

FABRICATION OF INGAASP QW SOA FOR SILICON PHOTONICS

DESIGN AND FABRICATION OF INGAASP
QUANTUM-WELL SEMICONDUCTOR OPTICAL
AMPLIFIERS FOR INTEGRATION WITH SILICON
PHOTONICS

BY MATTHEW VUKOVIC, B. ENG.

A Thesis Submitted to the School of Graduate Studies in Partial Fulfilment of the
Requirements for the Degree Master of Applied Science

McMaster University © Copyright by Matthew Vukovic, April 2020

McMaster University – Hamilton, Ontario, Canada – April 2020

M.A.Sc. – Engineering Physics

TITLE: Design and Fabrication of InGaAsP Quantum-Well Semiconductor Optical Amplifiers for Integration with Silicon Photonics

AUTHOR: Matthew Vukovic, B. Eng. (McMaster University)

SUPERVISORS: Dr. A.P. Knights, Dr. R.R. LaPierre

NUMBER OF PAGES: xvi, 121

Abstract

Silicon photonics provides an environmentally sustainable pathway to a more robust data infrastructure. To compensate for optical power losses, methods of amplification are required; specifically, amplifiers that can fit in a small footprint for applications in data centres. Semiconductor optical amplifiers (SOA) provide such a solution, and can be fabricated using III-V ternary or quaternary materials to enhance optical signals through a device on the scale of most CMOS components.

This research sought to fabricate an InGaAsP multiple quantum well semiconductor optical amplifier using the facilities in McMaster University's Centre for Emerging Device Technologies (CEDT). A ridge waveguide laser diode was first fabricated and validated, then altered by applying an anti-reflective coating to the waveguide facets to suppress reflections in the Fabry-Perot cavity in an attempt to create an SOA. The design process and fabrication methodology are explained, including an analysis of failed methodologies. Characterization measurement techniques are then detailed for the fabricated devices. Finally, the performance of the devices is presented, and future steps are suggested for improving the fabrication process to enhance device characteristics. The fabricated laser diodes produced an output power in excess of 20 mW at a peak wavelength near 1580 nm. The subsequently coated devices proved difficult to measure, displaying a maximum of 0 dB or 1 dB gain when checked for amplification, with suspicions that output loss (and therefore gain) was higher than measured. The coated devices exhibited gain saturation between -10 and 0 dBm of input power. Owing to the shapes of their characteristic curves, it was determined that SOA devices were successfully created.

Acknowledgements

As much as I would love to take all the credit for myself, I wouldn't be able to live with the lie; besides, no one would believe me!

Dr. Andy Knights, my supervisor, without whom I wouldn't be writing this, working a corporate job and wishing I had at least given a Master's degree a shot. Your combination of endless support and relentless pursuit of progress was the impetus I needed for a project that could at times, literally, be grating.

Dr. Ray Lapierre, my co-supervisor, for the advice over the years of my degree and creative insight into projects that unfortunately never saw the light of day.

Dr. Ross Anthony, Dr. David Hagan and Dr. Ye Mengyuan, thank you for all the time you put in to help me with setups or to correct my blatant oversights.

Jeremy Miller, for being willing to push film deposition runs through on the weekends to help me meet my deadlines and finally accomplish the project goals, as well as all of the thin film and ellipsometry help. Here's hoping you soon move to your end step.

Connor Wong, for all the long nights we sat waiting for the PFIB to destroy my waveguides, and for keeping me sane in the hotel office space.

Dr. Nebile Isik, for all of the time spent trying to get the MBE to grow nanowires that proved simply impossible to make, your perseverance in the face of futility is beyond appreciated.

Dr. Zhilin Peng, for trusting me to run the TAB SEM-FIB and providing excellent conversation while waiting for 15-minute milling sessions to finish.

Dr. Shahram Tavakoli, for transforming from an intimidating technician to a trusted colleague and friend, thank you for the MBE runs, the cleaving tool and the sum of all the small things you helped me with over the years.

Doris Stevanovic, for the clean room training and bringing levity to some otherwise boring fabrication sessions, and all the small favours throughout the years.

Simon McNamee and Peter Jonasson for providing tools when I needed them, 3D prints, fume hoods (sorry Simon) and great conversations.

To all the folks at RANOVUS who I had the pleasure to meet over my short four months there, and especially Dr. Zhao Wang for continuing to assist me even in absentia through his old scripts and Matlab license.

To my parents, for their unending love and support, although after 25 years I don't know why they would still put up with me.

Finally, to all the other friends I made who became family along the way, and helped shape me into the person I am today: you know who you are.

Table of Contents

1	Introduction.....	1
1.1	Proposed solutions.....	3
1.2	Internship:.....	5
2	Background & Theory	6
2.1	Diode Physics.....	6
2.2	Light Emission	9
2.2.1	Spontaneous emission.....	9
2.2.2	Stimulated emission	11
2.3	Converting from Laser to Optical Amplifier.....	14
2.4	Distributed Bragg Reflector	18
2.5	Fabrication Mechanisms	18
2.5.1	Molecular Beam Epitaxy (MBE).....	19
2.5.2	Etching and Lithography.....	20
2.6	Characterization	21
2.7	Electro-absorption Modulation	22
3	Design & Fabrication.....	23
3.1	Material	23
3.2	Mask.....	26
3.3	Device Fabrication Overview.....	30
3.4	Clean Room Fabrication.....	31
3.5	Device Separation and Faceting.....	37
3.5.1	Dicing and FIB.....	37
3.5.2	Cleaving and Thinning.....	40

3.6	Bonding	42
3.7	Anti-reflection Coating	43
3.8	Gratings	45
4	Characterization.....	48
4.1	Electrical Characteristics and Optical Power	49
4.2	Optical Spectra	51
4.3	Band Edge and Amplification	52
5	Performance and Analysis	57
5.1	Simulation	57
5.2	Anti-reflective (AR) Coating Deposition.....	57
5.3	Optical Power and Electrical Characteristics	60
5.3.1	LI Characteristics After AR Coating	66
5.4	Optical Spectra	69
5.4.1	Spectra After AR Film Deposition	75
5.5	Band Edge Absorption	76
5.6	Amplification	77
6	Conclusion.....	90
6.1	Future work	91
6.2	Skills acquired	93
7	References.....	94
	Appendix A: III-V Laser Diode/SOA Fabrication Guide.....	97
	Introduction	97
	Fabrication Procedure	99
	Appendix B: Device Fabrication Matrix.....	118

Appendix C: Amplification Sweep Matlab Code 119

List of Figures

Figure 2.1: PN junction energy band diagram, including (a) equilibrium, (b) forward and (c) reverse bias conditions [9].	7
Figure 2.2: Full I-V curve for a diode, displaying forward bias, reverse bias and breakdown behaviour.	8
Figure 2.3: (a) photon absorption, (b) spontaneous emission, and (c) stimulated emission for a hypothetical 2-level laser.	10
Figure 2.4: Spectrum from one of the fabricated laser diodes, with the peak power around 1580 nm.	13
Figure 2.5: Comparison of 3D and 2D density of states in semiconductors [19].	14
Figure 2.6: Thorlabs C-Band SOA output spectrum [20].	15
Figure 2.7: SOA gain vs output power for the same Thorlabs device as Figure 2.6 [20]. Note the steep gain drop-off around 7 dBm, after constant gain. The gain vs. input power curve has the same shape, but with different values on the x-axis.	18
Figure 2.8: Gas-source Molecular Beam Epitaxy schematic [Dr. Vegar Ottessen].	20
Figure 2.9: Thorlabs S145C integrating sphere optical responsivity [23].	22
Figure 3.1: Four of the alignment marks from the original mask on device A1, from different areas of the pattern.	27
Figure 3.2: Device A2, SEM image of a waveguide facet after dicing through the chip and waveguide.	27
Figure 3.3: Mask 1 key features included (a) dicing lanes, (b) cavity lengths and labels, (c) alignment marks, and (d) passive waveguide ends.	28
Figure 3.4: Mask 2 key changes included (a) cleaving lanes, (b) alphanumeric labels, (c) full-length metal, and (d) DBR grating accommodations.	30
Figure 3.5: (a) Schematic cross-section of final device [25], (b) Facet of device B9 after cleaving.	31
Figure 3.6: Delaminated photoresist from device A4.	32

Figure 3.7: Device B6 post-exposure photoresist discrepancies (a) initially, (b) after UV ozone clean.....	32
Figure 3.8: Device A3 alignment marks on (a) SiO _x hard mask pattern, (b) III-V material after etch and oxide removal.....	34
Figure 3.9: Etched angled facet of a waveguide on device A3.....	35
Figure 3.10: Cross-section of device A4. Note the metal thickness at the top two corners of the waveguide.	36
Figure 3.11: Progression of FIB facet milling currents on device A3 (a) initial, (b) 3 nA, (c) 100 pA. Note the reduced width of curtaining lines from the waveguide edge between (b) and (c).	39
Figure 3.12: Facet after milling by PFIB. Note the degradation of the top edge and corners of the waveguide.	40
Figure 3.13: Final device after mounting on a copper backer with silver paint.	43
Figure 3.14: Reflectance spectrum generated by the Filmetrics calculator [26] for the chosen dielectric stack.	44
Figure 3.15: Schematic of AR deposition mask in grey, covering the sample in brown...45	
Figure 3.16: SEM image of AR film coating over the ends of the waveguides.45	
Figure 3.17: Series of tests as performed on one of the B6 waveguides. Note the high variance between adjacent cuts, which had small or no changes made to beam parameters.46	
Figure 3.18: Images of some promising results achieved by the PFIB for DBR gratings in device B6.47	
Figure 4.1: LIV setup, including integrating sphere, probes, heat sink, lab jack and lateral translation stage.49	
Figure 4.2: Heat sink configuration used for fiber coupling to both ends of the device....50	
Figure 4.3: Configuration for optical spectrum analysis, the same as Figure 4.1 but with a coupled fiber instead of the integrating sphere.52	
Figure 4.4: (a) Photograph of the amplification setup. (b) Schematic of the setup including system losses.53	

Figure 5.1: (a) to (d) depict the TE₀ modes for 2 μm to 5 μm waveguide widths, while (e) to (h) show TE₁ modes. Note that only the 4 μm and 5 μm guides have real TE₁ modes. 58

Figure 5.2: AR film under Ga-FIB SEM, as deposited on device A3. (a) overhead view of film, with masked contacts visible, (b) overhead view of left facet of A3-21 (affected waveguide), (c) closer, angled image of A3-21’s left facet, (d) angled image of A3-21’s right facet.59

Figure 5.3: AR film under PFIB-SEM, as deposited on device B6. (a) evidence of film wrinkling next to one waveguide, (b) PFIB cross-section of a coated waveguide, noting the lighter colour is the metal contact, and the thicker layer on top is the AR film, (c) PFIB cross section of the chip facet, noting that the film runs over the facet to coat it, (d) coated facet of B6-17, one of the amplifying structures.60

Figure 5.4: Optical power vs. current (LI) curves for A2-21 at different temperatures. Linear scale is used to show the discrepancy in threshold current more clearly.61

Figure 5.5: Slope efficiency of A2-21 over different temperatures.61

Figure 5.6: Full forward and reverse bias IV curve for B6-M-33, a 5 μm wide device. ...62

Figure 5.7: IV comparison of 5 μm structures on A3 and B6. Threshold voltage remains the same but slope differs.63

Figure 5.8: Comparison of LI characteristics for A3 cavities of varying lengths and widths.64

Figure 5.9: Comparison of LI characteristics for B6 devices of varying widths with (a) a passive region and, (b) no passive regions. The dashed lines mark the threshold current values.65

Figure 5.10: Comparison of LI curves for coated and uncoated A3 devices.66

Figure 5.11: B6-L-12 LI curves for the coated and uncoated device. The waveguide width was 2 μm.67

Figure 5.12: B6-L-26 LI curves for the coated and uncoated device. The waveguide width was 4 μm.68

Figure 5.13: B6-L-17 LI curves for the coated and uncoated device. The waveguide width was 5 μm. No change is seen in the output power levels.68

Figure 5.14: B6-L-22 LI curves for the coated and uncoated device. The waveguide width was 4 μm . The output power level increases after coating.	69
Figure 5.15: Optical spectrum of laser A2-21 held at 20 $^{\circ}\text{C}$ with a 300 mA pump current. Note the ripples in the tails of the spectrum.	70
Figure 5.16: Comparison of optical spectra of A2-21 at different temperature (25 $^{\circ}\text{C}$) and pump current (400 mA) conditions.	70
Figure 5.17: Spectrum of A3-L1(2)-1, a 5 μm device at 20 $^{\circ}\text{C}$ and 200 mA. Peak width is approximately 10 nm.	71
Figure 5.18: Spectrum of A3-L1(2)-1 just before reaching the lasing threshold, at 120 mA pump current. Note the wide range of wavelengths over which power is distributed.	72
Figure 5.19: Spectrum of A3-L1(2)-6, a 4 μm device at 200 mA. Note the peak width of around 5 nm.	72
Figure 5.20: Evolution of the A3-L1(2)-1 spectrum with increasing pump current of (a) 300 mA, (b) 320 mA, (c) 420 mA, and (d) 500 mA. Note the changes in shape, especially from Figure 5.17, and the red shift of the peak wavelength.	73
Figure 5.21: Comparison of B6-L-18 devices at 200 mA pump current, with widths of (a) 5 μm , (b) 4 μm , (c) 3 μm , and (d) 2 μm	74
Figure 5.22: Spectra of B6-L-18 pumped at 50 mA (a) after AR coating, (b) with a 1563 nm light source at the input.	75
Figure 5.23: Spectrum of B6-L-12 at 200 mA before and after AR coating. Note the elimination of lasing behaviour in the coated spectrum.	76
Figure 5.24: Absorption measurements in reverse bias of B6-M-33 as part of the EAM project. Note the onset of absorption around 1580 nm for 0 V and 1 V bias.	77
Figure 5.25: Relation of laser power to laser wavelength from the Agilent 8164A output port. Series labels refer to the input power as set on the Agilent.	78
Figure 5.26: Output power and input power vs. wavelength for B6-L-17 at 50 mA. The legend refers to the power setpoint, not the real input power. Note the apparent peak near 1595 nm.	79

Figure 5.27: Output power vs. wavelength for B6-L-04 at 180 mA. Note the apparent peak near 1580 nm.....	80
Figure 5.28: Output power vs. wavelength for B6-L-08 at 180 mA. Note the apparent peak near 1585 nm.....	81
Figure 5.29: Gain vs. input power of B6-L-08 with 1580 nm input and pumped at 200 mA. Each curve represents a different digital filter width around the 1580 nm peak.	82
Figure 5.30: Gain vs. output power of B6-L-08 with 1580 nm input and pumped at 200 mA.	83
Figure 5.31: OSA spectrum of B6-L-08 with various levels of input power at 1580 nm. Note the oscillations in the signal from back-reflections, and magnitude of the smallest amplified peak.....	83
Figure 5.32: Optical filter inserted into the line with a tuned bandwidth of 4 nm, 1558 to 1562 nm.	84
Figure 5.33: Gain vs. input power of B6-L-04 with 1560 nm input and pumped at various current levels.	85
Figure 5.34: LI curves for the uncoated and coated B6-L-04 device. Note the maximum of the coated LI curve at 160 mA.....	85
Figure 5.35: Gain vs. output power of B6-L-04 with 1560 nm input at various current levels.	86
Figure 5.36: Output power vs. input power of B6-L-04 with 1560 nm input at various current levels. Note how the data deviates from the linear trend, then self-corrects.	87
Figure 5.37: Gain vs. input power of B6-L-08 with 1580 nm input at 200 mA. Filter set to 1 nm width, and loss increased by 5 dB over Figure 5.29.	88
Figure 5.38: Facet coupling loss vs. pump current for B6-L-04. Values were derived by subtracting measurements of the ASE through the fiber from measurements directly into the integrating sphere.	88

List of Tables

Table 3.1: Epitaxial growth structure of CEDT material, Growth A.....	24
Table 3.2: Epitaxial growth structure of Landmark Optoelectronics material, Growth B.....	25

List of Abbreviations

AC	alternating current
APC	angled physical contact
AR	anti-reflective
ASE	amplified spontaneous emission
CCEM	Canadian Centre for Electron Microscopy
CEDT	Centre for Emerging Device Technologies
CMOS	complementary metal-oxide-semiconductor
DBR	distributed Bragg reflector
DI	de-ionized (water)
EAM	electro-absorption modulation/modulator
FIB	focused ion beam
HF	hydrofluoric acid
III-V	composite materials composed of Group 3 and Group 5 elements
InGaAs	indium gallium arsenide
InGaAsP	indium gallium arsenic phosphide
InP	indium phosphide
IV	current-voltage characteristic
LI	optical power-current characteristic
LIV	electrical and optical power characteristics
MBE	molecular beam epitaxy
MOCVD	metal-organic chemical vapour deposition
MQW	multiple quantum well
OPM	optical power meter
OSA	optical spectrum analyzer
PFIB	plasma focused ion beam
PL	photolithography
QCSE	quantum-confined Stark effect

QW	quantum well
SiC	silicon carbide
SiO _x	silicon oxide
SiP	silicon photonics
SOA	semiconductor optical amplifier
SOI	silicon-on-insulator
TEC	thermoelectric cooler

1 Introduction

With the expansion of Internet connectivity both geographically and across everyday devices, a crisis is emerging with respect to bandwidth availability. The growth of communication infrastructure in developing countries, the “Internet of Things” trend amongst appliances, and the growing demand for media streaming are all contributing factors to an expected average bandwidth usage increase of 27.4% per year [1]. This growth rate puts enormous pressure on the existing data centre infrastructure, and will likely prove unsustainable (both functionally and environmentally) if no disruptive measures are taken. It is this issue that has led to increased research into silicon photonics (SiP).

