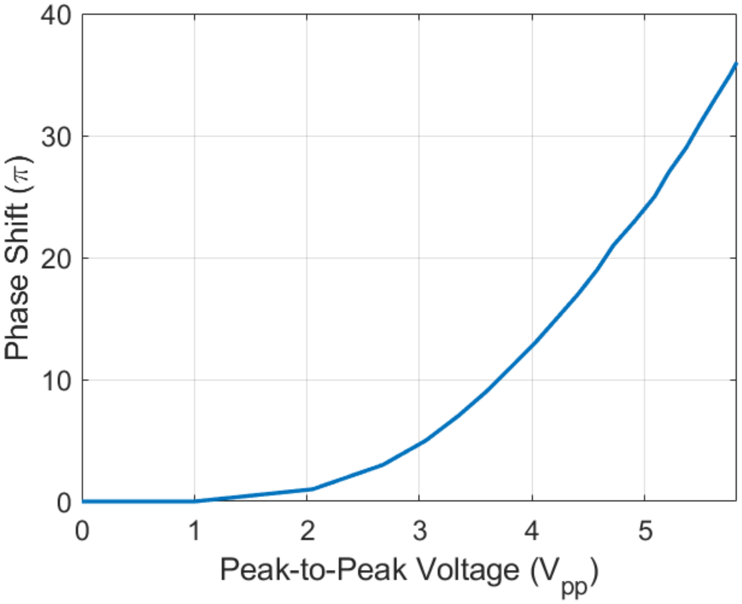

Figure 35: Experimentally Measured Phase Shift as a Function of Applied Peak-to-peak Voltage, Inferred from MZI Power Measurement

3.3 Conclusion

The birefringence of liquid crystal can be leveraged to design an integrated visible-light phase modulator. Many design parameters can be manipulated to improve the performance of the device. For example, the silicon-dioxide gap separating the waveguide and the liquid-crystal region can be minimized, the waveguide can be made narrower, or the waveguide can be revealed to the liquid-crystal region. These changes cause the waveguide mode to interact more with the liquid crystal and, hence, be impacted more by a change in the refractive index of the liquid crystal. This increased interaction results in larger phase shifts for a set shifter length. The device was fabricated with an integrated MZI and experimental measurements were taken that match closely to the simulated results. Overall, the liquid-crystal phase shifter enables compact and low-power integrated visible-light phase modulation.

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

4 LIQUID-CRYSTAL VARIABLE-TAP AMPLITUDE MODULATOR

In this chapter, an integrated visible-light variable-tap amplitude modulator is presented, the design process is outlined, and experimental results are shown and discussed. As discussed in Sec. 2.1, integrated visible-light amplitude modulation is a challenge due to the low thermo-optic coefficient and lack of electro-optic properties of silicon nitride. As shown in Sec. 3.2, amplitude modulation can be achieved by integrating a phase modulator into a MZI; however, this configuration has a large form factor since it requires a splitter at both the input and output of the device, phase shifters in each arm of the interferometer, and transitions connecting the liquid-crystal region to the splitter. Hence, a more compact device is necessary to more practically implement integrated visible-light amplitude modulation.

The variable-tap amplitude-modulator device described here leverages the birefringence of the liquid crystal to vary the coupling coefficient (which depends on mode overlap and propagation constant) between two waveguides to actively tune the amount of power coupled from one waveguide to another. The design process consists of choosing the proper waveguide widths and corresponding coupler length to achieve maximum variation in amplitude modulation. After the design process is explained via simulated results, experimental results are then shown and discussed.

4.1 Liquid-Crystal Variable-Tap Amplitude-Modulator Theory and Design

4.1.1 Device Structure

The integrated liquid-crystal visible-light variable-tap amplitude modulator consists of two vertically-stacked 160-nm-thick silicon-nitride waveguides, below a silicon-dioxide trench filled with liquid crystal, with integrated metal electrodes on either side of the liquid crystal, as shown in Figure 36. The upper waveguide, directly underneath the liquid-crystal region, is called the bus waveguide and the lower waveguide, situated 410 nm below the bus waveguide, is called the tap waveguide. Initially, the input light propagates in the bus waveguide and is then coupled into the tap waveguide. The coupling region has a set coupler length, defined in Figure 36a, which is the length of the coupler for which the two waveguides are vertically stacked, before the tap waveguide curves away.

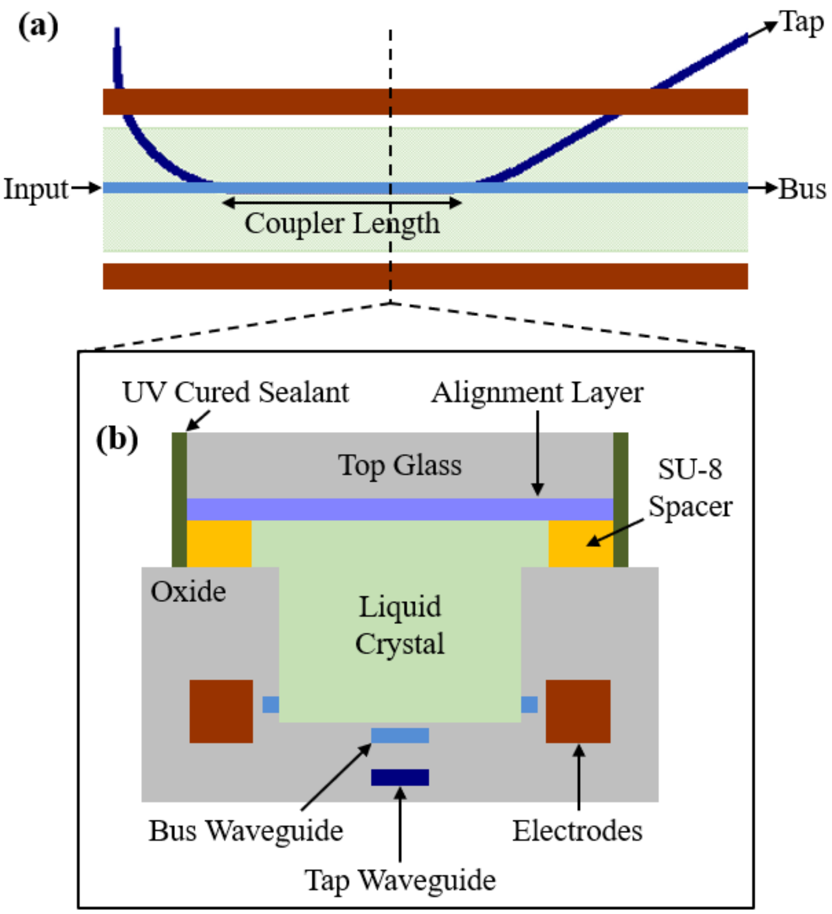

Figure 36: Simplified Top-view Schematic of the Liquid-crystal Variable-tap Amplitude Modulator (a) and (b) Cross Section of the Coupling Region After Packaging

4.1.2 Device Operation

In general, two waveguides have a coupling coefficient that dictates how light is coupled in between the two waveguides. This coupling coefficient depends on many variables, including the dimensions of both waveguides, the material composition of the waveguides and cladding, and the distance separating the two waveguides. Instead of describing the coupling coefficient via structural parameters, such as dimensions and material properties, the coupling coefficient can also be described based on the mode overlap and difference in propagation constants between two waveguides [39].

In this liquid-crystal variable-tap device, the amplitude of the light coupled from the bus waveguide into the tap waveguide depends on the coupling coefficient between the two waveguides. Since the bus waveguide is situated directly underneath the liquid-crystal trench, the confinement and propagation constant of the bus-waveguide mode are highly dependent on the refractive index of the liquid crystal. On the other hand, since the tap waveguide is situated farther away from the liquid-crystal region, the mode in the tap waveguide is not significantly affected by a change in the liquid-crystal refractive index. As the liquid-crystal refractive index is varied, the effective index of the tap-waveguide mode is not impacted, as shown in Figure 37.

However, because only the bus waveguide is significantly impacted by the liquid crystal, the coupling coefficient between the two waveguides changes as the liquid-crystal refractive index is tuned.

Similar to the phase modulator described in Sec. 3.1, the liquid-crystal refractive index can be actively varied by applying an electric field across the liquid-crystal region via the integrated electrodes, shown in Figure 36. As the liquid-crystal refractive index increases, the mode in the bus waveguide becomes less confined and is pulled up into the liquid-crystal region. This results in less mode overlap with the tap waveguide, as well as a change in effective refractive index of the bus-waveguide mode, while the effective refractive index of the tap-waveguide mode remains unchanged. This means that there is a change in propagation constant for the bus waveguide mode, but not for the tap waveguide, which results in an overall change in relative propagation constant between the two waveguides. This change in relative propagation constant and the change in mode overlap between the two waveguides impacts the coupling coefficient between the two waveguides and, hence, results in amplitude modulation of the light coupled from the bus waveguide to the tap waveguide.

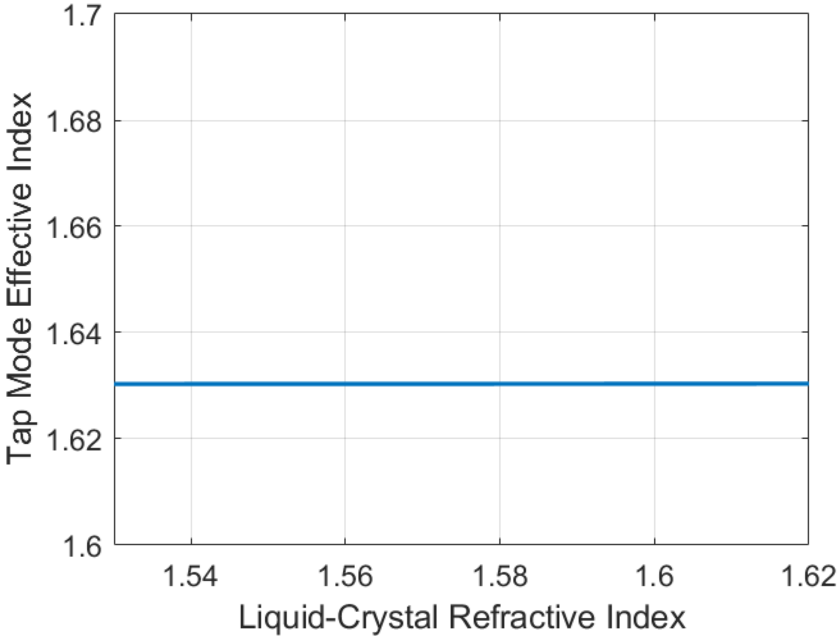

Figure 37: Simulated Effective Refractive Index of the Tap Waveguide as a Function of Liquid-crystal Refractive Index, Showing that the Index is Effectively Unaffected by Changes in the Liquid-crystal index

4.1.3 Device Design

As mentioned above, there are many parameters that impact the coupling between the two waveguides. Structurally, the waveguide material is set to be silicon nitride, the waveguide layer thicknesses are set to be 160 nm, and the vertical offset between the two waveguides is set to be 410 nm. Hence, the widths of the bus and tap waveguides need to be designed properly, as well as the coupler length.

One way to design the amplitude modulator is to first set the bus-waveguide width to a certain dimension, because it is the input of the device and will depend on the width of the prior device if the modulator is integrated into a larger system. A bus-waveguide width of 320 nm is chosen for the following simulations and discussion. Then, to ensure optimal modulation for a given bus-waveguide width, the coupler length and tap-waveguide width must be chosen.

First, the coupler length is chosen. Since the bus-waveguide mode pulls up into the liquid-crystal region as the liquid-crystal refractive index is increased, the mode overlap between the bus and tap waveguides decreases at the higher liquid-crystal refractive index. To take advantage of this phenomenon, the coupler length is chosen such that no light is coupled into the tap at the high liquid-crystal refractive index. To determine the appropriate length, the tap power versus coupler length is simulated for a bus-waveguide width of 320 nm and a variety of tap-waveguide widths, as shown in Figure 38. For example, for a 390-nm-wide tap waveguide, the appropriate coupler length is 17 μm.

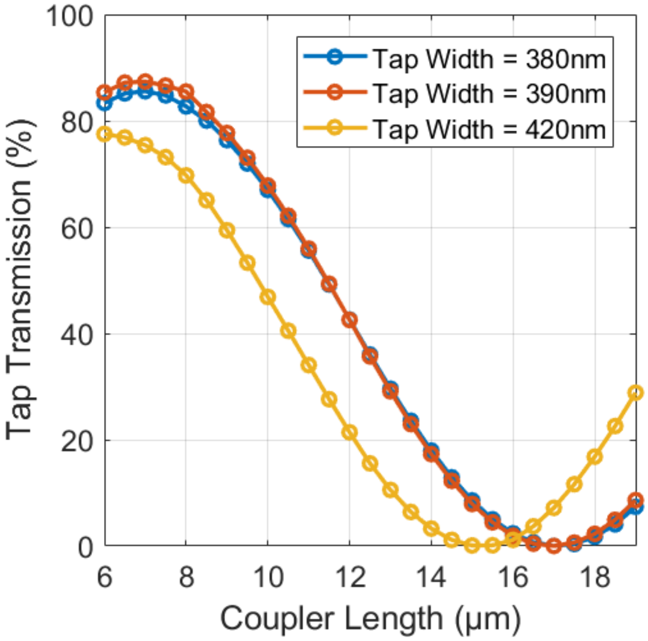

Figure 38: Tap Transmission Versus Coupler Length at the High Liquid-crystal Refractive Index of 1.62, a Bus-waveguide Width of 320 nm, and Various Tap-waveguide Widths

Second, the tap-waveguide width is chosen to maximize the amount of amplitude variation. To determine the optimal tap-waveguide width, the tap power versus liquid-crystal refractive index is simulated for each tap-waveguide width and corresponding coupler length, as shown in Figure 39. This determines the overall change in tap transmission. For example, for a bus width of 320 nm, the optimal tap width is found to be 390 nm, which results in amplitude variation from 60% at the low liquid-crystal refractive index to 0% at the high liquid-crystal refractive index.

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

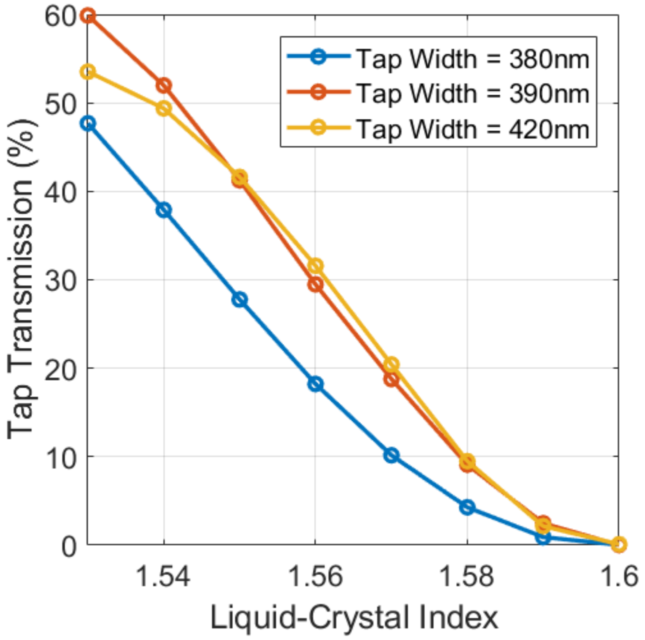

Figure 39: Tap Transmission Versus Liquid-crystal Refractive Index for a Bus-waveguide Width of 320 nm, Various Tap-waveguide Widths, and Appropriately Chosen Coupler Lengths

The above design process is shown for only three separate tap-waveguide widths for simplicity and to more clearly demonstrate the design process; however, the actual design process consists of a much more extensive sweep of waveguide widths. The tap power is simulated over a range of 16 different tap-waveguide widths at the high liquid-crystal refractive index of 1.62, as shown in Figure 40, to determine the proper coupler length to achieve zero coupling at the high index. The proper coupler length for each tap-waveguide width is then calculated from Figure 40, as shown in Figure 41.

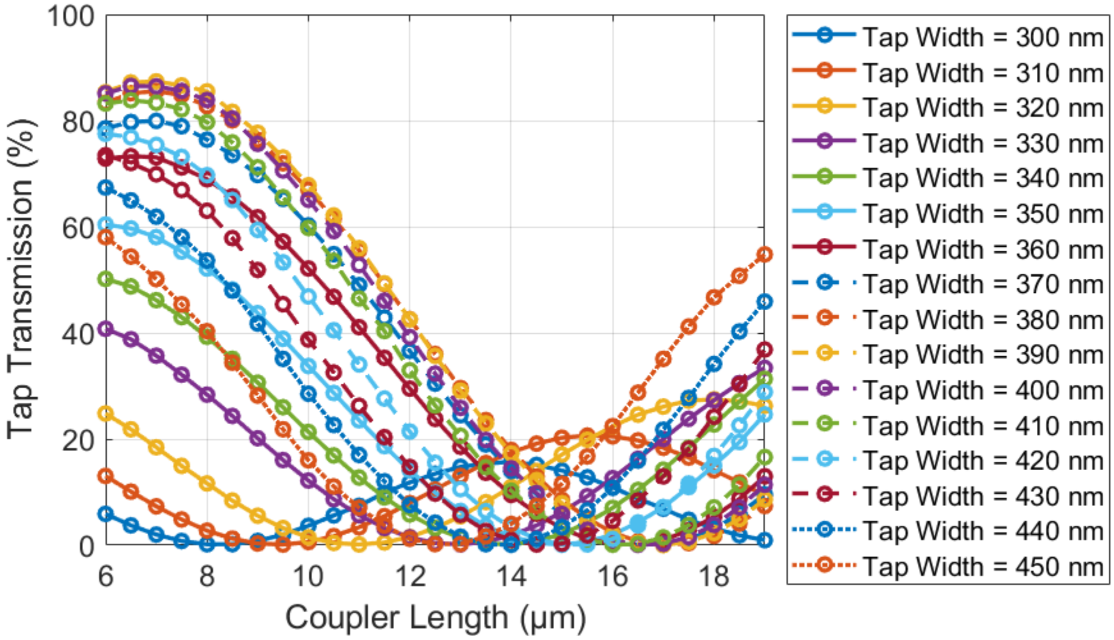

Figure 40: Tap Transmission Versus Coupler Length for a Liquid-crystal Refractive Index of 1.62, a Bus-waveguide Width of 320 nm, and an Extensive Range of Tap-waveguide Widths

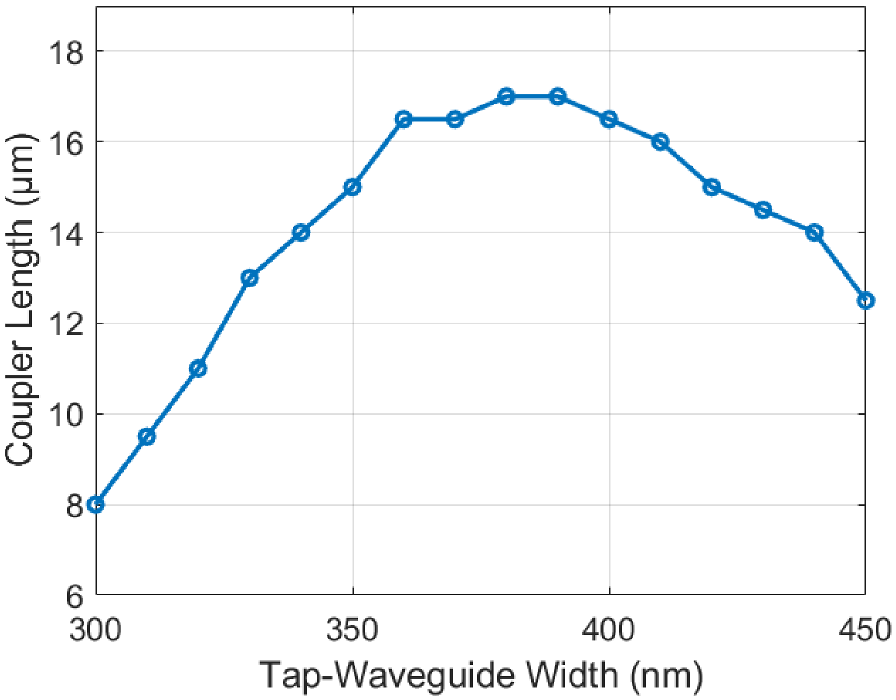

Figure 41: Calculated Coupler Length Required to Achieve Zero Coupling at the High Liquid-crystal Refractive Index for Various Tap-waveguide Widths and a Bus-waveguide Width of 320 nm

Then, the tap power is simulated over the range of 16 different tap-waveguide widths at the low liquid-crystal refractive index of 1.53, as shown in Figure 42, to determine the amplitude modulation range for each tap width. The variation in tap coupling over the entire liquid-crystal refractive index range for the various tap-waveguide widths (with the proper coupler length for each tap-waveguide width) is then calculated from Figure 40 and Figure 42.

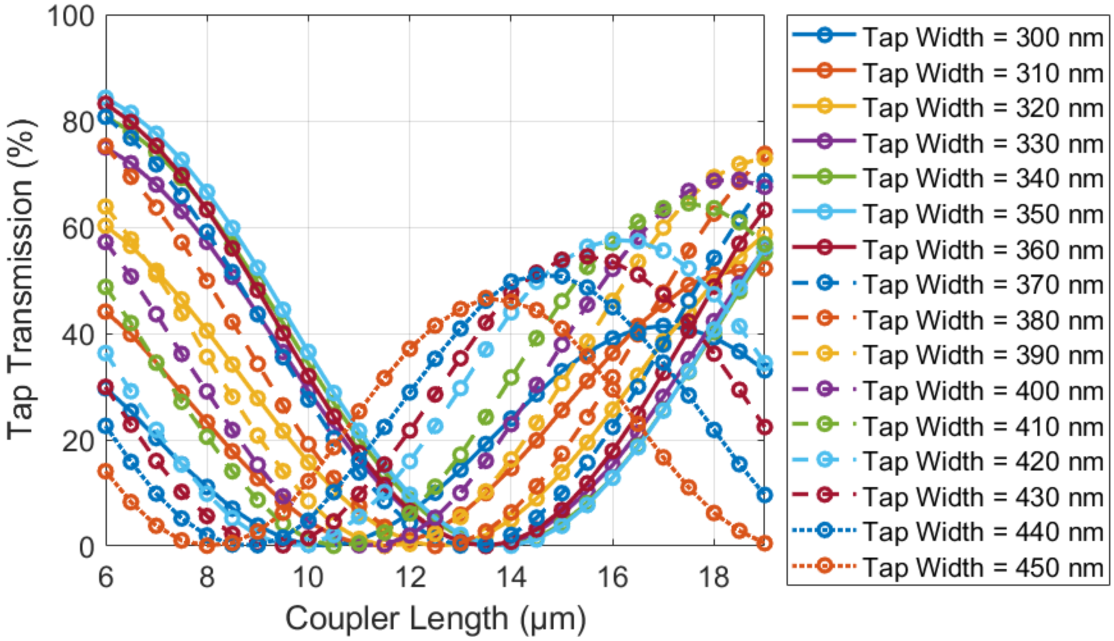

Figure 42: Tap Transmission Versus Coupler Length for a Liquid-crystal Refractive Index of 1.53, a Bus-waveguide Width of 320 nm, and an Extensive Range of Tap-waveguide Widths

The change in tap transmission as a function of tap-waveguide width is shown in Figure 43. This allows for the proper tap-waveguide width to be chosen to ensure optimal amplitude modulation to the tap waveguide.

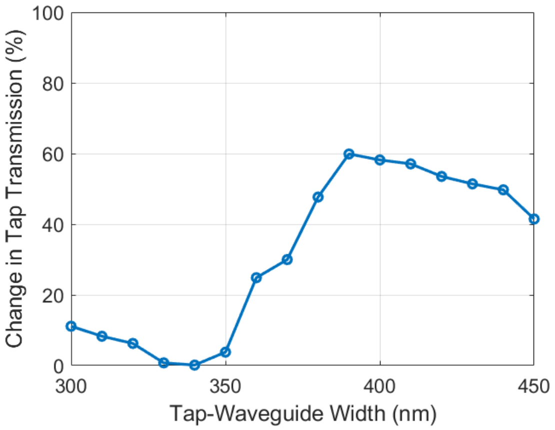

Figure 43: Change in Tap-waveguide Power Transmission Over the Maximal Range of Liquid-crystal Refractive Index for Various Tap-waveguide Widths and a Bus-waveguide Width of 320 nm

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

4.2 Liquid-Crystal Variable-Tap Amplitude Modulator Experimental Setup and Results

The liquid-crystal variable-tap amplitude modulator is fabricated and packaged via the process outlined in Sec. 2.5. A micrograph of the fabricated amplitude-modulator test structure is shown in Figure 44.

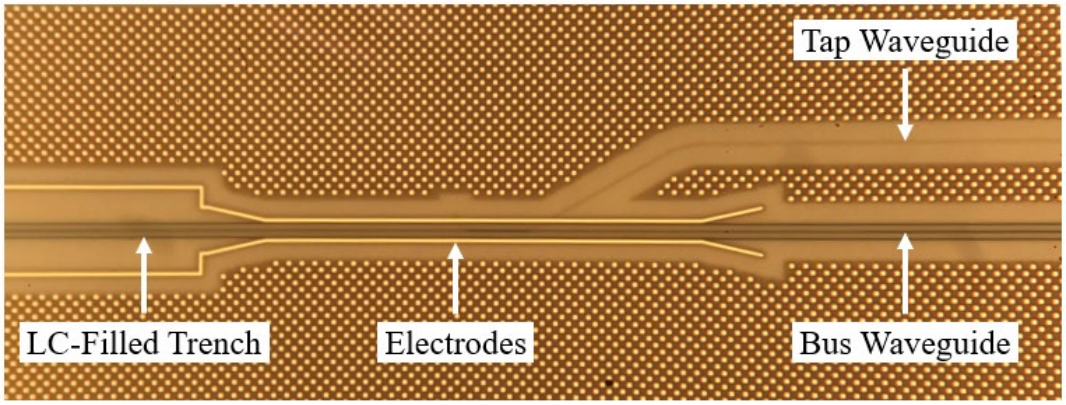

Figure 44: Micrograph of a Fabricated Liquid-crystal Variable-tap Amplitude Modulator

Similar to the experimental setup for testing the phase modulator, shown in Figure 32 and Figure 33, light from a He-Ne laser is coupled into a fiber via bulk optics lenses and mirrors, the fiber is routed through a polarization controller to ensure the light is transverse-electrically polarized, and then the light is coupled from the fiber onto the packaged chip via an edge coupler. As shown in Figure 45, the light is initially coupled into the lower waveguide layer and then adiabatically transitioned via an escalator to the upper bus waveguide under the liquid-crystal region, to eliminate the loss and reflection that would occur if the light in the bus waveguide suddenly approached the liquid-crystal interface, as discussed in Sec. 3.2. The light is then coupled from the bus waveguide to the tap waveguide for a set coupler length and, finally, coupled off chip via an edge coupler to an optical fiber that is connected to a power meter. Additionally, the light remaining in the bus waveguide is adiabatically transitioned to the lower waveguide layer via an escalator and also output to a fiber via an edge coupler, enabling measurement of the light that remains in the bus waveguide, as well as the light that is coupled into the tap waveguide.

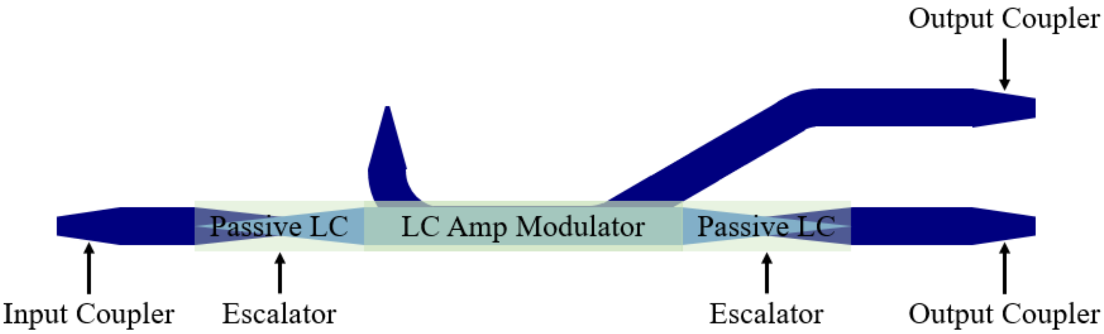

Figure 45: Diagram of the Variable-tap Amplitude-modulator Test Structure

Not to scale

A function generator is used to apply a 10-kHz square wave across the integrated electrodes on either side of the liquid-crystal region. The peak voltage of the square wave is varied to sweep the strength of the resulting electric field across the liquid-crystal region. As the peak-to-peak voltage is increased, the liquid-crystal refractive index increases and the coupling coefficient between the bus and tap waveguides changes, which results in amplitude modulation at the output of both the bus and tap waveguides. The experimentally-measured power in the tap waveguide is seen to increase in Figure 46a, as the experimentally-measured power decreases in the bus waveguide, as shown in Figure 47b. As expected, as the peak voltage of the square wave is increased, the power in the tap and bus change.

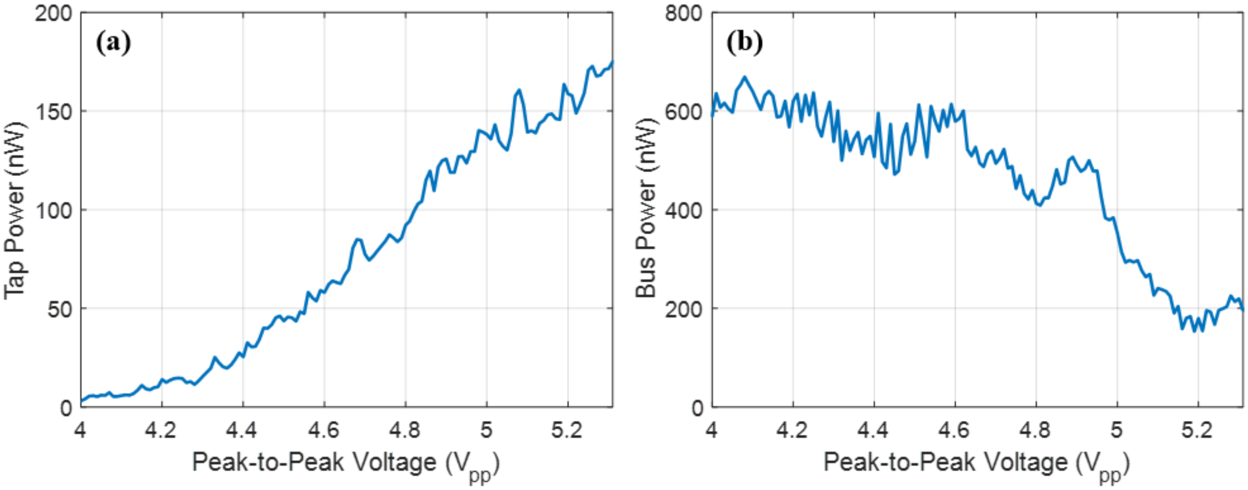

Figure 46: Experimental Results of the Power in (a) the Tap Waveguide and (b) the Bus Waveguide, as the Peak-to-peak Voltage of the 10-kHz Square Wave Applied Across the Liquid-crystal Region is Varied

Furthermore, the 10-kHz square wave can be modulated by a sinusoidal envelope signal to modulate the power in the tap waveguide. As the peak voltage is modulated by the envelope signal, the liquid-crystal refractive index is also modulated, which in turn varies the coupling coefficient between the bus and tap waveguides. This variation then results in amplitude modulation of the power in the tap waveguide. The power in the tap waveguide is measured as the peak voltage of the square wave is modulated with a 1-Hz sinusoidal signal. The power in the tap indeed modulates over time as the applied peak voltage is modulated, as shown in Figure 47.

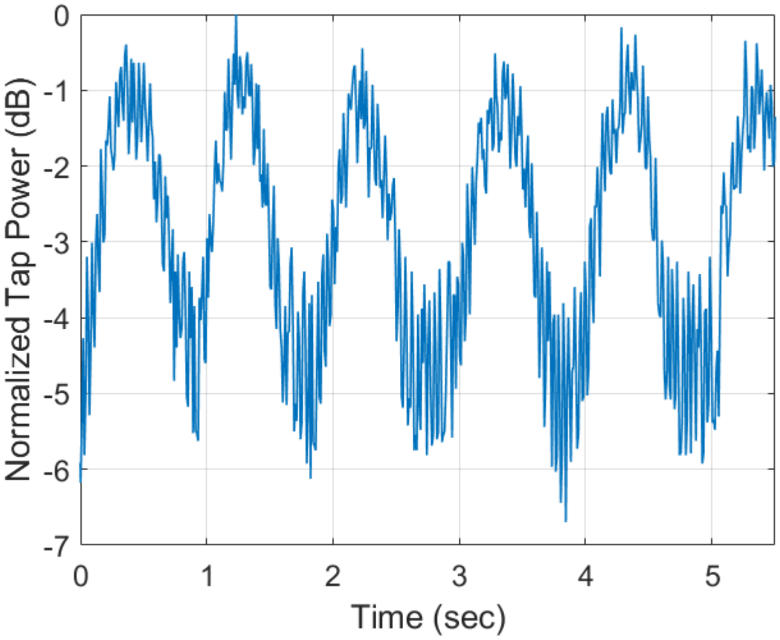

Figure 47: Experimental Results Showing Modulated Tap Power Output Under 1-Hz Sinusoidal Amplitude Modulation

The frequency of the sinusoidal envelope signal can be varied to demonstrate the variable switching speed of the liquid-crystal media. As the frequency of the envelope signal is increased, the frequency of the power modulation in the tap waveguide increases accordingly. Amplitude modulation up to 30 Hz is shown in Figure 48. When the envelope signal is modulating at 30 Hz, the liquid crystal is switching at a speed of 33 ms.

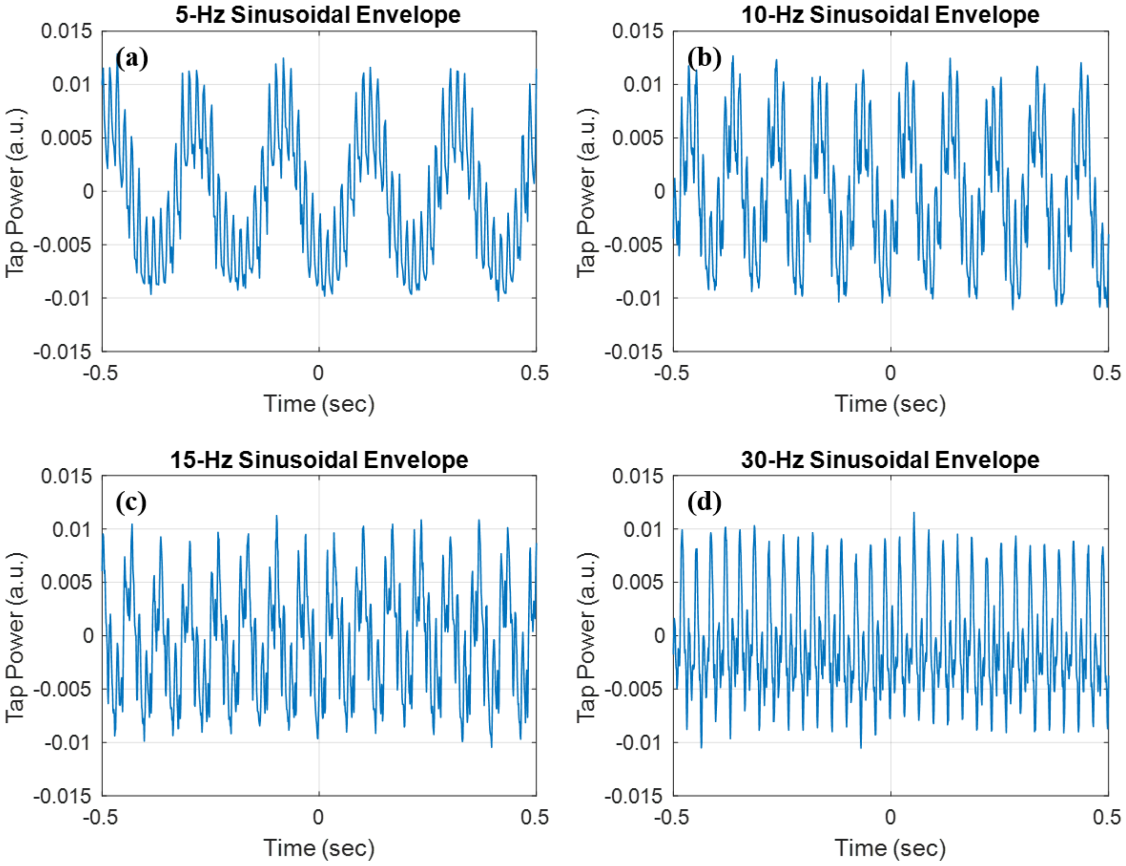

Figure 48: Experimental Results Showing Modulated Tap Power Output Under (a) 5-Hz, (b) 10-Hz, (c) 15-Hz, and (d) 30-Hz Sinusoidal Amplitude Modulation

4.3 Conclusion

The liquid-crystal birefringence can be used to vary the coupling coefficient between two waveguides to create an integrated visible-light variable-tap amplitude modulator. By actively tuning the refractive index of the liquid-crystal region above a bus waveguide, the mode overlap and propagation constant between this waveguide mode and a second tap waveguide mode can be actively tuned. The device is designed by first setting the bus-waveguide width and choosing the proper coupler length to achieve zero coupling into the tap at the high liquid-crystal refractive index. Then, the tap-waveguide width is chosen to maximize the variation in power coupled into the tap across the liquid-crystal index range. Using this design process, the device was simulated, fabricated, packaged, and experimentally tested. Experimentally, the liquid-crystal refractive index is tuned via a 10-kHz square wave with variable peak voltage, which is applied via integrated electrodes. The variation in peak voltage results in variation in the power coupled into the tap waveguide. Furthermore, the amplitude of the signal applied across the liquid-crystal region can be modulated via an envelope signal to modulate the power coupled into the tap waveguide, at varying frequencies. Overall, this small-form-factor variable-tap device is a compact and low-power solution to integrated visible-light amplitude modulation and will enable high-density integrated visible-light systems.

5 INTEGRATED VISIBLE-LIGHT LIQUID-CRYSTAL-BASED INTEGRATED OPTICAL PHASED ARRAYS

5.1 Introduction

Integrated optical phased arrays [40-42] have emerged as a promising technology for many applications, such as light detection and ranging (LiDAR) and free-space optical data communications, due to their ability to manipulate and dynamically steer free-space light in a compact form factor, at low costs, and in a non-mechanical way. However, optical phased array demonstrations to date have primarily focused on the infrared wavelength regime, although there are many potential wide-reaching applications of optical phased arrays that require visible light operation, including dynamic displays and projection systems (such as the VIPER display discussed in Sec. 6), underwater optical communications, LiDAR, and optogenetics.