The main principle behind SiP is the use of photons to carry information as opposed to electrons, on a low-cost, monolithic platform. There are numerous ways in which this provides benefits to data processing systems. Perhaps most obviously, photons are capable of achieving higher speeds than electrons, which can lead to higher data throughput rates. More importantly, SiP provides an advantage in the overall bandwidth available to the data transfer system. While electronic systems could, hypothetically, achieve speeds above the 10 Gbit/s threshold, a number of physical and practical barriers stand in the way. As the speed of oscillation of electrical signals increases, induced current effects from Faraday’s Law begin to become significant which lead to deterioration of signal quality as noise is introduced due to electromagnetic crosstalk. Additionally, higher electrical switching speeds increase the resistance of conductors from the skin effect derived from the Lorentz force [8], leading to heating which further increases resistance. Furthermore, the practical limitations of electronic circuit board fabrication are being reached, with each new generation of devices demanding a smaller footprint for higher computational power. Combatting all of this, SiP has minimal crosstalk between adequately spaced waveguides, no electromagnetic induction effects due to carrier travel, and minimal effect on the physical properties of the carrier structures during operation. This allows speeds to be increased until limited by electrical modulation of the optical sources, and provides greater signal stability over both short and long distances. Optical signals can also be encoded onto

different polarizations and wavelength bands by multiplexing, allowing for a higher signal density and simultaneous bidirectional communication across the same line. The network speed advantages of photonic communication have already led some internet service providers (ISP's) to replace their existing electrical transmission line infrastructure with fiber optic cables [2].

Another way in which SiP prevails over conventional electronics is in its environmental impact. Whereas the movement of electrons through solid material produces significant amounts of heat, and therefore requires further cooling to prevent the system from overheating, photons travelling through passive solid media produce minimal heat. Instead of cooling the entire chip or system as in conventional electronics, the primary amount of heat removal in SiP systems takes place at the laser or amplifier components which produce or boost optical signals. While there is presently no study that directly compares the thermal performance of both systems, it is this researcher's opinion that the thermal footprint of a SiP system would be significantly smaller than an electronic system, helping to mitigate the current global environmental threat posed by the proliferation of data centres [3,4].

However, SiP is not without its flaws. One of the largest issues, and the origin of the problem that this research seeks to address, is the large optical power losses incurred during signal modulation [5]. These losses are great enough that they limit widespread adoption of SiP. The most direct solution to this is to amplify the optical signal at the output of the modulator.

The erbium-doped fiber amplifier (EDFA) is one existing, and prolific, technique used to amplify optical signals. Briefly, these devices feature silica optical fibers that have been doped with erbium ions which, in combination with a separate laser source, amplify the signal travelling in the fiber. These devices provide a large gain in excess of 40 dB which is independent of signal polarization [6]. EDFA's also have low interfacial losses, as amplification occurs directly in the fiber. These amplifiers also display immunity to signal

crosstalk. However, due to the principles on which rare-earth amplifiers operate, they are incapable of switching at high frequencies as the carrier lifetime of the excited energy state is on the order of milliseconds [7]. Additionally, EDFA's are not suitable for amplifying signals with wavelengths outside of a narrow range (typically ~1530-1560 nm), due to the emission spectra of the excited erbium ions [7]. Limitations of fiber bend radius also mean that EDFA's must have a large physical footprint.

Another amplifying structure is the semiconductor optical amplifier (SOA), which uses the same principles as a diode laser to amplify optical signals within an electrically pumped semiconductor gain medium. III-V semiconductor materials are typically used to construct SOA's due to their direct band gap; silicon's indirect band gap makes the required energy absorption and emission characteristics very difficult to attain. Due to the short lifetimes of carriers in direct band-gap III-V materials, there is a much higher response time to input stimuli, allowing for a faster rate of switching. Also, as all amplification happens directly in the SOA device, the design can be made much more compact than conventional EDFA's. With respect to fabrication, SOA's can be made using wafer processing techniques, allowing many devices to be manufactured simultaneously for a relatively low cost compared to EDFA's; however, issues arise when attempting direct growth of a III-V device on a silicon-on-insulator (SOI) substrate, as the difference in lattice constants of the two materials is large enough to cause significant defect propagation. SOA's also suffer from high sensitivity to polarization, higher noise than EDFA's, and are susceptible to crosstalk between modulated signals within the waveguide [7]. Furthermore, optical losses occur at the facets of the SOA, where the amplifying waveguide couples to an input fiber or waveguide.

1.1 Proposed solutions

The overarching goal of the research project was the fabrication of a semiconductor device which could be used to amplify optical signals on-chip. The signals would lie approximately in the C-band for optical communication (1530-1565 nm). Novel techniques

were explored, and ultimately the project sought to prove the capability of the McMaster facilities to produce a working SOA device.

Several potential solutions were initially proposed and pursued for this research project. One approach was the novel solution of growing III-V indium gallium arsenide (InGaAs) nanowires directly on (100) SOI waveguides, whereby the signal mode travelling in the SOI waveguide would evanescently couple into the nanowires. The nanowires would be designed to amplify at the signal wavelength, creating an integrated on-chip amplification method. Nanowires were chosen specifically because the interfacial area of the III-V nanowire and SOI substrate would be small enough so that defect propagation due to lattice mismatch would be avoided. An electron beam lithography mask was designed in order to pattern a test piece of SOI in preparation for the nanowire growth, and several test growths were performed. However, due to the limitations of the gas-source molecular beam epitaxy (MBE) tool in the Centre for Emerging Device Technologies (CEDT) at McMaster, it was impossible to achieve the flux ratios between the arsenic, indium and gallium required to produce nanowires at the working wavelength. The project was to be attempted using metal-organic chemical vapour deposition (MOCVD), but due to administrative issues with the tool this approach was also forced to be discontinued.

The eventual approach turned out to be the most straightforward. An InGaAsP quantum well (QW) diode laser would be fabricated using the CEDT facilities, and characterized. Then, an anti-reflective (AR) coating would be deposited on the facets of the laser devices to reduce the reflectivity of the facets, creating an SOA using the laser cavity. As such, this research sought to fabricate an SOA device which can be fiber-coupled or free-space coupled to the facet of an existing SiP system. In the following chapters, the theoretical principles behind the operation of the SOA devices are outlined, followed by a description of the design and fabrication process used to manufacture the SOA devices. The characterization techniques are then discussed, followed by an analysis of the results collected from measuring the final devices, including both a laser and SOA device.

Overall, the contributions of this thesis are as follows:

- Revival and elaboration of a laser diode fabrication process to manufacture III-V devices in the CEDT at McMaster, including a detailed process manual,
- Understanding the viability of focused ion beam milling as an option for faceting waveguides and laser devices,
- Fabrication of InGaAsP multiple quantum well laser diodes with central wavelength of 1585 nm and output power in excess of 20 mW,
- Utility of a top-down deposition of silicon oxide-nitride anti-reflection coating,
- Fabrication of SOA devices by depositing anti-reflection coating on the previously fabricated laser diodes.

1.2 Internship:

The period of September 1, 2018 to December 21, 2018 was spent at RANOVUS Inc. in Ottawa, Ontario. The focus of this internship was on the characterization of similar (albeit higher calibre and quality) devices to the SOAs of this Master's work. A great deal of experience was gained with measurement and characterization techniques, as well as realistic SOA device operation parameters. The internship also emphasized the importance of SOA devices in the maturing field of silicon photonic systems. Due to the constraints of a non-disclosure agreement signed with the company, no further details on the work performed at RANOVUS Inc. may be detailed in this document.

2 Background & Theory

In this section, the principles behind the operation of a laser diode, as well as the subsequent changes made to it to produce an amplifier, are outlined. First, optical and electrical properties of the ridge waveguide quantum well laser provide insight into the device physics. Then, the fabrication and characterization techniques used throughout the research project are described.

2.1 Diode Physics

The basis of the laser diode, and by extension the semiconductor optical amplifier (SOA), is the PN junction. The PN junction is formed by creating two adjacent, oppositely doped regions: one doped with acceptor impurities (P) and one with donor impurities (N). Acceptor impurities (such as beryllium) create an excess of holes in the doped material, whereas donor impurities (such as phosphorus) provide an excess of electrons. The imbalance in charge carriers on each side results in the identification of majority and minority carriers, depending on which are in excess: on the P side, holes are majority and electrons are the minority, with the opposite being true for the N side. The resulting junction structure permits, among other things, one-way current flow under nominal operating conditions, and is the basis of the semiconductor diode.

Figure 2.1 describes the energy band structure of the junction. The region between the P- and N-doped volumes is referred to as the depletion region, wherein there are no free charge carriers at equilibrium. Due to the difference in energy levels of the P and N regions, a built-in voltage exists which serves to move charge carriers across the depletion region via the creation of an electric field. This field pushes minority carriers from one side of the junction to the other, i.e. holes from the N side to the P side, a phenomenon that is known as drift current.

The difference in concentrations of holes and electrons on each side of the junction leads to an opposing current to drift current. This current arises as majority carriers diffuse across the depletion region towards a volume of lower carrier concentration, moving holes from the P side to the N side or electrons from the N side to the P side. The unbiased system reaches equilibrium when the drift and diffusion currents are equal in magnitude, and does so almost instantaneously.




Figure 2.1: PN junction energy band diagram, including (a) equilibrium, (b) forward and (c) reverse bias conditions [9].

The application of an external electric field to the device, often via a voltage source, changes the shape of the depletion region and upsets the behaviour of the charge carriers from the equilibrium condition. When a positive voltage is applied to the P side of the device relative to the N side, the device is said to be in forward bias, as in Figure 2.1b. The superposition of the applied field with the built-in electric field results in a lower difference in energies of the two sides (as the two fields oppose each other), and reduces the width of the depletion region. This allows diffusion current to dominate, as the built-in field is reduced but the difference in carrier concentrations remains the same. Since the diffusion current arises from the flow of the majority carriers, a PN junction with an adequate forward bias effectively operates as a short circuit, with the current rising exponentially with applied voltage.

When a negative voltage is applied to the P side relative to the N side, the device is said to be in reverse bias, as seen in Figure 2.1c. The superposition of the internal and external fields this time leads to an increased difference in energy between the two sides and a wider

depletion region, further inhibiting majority carrier flow and diffusion current. Drift current dominates, and since it is composed of minority carriers a small but largely insignificant current is produced called the leakage current; the leakage current is so small that the net current in the reverse-biased device is determined by thermally generated drift currents [10]. At large enough reverse bias voltages the device operates in breakdown, where the energy difference between the two sides of the junction becomes large enough that the valence band energy of the P side exceeds the conduction band energy of the N side. This leads to a massive reverse current through the diode and often device failure, as can be seen in Figure 2.2.

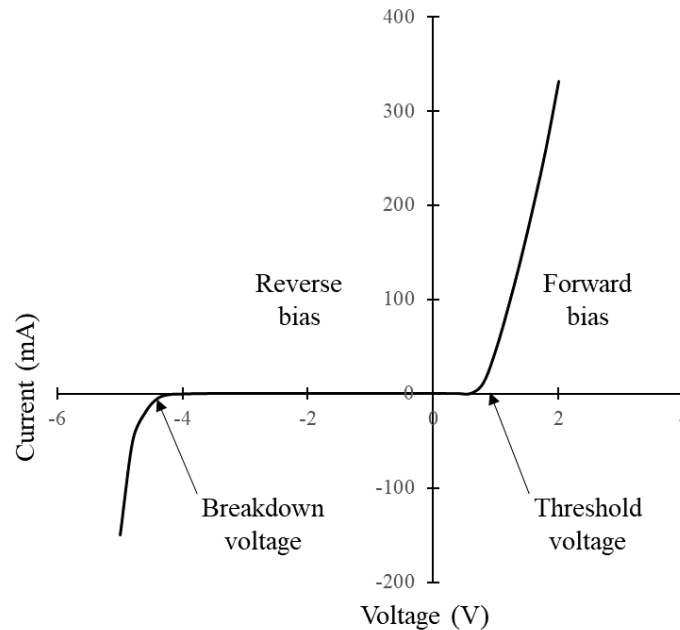


Figure 2.2: Full I-V curve for a diode, displaying forward bias, reverse bias and breakdown behaviour.

The IV curve of a diode like the one in Figure 2.2 is easy to recognize experimentally due to the exponential increase in current with applied voltage once a threshold voltage has been achieved. This becomes especially important when validating the operation of a newly fabricated laser diode: a simple “IV curve” test of the device will indicate whether a proper diode structure has been created. If the curve through forward bias is linear instead of

exponential, it is likely that some error in fabrication or some short circuit has caused the device to become ohmic.

In the case of the devices fabricated during this research project, a PIN diode is formed, where an intrinsic (undoped) region of semiconductor material lies between the P- and N-doped regions. The intrinsic region in this case is a series of undoped quantum well and barrier layers, the effect of which will be discussed later in this chapter. Inserting an intrinsic region between the P and N regions effectively establishes the depletion width of the diode as the width of the intrinsic layer [11]. Higher dopant levels are also introduced at the top and rear surfaces to facilitate ohmic contacts with probes at either end of the device.

2.2 Light Emission

In forward bias, an electric field is induced across the PIN diode and holes and electrons are excited into elevated energy states and injected across the depleted region from the P- and N-doped regions, respectively. When these carriers interact they recombine and release energy. This recombination occurs in the depletion region of device, in our case, the intrinsic region [12]. There are two major classes of recombination which occur: radiative, which produces a photon carrying the recombination energy, and non-radiative, which produces a phonon resulting in thermal energy. As we are interested in photon emission, the focus of this section will be on radiative recombination mechanisms. Radiative recombination is most often found in materials with a direct band gap, as the most favourable energy decay of an excited carrier does not need to emit a phonon to adjust its momentum.

2.2.1 Spontaneous emission

Spontaneous emission is the primary form of radiative recombination. As the name suggests, an excited carrier will drop from the state in the conduction band to a lower energy state in the valence band without external stimulus, emitting a photon as seen in Figure

2.3b. The energy of the emitted photon is equal to the energy difference between the initial and final energy of the excited carrier, as

$$E_{\text{photon}} = E_i - E_f = h\nu_{\text{photon}}$$

where ν_{photon} is the photon frequency and h is the Planck constant. As the carrier will most often be dropping from an excited state to the ground state, the energy of the photon will normally equal or exceed the material's band gap energy, E_g . A photon generated by spontaneous emission has no correlation to other emitted photons, and is produced with a random phase, polarization and propagation direction.

As the excited states necessary to create spontaneous emission do not exist in any significant quantity at thermal equilibrium, energy must be injected into the system. In rare-earth devices such as the EDFA, this is done optically using an external laser of a specific wavelength to excite carriers. For semiconductor devices such as the ones in this thesis, an external electric field is applied to inject carriers into the depletion region. The diode structure enables current flow through the device, leading to an electron-hole pair generation rate which is related to the applied current. The rate at which these excited states naturally decay is the carrier lifetime, which is typically in the nanosecond regime for direct band gap semiconductors [13].




Figure 2.3: (a) photon absorption, (b) spontaneous emission, and (c) stimulated emission for a hypothetical 2-level laser.

2.2.2 Stimulated emission

When a photon passes through a depletion region containing injected electron-hole pairs, there are two interactions that may take place. If the photon has an energy greater than the material band gap, then it may be absorbed by an unexcited carrier in the valence band, boosting that carrier to the conduction band as seen in Figure 2.3a. The photon may also cause stimulated emission, whereby the original photon encounters an electron-hole pair and induces the pair to undergo radiative recombination in such a way as to produce a photon with the same energy, phase, polarization and propagation direction as the original photon, effectively duplicating it, as seen in Figure 2.3c.

Stimulated emission is a much less likely phenomenon than spontaneous emission. For stimulated emission to dominate, a significant amount of external pumping of the system must take place [14], enough such that most of the states at the bottom of the conduction band are filled with electrons, and most of the states at the top of the valence band are filled with holes. This is known as population inversion. Once population inversion has been achieved, some photons will be spontaneously emitted as electron-hole pairs recombine, which go on to stimulate emissions from other pairs. This is called amplified spontaneous emission (ASE), and is the immediate precursor to achieving lasing in the device.