Recently, visible-light integrated optical phased arrays have been demonstrated in silicon-nitride platforms [7,8,43]; however, these systems have been limited to passive demonstrations. Since silicon nitride has a low thermo-optic coefficient and does not exhibit any significant electro-optic properties, integrated phase tuning at visible wavelengths is a challenge. As a solution, nematic liquid crystals, with strong birefringence in the visible spectrum, can be integrated into silicon-nitride platforms to enable visible-light phase modulation with low powers and short lengths (as discussed in Sec. 3).

In this section, liquid-crystal-based integrated optical phased arrays are proposed and experimentally demonstrated for the first time as a method for low-power and compact visible-light beam steering. A cascaded integrated optical phased array architecture with a liquid-crystal-based phase-shifting region is developed and used to experimentally demonstrate beam steering at a 632.8 nm wavelength with a 0.7° × 2.3° power full width at half maximum and 10.5° steering range within ±3.5 V.

5.2 Architecture

As a proof of concept, a visible-light integrated optical phased array was designed, fabricated in a CMOS-compatible foundry process at SUNY Poly, and packaged with liquid crystals using a back-end fabrication process at MIT. The integrated phased array consists of a silicon-nitride-based cascaded-phase-shifter architecture that linearly controls the relative phase applied to an array of antennas, as shown in Figure 49a.

At the input, an on-chip inverse-taper edge coupler couples light from an off-chip laser into an on-chip single-mode silicon-nitride waveguide. A 100-μm-long escalator device (an adiabatic layer transition structure) then couples the input light from the single-mode waveguide into the liquid-crystal-based phase-shifter bus region, as shown in Figure 49b. Next, evanescent tap couplers, placed with a pitch of 20 μm and with increasing coupling lengths, uniformly distribute the light from the bus region to 16 grating-based 400-μm-long antennas with a 2 μm pitch, shown in Figure 49c-d.

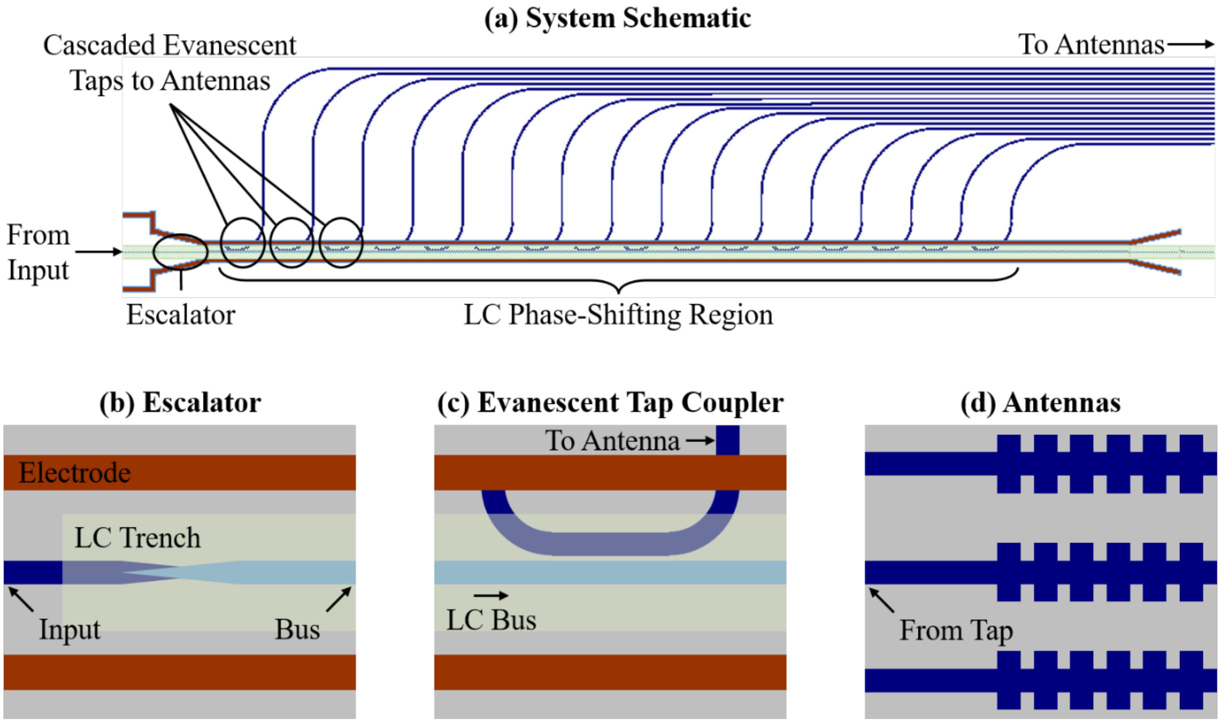

Figure 49: Partial Schematic of the Liquid-crystal-based Integrated Optical Phased Array Showing Major Components (a), Simplified Schematics of the (b) Layer-transition Escalator, (c) Evanescent Tap Coupler, and (d) Grating-based Antennas used in the Liquid-crystal-based Integrated Optical Phased Array

To enable one-dimensional far-field beam steering, the system utilizes the birefringence of liquid-crystal media to enable cascaded phase control to the array of antennas. As previously discussed, in a nematic liquid-crystal medium, the index of refraction varies based on the orientation of the liquid-crystal molecules. Thus, by applying an electric field across the liquid-crystal region to orient the molecules in the direction of the applied field, the index of the liquid-crystal media can be actively tuned, resulting in a linear phase shift to the antennas. To enable this functionality, the liquid-crystal-based phase-shifting region consists of a silicon-nitride waveguide to weakly confine and guide the light, liquid crystal deposited into an oxide trench above the waveguide to enable strong interaction between the optical mode and the liquid-crystal media, metal electrodes for applying an electric field across the liquid-crystal region, and a top glass chip with a mechanical alignment layer to anchor the liquid-crystal molecules. Additional details on the liquid-crystal-based phase shifter are provided in Sec. 3.

5.3 EXPERIMENTAL Results

To characterize the fabricated array, a 632.8-nm He-Ne laser was coupled onto the chip and an optical system was used to image the far field of the phased array onto a visible-light camera. A photograph of the experimental setup is shown in Figure 50, and the experimentally measured far-field pattern and cross sections of the far-field main lobe are shown in Figure 51a-b. As expected, the array forms a beam in the far field with a 0.7° × 2.3° power full width at half maximum, 8-dB sidelobe suppression, and second-order grating lobes at ±28°. Next, electronic probes were used to apply a 10-kHz square wave with variable peak-to-peak voltage across the electrodes of the liquid-crystal-based phase-shifting region to steer the beam in the array dimension (θ). As shown in Figure 51c, the system enables $10.5°$ of visible-light beam steering within ± 3.5 V.

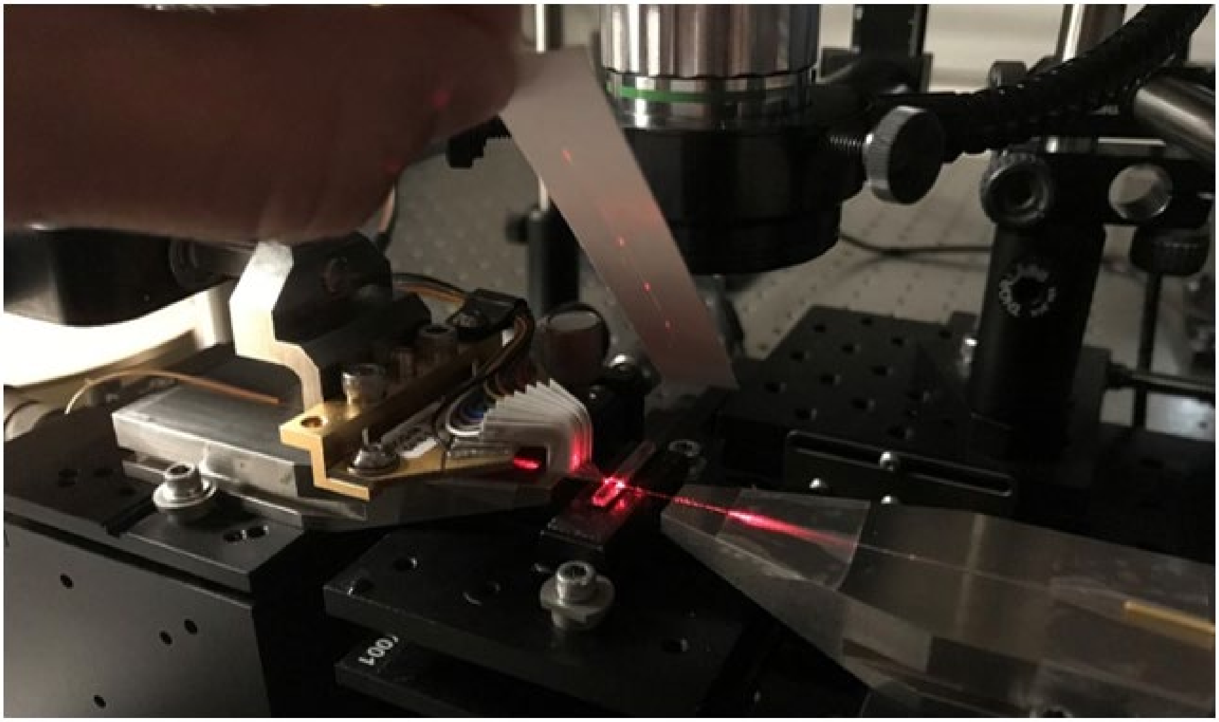

Figure 50: Photograph of the Packaged Optical-phased-Array Chip, Experimental Setup, and Output Radiated Light

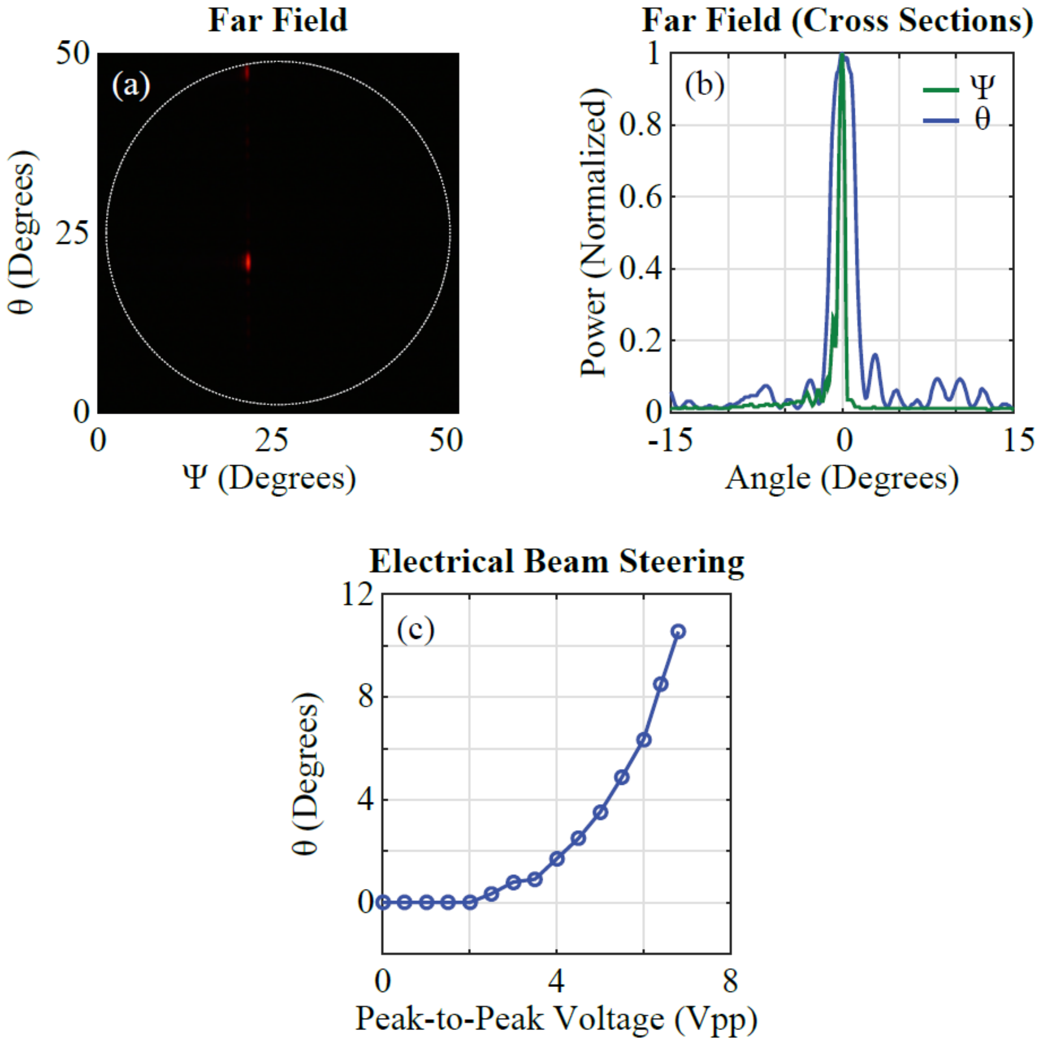

Figure 51: Measured Far Field Above the Chip Showing the Main Lobe of the Phased Array (a), (b) intensity Cross Sections of the Far-field Main Lobe in the Array Dimension (θ) and Antenna Dimension (Ψ). (c) Experimental Results Showing beam steering in the array Dimension (θ) Versus Applied Peak-to-peak Voltage

5.4 Conclusion

This work presents the first proposal and demonstration of liquid-crystal-based integrated optical phased arrays. A cascaded integrated optical phased array architecture with a liquid-crystal-based phase-shifting region was developed and used to experimentally demonstrate visible-light beam steering with a 0.7° × 2.3° power full width at half maximum and 10.5° steering range within ±3.5 V. This system has many important applications, ranging from optogenetics to underwater optical communications and LiDAR [13]. In Sec. 6, this liquid-crystal-based integrated optical phased array will be applied to the VIPER display to enable a dynamic visible-light near-eye holographic display for augmented-reality applications.

6 INTEGRATED-PHOTONICS-BASED HOLOGRAPHIC DISPLAYS FOR AUGMENTED REALITY

6.1 Introduction

In many situations, including military operations and medical procedures, access to real-time information can be a key determinant for success. Traditionally, this information has been displayed in real time using head-down or head-up displays. However, recently, there have been extensive efforts in developing head-mounted displays (HMDs) that are capable of relaying information directly in the user's field-of-view (FOV). These head-mounted displays enable the user to remain engaged with their surroundings while referencing information to real-world objects and events for an augmented-reality experience.

Typical commercially available augmented-reality head-mounted displays employ an optical relay system for each eye, wherein an image produced by a microdisplay is magnified using a system of lenses to generate an image superimposed on the external scene at a single virtual focal plane in the user's FOV [44], as shown in Figure 52a. However, the bulk-optics components utilized in these typical head-mounted displays result in large, heavy, and indiscreet head-mounted display solutions. Additionally, head-mounted displays typically employ low-luminance microdisplays (approximately 1000 cd/m^2), which render the systems inadequate for use in ambient daylight conditions, and optical relay systems which limit the user’s FOVs (limited to $< 40°$ compared to the 60° near-peripheral FOV of the human eye). Finally, typical head-mounted displays magnify the microdisplay image such that it appears at a single virtual focal plane (they are not capable of producing holographic images with full depth cues); this lack of depth information results in users experiencing eyestrain and headaches that limit long-term and wide-spread use of these displays (an effect known as the vergence-accommodation conflict). Although there have been a number of recent proposals and initial passive demonstrations of near-eye displays that utilize holographic image projection to emit full phase fronts and resolve the vergence-accommodation conflict [45,46], there is still a growing need for a dynamic, discrete, mobile, large-FOV, high-brightness augmented-reality head-mounted display with full binocular and monocular depth cues.

Figure 52: Simplified Diagram of (a) A Typical HMD Approach Using an Optical Relay System Versus (b) the Direct-view near-eye VIPER Approach

In this section, VIPER (Visible Integrated-Photonics Enhanced-Reality), a novel integrated-photonics-based visible-light near-eye holographic display, is proposed and experimentally demonstrated as a scalable solution to address this need. The VIPER display consists of a single discreet transparent chip, fabricated in a 300-mm-wafer foundry process, that sits directly in front of the user's eye, as shown in Figure 53, to enable a direct-view near-eye display approach, as shown in Figure 52b. The display is comprised of a grid of on-chip visible-light optical phased arrays, as shown in Figure 53a, that are encoded to emit light with the appropriate amplitudes and phases such that a virtual holographic image is formed that only the user can see, as shown in Figure 54b. It presents a highly-discreet and fully-holographic solution for the next generation of augmented-reality displays.

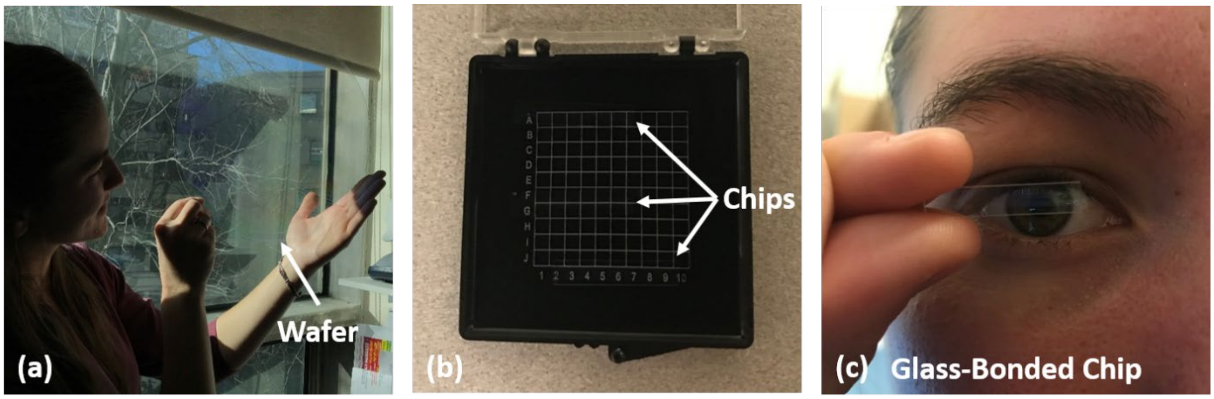

Figure 53: Photographs Showing (a) A 300-mm-Diameter Glass-bonded VIPER Wafer, (b) Three Glass-bonded VIPER Chips, and (c) A Glass-bonded VIPER Chip in the Near-eye Modality

(a) VIPER Schematic

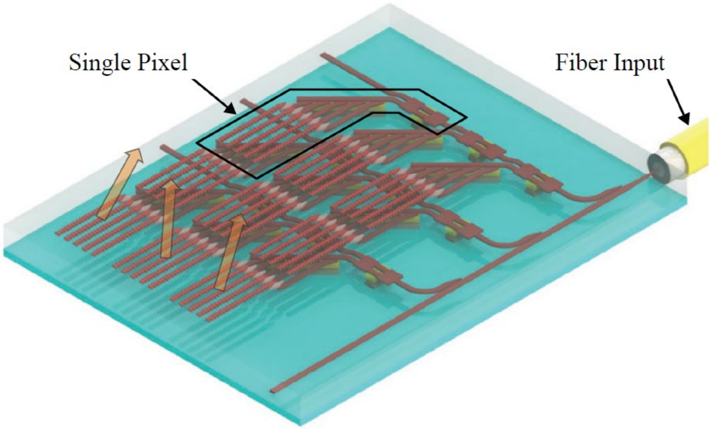

(b) VIPER Viewing Configuration

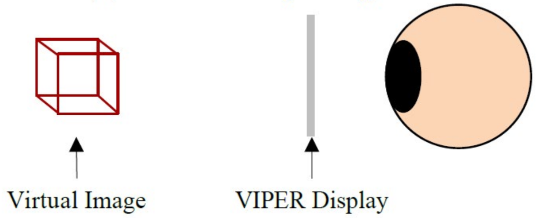

Figure 54: Simplified Schematic of the VIPER Display Showing the Input Optical Fiber, Active Distribution Network, and Grid of Optical-phased-array-based Pixels (a) and (b) Viewing Configuration for the VIPER Display Showing the Virtual HOLOGRAPHIC IMage Formed Behind the Display

6.2 Passive Architecture and Experimental Results

As an initial proof of concept demonstration, a passive VIPER display was designed at MIT and fabricated in a CMOS-compatible foundry process at SUNY Poly. The display is based on a grid of coherent visible-light integrated optical phased arrays that act as the pixels in the display, as shown in Figure 55. At the input, an on-chip inverse-taper edge coupler couples light from an off-chip laser into an on-chip single-mode silicon-nitride waveguide. A 5-stage multi-mode-interference splitter tree evenly distributes the input power to 32 rows with a final pitch of 32 µm. On each row, 32 compact optical-phased-array-based pixels are placed with a pixel pitch of 32 µm. Each pixel consists of (1) a phase taper structure [47] on the row waveguide to encode the absolute phase of the light emitted by each pixel, (2) an evanescent tap to couple light from the row waveguide to the pixel bus based on the desired pixel amplitude encoding, and (3) evanescent taps with increasing coupling lengths to uniformly distribute light from the pixel bus to 6 grating-based antennas with spatial offsets [8] to enable a linear phase gradient encoding for each pixel.