To initiate lasing, it is necessary to achieve an increased flux of photons in the system compared to the ASE condition. This is possible if an external laser source is coupled to a device and injects photons into the active region; however, for a self-contained system the most common approach is to use mirrors at the end facets of the device. In this case the end facets of the active region are fabricated as parallel, planar mirrors, forming a Fabry-Perot cavity. Photons that are produced in the cavity with propagation vectors normal to the planes of the facets will be reflected back and forth, stimulating further emissions of duplicate photons with the same direction of propagation. Photons without the appropriate propagation vector fail to resonate and eventually escape the cavity through the sidewall, or by being reabsorbed.

Mathematically, the steady-state cavity can be modelled by the transition rate equation as [15],

$$B_{12}N_1\rho(\nu_0) = A_{21}N_2 + B_{21}N_2\rho(\nu_0)$$

Here, N_1 and N_2 represent the electron concentrations at levels E_1 and E_2 , respectively, as defined in Figure 2.3. B_{12} , B_{21} and A_{21} are known as the Einstein coefficients of the system. $\rho(\nu_0)$ is the optical field strength, the total energy in the field per unit volume, for the stimulating photon frequency ν_0 . The B_{12} term represents the absorption rate of photons in the active region. The A_{21} term represents the spontaneous emission rate and the B_{21} term is the rate of stimulated emission. As lasing will be achieved when the stimulated emission rate is the dominant term, ratios can be taken between it and the absorption and spontaneous emission rates. Against the absorption rate it is found that,

$$\frac{B_{21}N_2\rho(\nu_0)}{B_{12}N_1\rho(\nu_0)} = \frac{B_{21}N_2}{B_{12}N_1}$$

So, when N_2 is much larger than N_1 , such as in the case of population inversion, the stimulated emission term will dominate over absorption. Against the spontaneous emission rate it is found that,

$$\frac{B_{21}N_2\rho(\nu_0)}{A_{21}N_2} = \frac{B_{21}}{A_{21}}\rho(\nu_0)$$

So, when the optical field strength is large, such as when a large flux of photons is present in the active region, stimulated emission will dominate over spontaneous emission. Laser is an acronym for Light Amplification by Stimulated Emission of Radiation; once the critical flux of photons is achieved through resonant amplification the device can be said to be lasing.

Lasing is observed on an optical spectrum as a sharp, intense peak around the band gap wavelength of the device, as seen in Figure 2.4. The laser device is identifiable from a normal light source in that it produces coherent, collimated and monochromatic light [16].




Figure 2.4: Spectrum from one of the fabricated laser diodes, with the peak power around 1580 nm.

The active region of the device used in this research is composed of multiple quantum wells (MQW) as opposed to a bulk material or a double heterostructure. The quantum well layer is constructed as a thin (< 10 nm) layer sandwiched between barrier layers which act to confine carriers to the well layer energetically by having a significantly larger band gap energy. This leads to two-dimensional confinement of carriers and has some notable effects on device operation. The narrowness of the quantum well limits the number of states available to excited carriers, discretizing the energy levels available and transforming the density of states in each band from a square root curve (as in the bulk) to a step function, as seen in Figure 2.5. This, in turn, reduces the number of carriers necessary to achieve population inversion in the quantum well layer [17]. It is also noted from Figure 2.5 that the lowest available energy state for carriers in a QW structure is just above the band gap energy of the material. As the number of carriers in the region is reduced the free absorption

coefficient, which affects optical absorption in the material, is also reduced, meaning that a lower level of stimulated emission is required to overcome absorption losses, and optical gain in the system increases [17]. When multiple quantum wells are subjected to sufficient pump current levels, the gain produced becomes N times the gain of a single quantum well, where N is the number of wells in the active region [18].




Figure 2.5: Comparison of 3D and 2D density of states in semiconductors [19].

2.3 Converting from Laser to Optical Amplifier

The conversion from a semiconductor laser diode to a semiconductor optical amplifier is straightforward. The primary difference between the two devices is that one forms an internal reflective cavity and the other is a single-pass device. Thus, by applying an anti-reflection coating to the facets of the laser, reflectivity may be reduced enough such that the only power emitted is ASE, as the photon flux lasing condition cannot be met. An example of an expected output spectrum of an SOA can be seen in Figure 2.6. For the case of the SOA, a longer active region is desired, as it provides more opportunities for incoming photons to interact with electron-hole pairs and stimulate emission.




Figure 2.6: Thorlabs C-Band SOA output spectrum [20].

The anti-reflection coating is designed to remove reflectivity for the wavelength produced by the active region. Cancellation of the reflected waveform requires deposition of a film of a precise refractive index n_f on each facet, where n_f is the geometric mean of the waveguide and air indices such that,

$$n_f = \sqrt{n_{air} n_{waveguide}}$$

However, in cases where no film exists with the appropriate refractive index, a multi-layer coating can be implemented, such as a silicon oxide-nitride stack. By varying the thicknesses of each layer in the stack, a maximum transmittance can be achieved at the desired wavelength.

As the SOA requires an external optical source to be of use, more attention must be paid to how the light is coupled into and out of the device. For the devices in this thesis, a ridge waveguide structure is implemented. This ridge is composed of InP, and lies on top of the InGaAsP MQW active region. The input signal is coupled into the waveguide, which is clad in material of lower refractive index. Due to the overlap of the signal mode and the mode supported by the active region, the signal is evanescently coupled into the MQW region of the device, where it becomes amplified. This amplified signal then couples back

into the waveguide and exits the device via the output facet. Any light reflected back into the system is detrimental as it introduces a large amount of noise; as such, extra steps are often taken to reduce the possibility of back-reflections such as integration of optical isolators and angled physical contact (APC) optical fibers.

It is worth noting that SOA devices such as the ones in this thesis are highly sensitive to input signal polarization. If the input signal polarization does not match the polarization of the gain region ASE, then amplification is severely hindered as the transverse fields will not optimally superimpose. Polarization restrictions are primarily due to the geometry of the device. The geometry of the waveguide of the SOA preferentially supports the TE optical mode [21].

Ultimately, the indication that an SOA has been created is the presence of internal optical gain. As the name implies, the power at the output facet should be greater than the coupled signal power at the input facet. This requires detailed loss measurements throughout the characterization system, as losses from measurement tools and polarization controllers are accrued in addition to interfacial losses where light from the fibers is coupled to/from the waveguide facets. Assuming that lasing behaviour has been suppressed by an anti-reflection (AR) coating, and that the spectrum of the device is that associated with ASE, the maximum gain should occur at the point of maximum power on a power-current (LI) curve measured from the device without optical input. At this point, the injected current causes the highest level of radiative recombination activity from carriers in the cavity before, at higher currents, thermal (non-radiative) recombination effects begin to dominate, resulting in lower optical power. Thus, the gain trend of the device should see an increase with pump current until the maximum ASE output power is achieved, then a reduction with increasing current thereafter.

With the optimal pump current determined, the gain of the device can be characterized using sweeps across input power and wavelength. The peak amplifying wavelength is

related to the peak wavelength of the device's ASE curve. In the case of Figure 2.6, the highest gain should occur at 1550 nm. The measured internal gain at a specific wavelength and pump current can then be plotted against the input power or the output power of the device. The gain of commercial devices is typically characterized using the output power, as in Figure 2.7. The gain vs. input power curve retains the same shape as this curve, but with several dB of difference on the x-axis scale. There are two notable regimes to the curves: the constant gain regime (in the case of Figure 2.7, the gain left of 7 dBm output power), and the saturation regime. The constant gain regime, sometimes referred to as the small-signal gain, is where the SOA will take any given input power and produce an output which is amplified by a constant gain. Here the rate of excited state generation by carrier injection is greater than or equal to the rate of stimulation or quenching. For Figure 2.7, the small signal gain is 15.75 dB. However, it is impossible for there to be an infinite number of inverted carriers in the conduction band of the SOA. As the input signal power is increased, the flux of photons becomes so great that stimulated emission occurs at a rate which outpaces the rate of excited state generation. If the carrier injection rate is increased to compensate for this, eventually a point is reached where the heat generated by carrier injection causes a significant rate of thermal recombination which also contests the excited state generation rate. As such, as the input power is driven higher, the output power becomes limited and the gain cannot stay constant. The output power limit in Figure 2.7 is between 14 and 15 dBm. As the input power is increased but the device cannot produce higher output power, the gain is seen to rapidly decrease.




Figure 2.7: SOA gain vs output power for the same Thorlabs device as Figure 2.6 [20]. Note the steep gain drop-off around 7 dBm, after constant gain. The gain vs. input power curve has the same shape, but with different values on the x-axis.

2.4 Distributed Bragg Reflector

A brief part of this project was dedicated to attempting to produce distributed Bragg reflection (DBR) gratings on the ends of the fabricated device waveguides. To do so, a grating of sufficient length is etched into the end of the device waveguide. The pitch and index of the grating are designed such that a narrow wavelength range is reflected back into the cavity constructively, while all other wavelengths become attenuated. The result is a laser which produces an intense peak at a single wavelength [22].

2.5 Fabrication Mechanisms

With the core theory behind laser diode operation covered, the next step becomes translating the device from theory to reality. In the case of the devices covered herein, quaternary devices were constructed from InGaAsP via MBE, and subsequently processed in CEDT clean room facilities to create the waveguides. This section seeks to clarify the operational and physical principles behind certain steps of the fabrication procedure, which is covered fully in Chapter 3: Design & Fabrication.

2.5.1 Molecular Beam Epitaxy (MBE)

The particular system most commonly used for semiconductor growth by the CEDT is gas-source MBE. As suggested by its name, the MBE grows epitaxial layers of material on a substrate, atomic plane by atomic plane. The growth rate is approximately 1 $\mu\text{m}/\text{hour}$. In this particular tool, gaseous forms of indium, gallium, arsenic, phosphorus and aluminum are created in the effusion cells and released into the reaction chamber, as illustrated in Figure 2.8. These gases are controlled via flow rate and temperature to determine their reactivity with the substrate. Beryllium and silicon can also be introduced as P- and N-type dopants during the construction of each layer. The material being grown will match the crystallographic orientation of the substrate, limiting the possible structures that can be created unless a different substrate orientation is selected. These limitations are further exacerbated by the need to “lattice match” the individual layers to minimize lattice strain and defect density. When growing layers of a quaternary material such as InGaAsP, the proportions of each element must be closely regulated to ensure minimal lattice strain while also maintaining the correct band gap energy for the layer. One major disadvantage of the gas-source MBE is its inability to create certain molar ratios of III-V materials, limiting the compositions of materials that can be grown. This limit on the III-V flux ratio arises because of hydrogen gas production resulting from cracking arsine (AsH_3) molecules.




Figure 2.8: Gas-source Molecular Beam Epitaxy schematic [Dr. Vegar Ottessen].

2.5.2 Etching and Lithography

There are a few points to highlight with respect to the wet bench work required to process the III-V material. Of particular interest are the mechanisms involved with wet etching InP and the resolution of the photolithography system used.

The etchant used to remove InP material was a mixture of hydrochloric acid and phosphoric acid. The etch is anisotropic, and etches one InP plane slower than the rest. As a result, the waveguides must be properly aligned to achieve vertical sidewalls instead of sloped sidewalls (leading to better optical confinement). The etch is also designed not to etch InGaAsP, which is added to the structure as an etch stop layer below the InP waveguide.

The lithography system in the CEDT is a contact lithography system. This has the user bring their prepared sample into contact with a chromium mask, through which a pattern is laid into photoresist via exposure to ultraviolet light. The resolution limit of the system

used was approximately 1 μm due to diffraction of the ultraviolet light from the exposure source as it propagated through the photoresist.

2.6 Characterization

Once the devices were fabricated, measurements were taken using various configurations of tapered fibers, micro-positioning stages, large area photodetectors and electrical probes. Further details on the characterization techniques used may be found in Chapter 4: Characterization.

There are multiple optical loss mechanisms which present themselves when measuring the light at the facet of the device. The facet itself may not have a smooth enough surface, resulting in scattering of light from the output. Dirt may simply have landed on the facet during the measurement process. The fiber coupled to the input or output of the device may also present issues. The spot size of the fiber, if significantly mismatched from the spot size of the waveguide, will result in coupling loss. Stress applied to the fiber will cause physical defects which perturb the light enough to cause noticeable loss. Furthermore, due to the polarization sensitivity of the SOA, any movement of the input fibers during measurement will require readjustment of the polarization controller to ensure maximum output power. It is also prudent not to rule out human error when aligning the fiber to the facet, as some apparent losses can be rectified by careful repositioning of the fiber with respect to the device.

Lastly, the measurement equipment has a response curve that must be accounted for before making assumptions about the measured power values. A Thorlabs Integrating Sphere was used to take optical power measurements for the laser and SOA devices in this thesis. The responsivity curve of the sphere can be seen in Figure 2.9. It is noted that there is minimal variation in the responsivity for the 1500 to 1600nm range over which this thesis was conducted.




Figure 2.9: Thorlabs S145C integrating sphere optical responsivity [23].

2.7 Electro-absorption Modulation

During the project, an opportunity with an industry partner arose to fabricate electro-absorption modulators (EAM) through the work of Dr. Ross Anthony, a post-doctoral fellow in the Knights group. In order to verify the fabrication principles employed, one of the laser devices in this research was tested as if it were an EAM device.

A reverse-bias voltage applied across the MQW section of the device will cause the creation of confined excitons within the wells. The applied field reduces the binding energy of the electron-hole pair in the exciton, but the energy barriers of the well keep the pair confined, leading to the presentation of exciton absorption peaks at the energy levels associated with each quantum well state. This is called the Quantum Confined Stark Effect (QCSE), and creates a change in the optical absorption of the device over a narrow wavelength range [24]. By sweeping the input wavelength of light coupled into the reverse-biased EAM, the band edges for absorption can be determined, with redshifting of the band edge with increasing reverse bias. As such, passing light through an EAM allows for voltage-controlled attenuation of the beam.

3 Design & Fabrication

The initial design and fabrication process was based on instructions from the lab manual for the CPFR Photonics Fabrication graduate short course [25]. While the manual laid out the basic theory and fabrication steps for creating an InGaAsP laser diode, there was no one remaining at McMaster who had fully completed the entire process. As such, the majority of this thesis project was dedicated to recreating, improving and modernizing the procedures in the manual, and eventually adding to them to transform the laser diode into an SOA. The fully detailed process distilled from this thesis project is presented in Appendix A: III-V Laser Diode/SOA Fabrication Guide. A breakdown of the devices fabricated and the various process iterations they went through is presented in Appendix B: Device Fabrication Matrix.

3.1 Material

The devices were constructed by processing InP wafers with epitaxially-grown layers of InGaAsP, InP and InGaAs. The first attempts at fabrication were performed with a 2-inch wafer grown by Dr. Shahram Tavakoli using the MBE in the CEDT at McMaster. The structure of these devices was based on the layout in [25], but due to various process variables the grown structure was modified and is illustrated as grown in Table 3.1. Validation of the band gap energy of each layer was performed by Dr. Tavakoli on a series of calibration samples using photoluminescence testing. Henceforth, this material shall be referred to as Growth A.

Table 3.1: Epitaxial growth structure of CEDT material, Growth A.

Composition	Doping (cm ⁻³)	Purpose	Layer Thickness
In _{0.531} Ga _{0.469} As	P ⁺ = 1x10 ¹⁹ , Be	Contact Layer	200 nm
InP	p = 1x10 ¹⁸ , Be	Ridge	1.4 μm
In _{0.758} Ga _{0.242} As _{0.525} P _{0.475}	1.24Q, p = 1x10 ¹⁸ , Be	Etch Stop	10 nm
InP	p = 5x10 ¹⁷ , Be		200 nm
In _{0.819} Ga _{0.181} As _{0.395} P _{0.605}	1.15Q, undoped	Confinement	80 nm
In _{0.758} Ga _{0.242} As _{0.525} P _{0.475}	1.24Q, undoped	Barrier	70 nm
In _{0.758} Ga _{0.242} As _{0.85} P _{0.15}	1.58Q, undoped	Quantum Well	5 nm
In _{0.758} Ga _{0.242} As _{0.525} P _{0.475}	1.24Q, undoped	Barrier	10 nm
In _{0.758} Ga _{0.242} As _{0.85} P _{0.15}	1.58Q, undoped	Quantum Well	5 nm
In _{0.758} Ga _{0.242} As _{0.525} P _{0.475}	1.24Q, undoped	Barrier	10 nm
In _{0.758} Ga _{0.242} As _{0.85} P _{0.15}	1.58Q, undoped	Quantum Well	5 nm
In _{0.758} Ga _{0.242} As _{0.525} P _{0.475}	1.24Q, undoped	Barrier	70 nm
In _{0.819} Ga _{0.181} As _{0.395} P _{0.605}	1.15Q, n = 5x10 ¹⁷ , Si	Confinement	80 nm
InP	n = 1x10 ¹⁸ , Si	Buffer	0.5 μm
n ⁺ -InP		Substrate	500 μm

The second wafer used for processing was a 3-inch wafer grown by Landmark Optoelectronics in Taiwan, obtained with support from CMC Microsystems. These were procured by Dr. Ye Mengyuan, and their structure is laid out in Table 3.2. The growth structure of these devices was designed by Dr. Mengyuan. Henceforth, this material shall be referred to as Growth B.