(a) System Schematic

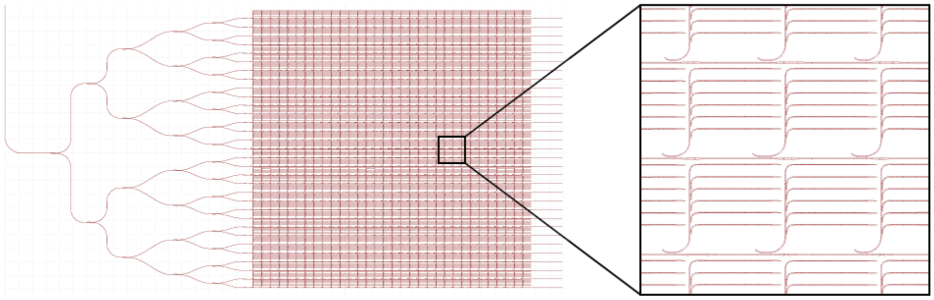

(b) Pixel Schematic

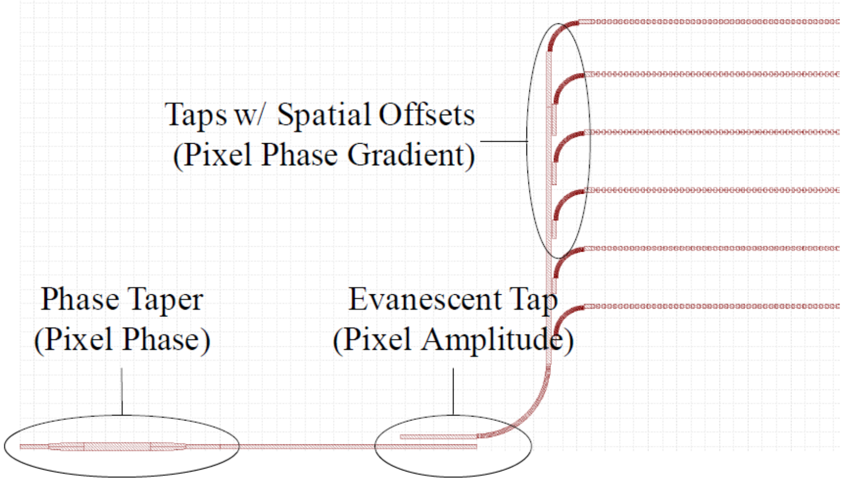

Figure 55: Schematic of the Passive VIPER Display with 32 × 32 pixels, 32 µm Pixel Pitch, 6 Antennas Per a Pixel, and 4 µm Antenna Pitch (a) and (b) Schematic of a Single Optical-phased-array-based Pixel of the Passive VIPER Display Showing the Phase Taper for Pixel Absolute Phase Encoding, Evanescent Tap for Pixel Amplitude Encoding, and Pixel-to-antenna Taps with Varying Spatial Offsets for Pixel Phase Gradient Encoding

To generate an example virtual image of a wire-frame cube using the passive VIPER display, a holographic encoding procedure was used to determine the absolute phase, amplitude, and phase gradient encodings for each pixel. Specifically, the holographic phase and amplitude distributions necessary for generating the desired image on the retinal plane were closely approximated by discretizing these distributions into local one-dimensional phase gradients with arbitrary amplitudes and absolute phases. These discretized phase and amplitude distributions were then iteratively optimized using the Gerchberg–Saxton algorithm [48] to accurately generate the desired virtual image. The resulting image and corresponding amplitude, absolute phase, and phase gradient encodings — simulated assuming a 632.8-nm operating wavelength, 1-m virtual object distance, 12-mm eye relief, and 20-mm human eye focal length — are shown in Figure 56.

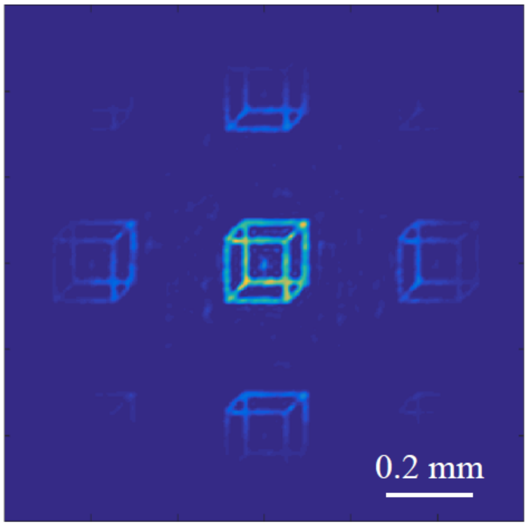
(a) Simulated Image

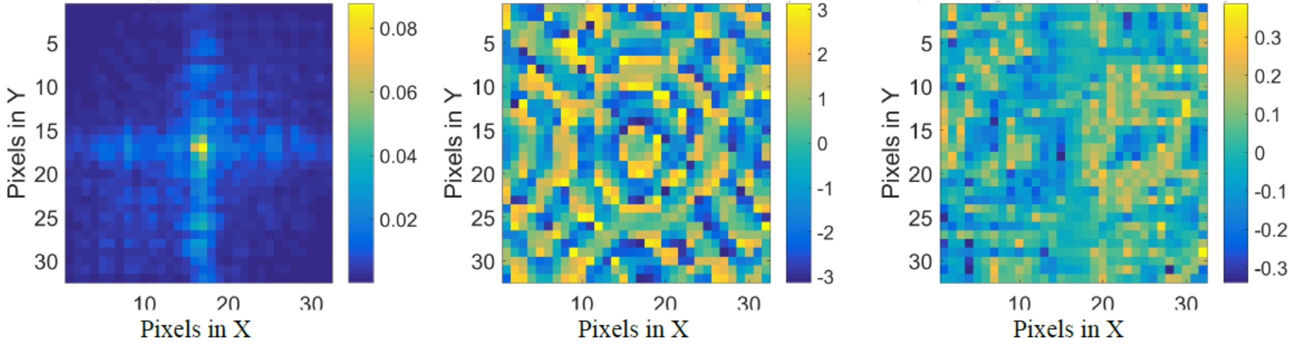
(b) Amplitude (Tap Fraction)

(c) Absolute Phase (Radians)

(d) Phase Gradient (Radians/Antenna)

Figure 56: Simulation of the Virtual Image Projected by the Passive VIPER Display and Corresponding (a), (b) Amplitude, (c) Absolute Phase, and (d) Phase Gradient Encodings, Assuming a 632.8 nm Operating Wavelength, 1 m Virtual Object Distance, 20 mm Human Eye Focal Length, and 12 mm Eye Relief

To characterize the fabricated passive VIPER display, a 632.8-nm He-Ne laser was coupled onto the chip and an optical imaging system, consisting of a 20-mm-focal-length lens and a visible-light camera, was used to emulate the functionality of the human eye (photographs of the experimental setup and fabricated chip are shown in Figure 57a-b). As expected, the display generates the desired wire-frame image at a single focal plane with a virtual object distance around 1 m and an emulated eye relief around 12 mm, as shown in Figure 57c.

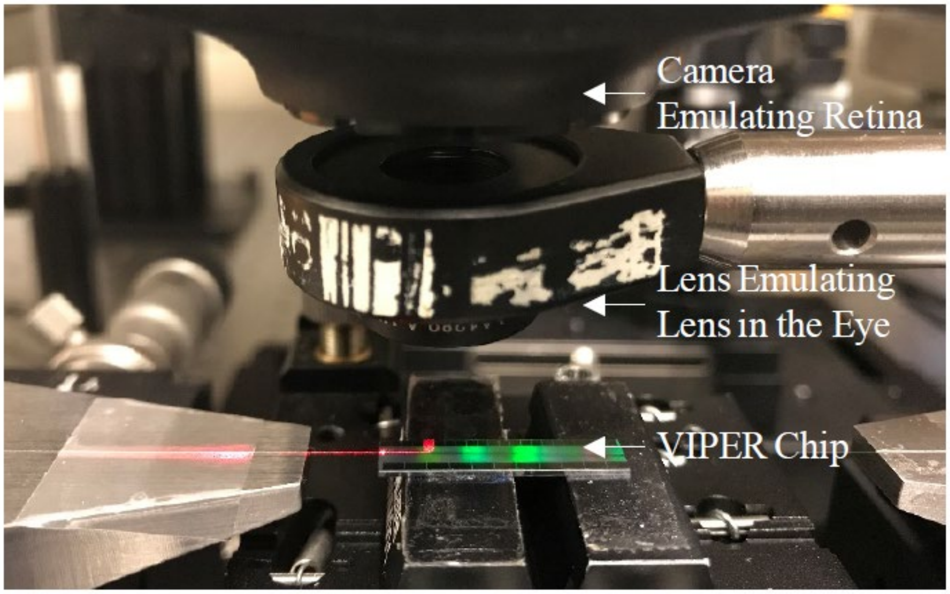
(a) Experimental Setup

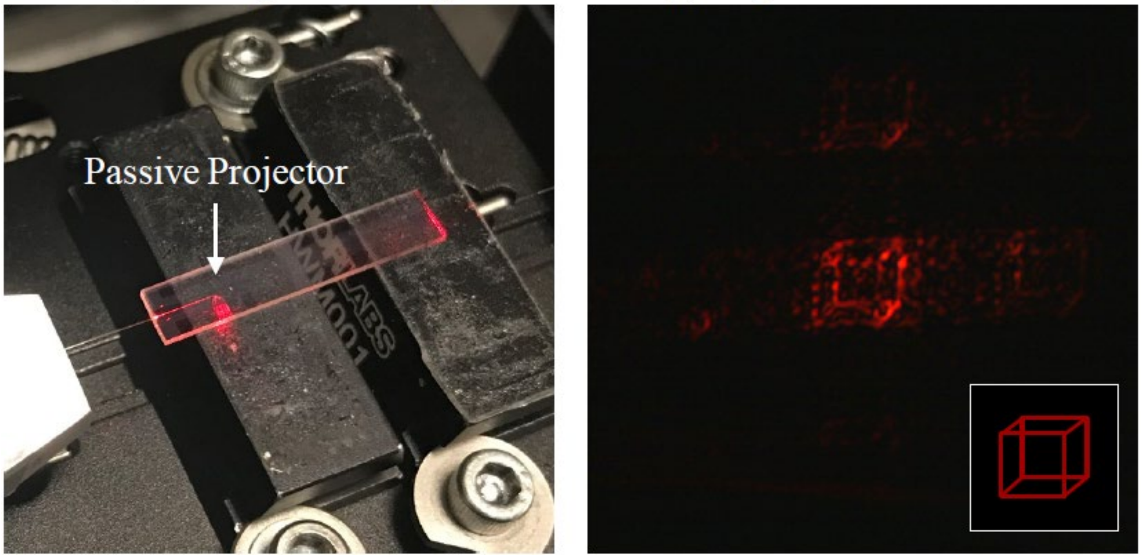
(b) Transparent Chip

(c) Experimental Results

Figure 57: Photograph of (a) the VIPER Characterization Setup Showing the Input Optical Fiber, Photonic Chip, Lens Emulating the Lens in the Eye, and Camera Emulating the Retina, and (b) the Transparent Holographic-display Chip and (c) Experimental Measurement of the Virtual Image Projected by the Passive VIPER Display with a ~1-m Virtual Object Distance, 20-mm Focal Length Lens, and 12-mm Eye Relief

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

6.3 Active Architecture and Experimental Results

Next, an active video version of the VIPER display was developed by leveraging the integrated liquid-crystal-based components discussed in Sec. 3-5. Similar to the passive VIPER display, the active display is based on a grid of coherent visible-light integrated optical phased arrays that act as the pixels in the display, as shown in Figure 58; however, in the active display, the amplitude and phase encodings for each pixel are dynamically controlled using liquid-crystal-based components.

]

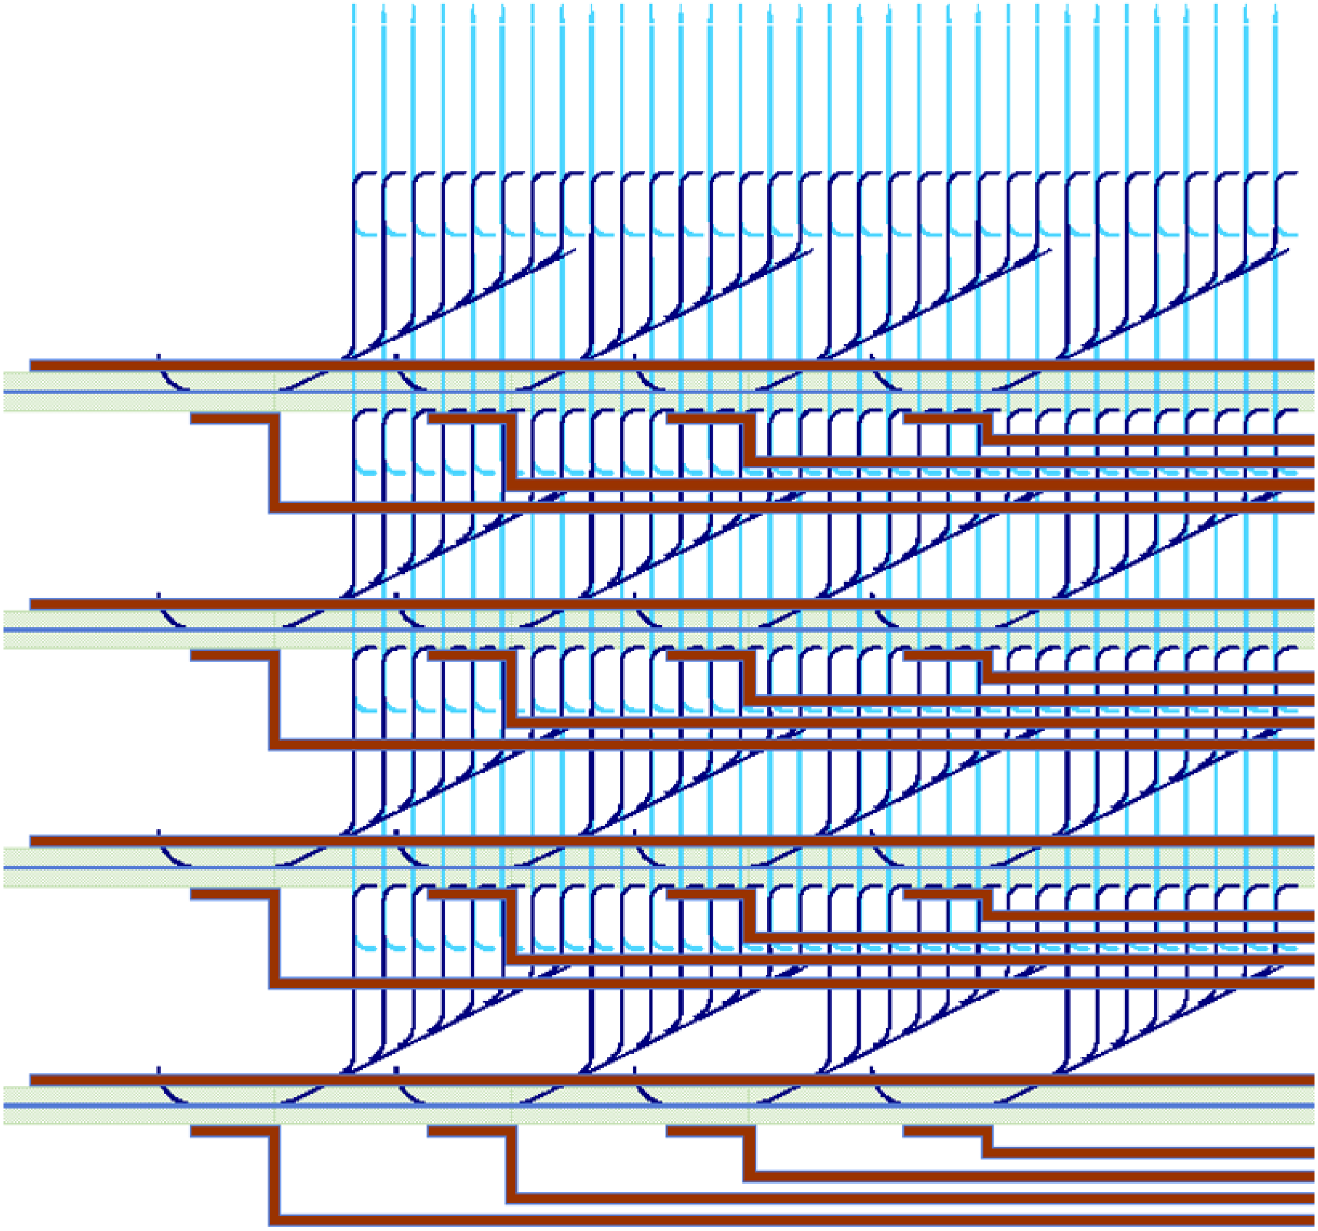

Figure 58: Partial Schematic of the Active VIPER Display with 4 × 4 pixels, 32 µm Pixel Pitch, 8 Antennas Per a Pixel, and 4 µm Antenna Pitch

At the input, an on-chip inverse-taper edge coupler couples light from an off-chip laser into an on-chip single-mode silicon-nitride waveguide. A 2-stage multimode interference splitter tree evenly distributes the input power to 4 rows with a final pitch of 32 µm. On each row, 4 compact optical-phased-array-based pixels are placed with a pixel pitch of 32 µm for a total of 16 pixels in the display.

Figure 59 shows a schematic of one of the pixels in the active VIPER display. Each pixel consists of (1) a liquid-crystal-based phase shifter [2] on the row waveguide to modulate the absolute phase of the light emitted by each pixel, (2) a liquid-crystal-based variable tap [4] to modulate the amplitude of light coupled from the row waveguide to each pixel bus, and (3) a liquid-crystal-based pixel bus with compact cascaded pixel-bus-to-antenna taps to distribute the light from the pixel bus to 8 grating-based antennas and modulate the linear phase gradient of the light emitted by each pixel [5].