Table 3.2: Epitaxial growth structure of Landmark Optoelectronics material, Growth B.

Composition	Doping (cm ⁻³)	Purpose	Layer Thickness
InP	Undoped	Capping	10 nm
In _{.53} Ga _{.47} As	P ⁺ = 1x10 ¹⁹	Contact Layer	200 nm
InP	p = 5x10 ¹⁷	Ridge	1.5 μm
In _{.7322} Ga _{.2678} As _{.581} P _{.419}	Undoped	Confinement	105 nm
In _{.53} Ga _{.47} As	Undoped	Quantum Well	8 nm
In _{.7322} Ga _{.2678} As _{.581} P _{.419}	Undoped	Barrier	10 nm
In _{.53} Ga _{.47} As	Undoped	Quantum Well	8 nm
In _{.7322} Ga _{.2678} As _{.581} P _{.419}	Undoped	Barrier	10 nm
In _{.53} Ga _{.47} As	Undoped	Quantum Well	8 nm
In _{.7322} Ga _{.2678} As _{.581} P _{.419}	Undoped	Barrier	10 nm
In _{.53} Ga _{.47} As	Undoped	Quantum Well	8 nm
In _{.7322} Ga _{.2678} As _{.581} P _{.419}	Undoped	Barrier	10 nm
In _{.53} Ga _{.47} As	Undoped	Quantum Well	8 nm
In _{.7322} Ga _{.2678} As _{.581} P _{.419}	Undoped	Barrier	10 nm
In _{.53} Ga _{.47} As	Undoped	Quantum Well	8 nm
In _{.7322} Ga _{.2678} As _{.581} P _{.419}	Undoped	Confinement	105 nm
InP	n = 5x10 ¹⁷	Buffer	0.5 μm
n ⁺ -InP	n ⁺ = 2x10 ¹⁸	Substrate	600 μm

There are a few key differences to note between Growth A and Growth B. Most glaring is in the quantum wells. Growth A has 3 quantum wells which are composed of In_{.758}Ga_{.242}As_{.85}P_{.15} and are 5 nm wide each. Growth B, on the other hand, has 6 quantum wells composed of In_{.53}Ga_{.47}As which are 8 nm wide each. Furthermore, the confinement layers of Growth A are doped while those layers in Growth B are undoped. There is also a thin capping layer of InP on Growth B, which is to protect the device surface from damage before fabrication. Finally, something that cannot be seen from Table 3.1 and Table 3.2 is that the rear side of Growth B was mirror-polished, leading to extra precautions during the fabrication process to ensure the researcher was working on the correct side of the material.

A total of nine devices were put through fabrication procedures, in variations of the process: A1, A2, A3, A4, B2, B3, B6, B8 and B9.

3.2 Mask

There are three fundamental layers to the mask used to fabricate the laser diodes: mesa (MESA), via (VIA) and metal (MET). The patterns are transferred from the mask by contact lithography onto S1808 photoresist. The mesa layer creates the InP ridges for the waveguides, the via layer etches a hole through a barrier oxide to allow electrical contact with the waveguide, and the metal layer protects certain areas of the chip during metal deposition to create the contact pads.

The first lithography mask used was found in the CEDT clean room from a former user in 1994. The mask had been designed with the purpose of creating III-V laser diodes, although whether this was for course work or a research project was unknown. While this mask was useful in the initial stages of process validation and practicing handling III-V material, there were ultimately a number of significant issues which warranted the design of a new mask. First, due to its age, the chromium of the mask pattern was beginning to flake off and leaving debris on the sample during processing. The alignment marks were also poorly designed, placed all over the mask in different configurations and without any proper parallax indicators. Four of these marks can be seen in Figure 3.1. This made proper alignment of the vias very difficult to perform, resulting in every 2 μm -wide and most 3 μm -wide waveguides failing. Additionally, the original mask was designed so that the sample could be cleaved into laser bars once clean room processing was finished. However, as will be discussed later in this chapter, the next attempts to separate devices were with a dicing saw instead, requiring the addition of dicing lanes. Without dicing lanes, pieces of the waveguide were ripped off by the saw upon separation, leading to more work to clean up the facets, as can be seen in Figure 3.2. Finally, there were no numerical indicators for device locations on the chip, unnecessarily complicating measurements. This mask, which will be referred to as Mask 0, was used to fabricate Devices A1 and A2.




Figure 3.1: Four of the alignment marks from the original mask on device A1, from different areas of the pattern.




Figure 3.2: Device A2, SEM image of a waveguide facet after dicing through the chip and waveguide.

Mask 1 was designed with the intention of dicing the chips that used it. It was designed in KLayout and ordered from Advance Reproductions as a chromium-on-quartz mask of the appropriate dimensions for the mask aligner in the CEDT clean room. Dicing lanes like in Figure 3.3a were added to the mask by creating intermittent breaks in the waveguides and avoiding metal deposition in a row spanning the chip, with a width of 30 μm calibrated to the kerf of the dicing saw. The addition of these lanes also required the cavity lengths of each device to be specified ahead of time, and labelled in the metal layer as in Figure 3.3b.

Uniform cross-type alignment marks with parallax correction were placed in each corner of the mask design such as those in Figure 3.3c. Finally, due to an error in the design, metal was not deposited all the way to the end of every waveguide; the initial intention was to have half of the devices with passive end sections to compare performance with fully active devices, which can be seen by the overlap of the mask layers in Figure 3.3d. Mask 1 was used to fabricate Devices A3, A4 and B9.

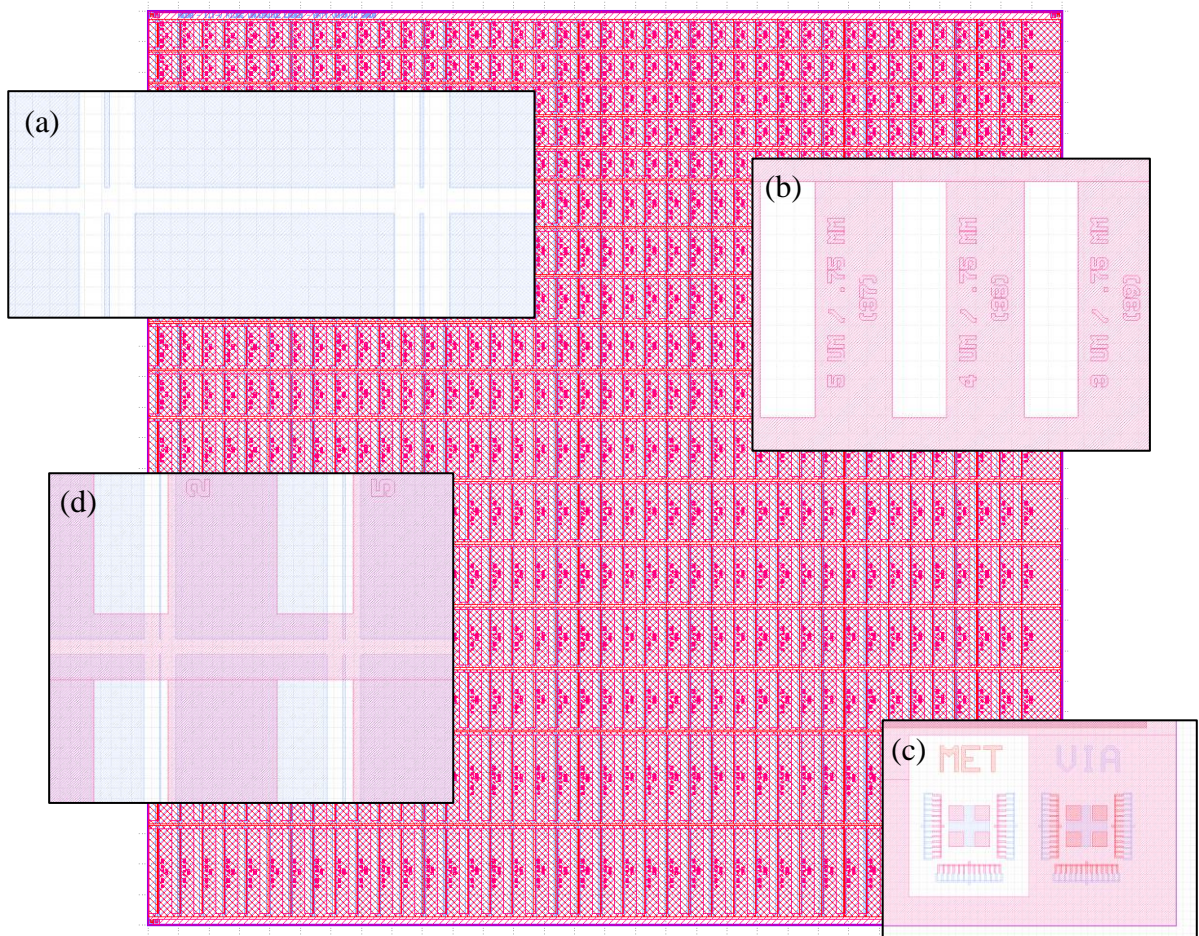


Figure 3.3: Mask 1 key features included (a) dicing lanes, (b) cavity lengths and labels, (c) alignment marks, and (d) passive waveguide ends.

Mask 2 was also designed in KLayout and sourced from Advance Reproductions. The primary reason for designing this mask was the decision to move back to cleaving devices apart instead of dicing. Thus, the dicing lanes were removed from the design. However, to

reduce stress on the structure when cleaving and to indicate where to cleave, paths were cut through the MESA layer to show cleaving lanes, along with notches in the side of each mesa showing the intended fracture line, as can be seen in Figure 3.4a. While this appears similar to Figure 3.3a, note that the waveguides are now continuous across the lane in Mask 2. Labels were changed to be alphanumeric, indicating the exact position on a chip that the device was located, shown in Figure 3.4b. The changes to the alignment marks in Mask 1 led to improved results over Mask 0; however, the minimum feature size of the alignment marks on Mask 1 was 1 μm , which proved difficult to resolve using the CEDT mask aligner along with causing etching issues. Accordingly, there were still issues with properly aligning the via over the 2 μm waveguides. To rectify this, the alignment marks were doubled in size when designing Mask 2. It was also discovered through this researcher's work at RANOVUS Inc. that passive regions in the III-V waveguide lead to significant losses in the system from absorption effects, so Mask 2 ensured that metal was deposited across the entire length of each waveguide, as shown in Figure 3.4c. A section of the chip was also designed for the addition of DBR gratings; as such, one corner of the chip was modified to accommodate them, as seen in Figure 3.4d.

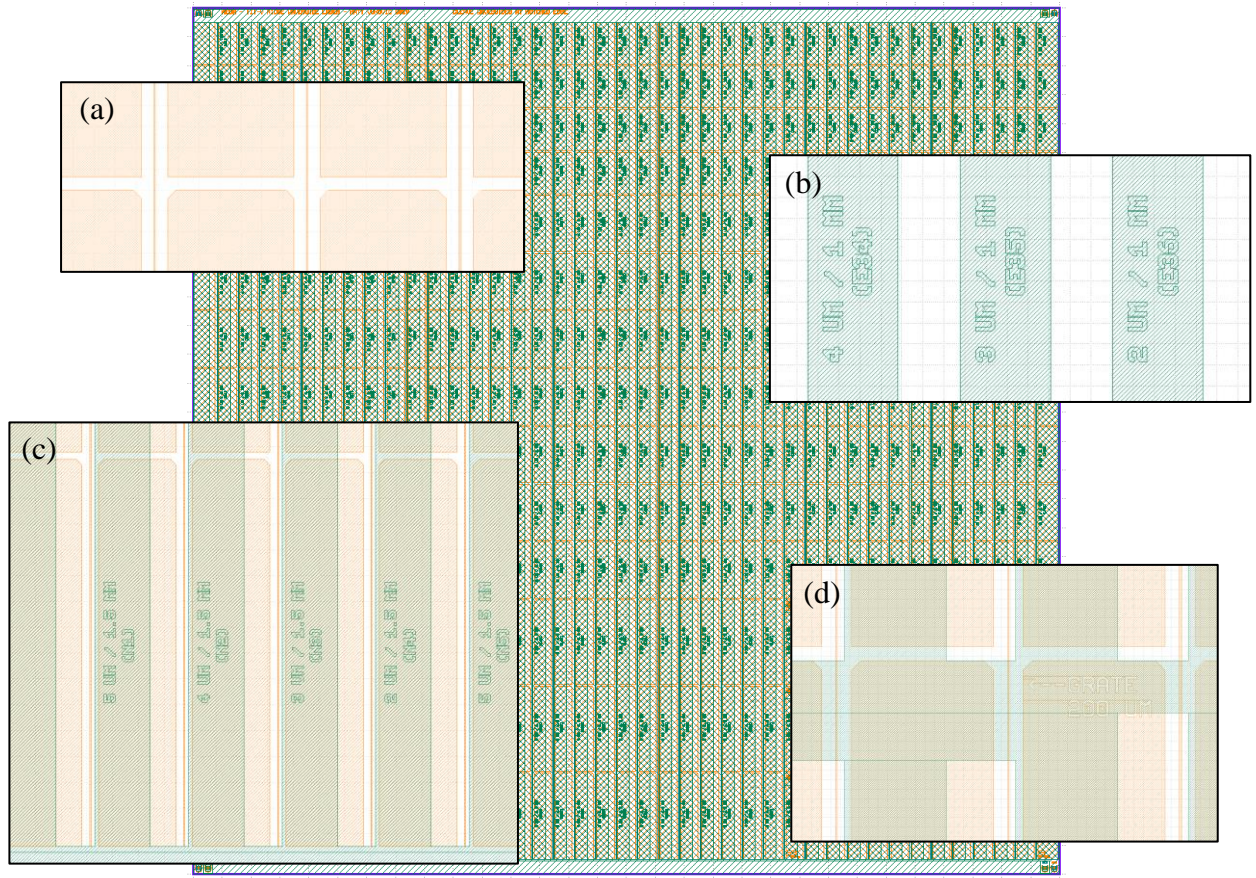


Figure 3.4: Mask 2 key changes included (a) cleaving lanes, (b) alphanumeric labels, (c) full-length metal, and (d) DBR grating accommodations.

3.3 Device Fabrication Overview

For the full process used to fabricate the devices, refer to Appendix A. The schematic and actual cross-sections of the final device can be seen in Figure 3.5. The final process for fabricating the devices is, briefly, as follows:

1. Clean the sample and deposit 100 nm of silicon oxide (SiOx) as a mask oxide.
2. Pattern the MESA mask, then etch through it to create the InP ridge waveguides.
3. Remove the mask oxide, then deposit 100 nm of SiOx as a barrier oxide.
4. Pattern the VIA mask, then etch through it to create the vias on top of the waveguides.
5. Strip the photoresist leftover from the VIA mask, then pattern the MET mask.
6. Deposit a Ti/Pt/Au contact on the top of the chip.

7. Thin the chip, then deposit a Ni/Ge/Au rear contact.
8. Anneal the metal contacts into the device.
9. Cleave device bars.
10. Affix bars to a copper backing piece using silver paint.

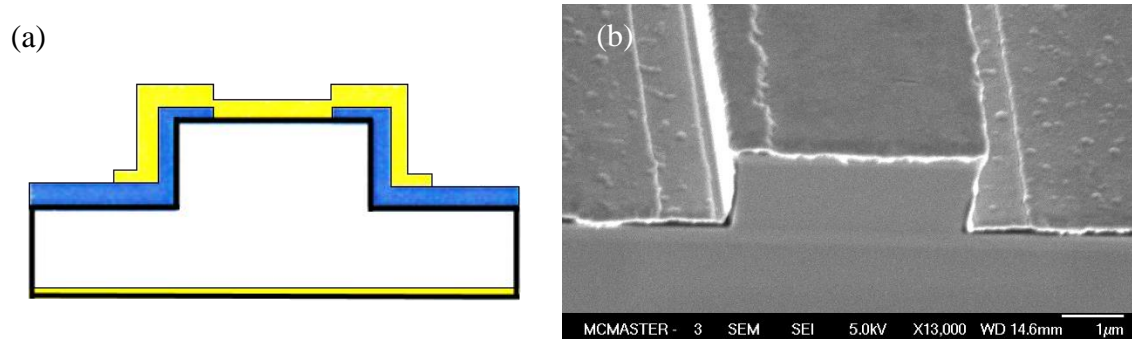


Figure 3.5: (a) Schematic cross-section of final device [25], (b) Facet of device B9 after cleaving.

3.4 Clean Room Fabrication

This section will focus on the failures and troubleshooting that occurred during development of the clean room process, from the initially grown material until the devices were separated into bars.

There were some significant issues faced with photoresist quality during the fabrication processes, and it is hypothesized that they were due to the presence of organic molecules and residue on the surface of the devices. This residue could be left behind by resilient photoresist, atmospheric exposure or, in some suspected cases, by the clean room CVD tool. Organic poisoning of the process presented itself as a problem at two major points: spincoating and post-exposure. If organic molecules were present on the surface before spincoating, the film would delaminate from the sample as in Figure 3.6. Later, when working with Growth B samples, the presence of organics on the surface would cause beads in the photoresist mask after exposure under the mask aligner, as seen in Figure 3.7a. There were several theories as to what could have been causing the issues, from the quality of the chemicals to physical discrepancies to the grown material itself. But after many devices

and calibration samples it became clear that organics were the root cause, especially once some potential solutions were tested.




Figure 3.6: Delaminated photoresist from device A4.

The hallmark of organic contamination ended up being the hydrophilicity of the sample surface. As the photoresist was always spun onto a layer of SiO_x, and SiO_x is naturally hydrophilic, it was observed that samples with organic residuals exhibited partially or fully hydrophobic surfaces instead. This was later tested at every stage by dipping the sample into a shallow dish of de-ionized (DI) water and checking if the water beaded off the surface or stayed on as a uniform layer. If the sample was hydrophilic, processing could continue.




Figure 3.7: Device B6 post-exposure photoresist discrepancies (a) initially, (b) after UV ozone clean.

The first solution attempted was to use a series of non-polar solvent baths to remove as much of the residue as possible. This involved acetone, methanol, isopropyl alcohol and ethanol. While this worked in some instances, it was considered too inconsistent to be reliable. Another solution was to perform a piranha etch, using hot sulfuric acid and hydrogen peroxide on the sample to remove any organics. The piranha etch is often employed in silicon-based devices as an organics clean. This was tried on devices A2 and A3 when they experienced delamination issues, with SiO_x films on the surface. While the piranha etch worked and restored hydrophilicity, the process was arduous and hazardous. Furthermore, the sulfuric acid acts as an etchant of InGaAs, the top contact layer of the device. As such, the piranha etch was unable to be employed if there was organic residue when the vias had been opened up over the waveguides (when spincoating photoresist to form the MET mask). A flash hydrofluoric acid (HF) etch was also considered and tested, again leading to positive results. The quick (2 to 3 seconds) dip in HF would remove a very thin layer of SiO_x, getting rid of any organics clinging to the surface in the process. However, this also had the effect of eroding some of the SiO_x surface, and again proved unusable when organic contamination presented and the vias had been opened, as the finer features were erased or over-broadened. Additionally, it was observed that the HF could unevenly etch the SiO_x surface, leading to height variations and differently degrading the photoresist film quality.

Eventually, the initial steps of the whole process were revisited, where the UV Ozone chamber is used to perform an initial organics clean on the bare III-V surface before deposition of the first oxide layer. This step was replicated when organic residue was encountered, and was found to be highly effective, resulting in fully hydrophilic surfaces after every application. Additionally, since this technique did not involve etching of any existing material layers it could be applied once the vias were cut into the oxide as well. Using the UV Ozone became the go-to method of removing organic contamination from samples. Once the UV Ozone treatment had been applied and the photoresist exposed again,

the beads in the Growth B photoresist were also significantly mitigated as seen in Figure 3.7b.

Another obstacle that was encountered resulted from the anisotropy of the $\text{HCl}:\text{H}_3\text{PO}_4$ InP etch. As the minimum feature size of the alignment marks for Mask 1 were $1\ \mu\text{m}$ and the waveguides needed to be etched $1.5\ \mu\text{m}$ deep, the timing of the InP etch led to full or partial erasure of the alignment marks as seen in Figure 3.8. Mask 2 rectified this by increasing the minimum size of the alignment marks to $2\ \mu\text{m}$. The anisotropy also caused issues with the exposed ends of the waveguides. InP etchant would undercut the SiO_x mask oxide and create pointed facets instead of flat faces, as seen in Figure 3.9. This was undesirable as it increased the amount of work required to create a Fabry-Perot cavity with either the FIB tool or a cleave. This angled facet issue was again mostly a problem for Mask 1 devices, where there were more exposed waveguide ends due to the interruptions for dicing lanes; the problem was again mostly fixed by changes made for Mask 2. It is interesting to note, however, that by controlling this angled facet etch it may be possible to wet etch tapers into an InP waveguide in future projects.




Figure 3.8: Device A3 alignment marks on (a) SiO_x hard mask pattern, (b) III-V material after etch and oxide removal.




Figure 3.9: Etched angled facet of a waveguide on device A3.

The final issues encountered during clean room processing were related to the metal contacts deposited by the e-beam evaporator. The gold of the top metal contact was to be deposited at a 30° and -30° angle to the chip face, as stated in [25]. Through human error, this step was overlooked once, and the reasoning behind the angled deposition became clear. If metal was only deposited directly normal to the face of the chip, not enough metal coated the sidewalls of each waveguide to form a cohesive contact over the edge of the waveguide to the contact pad, as can be seen in Figure 3.10. This meant that electrical pumping of the waveguide was only possible by landing a probe directly on top of the waveguide, almost guaranteeing physical damage to the device. Furthermore, the weak contact would also concentrate the pump current over a smaller area, leading to a higher current density and more heat buildup, eventually causing device failure by overheating.




Figure 3.10: Cross-section of device A4. Note the metal thickness at the top two corners of the waveguide.

The rear contact deposition was more straightforward than top contact deposition, as it only required the metal to be deposited normal to the device surface. However, the devices had been thinned when it came time to deposit the rear contact, and were much more fragile. Since the stage clips for the e-beam evaporator were able to snap normal thicknesses of III-V wafer, a solution using a carrier wafer was devised. 1827 photoresist was selected due to its viscosity, drawn into a syringe with a filter, and brushed onto a silicon wafer as a thin coat. The thinned device was then dropped onto this photoresist layer and lightly depressed to adhere, then placed on a hot plate to cure. This proved to be an effective, low-risk method for bonding the device to a carrier. During the first attempt at deposition using this carrier technique, the low pressure of the e-beam system caused trapped air bubbles in the photoresist to expand, lifting the device off the carrier wafer and suspending it by a thin-walled bubble of photoresist. On a second attempt, as much as possible was done to remove air bubbles underneath the device, applying light pressure and shifting the chip around, and the only bubbles seen appeared at the peripheries of the chip. Once these were removed with tweezers, the deposition could be performed successfully. The devices were then annealed in a rapid thermal annealer and ready for cleaving.

3.5 Device Separation and Faceting

The initial approach suggested by [25] was to thin and cleave the laser bars apart, although no specific directions were given. First attempts at thinning in the CCEM and cleaving in the CEDT were tedious, and it became clear that a large amount of practice would be required to make the techniques workable. As such, the thinning and cleaving approach was abandoned in favour of dicing the chips apart in the CEDT, then using a gallium focused ion beam (FIB) to polish the facets.

3.5.1 Dicing and FIB

The dicing approach initially came to light as other members of the McMaster Silicon Photonics Research Group were working with the dicing saw to separate silicon wafers into smaller pieces for sample prep. Special blades for cutting III-V material were procured from Disco, with a finer grit than those used to cut silicon due to the fragility of the III-V material. The dicing saw used was a MicroAce model, run by Dr. Chang-qing Xu. The silicon wafers could be cut into clean squares or rows with the saw; however, even with the special blades, the III-V material fractured during dicing every time. This sometimes led to the loss of chunks of wafer which went flying off the saw bed, and also sprayed adjacent pieces of the device with debris that often damaged the quality of the top surfaces. From Figure 3.2 it can also be seen that the facet is left grooved by the diamond grit of the saw blade, and needed to be polished by FIB to form proper facets. The major advantages of the process were in the ease of operation, which meant that no physical skills were required to process the material, including keeping the material at its original thickness. Although the reduced yield was a factor, it was ultimately the desire to thin the material for electrical characteristics that motivated the switch back to a thinning and cleaving approach. The dicing and FIB approach was applied to devices A1, A2 and A3.

Once the diced devices were bonded to a copper holder, they were taken to the FIB to be faceted. The gallium FIB tool was run by Dr. Zhilin Peng, owned by Dr. John Luxat. The tool uses a focused beam of gallium ions to ablate material from the sample. The goal in

this project was to use the FIB to improve the quality of the diced facets, in order to generate a mirror finish and create a Fabry-Perot cavity in the device to enable lasing. By changing the ion beam current it was possible to affect the quality of the polished surface; generally, a higher beam current produced a rougher surface, with the benefit of faster material ablation. A progression of currents was employed, depending on the quality of the surface, as seen in Figure 3.11. To remove the bulk of the damaged surface, a 3 nA beam current was applied to the area of interest. After this initial cut, the surface was inspected under the built-in SEM and it was subjectively determined whether or not more polishing was required, based on surface roughness. One cause of surface roughness at this point was curtaining, which is caused by diffraction of the ion beam by defects or changes in material at the surface of the sample. For instance, the edges of the waveguide cause a disturbance in the polished surface, as seen in Figure 3.11. If the facet roughness was deemed to be too high after the initial cut, a 100 pA current was used to polish the facet; this 100 pA current would always result in a high-quality facet, at the expense of a long milling time. After testing a number of FIB'd devices, it was discovered that the impact of curtaining on the facets was not as significant as first thought, as devices with curtaining could still exhibit lasing behaviour. This sped up future FIB sessions as fewer cuts needed to be made to produce adequate facets.




Figure 3.11: Progression of FIB facet milling currents on device A3 (a) initial, (b) 3 nA, (c) 100 pA. Note the reduced width of curtaining lines from the waveguide edge between (b) and (c).

While the FIB proved effective in producing high quality facets, even on waveguides that had been heavily damaged, it still had problems. One issue stemmed from visually lining up the FIB cut to make the facet normal to the axis of the waveguide. If the facet was offset by too great of an angle it was possible for lasing behaviour to be affected as the facet would not function well enough as a mirror. In order to align the FIB aperture perpendicular to the device surface, the stage would be tilted at a 52° angle in the tool. However, any unevenness in the copper backer or the carbon tape holding the sample to the stage would require adjustment of this tilt angle. This could be further complicated by the shape of the dicing saw blade. One side of the blade would cut perpendicular to the face of the chip, while the other would cut at a slight angle; this was necessary to create a fine cutting edge. The facet of the chip which was cut with the angled side of the saw blade would become difficult to line up under the FIB beam, as some intuitive visual corrections would need to

be made with no guarantees that the facet would be perfectly normal to the waveguide axis. Lastly, faceting was also attempted using a plasma focused ion beam (PFIB), but the Gaussian distribution of the beam power was too wide, affecting the precision of the cut as seen in Figure 3.12. The partially eroded waveguide facets in this case would lead to more loss and would likely be unable to lase. As such, the PFIB was abandoned as a method of faceting the waveguides.




Figure 3.12: Facet after milling by PFIB. Note the degradation of the top edge and corners of the waveguide.

3.5.2 Cleaving and Thinning

The thinning and cleaving process was revisited for several reasons. First, thinner devices were hypothesized to operate on a lower threshold current and voltage, making them more electrically efficient and potentially raising their peak power. Second, a thinner device would be easier to affect with a heat sink, leading to fewer ruined devices due to overheating. Third, if thinning and cleaving could be effectively employed, the device yield would be higher than for the dicing saw purely based on the number of devices to survive the separation process. Fourth, it was desired to check if the facet quality of a thinned and cleaved device was adequate for lasing, or if FIB was truly required to fabricate laser diodes with this process.

Thinning was carried out in the CCEM prep lab, using a lapping chuck sourced from South Bay. The sample would be glued onto a metal puck with crystal bond, then attached to the chuck. The chuck would then use a micrometer to set the target thickness, and the sample would be ground down using a series of progressively finer grits of silicon carbide (SiC) paper. The bulk of the undesired material would be removed with 600 ANSI grit paper, then polished to a near-mirror finish with 1200 ANSI grit paper once the target thickness was almost achieved. The thickness would be monitored throughout using a micrometer. When the target thickness and finish were achieved, the sample puck was placed on a hot plate to remove the thinned devices, then the sample was cleaned in acetone to prepare for rear contact deposition.

There were numerous issues with the thinning process. The largest impediment was a lack of practice with the tools. There were certain aspects of the process such as when to stop or continue grinding, how hard to land the micrometer probe and the movements of the chuck on the lapping paper that were skills that needed to be acquired. Complicating things further was the imprecision of the chuck itself, which would indicate that the target thickness had been reached when, upon measurement, it had not. This could have been remedied with a new lapping chuck, but South Bay Technology Inc. was no longer in business and no vendors of similar equipment could be found. When finished thinning, and removing the sample from the puck, cracks sometimes propagated in the sample causing fracture. One cause of this was tapping too hard on the chip with the micrometer probe; another cause was grinding too deep with the 600 grit SiC paper, causing deep scratches in the back of the chip. The difference in the thermal expansion rates of the metal puck, crystal bond and semiconductor sample proved an issue as well. Ramping up the temperature of the hot plate too quickly would cause fracture in the sample, so a standard procedure of ramping up by 10 °C over one minute intervals was devised. Once thinning was finished, the devices were scribed by a tool in the CEDT and cleaved by hand. Again, this required practice to figure out the correct vectors and magnitudes of force to apply to the scribed chip to ensure a

straight cleave down the desired atomic plane. Incorrect application of force would cleave only a small, unusable chunk off the device instead of a bar.

An attempt was made to circumvent the manual thinning process by using an auto-polisher. Such a tool would have made it possible to thin up to six samples simultaneously and within minutes. Unfortunately, it proved difficult to control the thickness of the samples lapped by the tool as it was intended for use with materials such as steel. Ultimately, the effort was abandoned due to safety concerns over off-gassing of phosphine during the lapping process, and the lack of an adequately ventilated environment in which to use the auto-polisher. Such a tool could be worth revisiting if a proper fume extraction system could be implemented.

3.6 Bonding

With the device bars thinned, cleaved, contacted and annealed, all that needed to be done to complete the testable device was to stick them to a carrier. The carrier needed to be thermally conductive and easy to shape to the size of the device bar, so copper sheet metal was chosen. To adhere the device to the copper while also allowing the back to be contacted, silver paint was used. This allowed the researcher to land a probe on a silver paint pad which would act as a common ground for every device on the chip. The only part of this process to optimize was ensuring that the copper backer was as flat as possible. A flat backer would ensure better rear contact, as well as making it easier to FIB the facets of those devices which went through that process. The copper was flattened using a hammer and the anvil of a vice. Once mounted, the devices looked like the one in Figure 3.13.




Figure 3.13: Final device after mounting on a copper backer with silver paint.

3.7 Anti-reflection Coating

The addition of an AR coating to the facets of the laser diodes would, in theory, transform them into SOA's. For relative ease of processing, a four-layer stack of SiOx and Si₃N₄ was selected as the AR film, as the CVD operators were well-acquainted with growing such films. Having measured the peak wavelength of the devices to be around 1580 nm, a starting idea of the film thicknesses of each layer could be calculated using the optical thickness of a quarter wavelength within each material, as:

$$t_{optical} = \frac{\frac{1}{4}\lambda}{n}$$

where n is the index of refraction of the layer and λ is the peak wavelength of the laser in air. The initial guesses were then plugged into an online calculator provided by Filmetrics [26], and from there the coating was optimized through trial and error, re-simulating with different layer thicknesses. The resulting film had Si₃N₄ layers of 53 nm and SiOx layers of 73 nm, giving the reflectance spectrum in Figure 3.14. As can be seen, the chosen film should have reduced reflectance to below 0.5% for light from 1500-1650 nm.




Figure 3.14: Reflectance spectrum generated by the Filmetrics calculator [26] for the chosen dielectric stack.

The film was grown using Dr. Peter Mascher's ECR-PECVD tool with the help of Jeremy Miller. A layer of Si_3N_4 layer was grown first on the substrate, due to its higher index, then a layer of SiO_x , then Si_3N_4 , then the final layer of SiO_x . Previous work with silicon-based substrates suggested creating a stack of devices and laying them so that the facets faced the incident plasma; however, such a technique was dubious for III-V materials due to their fragility. As such, it was decided to deposit the coating on top of the chip with the hope that enough would slip over the chip edges to coat the facets. It was known that III-V devices had successfully had their facets coated in AR films before, but the literature, while ripe with such devices, was bereft of any explanation as to the exact methods of deposition.