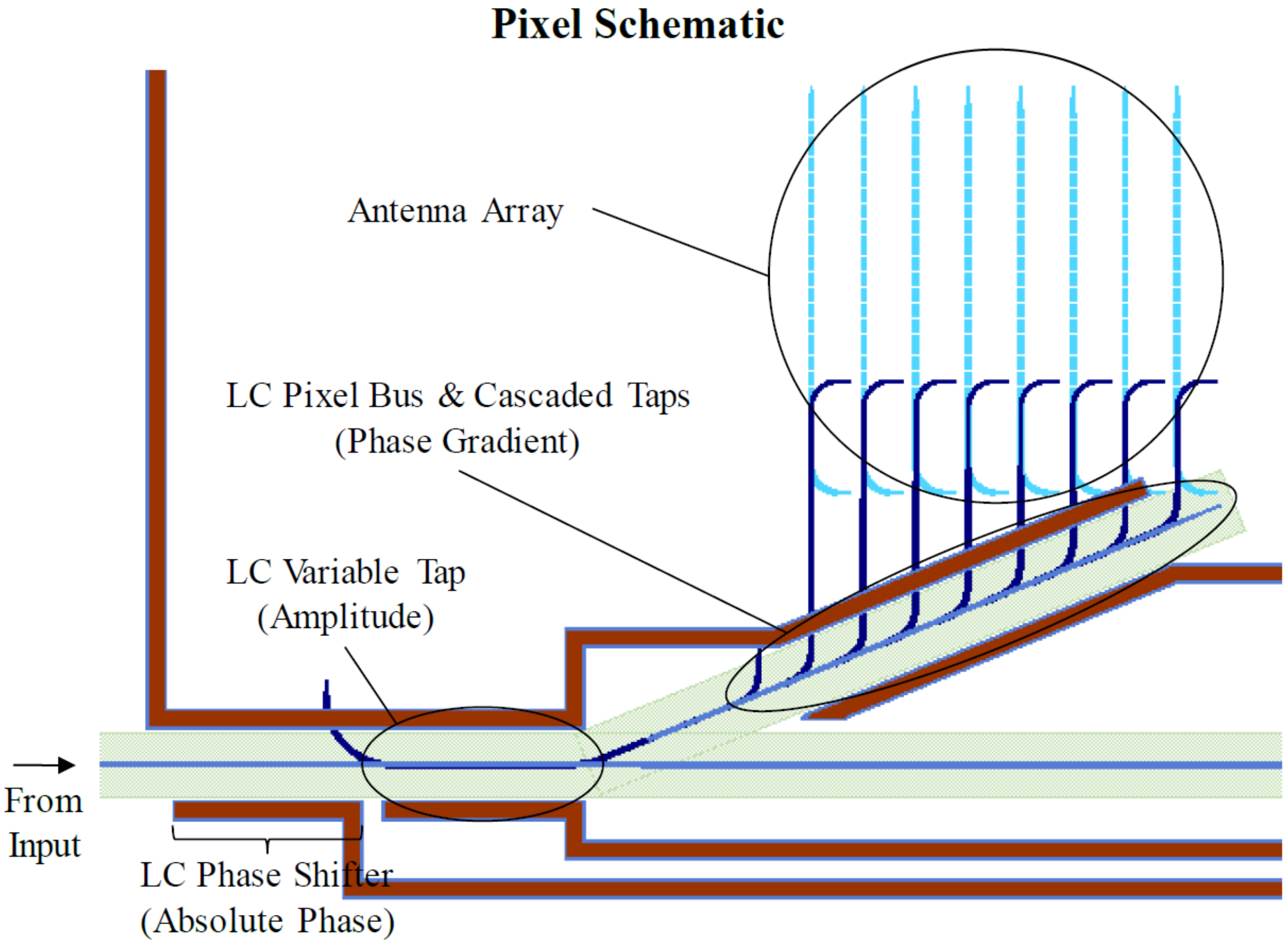

Figure 59: Schematic of a Single Optical-phased-array-based Pixel of the Active VIPER Display Showing Major Components, including the Liquid-crystal-based Phase Shifter for Pixel Absolute Phase Encoding, Liquid-crystal-based Variable Tap for Pixel Amplitude Encoding, and Liquid-crystal-based Pixel Bus with Compact Cascaded pixel-bus-to-Antenna Taps for Pixel Phase Gradient Encoding

To characterize the fabricated display, electronic probes were used to apply a 10-kHz square wave with variable peak-to-peak voltage to the active components in one of the pixels in the display. As shown in Figure 60a-b, amplitude modulation from the “off” state to the “on” state with a contrast ratio of approximately 100:1 was achieved within ±4.5 V and phase gradient modulation of 0.65π per antenna was achieved within ±4.75 V.

Finally, as an initial demonstration of the video functionality of the active VIPER display, the display was encoded to spell out the letters in the word “LIGHT”. An optical system was then used to image the near field of the chip as the display switched between letters, as shown in Figure 60c.

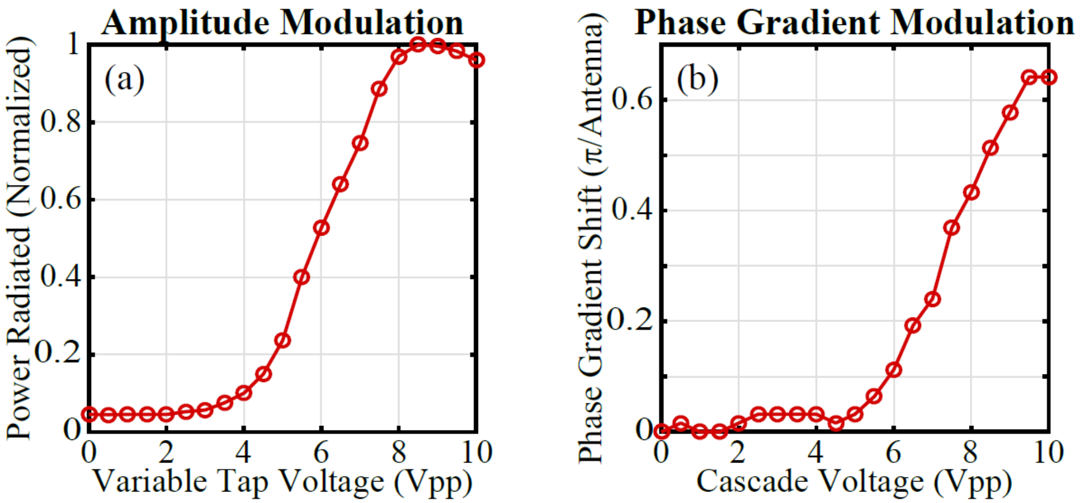

(c) Active Display Modulation

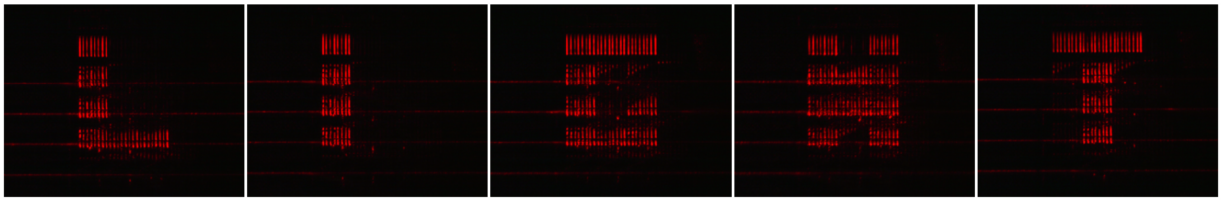

Figure 60: Experimental Results for a Single Active VIPER Pixel

Power radiated out of the pixel versus applied peak-to-peak variable tap voltage (a) and (b) pixel phase gradient shift versus applied peak-to-peak cascade voltage. (c) Experimental results showing five images of the near field of the multi-pixel active VIPER display; in this experiment, the display was encoded to spell out the letters in the word “LIGHT”

6.4 Conclusion

This work presents the first proposal and demonstration of an integrated-photonics-based visible-light near-eye holographic display. First, a novel transparent 300-mm-wafer foundry platform on glass for visible-light integrated photonics was demonstrated. Next, a passive phased-array-pixel-based architecture and a holographic image encoding methodology were developed and used to experimentally demonstrate virtual image projection of a wire-frame cube with a 32 × 32 pixel display located 1 m from the virtual object with 12 mm of eye relief. Finally, a corresponding active architecture consisting of cascaded compact active optical-phased-array-based pixels was developed and used to experimentally demonstrate initial active video display functionality.

In the future, a custom thin-film-transistor electronics backplane could be developed to enable integrated control of the active VIPER display and scaling up to a large-scale video-hologram demonstration. Furthermore, the architecture of the VIPER display could be further developed to support multi-color functionality. Finally, gain material could be integrated into the VIPER foundry platform to enable on-chip lasers for the VIPER display.

The VIPER system has the potential to enable the next generation of augmented-reality head-mounted displays — with highly-efficient components for daytime operation, compact form factors for discreet and mobile use, and vergence-accommodation-conflict resolution for long-term wear — with a wide range of applications, including military, medical, engineering, and gaming.

7 ACKNOWLEDGEMENTS

We would like to thank DARPA’s Microsystems Technology Office (MTO) and the Air Force Research Laboratory (AFRL) for supporting the VIPER project. We would especially like to thank Dr. Gordon Keeler (DARPA Program Manager) and Dr. Nicholas Usechak (AFRL Program Manager) for their continued support and helpful discussions.

This project and resulting accomplishments were made possible through the collaborative efforts of the VIPER team, including Jelena Notaros (MIT), Milica Notaros (MIT), Manan Raval (MIT), Michael Watts (MIT), Thomas Dyer (SUNY Poly), Christopher Baiocco (SUNY Poly), and Daniel Coleman (SUNY Poly).

8 REFERENCES

1. M. Notaros, J. Notaros, M. Raval, and M. R. Watts, “Integrated visible-light liquid-crystal-based phase modulator.” in preparation.
2. M. Notaros, M. Raval, J. Notaros, and M. R. Watts, “Integrated visible-light liquid-crystal phase modulator,” in *Frontiers in Optics*, pp. FW6B–5, Optical Society of America, 2018.
3. M. Notaros, J. Notaros, M. Raval, and M. R. Watts, “Integrated liquid-crystal-based variable-tap device for visible-light amplitude modulation.” in preparation.
4. M. Notaros, J. Notaros, M. Raval, and M. R. Watts, “Integrated visible-light liquid-crystal variable-tap amplitude modulator,” in *Integrated Photonics Research, Silicon and Nanophotonics*, pp. ITh2C–6, Optical Society of America, 2019.
5. J. Notaros, M. Notaros, M. Raval, and M. R. Watts, “Liquid-crystal-based visible-light integrated optical phased arrays,” in *CLEO: Science and Innovations*, pp. STu3O–3, Optical Society of America, 2019.
6. J. Notaros, M. Notaros, M. Raval, T. Dyer, C. Baiocco, and M. R. Watts, “Integrated-photonics-based holographic display for augmented reality.” in preparation.
7. J. Notaros, M. Raval, M. Notaros, and M. R. Watts, “Integrated-phased-array-based visible-light near-eye holographic projector,” in *CLEO: Science and Innovations*, pp. STu3O–4, Optical Society of America, 2019.
8. M. Raval, A. Yaacobi, and M. R. Watts, “Integrated visible light phased array system for autostereoscopic image projection,” *Optics Letters*, vol. 43, no. 15, pp. 21–24, 2018.
9. T. Agocs, T. Balogh, T. Forgacs, F. Bettio, E. Gobbetti, G. Zanetti, and E. Bouvier, “A large scale interactive holographic display,” in *Virtual Reality Conference*, p. 311, IEEE, 2006.
10. J. Hahn, H. Kim, Y. Lim, G. Park, and B. Lee, “Wide viewing angle dynamic holographic stereogram with a curved array of spatial light modulators,” *Optics Express*, vol. 16, no. 16, pp. 12372–12386, 2008.
11. T.-C. Wu, Y.-C. Chi, H.-Y. Wang, C.-T. Tsai, and G.-R. Lin, “Blue laser diode enables underwater communication at 12.4 gbps,” *Scientific Reports*, vol. 7, p. 40480, 2017.
12. S. K. Cool, K. Breyne, E. Meyer, S. C. D. Smedt, and N. N. Sanders, “Comparison of in vivo optical systems for bioluminescence and fluorescence imaging,” *J Fluoresc*, vol. 23, pp. 909–920, 2013.
13. E. Shim, Y. Chen, S. Masmanidis, and M. Li, “Multisite silicon neural probes with integrated silicon nitride waveguides and gratings for optogenetic applications,” *Scientific Reports*, vol. 6, p. 22693, 2016.
14. W. D. Sacher, X. Liu, I. F. Almog, A. Fomenko, T. Lordello, F.-D. Chen, H. Moradi-Chameh, A. Naderian, M. Chang, T. M. Fowler, T. A. Valiante, A. M. Lozano, L. C. Moreaux, J. K. S. Poon, and M. L. Roukes, “Nanophotonic neural probes for in vivo light sheet imaging,” in *Conference on Lasers and Electro-Optics*, p. SM4H.6, Optical Society of America, 2019.
15. A. Mohanty, Q. Li, M. A. Tadayon, G. Bhatt, E. Shim, X. Ji, J. Cardenas, S. A. Miller, A. Kepecs, and M. Lipson, “A Reconfigurable Nanophotonics Platform for Sub-Millisecond, Deep Brain Neural Stimulation,” *arXiv e-prints*, p. arXiv:1805.11663, 2018.
16. K. Shtyrkova, P. T. Callahan, N. Li, E. S. Magden, A. Ruocco, D. Vermeulen, F. X. Kärtner, M. R. Watts, and E. P. Ippen, “Integrated cmos-compatible qswitched mode-

locked lasers at 1900nm with an on-chip artificial saturable absorber," *Optics Express*, vol. 27, no. 3, pp. 3542–3556, 2019.

17. Y. Xing, T. Ako, J. P. George, D. Korn, H. Yu, P. Verheyen, M. Pantouvaki, G. Lepage, P. Absil, A. Ruocco, C. Koos, J. Leuthold, K. Neyts, J. Beeckman, and W. Bogaerts, "Digitally controlled phase shifter using an soi slot waveguide with liquid crystal infiltration," *IEEE Photonics Technology Letters*, vol. 27, no. 12, pp. 1269–1272, 2015.
18. H. Desmet, W. Bogaerts, A. Adamski, J. Beeckman, K. Neyts, and R. Baets, "Silicon-on-insulator optical waveguides with liquid crystal cladding for switching and tuning," *ECOCs*, vol. 3, pp. 430–431, 2003.
19. W. D. Cort, J. Beeckman, R. James, F. A. Fernández, R. Baets, and K. Neyts, "Tuning of silicon-on-insulator ring resonators with liquid crystal cladding using the longitudinal field component," *Optics Letters*, vol. 34, no. 13, pp. 2054–2056, 2009.
20. W. D. Cort, J. Beeckman, T. Claes, K. Neyts, and R. Baets, "Wide tuning of silicon-on-insulator ring resonators with a liquid crystal cladding," *Optics Letters*, vol. 36, no. 19, pp. 3876–3878, 2011.
21. A. D. Falco and G. Assanto, "Tunable wavelength-selective add–drop in liquid crystals on a silicon microresonator," *Optics Communications*, vol. 279, no. 1, pp. 210–213, 2007.
22. C. Vázquez, P. Lallana, J. Montalvo, J. Sánchez-Pena, A. D'Alessandro, and D. Donisi, "Switches and tunable filters based on ring resonators and liquid crystals," Proceedings of *SPIE*, vol. 6593, p. 65931F, 2007.
23. B. Maune, R. Lawson, C. Gunn, A. Scherer, and L. Dalton, "Electrically tunable ring resonators incorporating nematic liquid crystals as cladding layers," *Applied Physics Letters*, vol. 83, no. 23, pp. 4689–4691, 2003.
24. J. Ptasinski, S. W. Kim, L. Pang, I.-C. Khoo, and Y. Fainman, "Optical tuning of silicon photonic structures with nematic liquid crystal claddings," *Optics Letters*, vol. 38, no. 12, pp. 2008–2010, 2013.
25. Z. Zhang, Z. You, and D. Chu, "Fundamentals of phase-only liquid crystal on silicon (LCOS) devices," *Light: Science & Applications*, vol. 3, p. e213, 2014.
26. W. P. B. Jr. and L. A. Lei, "Advances in liquid crystal on silicon (LCOS) spatial light modulator technology," in *SPIE Display Technologies and Applications for Defense, Security, and Avionics*, vol. 8736, SPIE, 2013.
27. N. Collings, T. Davey, J. Christmas, D. Chu, and B. Crossland, "The applications and technology of phase-only liquid crystal on silicon devices," *Journal of Display Technology*, vol. 7, no. 3, pp. 112–119, 2011.
28. P. G. de Gennes and J. Prost, The Physics of Liquid Crystals. Oxford University Press, 1995.
29. A. Sneh and K. M. Johnson, "High-speed continuously tunable liquid crystal filter for wdm networks," *Journal of Lightwave Technology*, vol. 14, no. 6, pp. 1067– 1080, 1996.
30. J. Beeckman, K. Neyts, and P. J. M. Vanbrabant, "Liquid-crystal photonic applications," *Optical Engineering*, vol. 50, no. 8, pp. 1–18, 2011.
31. R. A. Alla, On the Control of Nematic Liquid Crystal Alignment. PhD thesis, University of Gothenburg, 2013.
32. H. Jansen, H. Gardeniers, M. de Boer, M. Elwenspoek, and J. Fluitman, "A survey on the reactive ion etching of silicon in microtechnology," *Journal of Micromechanics and Microengineering*, vol. 6, no. 1, pp. 14–28, 1996.

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

33. S. Tachi, K. Tsujimoto, and S. Okudaira, "Low-temperature reactive ion etching and microwave plasma etching of silicon," *Applied Physics Letters*, vol. 52, no. 8, pp. 616–618, 1988.
34. M. Elwenspoek, U. Lindberg, H. Kok, and L. Smith, "Wet chemical etching mechanism of silicon," in *IEEE Micro Electro Mechanical Systems An Investigation of Micro Structures, Sensors, Actuators, Machines and Robotic Systems*, pp. 223–228, IEEE, 1994.
35. A. J. van Roosmalen, "Review: dry etching of silicon oxide," *Vacuum*, vol. 34, no. 3–4, pp. 429–436, 1984.
36. V. R. Almeida, R. R. Panepucci, and M. Lipson, "Nanotaper for compact mode conversion," *Optics Letters*, vol. 28, no. 15, pp. 1302–1304, 2003.
37. J. Notaros, F. Pavanello, M. T.Wade, C. M. Gentry, A. Atabaki, L. Alloatti, R. J. Ram, and M. A. Popović, "Ultra-efficient cmos fiber-to-chip grating couplers," in *Optical Fiber Communication Conference*, p. M2I.5, Optical Society of America, 2016.
38. L. B. Soldano and E. C. M. Pennings, "Optical multi-mode interference devices based on self-imaging: principles and applications," *Journal of Lightwave Technology*, vol. 13, no. 4, pp. 615–627, 1995.
39. H. A. Haus, Waves and Fields in Optoelectronics. Prentice Hall, 1984.
40. F. Aflatouni, B. Abiri, A. Rekhi, and A. Hajimiri, "Nanophotonic projection system," *Optics Express*, vol. 23, no. 16, pp. 21012–21022, 2015.
41. D. N. Hutchison, J. Sun, J. K. Doylend, R. Kumar, J. Heck, W. Kim, C. T. Phare, A. Feshali, and H. Rong, "High-resolution aliasing-free optical beam steering," *Optica*, vol. 3, no. 8, pp. 887–890, 2016.
42. S. A. Miller, Y.-C. Chang, C. T. Phare, M. C. Shin, M. Zadka, S. P. Roberts, B. Stern, X. Ji, A. Mohanty, O. A. J. Gordillo, U. D. Dave, and M. Lipson, "Large-scale optical phased array using a low-power multi-pass silicon photonic platform," *Optica*, vol. 7, no. 1, pp. 3–6, 2020.
43. C. V. Poulton, M. J. Byrd, M. Raval, Z. Su, N. Li, E. Timurdogan, D. Coolbaugh, D. Vermeulen, and M. R. Watts, "Large-scale silicon nitride nanophotonic phased arrays at infrared and visible wavelengths," *Optics Letters*, vol. 42, no. 1, pp. 21–24, 2017.
44. B. Kress and T. Starner, "A review of head-mounted displays (HMD) technologies and applications for consumer electronics," in *Photonic Applications for Aerospace, Commercial, and Harsh Environments IV*, vol. 8720, p. 87200A, International Society for Optics and Photonics, 2013.
45. C. Martinez, V. Krotov, B. Meynard, and D. Fowler, "See-through holographic retinal projection display concept," *Optica*, vol. 5, no. 10, pp. 1200–1209, 2018.
46. G.-Y. Lee, J.-Y. Hong, S. Hwang, S. Moon, H. Kang, S. Jeon, H. Kim, J.- H. Jeong, and B. Lee, "Metasurface eyepiece for augmented reality," *Nature Communications*, vol. 9, no. 1, pp. 1–10, 2018.
47. J. Notaros, C. V. Poulton, M. J. Byrd, M. Raval, and M. R. Watts, "Integrated optical phased arrays for quasi-Bessel-beam generation," *Optics Letters*, vol. 42, no. 17, pp. 3510–3513, 2017.
48. J. Zhou, J. Sun, A. Yaacobi, C. V. Poulton, and M. Watts, "Design of 3D hologram emitting optical phased arrays," in *Integrated Photonics Research, Silicon and Nanophotonics*, p. IT4A.7, Optical Society of America, 2015.