As the coating was to be deposited on top of the chip, a mask needed to be created to protect the contacts so that the devices could still be electrically pumped afterwards. A piece of aluminum wire was flattened with a hammer and formed such that it would lie in contact with both the device and the deposition stage to which it was clipped, as in Figure 3.15. The thickness of the aluminum wire was chosen such that, when flattened, it covered $2/3$ of the device width, leaving the facets exposed. The result was a film that coated the device as in Figure 3.16.



Figure 3.15: Schematic of AR deposition mask in grey, covering the sample in brown.




Figure 3.16: SEM image of AR film coating over the ends of the waveguides.

3.8 Gratings

One sub-project was to etch DBR gratings into the ends of the laser diode waveguides to produce single line lasers. The project was attempted using the PFIB in the CCEM with the assistance of Connor Wong. The PFIB was selected over the regular gallium FIB because it was capable of programmable etch patterns. Values for grating pitch, length and depth were provided by Dr. Ye Mengyuan.

Initial cuts into fully fabricated material (metal deposited and annealed with barrier oxide underneath) proved difficult, as the metal would quickly backfill any gratings that the tool had cut. As such, device B6 was fabricated until right before the barrier oxide was to be deposited, and gratings were then attempted on the bare III-V waveguides. Beam current,

ion dose and dwell time were varied in attempts to create defined gratings, with the results visible in Figure 3.17. There were numerous failures, seemingly due to errors in the control program; identical beam settings sometimes produced wildly different results, either making a cut too shallow across the waveguide or ablating it entirely. It was this inconsistency that forced the project to be abandoned, even with some promising results such as those in Figure 3.18.




Figure 3.17: Series of tests as performed on one of the B6 waveguides. Note the high variance between adjacent cuts, which had small or no changes made to beam parameters.




Figure 3.18: Images of some promising results achieved by the PFIB for DBR gratings in device B6.

4 Characterization

This section will detail the measurement processes used to characterize the fabricated devices, both electrically and optically. The current and voltage response were measured alongside the output optical power of the laser diodes and SOA's. The spectral output was then measured for both types of devices. Finally, the amplification of the SOA devices was measured in the same fashion as band edge measurements performed on a laser diode as part of the EAM project.

The device, on heat sinks, was attached to a Thorlabs lab jack to allow for vertical adjustment, which was mounted on top of a lateral translation stage to easily move samples under the microscope's field of view. The entire setup was built on an optical bench. The devices were mounted on a copper heat sink with silicone thermal paste, which was screwed onto an aluminum block with a thermoelectric cooler (TEC) sandwiched between the two. The copper mount would dissipate heat from the chip, and the block would dissipate heat rejected from the TEC. A thermistor was mounted within the copper heat sink, not directly on the sample; as a result, it did not accurately measure the surface temperature of the device, but some lower temperature instead. Accordingly, the temperature setpoint was selected as 20 °C in an attempt to keep the actual surface temperature around 25 °C. Two configurations of the heat sink setup were used: one for the integrating sphere, and one for fiber coupling from both ends. The former can be seen in Figure 4.1, while the latter can be seen in Figure 4.2.




Figure 4.1: LIV setup, including integrating sphere, probes, heat sink, lab jack and lateral translation stage.

4.1 Electrical Characteristics and Optical Power

To test the electrical and optical power (LIV) characteristics, a current was applied through the device via the top and bottom contacts. The tests were current-controlled to measure the light-current (LI) relationship, and the current-voltage relationship (IV) was also recorded. This would define the threshold current at which lasing occurs in the device, as well as operating point conditions. Additionally, driving with voltage instead of current has the potential to damage the device. A Keithley 2400 acted as both the current source and voltmeter. A Thorlabs S145C integrating sphere with a responsivity centred around 1550 nm was used to measure output optical power, and a Lightwave LDC 3742 laser diode controller was used to regulate the temperature of the samples using the TEC and thermistor. The setup can be seen in its entirety in Figure 4.1.




Figure 4.2: Heat sink configuration used for fiber coupling to both ends of the device.

To test, one probe was landed on the silver paint of the copper carrier as a common rear contact. If this was the first time the device had been tested since mounting to a copper backer, the second probe would be landed on a different area of the silver paste to check for continuity. The second probe was then landed on the top contact of whichever device was being tested. For direct optical power measurements, the device bar was brought as close to the edge of the heat sink as possible and the integrating sphere was butted up to the aluminum block. The large detection area of the integrating sphere meant that it could collect almost all of the power emitted by the device under test, giving an accurate reading of the real output power without having to account for any interfacial losses. Once the temperature setpoint was reached, the Keithley current was set to pump 20 mA across the device with a 3.2 V voltage compliance limit, to check if the device was electrically viable. An open circuit at this point would indicate a blown device with no chance of working. If a voltage around 1 V was seen, the current was slowly ramped up in 20 mA steps, recording the voltage and optical power at each step. The results were then compiled in Microsoft Excel.

As the current was ramped up, it would be obvious to the observer if a device crossed the lasing threshold or if the device was not functioning either as a laser or a diode in general. Lasing behaviour was signified by a sharp rise in output optical power when increasing current by a step; this increase was often on the order of 10 or even 15 dB. Some devices, particularly ones with 2 μm waveguides, showed much smaller increases at threshold, and the lasing behaviour of such devices was normally only spotted when graphing the LI curve. Some devices only exhibited optical power that rose proportionally to the pump current and never rose significantly; in these cases, the device either failed to achieve population inversion or had inadequately mirrored facets to form a Fabry-Perot cavity. If the voltage across the device rose linearly with applied current, the device was ohmic and not functioning as a diode. This was caused by a short from the top to bottom contact, such as a rogue streak of silver paint on the facet or punch-through due to over-voltage.

Device failure during testing was easily spotted. Failed or “blown” devices would experience a rapid decrease in power and increase in voltage before failing to an open-circuit condition. This was most often experienced when testing at high current pump values and can be attributed to thermal failure as a result of current heating. Excessive heat could cause damage to the metal contacts or physical expansion. Lastly, if the probes were bumped, slipped, or otherwise removed from the contact pad during operation, a spark could jump across the contact as the system attempted to correct for a continuous current flow. This would cause catastrophic failure of the device, with a scorch mark easily visible on the surface of the contact pad.

4.2 Optical Spectra

The optical spectra of the devices were measured with a Thorlabs 203C optical spectrum analyzer (OSA). The spectrum data was recorded from the OSA using the Thorlabs OSA application on a connected computer. To perform these measurements the setup had to be changed slightly from the direct power configuration, to the arrangement shown in Figure

4.3. The integrating sphere was removed in favour of a tapered fiber which was mounted on a Thorlabs 3-axis stage and edge-coupled to the device facet.




Figure 4.3: Configuration for optical spectrum analysis, the same as Figure 4.1 but with a coupled fiber instead of the integrating sphere.

To couple the fiber, the device was pumped with 200 mA to reach an operating point that produced lasing or at least noticeable power in every device, without danger of blowing the device. The APC end of the fiber was connected to the integrating sphere and the fiber was adjusted using the stage to find the point of maximum coupled power. The difference between this power and the direct optical power from the device was taken as the interfacial coupling loss of the fiber to facet. The APC connector could then be swapped over to the OSA, where the spectrum was taken for a range of different pump currents.

4.3 Band Edge and Amplification

For amplification, an Agilent 8164A was used as an external laser source in conjunction with its internal optical power meter (OPM). A JDS Uniphase OAB EDFA was connected to the output of the Agilent to amplify the signal up to 15 dBm to sweep a broad range of input power, as the maximum attainable power of the Agilent on its own was below 5 dBm. The EDFA had a wavelength limit of 1528 to 1565 nm. A JDS Uniphase HA1 Optical

Attenuator was placed at the output of the EDFA, to enable measurements with input lower than -12 dBm, the lowest output power of the Agilent. A polarization controller was inserted between the attenuator and the input facet of the device, since amplification within the SOA is highly sensitive to polarization. Tapered fibers were edge-coupled to both ends of the device, with an input facet and output facet arbitrarily determined. Finally, an OzOptics tunable bandwidth filter was inserted between the output facet and the OPM to pass only the wavelength region of interest. The range of the filter was 1525 to 1570 nm. A photograph of this setup is shown in Figure 4.4a, with a schematic in Figure 4.4b. The amplification sweeps of the coated device were performed in forward bias, while the band edge measurements were in reverse bias for an uncoated device. Otherwise, the band edge setup was a simplified version of the amplification setup, without the filter, EDFA or attenuator. Dr. Ross Anthony performed the band edge measurement and the researcher performed the amplification measurements.

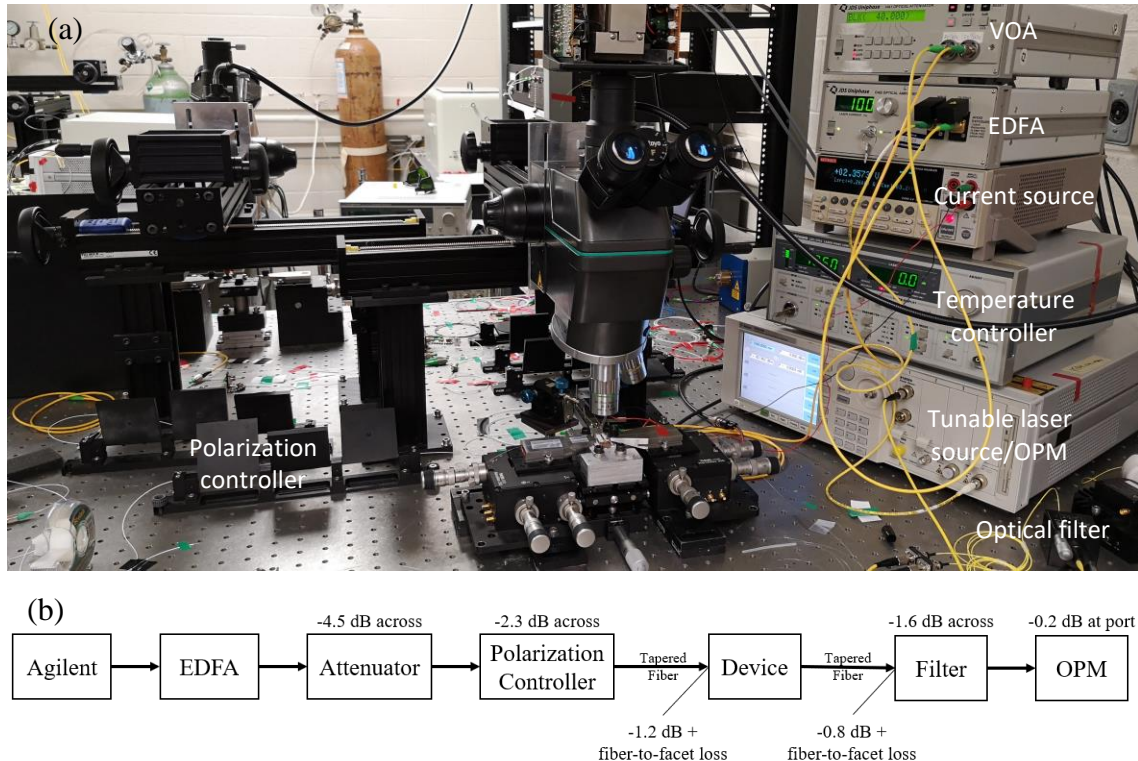


Figure 4.4: (a) Photograph of the amplification setup. (b) Schematic of the setup including system losses.

Before measurements began, system losses were calculated. For the band edge measurement, the system loss was characterized with a fiber-to-fiber power measurement. For the amplification measurements, the system output power curve was measured using the on-board sensor, and the integrating sphere was used to systematically determine the losses at each point of the setup. Since the chip output power was known, the fiber-to-facet coupling loss could be accurately determined by coupling the fiber to the chip and measuring the attenuated power at the fiber output for each device operating at 200mA, as before. Differences in the spot sizes of the tapered fibers and facets, as well as any facet defects meant that a fiber-to-fiber measurement was not necessarily representative of the system with the III-V device inserted. There were also losses across each tapered fiber, which were measured by applying a known power to the APC ferrule and sticking the tapered end into the integrating sphere cavity. The losses for the amplification setup are shown in Figure 4.4b.

The band edge measurements, as performed by Dr. Anthony, applied a reverse bias across the device from 0 to -5 V. At each bias level, the wavelength was swept from 1560 to 1640 nm with a power of 0 dBm from the Agilent laser source, and the output power measured. A sharp increase in output power at a given wavelength would indicate the onset of transparency in the material and thus indicate the band edge of the device. A laser diode from device B6 was used for these tests.

A coated B6 device was used to test amplification. The device was pumped with 200 mA and the tapered fibers were coupled to the facets by maximizing the measured output power. The current was then adjusted to the desired setpoint, and the fibers re-positioned for any changes in the optimum coupling point. Matlab code would then tell the Agilent to output at 0 dBm and 1580 nm to allow the researcher to adjust the polarization of the input laser to match the polarization of the waveguide. These Agilent settings provided an input signal that, by amplification in the guide, would noticeably alter the output power. 1580 nm was determined as an optimal wavelength for this after analyzing the band edge data and a

sweep of output power over input wavelength, as seen in Figure 5.27. Polarization matching of the signal and cavity could be found by maximizing the amplified power seen at the OPM by adjusting the paddles of the polarization controller. There was no observed change in the optimal input coupling point for the device into the fiber compared to the fiber into the device. The drift of the output coupling point was found to be negligible (less than 1 dB) over the sweep time unless the pump current was high, in which case thermal effects could cause the optimal point to drift. Once aligned, the rest of the Matlab code was executed. The Matlab code was roughly based on a script written by Dr. Zhao Wang, but significant changes were made based on this researcher's work at RANOVUS Inc. A few iterations were used for the various sweeps, but a representative script can be found in Appendix C: Amplification Sweep Matlab Code. The code performed a stepped wavelength sweep at various levels of input power or attenuation, recording the wavelength, input and output power to the device and true attenuation before concatenating the data vectors and writing to a CSV file.

The first set of sweeps was to determine the behaviour of the device over a range of input wavelengths. To avoid limitations on the input wavelengths, the EDFA and filter were not used for this measurement. Instead, the Agilent was set to sweep across 1510 to 1640 nm, first at -10 dBm for a range of 0 to 30 dB of attenuation, then from -10 to 4 dBm with 0 dB of attenuation. The Agilent was incapable of producing more than 4.5 dBm of output power. During the sweep, the OPM sensor and attenuator wavelengths were shifted to the laser wavelength.

From the first set of sweeps it was found that the best wavelength for amplification in the devices was 1580 nm. Further sweeps were performed at this wavelength for a wider range of input powers. As the filter and EDFA could not perform at 1580 nm, a different solution was devised. The optical spectra at each level of attenuation was recorded, and by weighting the measured levels of the spectrum against the total measured power (since loss at the

OSA input was unknown) the signal could be digitally filtered by isolating a band of values around the signal wavelength, and the gain curves then determined.

As a last step, the EDFA and filter were added to the system to check the broadest possible range of input power. To avoid operating at the fringes of their specified ranges, while also using a wavelength that was known to display some degree of amplification in the devices, 1560 nm was chosen for the input signal. The EDFA was set to maximum amplification, and the filter set to a 2 nm band pass around the signal peak (i.e. 1559 to 1561 nm). The input and output power were then recorded, along with ASE power.

After characterizing the total output power, spectrum and amplification, there were three distinct criteria that could be used to determine if a device was functioning as an SOA:

1. The output power of the device when pumped to 200 mA is in the range of -18 to -23 dBm, indicating **amplified spontaneous emission**. This often means population inversion has been achieved, but the photon flux in the cavity is not high enough to cause lasing anymore.
2. The optical spectrum displays **no lasing peaks**, but rather a broad, continuous distribution of power across a wide range of wavelengths.
3. Coupling of a laser signal to the device causes an appreciable **gain** in the output power in proportion to the signal power. The output power may not exceed the input power due to system losses, but this still indicates that the source signal is stimulating emissions in the cavity.

5 Performance and Analysis

5.1 Simulation

The structure of Growth A was simulated in RSoft to check which waveguides, if any, on the design would carry only a single mode. The simulation was set up with the help of Dr. David Hagan. The waveguide width was tested at 2 μm , 3 μm , 4 μm and 5 μm . It was discovered that the waveguides began to carry a single mode at 3 μm wide, with the 2 μm guide also carrying a single mode, and the 4 μm and 5 μm guides carried up to two modes. The simulated modes can be seen in Figure 5.1.

5.2 Anti-reflective (AR) Coating Deposition

Two coating runs were performed to deposit the four-layer silicon oxide-nitride stack. The first run was performed on device A3, and the second run was performed on device B6. The parameters for each growth were the same, but there were noticeable differences in the film quality when examined under the scanning electron microscope (SEM). The A3 film can be seen in Figure 5.2. Of the individual devices on the coated piece of A3, only one (A3-21) showed signs of adequate power attenuation. It is unclear what caused the “wrinkled” nature of the A3 film, however the anti-reflective properties of the coating were less effective than for the B6 film, as will be described later in this chapter. It is hypothesized that the wrinkling effect could have to do with the amount of environmental exposure before coating, as A3 was fully fabricated and idle for almost 5 months before it was coated, whereas B6 was only idle for 1 month. This leaves moisture and dust as prime suspects to consider. It is also noted that A3 had facets polished by focused ion beam (FIB) and was not thinned, compared to B6 which was thinned and cleaved.