APPENDIX A:
SUNY POLY PHOTONICS FABRICATION SUMMARY

8.1 Program Overview

This appendix covers the wafer fabrication services provided by SUNY Poly for DARPA's VIPER Program. The objective of the VIPER program was to develop visible-light integrated photonics technology for wearable augmented reality eyewear. MIT managed the program, provided the integrated photonics designs, performed specialized post-processing on the integrated photonics chips, and characterized the performance of the completed optical devices.

SUNY Poly fabricated the integrated photonics circuits at their 300-mm wafer development line in Albany, NY. The wafer fabrication phases of the VIPER program spanned August 2017 through November 2020. The program called for the fabrication of four fully integrated 300-mm product builds and a single-layer photonics product on 300-mm optical-grade fused silica wafers instead of silicon. The purpose of the fused silica substrates was to provide optical transparency through the completed product to meet the requirements of its end use. In the course of wafer processing, various enabling techniques we developed and implemented to overcome unanticipated fabrication challenges that arose.

The application wavelength for the VIPER test vehicles was 633 nm. This dictated the use of silicon-nitride ("nitride") for the optical components since it is transparent in the visible spectrum. The silicon nitride material used in this program was formed using plasma-enhanced chemical vapor deposition (PECVD). The components were clad in silicon dioxide ("oxide").

8.2 Visible-Light Integrated Photonics

8.2.1 Processes

Two base process flows were required for the fabrication of the four VIPER products. The first flow ("Cobra") had seven mask layers. The second flow ("Python") had the same seven Cobra layers plus an additional two layers.

The seven mask layers common to Cobra and Python consisted of four core photonics layers used for the key components of the photonics integrated circuits (PICs); one process-assist layer; and two layers to permit optical and electrical probing of the circuits.
The four core photonics mask levels were:

i) BN: the first of two 160-nm thick nitride waveguide levels,

ii) TN: the second of two 160-nm thick nitride waveguide levels,

iii) TS: a trench level for clearing out oxide cladding above selected waveguide sections, and

iv) M1: a metal level for forming electrodes on each side of the trenches.

The one process-assist level was:

v) SN: a nitride etch-stop level to aid in trench formation.

The two levels to enable optical and electrical probing were:

vi) PA: a pad opening level for clearing oxide from the metal pads so a voltage bias can be applied to the electrodes, and

vii) DI: a level for cleanly etching through the surface material around the dicing channels to smoothly cut the ends of the waveguides for the efficient edge coupling of light into them.

Both Cobra and Python used these seven mask levels. Cobra used just these seven and Python used these seven plus an additional two. The additional two layers in Python were optical "emitter" layers. The mask designations for these levels were ZN and FN.

8.2.1.1 Cobra

Figure 61 shows the cross-section for the Cobra layers. It shows the two 160-nm thick waveguide levels (BN and TN); the 160-nm thick trench reactive ion etching (RIE) etch-stop layer (SN); the 800-nm deep trench level (TS); and the 820-nm thick aluminum electrode level (M1). Oxide encases all of these structures as illustrated in yellow. Also shown here is a 2-μm thick layer of oxide used between the silicon substrate and the first waveguide layer. This oxide was to provide optical isolation between the waveguides and the supporting silicon substrate. The figure does not show the pad openings (PA) or chip edge-cut (DI) levels.

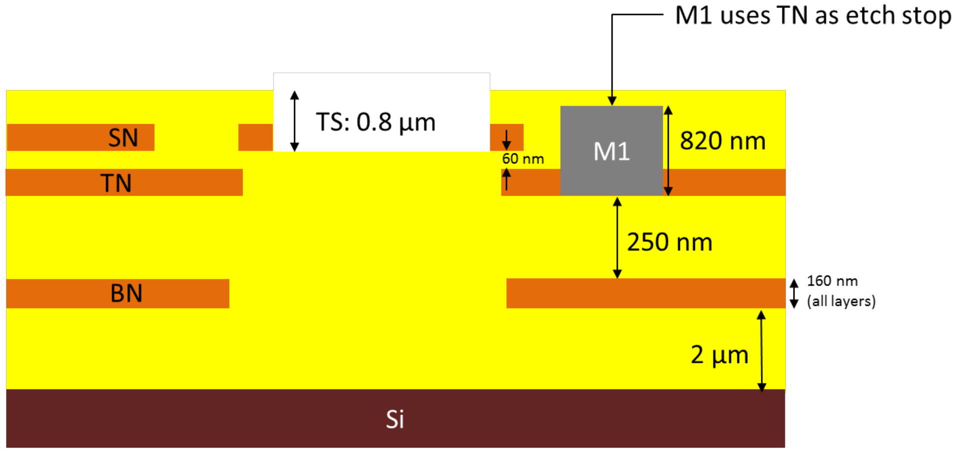

Figure 61: Cobra Stack Depiction

8.2.1.2 Python

In addition to the seven levels used in Cobra, Python had two additional "emitter" levels. Figure 62 shows the build cross-section for Python with eight of the nine levels shown (DI is not shown). The figure shows the two 60-nm thick emitter levels (ZN and TN); the two waveguide levels (BN and TN); the etch-stop layer (SN); the trench layer (TS); and the electrode layer

(M1). Again, all of these structures are encased in oxide as indicated in yellow. Also shown here is the 2-µm thick oxide layer between the substrate and the first emitter level to provide optical isolation of the optical components from the substrate. The absence of oxide above the M1 aluminum illustrates the pad opening level (PA). The figure does not show the chip edge-cut (DI) level.

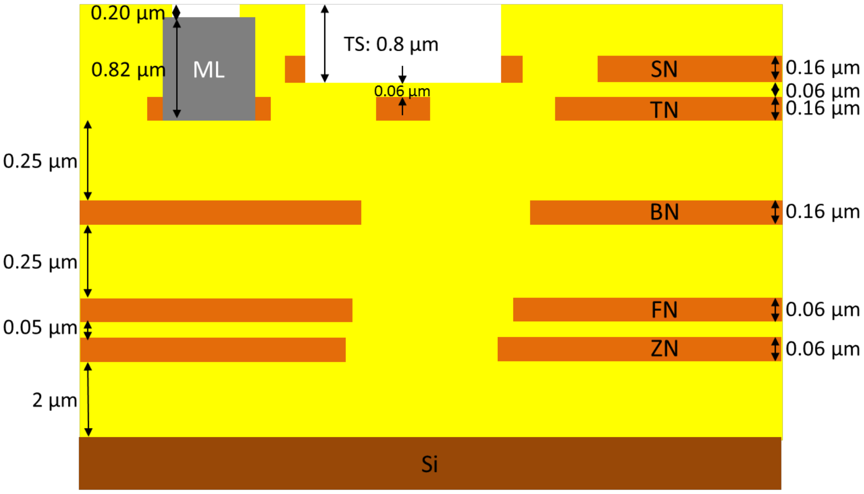

Figure 62: Python Stack Depiction with Embedded Cavity Waveguide

A variation on the core Python process was also used. In this variation, the TS trench was etched further down to a level even with the bottom of the TN waveguide instead of stopping at the SN etch-stop layer above the TN waveguide. The SN etch-stop layer was not needed and was removed in this build.

Figure 63 illustrates this alternative Python 1 "reveal" build. It shows the 1020-nm deep TS trench extended to the bottom of the TN waveguide. It also shows the absence of an SN etch-stop layer since it is not used in this version of Python.

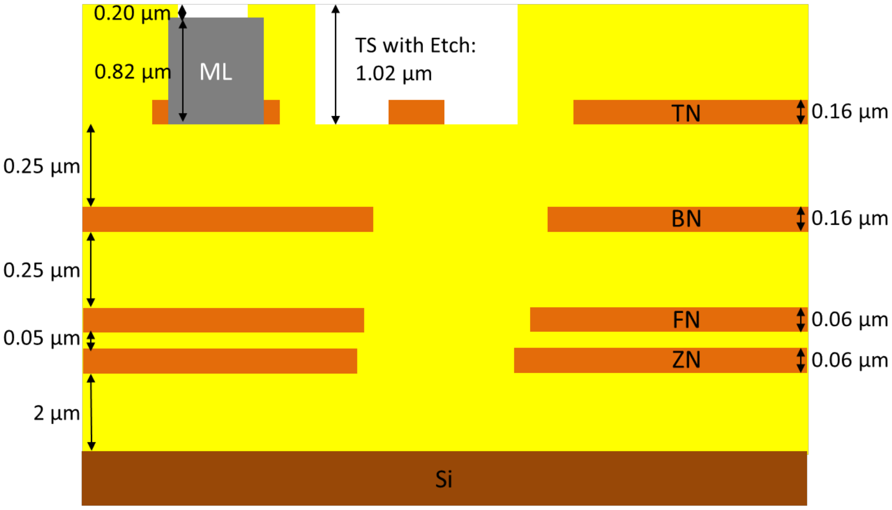

Figure 63: Python Stack Depiction with Revealed Cavity Waveguide

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

8.2.2 Products

The VIPER program required four integrated product builds: the first was a Cobra build ("Cobra 1") and the remaining three were Python builds ("Python 1", "Python 2", and "Python 3"). These four product builds were managed in two phases. The first three products (Cobra 1, Python 1, and Python 2) were managed in Phase 1 and the last product (Python 3) was managed in Phase 2.

For all of the products, MIT provided the designs and SUNY Poly merged the MIT design data with a standard frame that contained all the necessary alignment marks, overlay structures, critical dimension measurement structures, film-thickness measurement sites, and other standard marks and structures needed for processing and in-line monitoring.

Since all of the layouts contained a wide range of pattern densities, SUNY Poly also added dummy pattern to fill the un-patterned areas to smooth out the overall pattern density and ensure process uniformity for chemical mechanical planarization (CMP) and RIE processes.

8.2.2.1 Cobra

The single Cobra product was Cobra 1. Figure 64 illustrates the MIT Cobra 1 design. The BN waveguide level of the design is shown here without the SUNY Poly frame or fill pattern added.

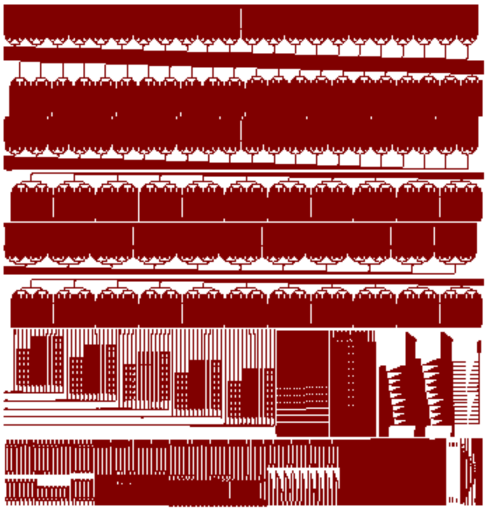

Figure 61: Cobra1 Design provided by MIT

8.2.2.2 Python

There were three Python products built for the program. Figure 65 illustrates the MIT Python 1, Python 2, and Python 3 designs. Again, only the BN waveguide levels are shown and the SUNY Poly frame and fill pattern are left out. All of the Python designs were similar in overall pattern density for the various levels.

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

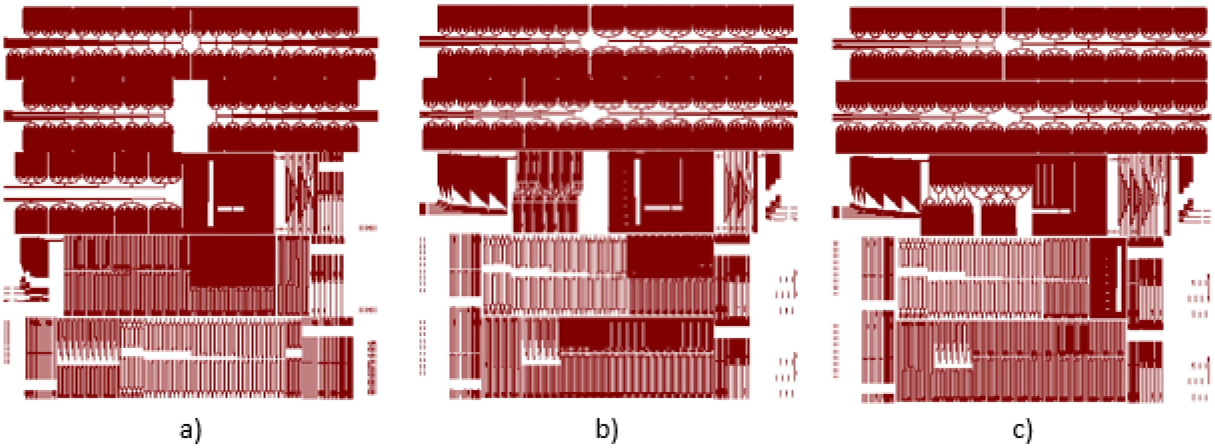

Figure 62: MIT Reticle Designs for a) Python1, b) Python2 and c) Python3

8.2.3 Wafer Processing

This section contains process descriptions and data collected in the course of processing the four VIPER integrated builds. Some processes were legacy photonics processes and others were newly developed specifically for VIPER. Since all of the product builds have a majority of process steps in common, no distinction is made here for the processing of the different products.

8.2.3.1 Waveguides

Figure 66 is a cross-section TEM showing waveguide nitride and cladding oxide layer thicknesses. The BN, TN, SN nitride layers are shown along with the intervening oxide layers and thickness measurements are superimposed on the image. This cross-section was taken through a film thickness measurement site where the films are measured in-line by ellipsometry.

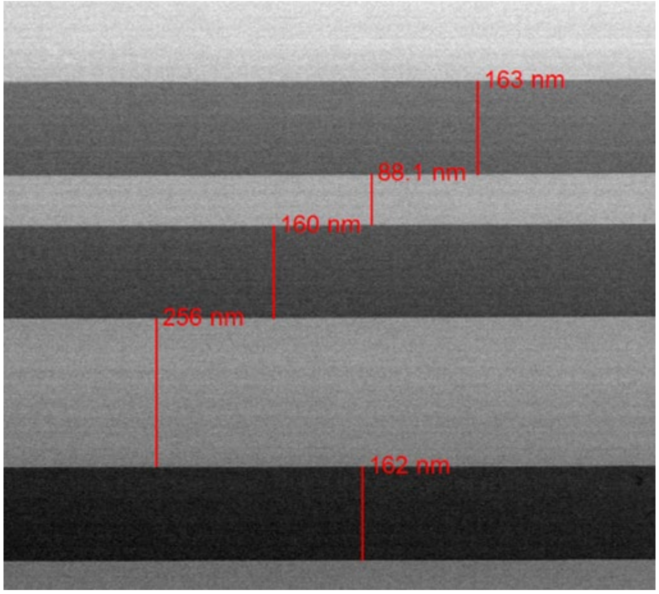

Figure 63: Cross-section TEM Showing Waveguide Nitride and Cladding Oxide Layer Thicknesses

The cross-section TEMs in Figure 67 and Figure 68 are a low magnification image of ZN, FN, BN, TN, SN, ML, and TS and a high magnification image of ZN, TN, BN, and TN while Figure 69 shows two locations of the BN layer patterned in waveguide arrays.

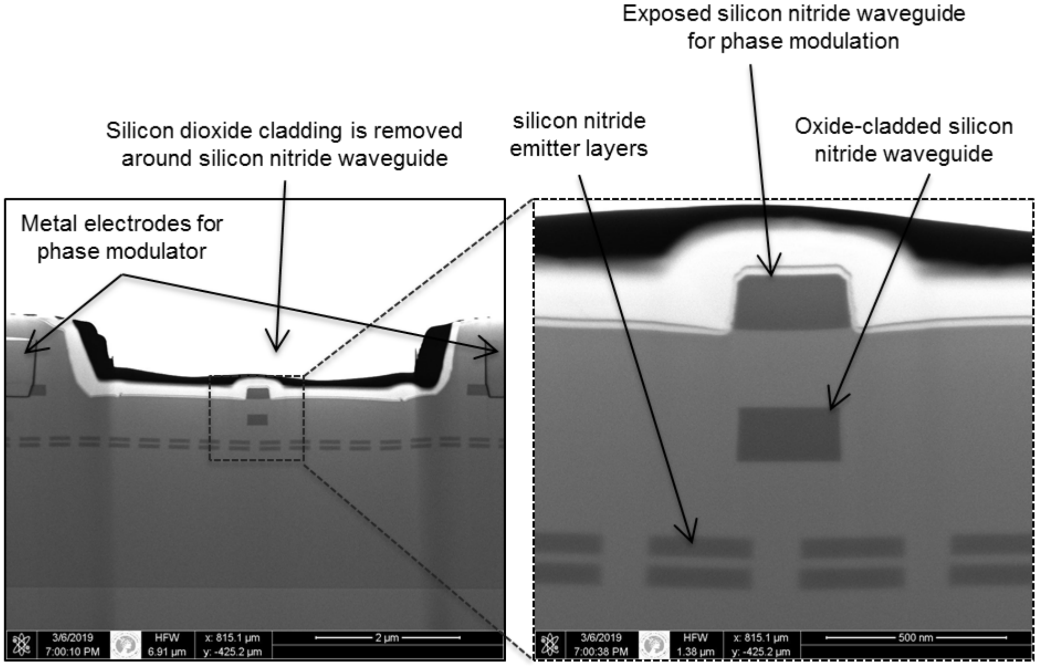

Figure 64: Low Magnification Image of ZN, FN, BN, TN, SN, ML, and TS and a High Magnification Image of ZN, TN, BN, and TN

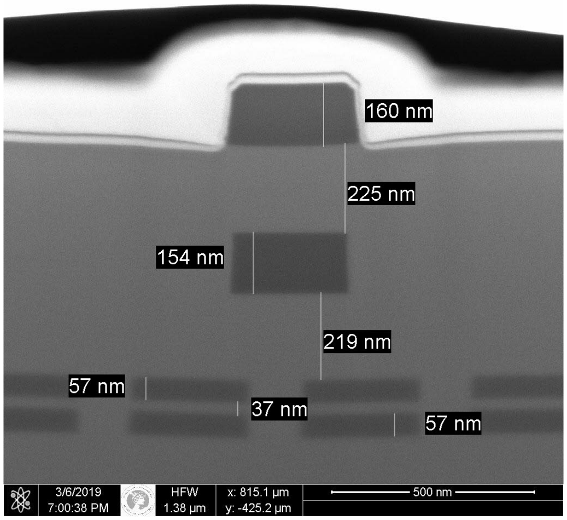

Figure 65: Close-up Image Again with Measurements of the Different Film Layers

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

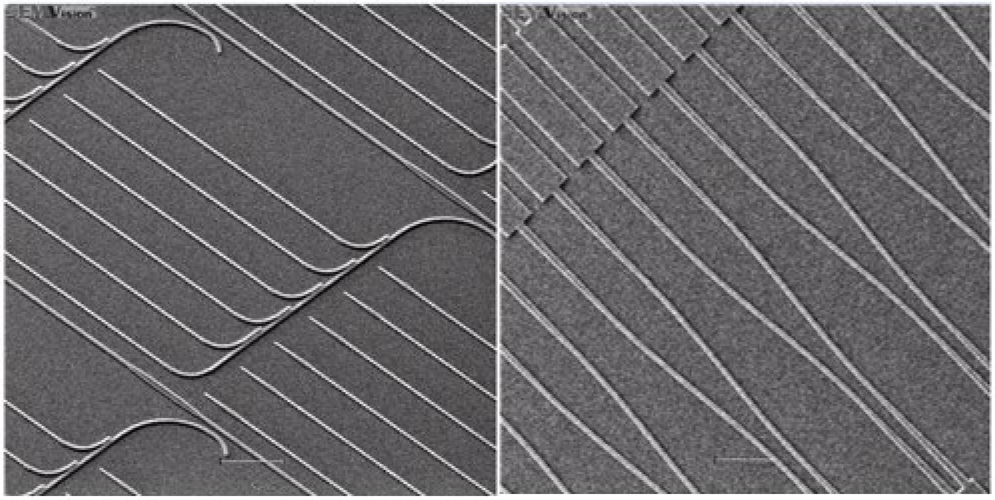

Figure 66: Tilted Inline Scanning Electron Microscope (SEM) Images of BN Waveguides

8.2.3.2 Trenches

Both versions of the TS trench are described below. The TS trench that stops at the SN etch-stop layer above the TN waveguide and keeps the waveguide buried is referred to as the "standard" trench. The version that extends further down and reveals the underlying TN waveguide is referred to as the "reveal" trench.