Figure 5.1: (a) to (d) depict the TE_0 modes for 2 μm to 5 μm waveguide widths, while (e) to (h) show TE_1 modes. Note that only the 4 μm and 5 μm guides have real TE_1 modes.




Figure 5.2: AR film under Ga-FIB SEM, as deposited on device A3. (a) overhead view of film, with masked contacts visible, (b) overhead view of left facet of A3-21 (affected waveguide), (c) closer, angled image of A3-21's left facet, (d) angled image of A3-21's right facet.

The B6 film turned out uniform, with signs of wrinkling presenting next to the waveguides, although to nowhere near the same extent as A3, which can be seen in Figure 5.3a. A comparable overhead view to Figure 5.2a can be found in Figure 3.16. Plasma focused ion beam (PFIB) SEM cross-sections of the device reveal that the AR film not only coats the top of the device, as in Figure 5.3b, but also the facet as per Figure 5.3c. Figure 5.3d shows the facet of B6-17, which was later proven to amplify, and shows an exceptionally clean facet after coating.




Figure 5.3: AR film under PFIB-SEM, as deposited on device B6. (a) evidence of film wrinkling next to one waveguide, (b) PFIB cross-section of a coated waveguide, noting the lighter colour is the metal contact, and the thicker layer on top is the AR film, (c) PFIB cross section of the chip facet, noting that the film runs over the facet to coat it, (d) coated facet of B6-17, one of the amplifying structures.

5.3 Optical Power and Electrical Characteristics

In general, devices which were fabricated using Mask 0 (A1 and A2) were too inconsistent to compare to devices made with Masks 1 and 2. No devices from A1 were found to lase. Only one device on A2, A2-21, was found to lase at a threshold current of 180 mA and an output power greater than 7 dBm. These threshold and power characteristics would be far outclassed by the devices from A3. Neither A1 nor A2 was thinned, and both had FIB-polished facets to create the cavity mirrors. A2 did, however, demonstrate the effect of temperature on the threshold value as shown in Figure 5.4. As the diode temperature is increased the intrinsic losses increase, so the threshold current also increases and power decreases [27]. The slope efficiency of the laser (output power over input current) decreased with temperature similarly to the output power, as seen in Figure 5.5. For future

devices, 20 °C was chosen as the setpoint as the thermoelectric cooler (TEC) had difficulty consistently driving to 15 °C. A lower setpoint temperature allowed for not only lower threshold current, but also reliable operation at higher pump current levels.




Figure 5.4: Optical power vs. current (LI) curves for A2-21 at different temperatures. Linear scale is used to show the discrepancy in threshold current more clearly.




Figure 5.5: Slope efficiency of A2-21 over different temperatures.

A full current-voltage (IV) sweep of one of the B6 devices was taken to check how far it could be pushed into reverse bias, to check the band edge properties which will be described later in this chapter. B6-M-33 was measured from -5 to 2 V, with the resulting curve shown in Figure 5.6. The diode behaviour of the device is clearly exhibited, with an exceptionally low dark current of around -20 nA and no indication of approaching breakdown even at -5 V bias. Comparing the IV curves for A3 and B6, as in Figure 5.7, the only major difference for similar waveguide structures is the slope of the IV curve. The change in voltage per unit current is larger for A3, indicating that it has a higher resistance than B6. This makes sense, as B6 was thinned while A3 was not, meaning A3 had more material between the top and bottom contacts. The voltage can also change easily depending on the quality of the electrical contact, and was seen to decrease over the course of some measurements as the contact burned into the device and improved the connection.




Figure 5.6: Full forward and reverse bias IV curve for B6-M-33, a 5 μm wide device.




Figure 5.7: IV comparison of 5 μm structures on A3 and B6. Threshold voltage remains the same but slope differs.

The effects of cavity length on the LI curve can be demonstrated with device A3, but not with B6. A3 was diced, meaning that, despite the trauma inflicted on the chip by the saw blade, it was possible to accurately control the cavity lengths of device bars. When cleaving a thinned chip such as B6, the brittleness of the III-V material allowed the cleave line to wander from the initial scribe mark, resulting in cavities of approximate, but not precise, length. In Figure 5.8, the LI curves of A3 devices with different cavity lengths and waveguide widths are compared. Only one A3 device with a waveguide width of 3 μm was found to lase. A trend emerges, showing that the longer cavity lengths have lower threshold current, and are also more resilient to thermal losses as pump current is increased. The former trend may be explained by an increased probability of stimulated emission as there are more excited states available along the cavity. The latter may be explained by a slightly larger surface area over which heat could dissipate; however, it is also possible that this was a result of slightly different setup conditions during testing. For varying widths of A3 devices, there does not seem to be a difference in the output power, at least between 4 μm and 5 μm waveguides. The maximum power achieved from the A3 devices was 13.62 dBm, or 23 mW.




Figure 5.8: Comparison of LI characteristics for A3 cavities of varying lengths and widths.

Unlike A3 devices, B6 devices were cleaved as full bars, and could thus be tested under the same conditions as one another, in one session. Using Mask 2, devices on one end of B6 contained passive regions initially intended for DBR gratings; however, in this case they serve to show differences between waveguides with passive regions and waveguides with metal across their entire ~ 1.5 mm length. As mentioned in Chapter 3: Design & Fabrication, A3 used Mask 1 and had passive regions on both ends of every waveguide. A comparison of various widths of B6 waveguides, both with and without the passive regions, can be seen in Figure 5.9. One of the biggest differences is that the threshold current is around 60 mA for devices with a passive region, whereas threshold current has decreased to 40 mA for fully active devices. There is also a slightly higher peak power for fully active devices, as well as a higher resilience to thermal losses at higher pump currents. These are both indicators that there is less intrinsic loss in the cavity, aligning with the hypothesis that the passive regions are absorbing photons. There was a wide performance variance among all B6 2 μm waveguides, so it is difficult to say if the fully active 2 μm waveguide truly outperforms the partially passive one by not dropping off in power as quickly.

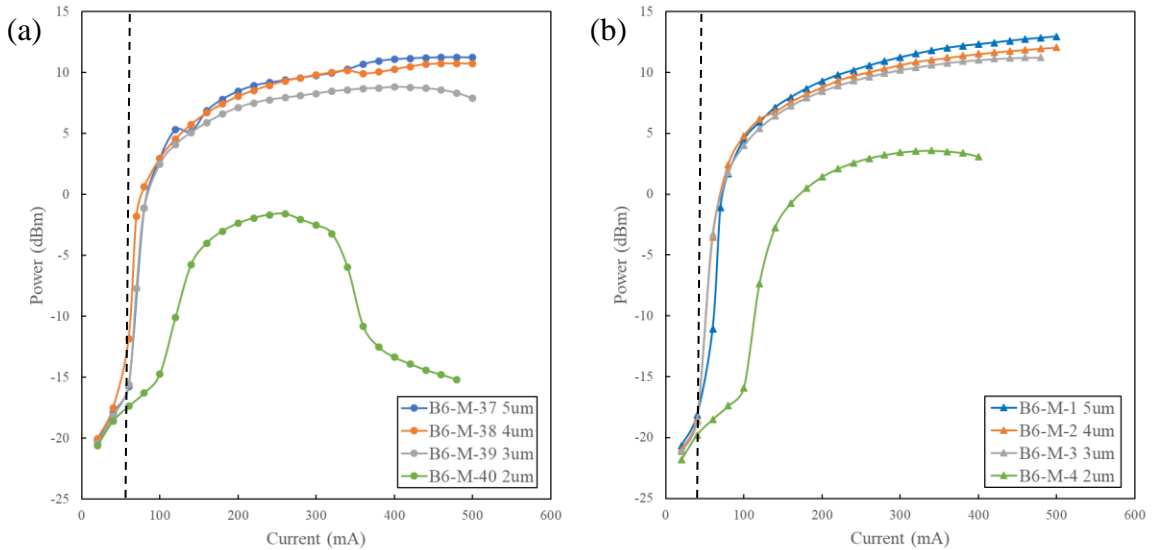


Figure 5.9: Comparison of LI characteristics for B6 devices of varying widths with (a) a passive region and, (b) no passive regions. The dashed lines mark the threshold current values.

The threshold current stays the same regardless of width for B6 devices, and the power slightly decreases with decreasing width. This is most likely due to a change during the fabrication process, as a similar decrease in power with decreasing width was not observed for A3 devices. One possibility is a difference in epitaxial structure, as A3 and B6 were made of different materials. The materials could have etched in different ways, leaving one with rougher waveguide side walls. Slight changes in the way the vias came out could also affect the power, as metal coating the sidewall of the waveguide instead of the top would change the electric field across the active region. The via etch showed variance over different fabrication runs due to changes in the quality of the deposited barrier oxide.

Comparing the LI characteristics of A3 and B6, the threshold current is clearly lower for B6 devices, which lased at 60 mA compared to 100 mA for A3. This indicates that there are more intrinsic losses in the A3 cavities, most likely due to the presence of passive regions at both ends of the waveguides. Despite the difference in threshold current, lasing power remained consistent between the two runs, indicating a similar quality of facet and active region, despite the difference in fabrication techniques (dice/FIB for A3 and thin/cleave for B6) and material composition. The increased precision of Mask 2 over Mask

1 also led to successful attempts at fabricating devices with waveguide widths of 3 μm and 2 μm on B6, compared to the single 3 μm laser from A3 and no successful 2 μm devices. The maximum power achieved from B6 laser devices was 13.02 dBm, or 20 mW.

5.3.1 LI Characteristics After AR Coating

After depositing the AR film on A3 and B6, the LI curves were remeasured to check if lasing had been suppressed. For A3, the power clearly dropped after coating, as seen in Figure 5.10. Only A3-21 and A3-22 were functional after coating. The power of A3-22 dipped slightly but the device continued to lase at the same pump current; however, the power of A3-21 dropped completely, down to levels expected for amplified spontaneous emission (ASE), and did not exhibit a lasing threshold current. This was a promising result, but before any spectrum or amplification measurements could be taken the device was rendered untestable by thermal paste. Even still, this gave promise to the crude masking method developed to shield the contacts of the devices while still allowing the AR coating to be deposited on the facets. The change in power at low pump current is most likely due to differences in the measurement setup.




Figure 5.10: Comparison of LI curves for coated and uncoated A3 devices.

There were inconsistent results from coating the B6 waveguides, although the coatings were ultimately successful on a handful of devices such as B6-L-12, shown in Figure 5.11. The successful coatings could be seen in the same way as A3-21, with a large drop in output power and no sign of lasing behaviour on the LI curve. Later amplification testing would confirm that these coatings had worked.




Figure 5.11: B6-L-12 LI curves for the coated and uncoated device. The waveguide width was 2 μm .

Some coated devices only exhibited a small drop in the output power compared to before, such as B6-L-26 in Figure 5.12. In this case, it is most likely that the coating covered the facet, but not in the desired thickness. This could have resulted in attenuation of the output instead of AR behaviour. Upon further examination, the threshold current also appears to have increased after coating, suggesting that the film has introduced more loss into the device, resulting in some AR behaviour. This would reaffirm the theory that the facets are improperly coated.




Figure 5.12: B6-L-26 LI curves for the coated and uncoated device. The waveguide width was 4 μm .

Devices such as B6-L-17, shown in Figure 5.13, exhibited virtually no change in LI characteristics after coating. This would suggest that no coating covered the end facets; however, later in this chapter it will be shown that such devices can still amplify if operated below their threshold current. A few devices such as B6-L-22 in Figure 5.14 displayed an increased output power after coating, without significant change to the threshold current. This is unintuitive, and may indicate pre-existing facet defects on the device that were smoothed out by the film deposition, creating better facet mirrors. Only 2 of the 25 devices measured displayed increased power after coating.




Figure 5.13: B6-L-17 LI curves for the coated and uncoated device. The waveguide width was 5 μm . No change is seen in the output power levels.




Figure 5.14: B6-L-22 LI curves for the coated and uncoated device. The waveguide width was 4 μm . The output power level increases after coating.

5.4 Optical Spectra

The one A2 device that was lasing (A2-21) had its spectrum taken across a range of pump currents and temperature setpoints. Figure 5.15 shows the spectrum as recorded at 20 $^{\circ}\text{C}$ with a 300 mA pump current; this device had a high threshold current around 200 mA. One may notice the multi-mode nature of the device, with many secondary peaks surrounding a clear central maximum near 1565 nm. It is also noted that there is significant ripple in the spectrum, potentially due to back reflections in the measurement setup, or due to the light being coupled to the optical spectrum analyzer (OSA) through free space instead of a tapered fiber.




Figure 5.15: Optical spectrum of laser A2-21 held at 20 °C with a 300 mA pump current. Note the ripples in the tails of the spectrum.

As the pump current and temperature setpoint were individually increased, similar effects on spectral performance were noted. In Figure 5.16, the conditions in Figure 5.15 are compared to a higher temperature setpoint at the same pump current level, and a higher pump level at the same temperature. Increasing either quantity pushes the output peak of the device to a slightly higher wavelength and holds output power slightly lower. This indicates that the two variables are affecting the same parameter of the device, with increased pump current likely replicating the thermal effects.




Figure 5.16: Comparison of optical spectra of A2-21 at different temperature (25 °C) and pump current (400 mA) conditions.

With device A3 and later, OSA measurements were performed with tapered fibers as opposed to free space coupling. A3 demonstrates more reasonable behaviour than A2, with no ripple and none of the comb-like structure. The spectrum of A3-L1(2)-1 (a 5 μm waveguide) is displayed in Figure 5.17 at 20 °C and 200 mA pump current.




Figure 5.17: Spectrum of A3-L1(2)-1, a 5 μm device at 20 °C and 200 mA. Peak width is approximately 10 nm.

The threshold current of A3-L1(2)-1 is 130 mA. A measurement of the spectrum at 120 mA, just before reaching threshold, reveals that the device is emitting strong ASE just prior to lasing, as suggested by theory. This spectrum is shown in Figure 5.18. The power is low at the output of the pre-threshold device, but also distributed over a broader range of wavelengths than the lasing device. This regime was later explored with B6 devices to test if amplification could be seen at low current in devices whose output powers were not fully attenuated by the AR film.




Figure 5.18: Spectrum of A3-L1(2)-1 just before reaching the lasing threshold, at 120 mA pump current. Note the wide range of wavelengths over which power is distributed.

From the simulations earlier in this chapter, it was seen that the 5 μm waveguide more readily supports multiple transverse modes compared to its narrower counterparts. This becomes evident when comparing the peak width of A3-L1(2)-1 to A3-L1(2)-6, a 4 μm device. The spectrum of A3-L1(2)-6 at a pump current of 200 mA is shown in Figure 5.19. The peak width of this device is only 5 nm compared to the 5 μm device, which had a peak width of around 10 nm at 200 mA of pump current in Figure 5.17. Despite the difference in peak width, both devices share the same central wavelength, around 1587 nm.




Figure 5.19: Spectrum of A3-L1(2)-6, a 4 μm device at 200 mA. Note the peak width of around 5 nm.

As the pump current was increased, the shape of the output spectrum evolved. Using A3-L1(2)-1 as a representative example, the peak wavelength would tend to increase with pump current, as seen in Figure 5.20. Additionally, mode hopping was observed with increasing current, resulting in sudden changes to the shape of the output spectrum or the peak wavelength. These changes should be obvious from Figure 5.20 as they were significant. Again, A3-L1(2) is being used as an example for this behaviour, as it was exhibited in every laser device tested including B6.