8.2.3.3 Standard

The series of images displayed in Figure 70 are tilt-SEM images of the standard trench. In the high magnification shot on the right, a high degree of sidewall roughness is evident and the SN etch-stop nitride is delineated from the oxide above it.

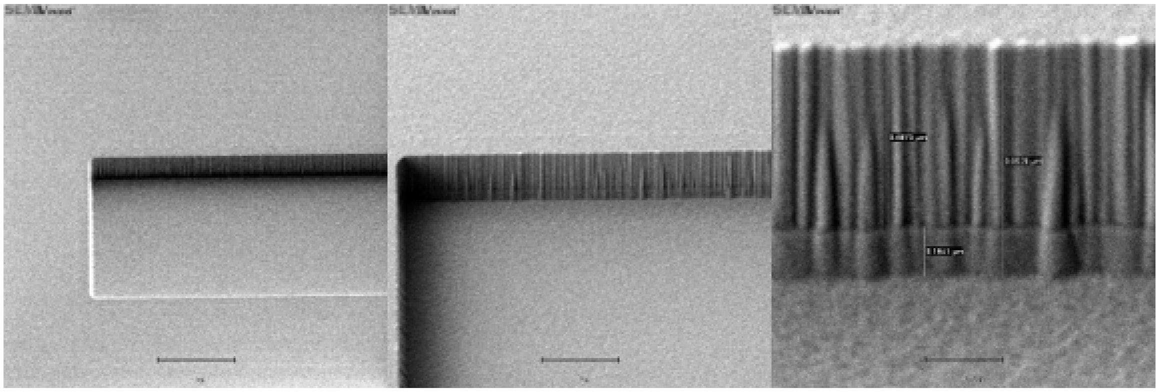

Figure 70: Tilted Inline Scanning Electron Microscope (SEM) Images of TS Trenches

Figure 71 and Figure 72 are some cross-section TEMs of the aluminum electrodes, TS trench, and TN waveguide under the trench. An early problem encountered in the program was incomplete etch of the SN etch-stop layer after the oxide was cleared and the nitride-open etch process completed. The challenge here was the difficult balance between clearing the entire thickness of nitride but not over etching which would erode the thin oxide layer between the bottom of the SN nitride and top of the TN waveguide. This problem was resolved by creating a more robust endpoint algorithm for ending the nitride etch process and instituting an energy-dispersive X-ray (EDX) inspection of the bottom of the trenches after RIE to catch any residual nitride as it occurred so corrections could be implemented quickly if needed.

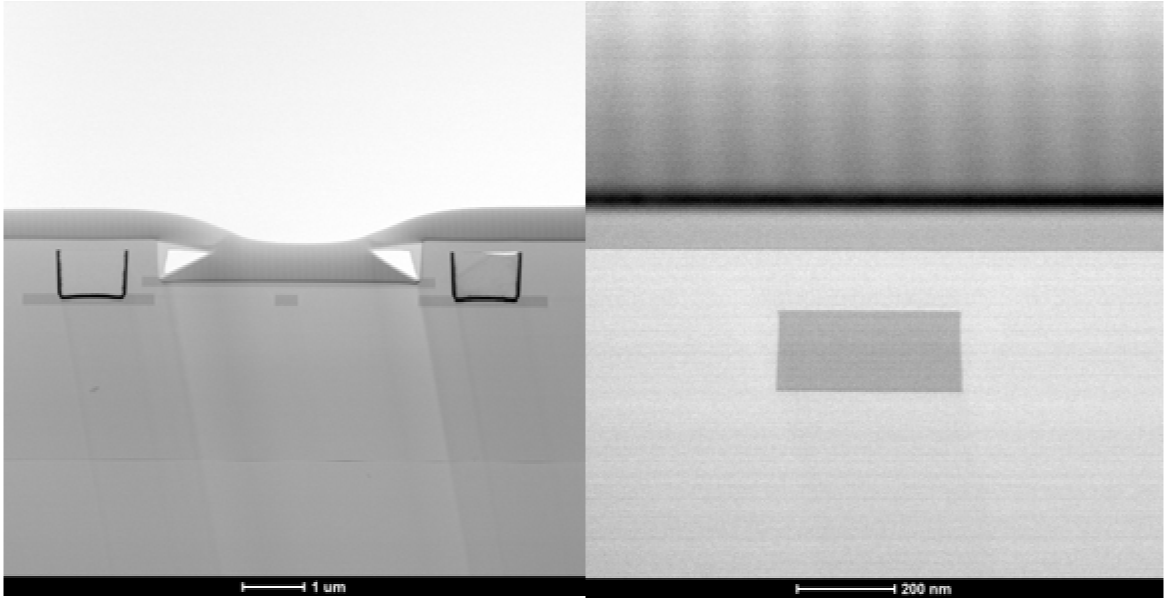

Figure 77: Cross-sectional TEM of the Standard TS Trench Over an Embedded TN Waveguide

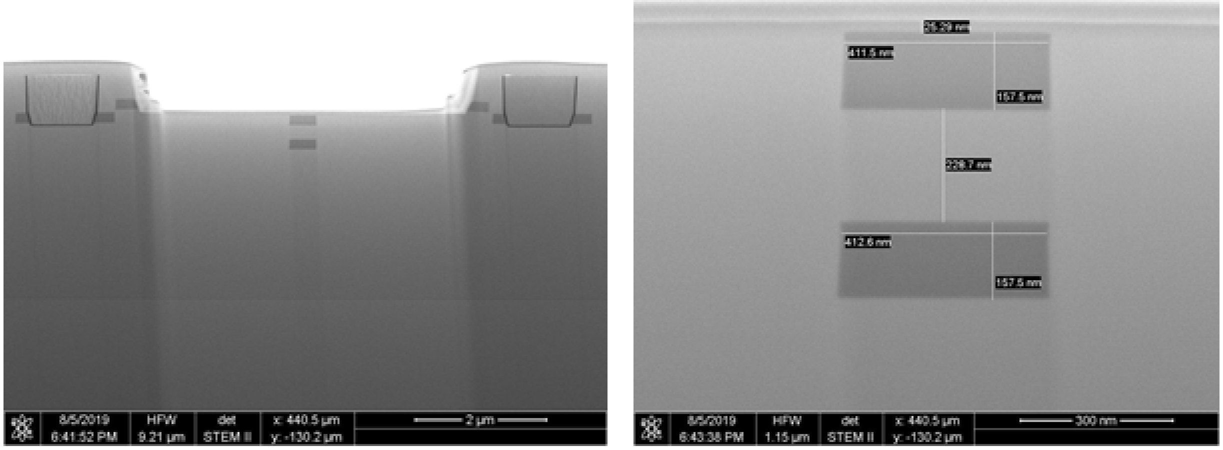

Figure 72: Cross-section TEM of the Standard TS Trench Opening Over Stacked TN and BN Waveguides

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

8.2.3.4 Reveal

For the reveal trench, the trench depth was extended to the bottom on the TN waveguide to fully expose the top and sides of the waveguide. The images displayed in Figure 73 and Figure 74 illustrate the TN waveguides after they were revealed with the TS trench reveal RIE.

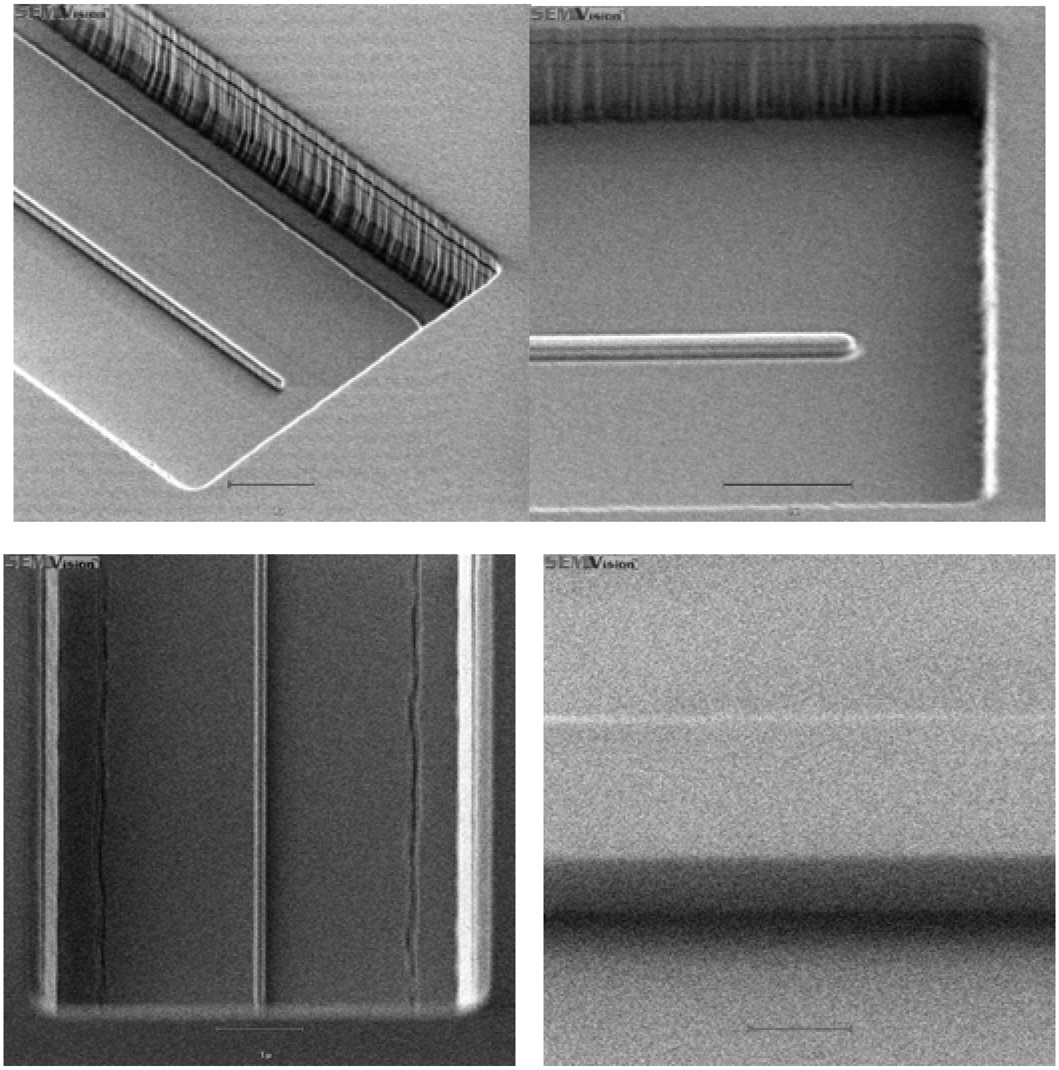

Figure 73: Various Inline Tilted SEM Images of a Revealed WG Inside of a TS Trench Cavity

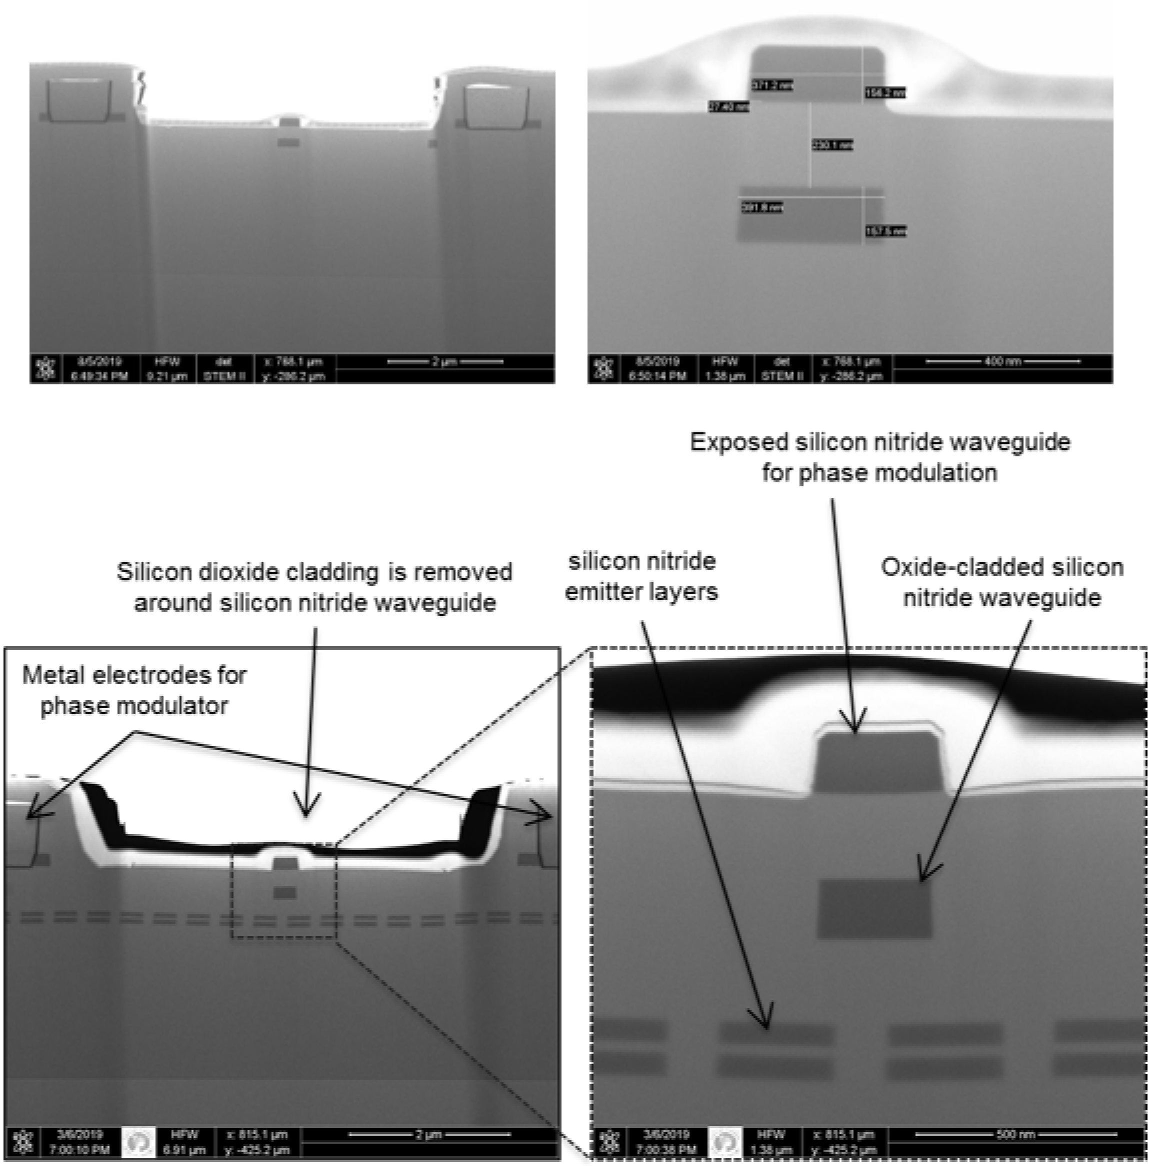

Figure 74: Cross-sectional TEMs of Revealed TN Waveguides Inside of a TS Trench Cavity

8.2.3.5 Electrodes

Electrodes were formed using a damascene aluminum process. Figure 75 illustrates the trench prior to electrode metallization and planarization. The M1 trench etch process is the same as the standard TS trench etch. In the case of M1, the TN nitride waveguide serves as an etch-stop analogous to SN in the standard TS RIE. Optical images of the metal electrodes on each side of the TS trench at the end of wafer processing are shown in Figure 76 and Figure 77. Figure 78 shows the passivation opening as well as the grain structure of the aluminum electrode.

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

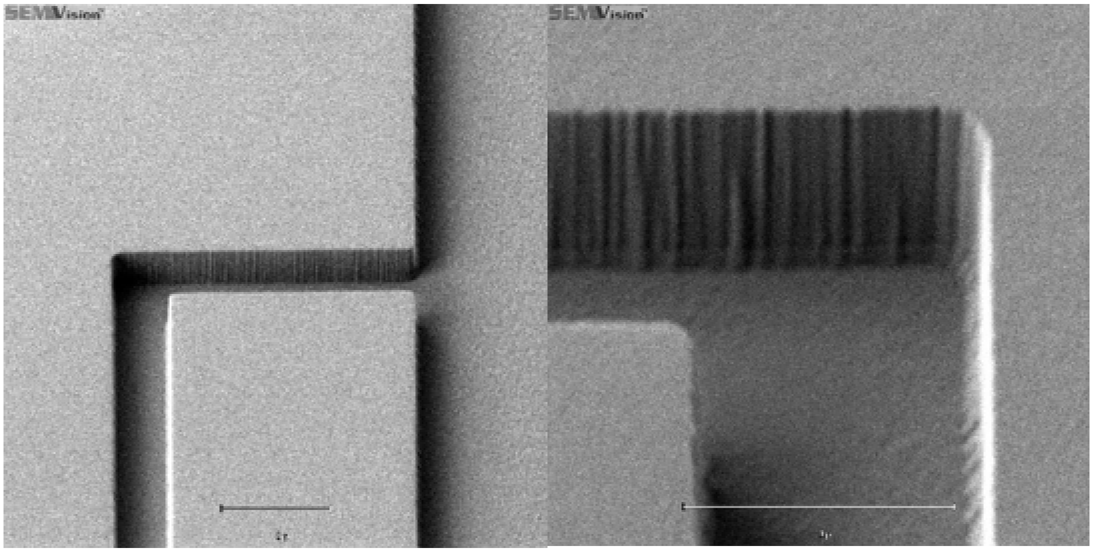

Figure 75: Tilted Inline SEM Images of the Metal Electrode Trench Prior to Metallization

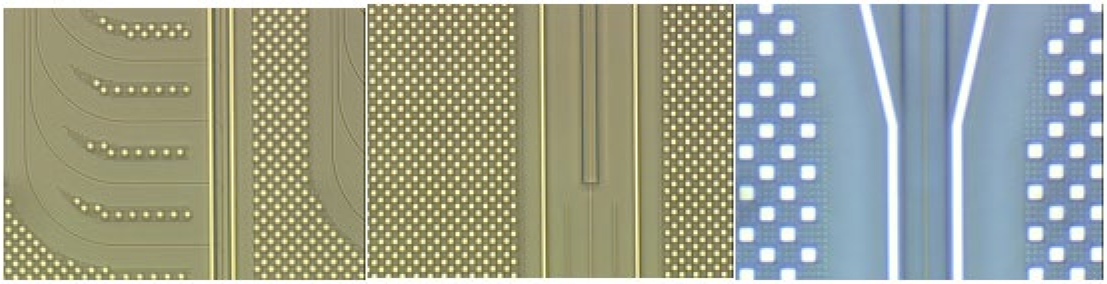

Figure 76: Optical Micrographs of the Electrodes and Dummy Fill Shapes

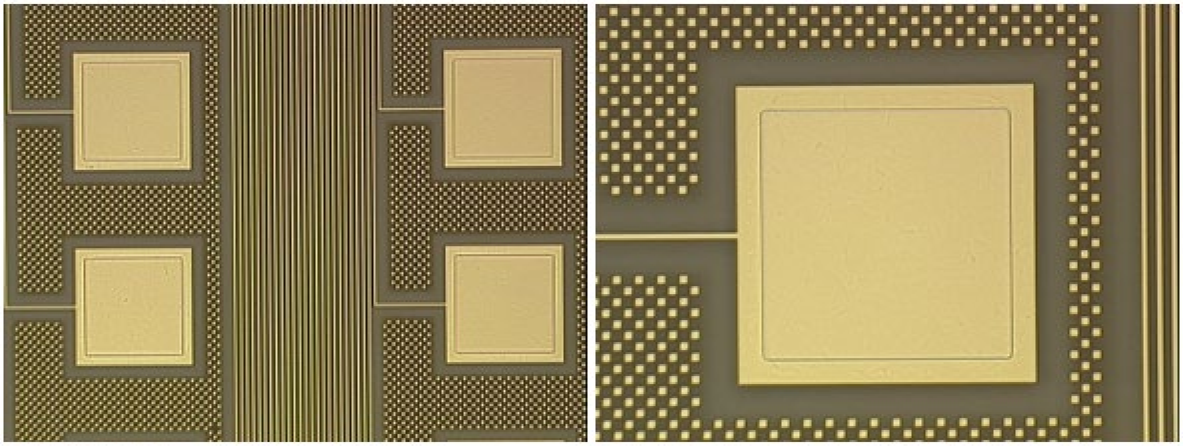

Figure 77: Optical Micrographs of Metal Probe / Bond Pads with the Liner/Passivation Opening (PA)

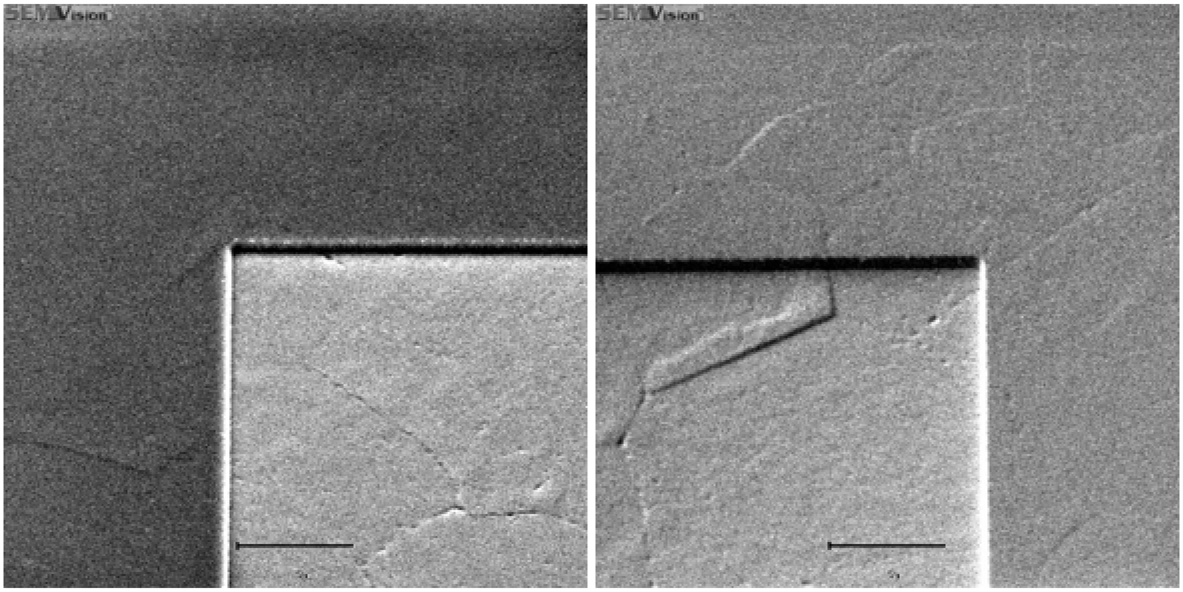

Figure 78: Inline Tilted SEM Images of the PA Opening Exposing the Aluminum Probe/Bonding Pad

8.2.3.6 Edge-Cut Channel

A dicing trench patterning level serves to define the dicing channels and provide a clean facet for coupling into the waveguide tapers. Figure 79-81 illustrate various trench geometries and taper regions. The protrusion at the tapers is attributed to a slightly different etch rate between oxide and nitride in the edge-cut etch process.

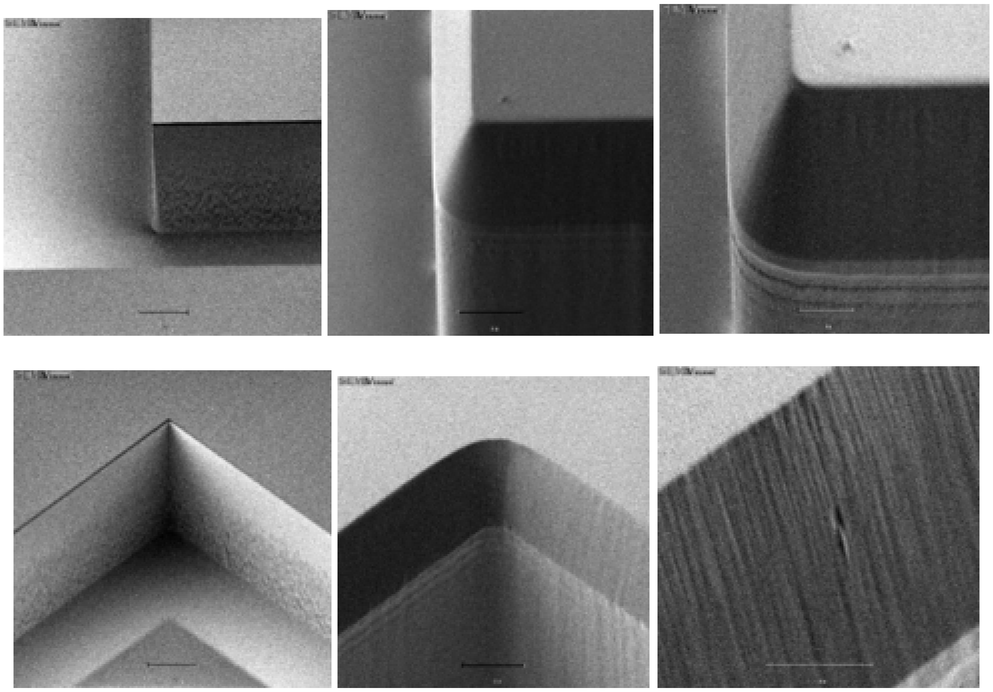

Figure 79: Inline Tilt-SEM Images of the DI Trenches

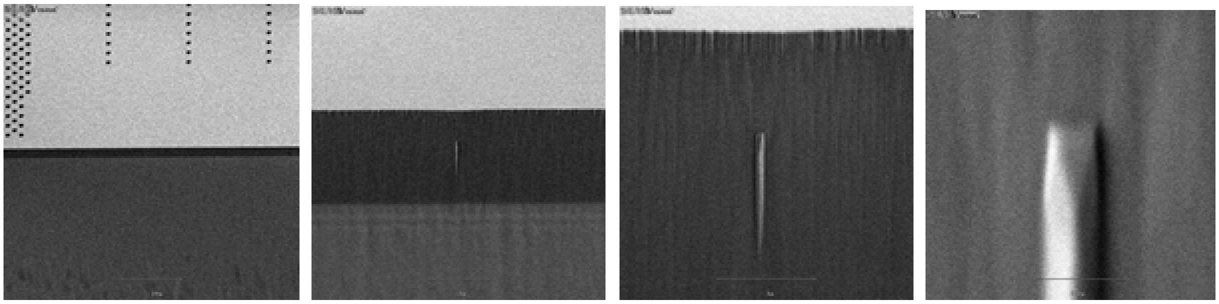

Figure 810: Inline Tilt-SEM Images of the SiN Taper / DI-Trench Intersection and Normal-Twist

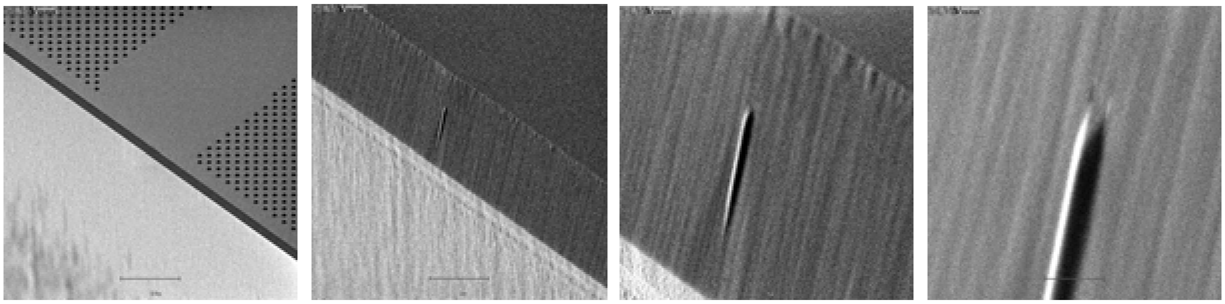

Figure 811: Inline Tilt-SEM Images of the SiN Taper / DI-Trench Intersection and with a 45-Degree Twist

8.3 Visible-light PICs on Fused Silica

The integrated photonic circuits in augmented reality (AR) eyewear must not obscure the vision of the user which presents a challenge when working with traditional plaforms. As a result, this work makes use of silicon nitride waveguides and silicon dioxide cladding which are inherently transparent in the visible spectrum. The silicon wafer substrates they are built on are not transparent at these wavelengths. The photonics layers in AR glasses need to be supported on transparent substrates to be useful.

The photonics layers could be fabricated on transparent substrates if the processing equipment could handle such substrates. Unfortunately, high volume wafer processing tools and their robot handlers are designed to work with opaque wafers. Various sensors used for wafer handling and detection reflect light off the wafer surface to detect the presence and positioning of the wafers. It would be an expensive and time-consuming task to retrofit all the needed wafer fabrication equipment to handle transparent substrates.

This problem can be circumvented by building all of the integrated photonics layers on silicon wafers and then transferring them to transparent substrates after the circuits are built. The key challenge for such a thin layer transfer process would be the mechanical handling of the very thin, delicate layer containing the photonics without damaging the structures inside.

8.3.1 Overview

Phase 3 of the VIPER program was an investigation of the use of optical-grade fused silica to replace silicon as the substrate holding the photonics circuits. In this phase, feasibility tests for handling and processing fused silica wafers were performed; processes for transferring the very thin photonics layers to fused silica without damaging the critical structures inside were developed; and successful demonstrations of the transfer of VIPER product layers to 300-mm fused silica wafers were accomplished.

8.3.2 Wafer-Handling Test

A process sequence for transferring the thin photonics layers from the silicon wafers to fused silica wafers was defined early in the program. This sequence required the use of a single 300-mm processing tool, a wafer bonder, to be able to handle and process the transparent wafers as the last step in the sequence. The ability of the bonder to handle the fused silica wafers was demonstrated early on and is depicted in Figure 82.

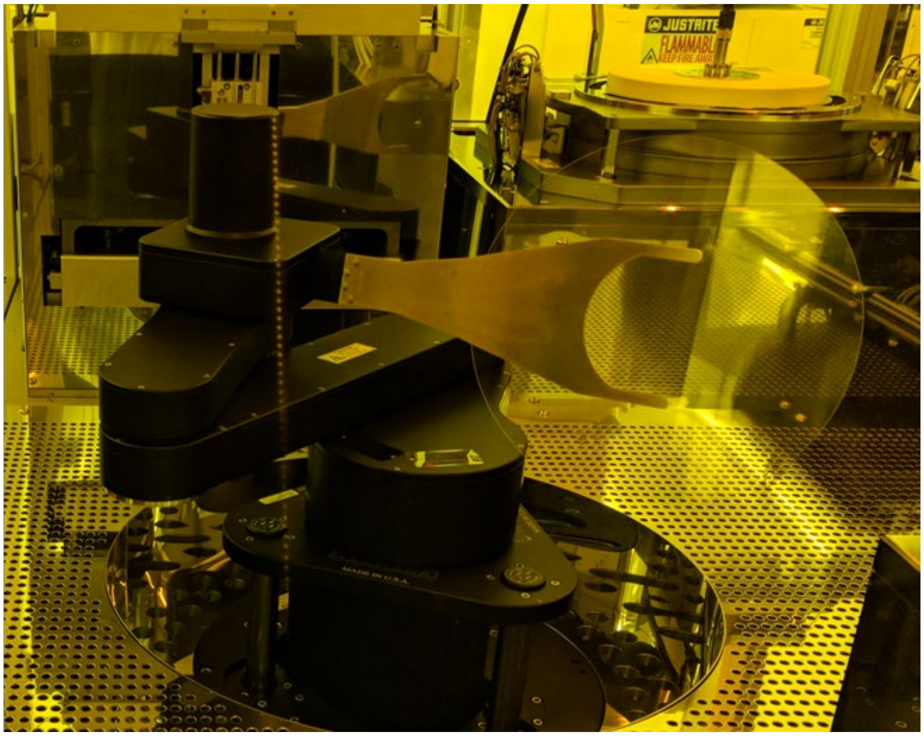

Figure 812: Initial Glass Wafer-handling Test for the 300-mm Wafer Bonder

8.3.3 VIPER PIC on Fused Silica

Ultimately, a process was developed for successfully transferring photonics layers to the fused silica wafers. The chip yield for the layer transfer process currently stands at about 75%. Figure 83 shows a 300-mm single-level Python 1 BN photonics product on a 300-mm grade fused silica wafer. The transferred layer is 4-µm thick. It consists of a 160-nm thick patterned BN waveguide layer sandwiched between two, 2-µm layers of oxide cladding.

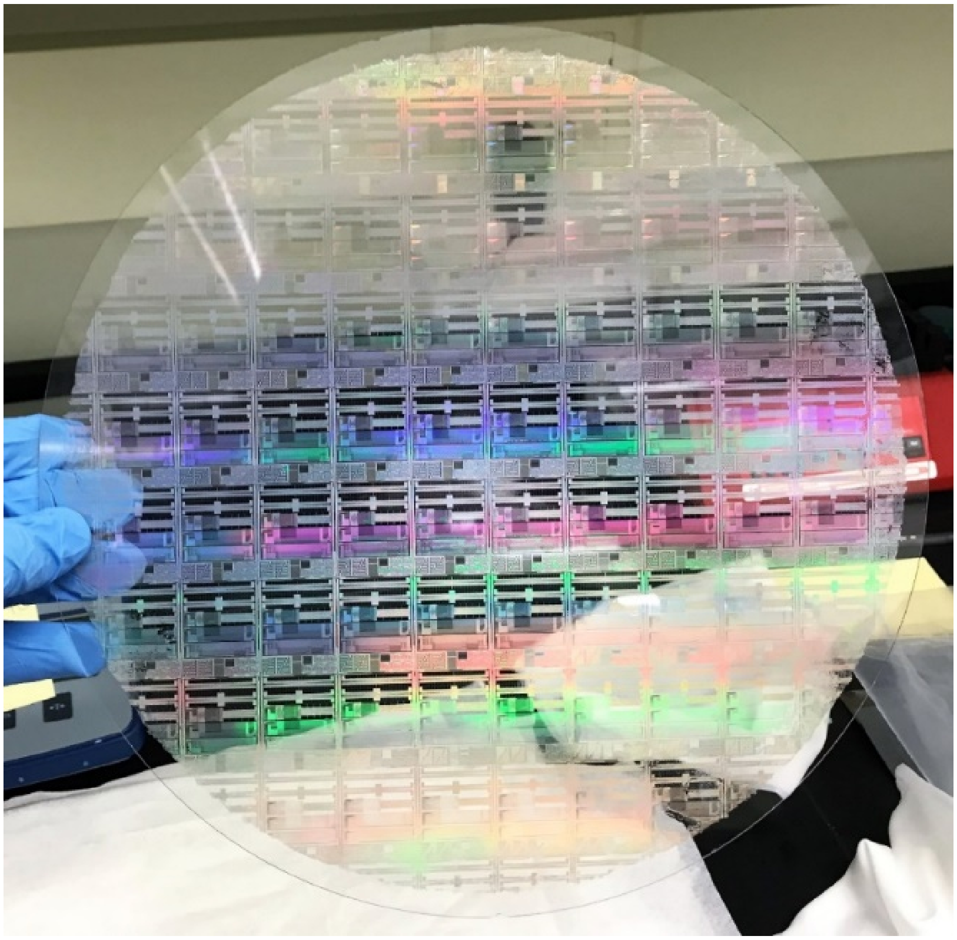

Figure 813: Python 1 BN Photonics Product on a 300mm Grade Fused Silica Wafer

LIST OF ABBREVIATIONS, ACRONYMS, AND SYMBOLS

ACRONYM	**DESCRIPTION**
5CB	4-cyano-4'-pentylbiphenyl
AR	Augmented Reality
CMOS	Complementary Metal-Oxide-Semiconductor
CMP	Chemical Mechanical Planarization
CNSE	College of Nanoscale Science and Engineering
EDX	Energy-Dispersive X-ray
FOV	Field-Of-View
He-Ne	Helium Neon
HMD	Head-Mounted Display
LC	Liquid Crystal
LiDAR	Light Detection and Ranging
MIT	Massachusetts Institute of Technology
MMI	Multimode Interference
MZI	Mach–Zehnder Interferometer
OPA	Optical Phased Array
PECVCD	Plasma-Enchanced Chemical Vapor Deposition
PIC	Photonic Integrated Circuit
RIE	Reactive Ion Etching
SCCM	Standard Cubic Centimeters Per Minute
SEM	Scanning Electron Microscopy
SiN	Silicon Nitride
SUNY	State University of New York
TEM	Transmission Electron Microscopy
UV	Ultraviolet
VIPER	visible Integrated Photonics Enhanced Reality