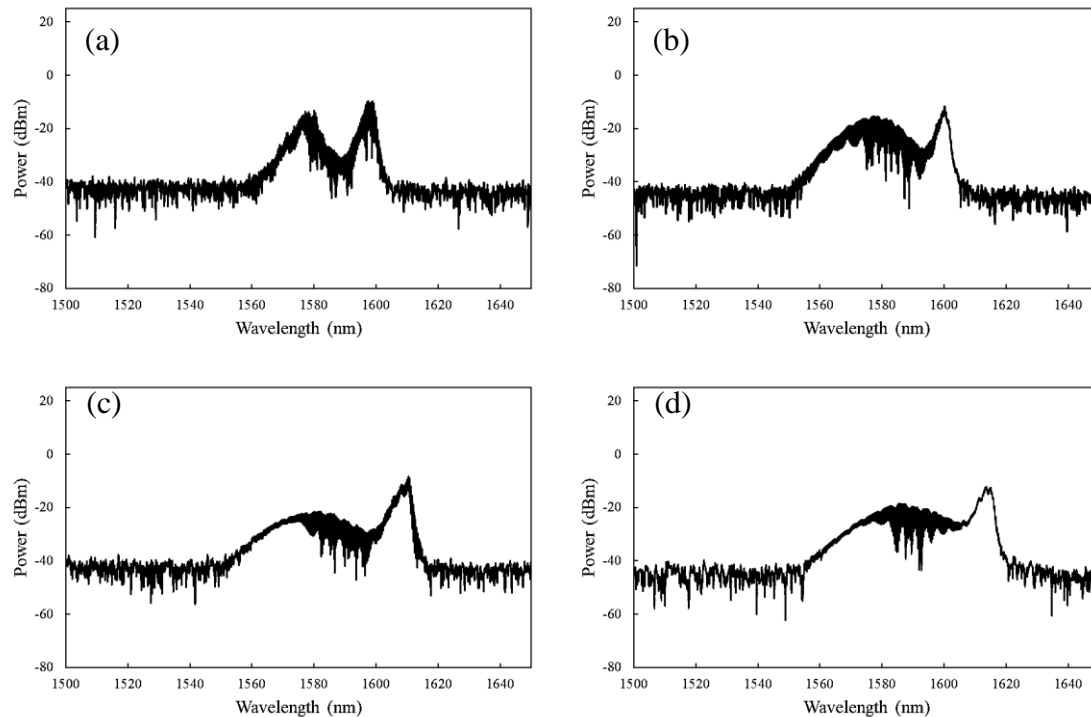


Figure 5.20: Evolution of the A3-L1(2)-1 spectrum with increasing pump current of (a) 300 mA, (b) 320 mA, (c) 420 mA, and (d) 500 mA. Note the changes in shape, especially from Figure 5.17, and the red shift of the peak wavelength.

Among the B6 devices there were variations when it came to the shapes of the spectra, but all devices shared the trends exhibited by A3: increasing peak wavelength with pump current and mode hopping which would change the shape of the spectrum. Figure 5.21 displays a device of each width (5 μm , 4 μm , 3 μm and 2 μm) at 200 mA pump current. No trends are obvious based on device width here, aside from a similar “shark fin” shape to

the lowest wavelength peak. That aside, the distributions of peak wavelengths and breadth of each peak are independent from one another. Such variance is present for the spectra across all B6 devices. This could be caused by errors in manufacturing the grown material, such as uneven growth of the active layers. The different peak wavelengths could also be related to the mode-carrying abilities of each waveguide, as the most confining device (2 μm) exhibits the lowest peak wavelength at 1570 nm instead of 1590 nm or even 1600 nm. The variance in the B6 devices makes it difficult to definitively say if there is a difference between the peak wavelengths of Growth A and Growth B material, as the A3 devices measured within the range of wavelengths produced by B6.

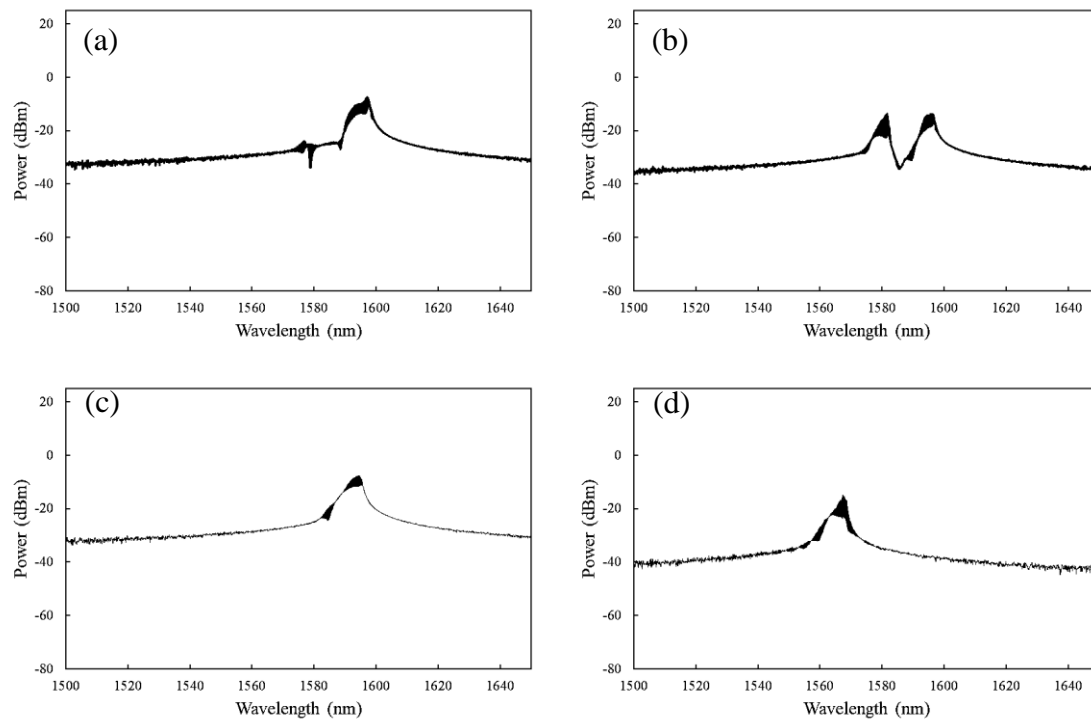


Figure 5.21: Comparison of B6-L-18 devices at 200 mA pump current, with widths of (a) 5 μm , (b) 4 μm , (c) 3 μm , and (d) 2 μm .

Pre-threshold behaviour in B6 devices was identical to A3 devices in that the spectrum showed the hallmarks of an ASE spectrum. Figure 5.22 displays the pre-threshold spectrum for B6-L-18 at 50 mA of pump current after coating, and the response when a 1563 nm laser is added to the input of the coated device. The output of B6-L-18 was not attenuated

by the AR coating. With optical input, a degree of amplification around the input signal can be seen which encouraged further testing at low pump current with a wider range of wavelengths.




Figure 5.22: Spectra of B6-L-18 pumped at 50 mA (a) after AR coating, (b) with a 1563 nm light source at the input.

5.4.1 Spectra After AR Film Deposition

After deposition, the devices were measured with the OSA to determine if lasing had been suppressed, and that the spectrum resembled something closer to an ASE curve. The coating of B6 slightly changed the spectra of most of the devices, causing some minor shifting of peak wavelength and spectrum shape. However, the most significant effect was on the 2 μm waveguides. Figure 5.23 shows the change effected on B6-L-12, which is representative of the two other 2 μm devices that showed attenuated output power (B6-L-04 and B6-L-08). The lasing activity in the device was not very strong to begin with, but

with the coating the device is unable to lase at all. Instead, the device teeters on the cusp of lasing, with its power distributed across a wide range of wavelengths indicative of ASE.




Figure 5.23: Spectrum of B6-L-12 at 200 mA before and after AR coating. Note the elimination of lasing behaviour in the coated spectrum.

5.5 Band Edge Absorption

As part of the electro-absorption modulator (EAM) project, B6-M-33 was put into varying levels of reverse bias and checked for absorption over a range of wavelengths, as seen in Figure 5.24. This provided insight into the wavelength at which absorption begins to occur in Growth B devices, indicating their approximate band gap. In this case, the onset appears around 1580 nm at 0 V bias, where power begins to transmit through the waveguide instead of being fully absorbed by it. Further reverse bias serves only to increase the wavelength at which absorption occurs.




Figure 5.24: Absorption measurements in reverse bias of B6-M-33 as part of the EAM project. Note the onset of absorption around 1580 nm for 0 V and 1 V bias.

5.6 Amplification

The optical gain of the SOA can be calculated by the following equation:

$$P_{in} - |Loss_{in}| - Sys_{in} + Gain = P_{out} + |Loss_{out}| + Sys_{out}$$

And, rearranging,

$$Gain = (P_{out} - P_{in}) + |Loss_{in}| + |Loss_{out}| + Sys_{in} + Sys_{out}$$

Where P_{in} is the power measured directly from the laser source in dBm, P_{out} is the output power measured at the OPM in dBm, $Loss_{in}$ is the input facet coupling loss in dB, $Loss_{out}$ is the output facet coupling loss in dB, Sys_{in} is the loss from the laser source to the tip of the input fiber in dB, and Sys_{out} is the loss from the output fiber ferrule to the OPM in dB. The output power is further corrected by subtracting the power of the ASE of the amplifier in mW as background noise. The gain is then the internal gain of the supposed SOA device in dB. The facet coupling losses differed between individually measured devices, likely

due to varying facet quality. Additionally, the laser source power had to be monitored, as it was limited in power depending on the operating wavelength; at 1510 nm or 1640 nm the maximum achievable power was -4.5 dBm, while at any wavelength the maximum power was 4.6 dBm. This input power to wavelength relation can be seen in Figure 5.25. As the two wavelengths of interest were 1560 nm and 1580 nm, 0 dBm was selected as the output power of the Agilent to maintain a well-known power level.




Figure 5.25: Relation of laser power to laser wavelength from the Agilent 8164A output port. Series labels refer to the input power as set on the Agilent.

The modified output power was measured for a range of input powers from -40 dBm to 5 dBm, before accounting for line or coupling losses. After accounting for losses, the input power range was closer to -53 dBm to -3 dBm. Devices with successful coatings as seen through OSA measurements were tested, as well as devices that had not seen lasing suppression after coating. The latter devices were tested at their pre-threshold current conditions, as with Figure 5.22. B6-L-17 was one such device, and the output power vs. wavelength is shown in Figure 5.26. The output power rises with input power and wavelength, showing a maximum around 1595 nm, although exact quantification is made difficult due to an underlying sawtooth-shaped ripple in the output power. This is a result

of back reflections in the system, caused by an imperfect AR coating on the facets. The same measurement was taken for devices B6-L-04 and B6-L-08 at 180 mA, the results of which are shown in Figure 5.27 and Figure 5.28, respectively. The peak output for B6-L-04 appears near 1580 nm, and the peak for B6-L-08 appears near 1585 nm. These peak wavelengths also correspond to the highest power on an output spectrum measured directly from the device.




Figure 5.26: Output power and input power vs. wavelength for B6-L-17 at 50 mA. The legend refers to the power setpoint, not the real input power. Note the apparent peak near 1595 nm.

The pump current levels for B6-L-04 and B6-L-08 were informed by the peak power on the LI curve of each device after coating. The peak power indicated the current with the most injected carrier activity in the cavity before thermal recombination effects began to quench excited carriers faster than they could be produced.




Figure 5.27: Output power vs. wavelength for B6-L-04 at 180 mA. Note the apparent peak near 1580 nm.

From the peak wavelength results of the wavelength sweeps and band edge absorption measurements, it was determined that sweeping at 1580 nm would yield the highest gain from the coated devices. However, due to the wavelength constraints of the equipment available, 1580 nm could only be swept over a limited range of power and could not be optically filtered. A filter was required to reduce background power introduced by the ASE of the amplifier. Instead, the output power was measured as usual, but the spectrum at each level of attenuation was also measured into the OSA. This allowed the signal to be digitally filtered and was used to analyze data for B6-L-08 at 200 mA. To achieve a wider range of input power and to optically filter the output signal, B6-L-04 was measured across a range of pump currents with an input wavelength of 1560 nm, with the knowledge that this would not result in the highest achievable gain.

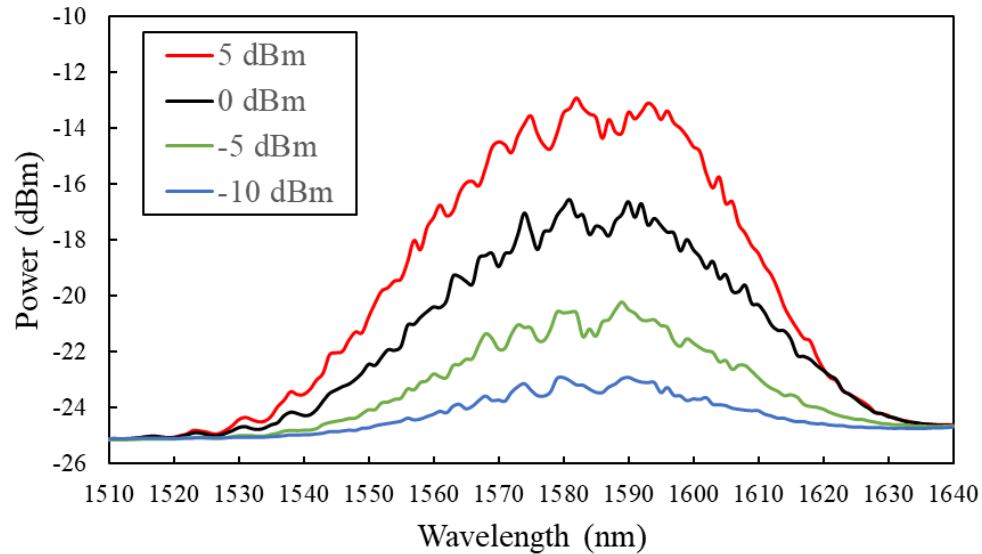


Figure 5.28: Output power vs. wavelength for B6-L-08 at 180 mA. Note the apparent peak near 1585 nm.

The gain of B6-L-08 at 1580 nm input and pump current of 200 mA is plotted against input power in Figure 5.29. The results without any filtering made no sense due to the presence of background noise in the output signal from the ASE of the amplifier. The output spectrum was collected with the OSA, then each individual power measurement at each wavelength point was weighted as a percentage of the total measured power. The filter width could then be set by eliminating all data outside of a specified wavelength range. The resolution of this technique was limited by the spacing of wavelength measurements from the OSA, and resulted in a filter with a 100% extinction ratio at the filter edges. Multiple filter widths were tested, as seen in Figure 5.29. As expected, expanding the filter increases the output power while input power remains constant, resulting in a gain that rises with increasing filter width. The maximum gain achieved by B6-L-08 is 0 dB, ignoring the low-power tail. Gain saturation is achieved at an input power of -20 dBm.




Figure 5.29: Gain vs. input power of B6-L-08 with 1580 nm input and pumped at 200 mA. Each curve represents a different digital filter width around the 1580 nm peak.

The gain of B6-L-08 at 1580 nm input and pump current of 200 mA is plotted against output power in Figure 5.30. Increasing the filter width once again results in higher gain values, as well as causing gain saturation to occur at higher levels of output power. The sharp increase of gain at lower output power is attributed to the output power not changing with increasing input power. At lower input power levels, the signal power peak barely rises above the ASE power, and has additional peaks which dip below the ASE level due partially to back-reflections in the system. The summation of these high and low peaks leads to an overall signal which measures close to the ASE power after digitally filtering. Figure 5.31 shows the spectral peaks of B6-L-08 to contrast the amplified power at varying levels of input power. The apparently flat baseline of the signal is the ASE of the amplifier.




Figure 5.30: Gain vs. output power of B6-L-08 with 1580 nm input and pumped at 200 mA.




Figure 5.31: OSA spectrum of B6-L-08 with various levels of input power at 1580 nm. Note the oscillations in the signal from back-reflections, and magnitude of the smallest amplified peak.

B6-L-04 was measured with an input wavelength of 1560 nm to make use of the erbium-doped fiber amplifier (EDFA) and optical filter available. The EDFA boosted the initial power out of the Agilent to 15 dBm, allowing for a wider range of input powers to be tested. The filter was set to a bandwidth of 4 nm around 1560 nm (1558 to 1562 nm), and can be seen in action in Figure 5.32. The extinction ratio of the filter can be calculated from Figure 5.32 as 11.1 dB/nm; for simplicity, the extinction ratio of the filter was deemed to be 100% at the edges of the target filtered range. The addition of the filter was intended to remove the need to digitally filter the data from OSA measurements.




Figure 5.32: Optical filter inserted into the line with a tuned bandwidth of 4 nm, 1558 to 1562 nm.

The gain of B6-L-04 at 1560 nm input and various pump currents is plotted against input power in Figure 5.33. Gain saturation is achieved between -5 and 0 dBm of input power, depending on the pump current level. The gain trends with pump current mostly in agreement with the LI curve of B6-L-04, shown in Figure 5.34. When the device is at its peak ASE output power the gain is also highest, in this case at 160 mA. It is noted that there are sudden drops in output power as input power is decreased, leading to the jagged appearance of the gain curve. Without the power drop it is reasonable to assume that the gain would increase to higher levels for 160 mA and 200 mA. The waveguide of B6-L-04 was 2 μm wide, meaning that mode hopping is an unlikely, but possible, explanation for

this behaviour, especially as it only presents itself for pump current levels which produce significant ASE power as seen from Figure 5.34.




Figure 5.33: Gain vs. input power of B6-L-04 with 1560 nm input and pumped at various current levels.




Figure 5.34: LI curves for the uncoated and coated B6-L-04 device. Note the maximum of the coated LI curve at 160 mA.

The gain of B6-L-04 at 1560 nm input and various pump currents is plotted against output power in Figure 5.35. The same low-power tail behaviour is observed here as in Figure 5.33. Gain saturation occurs between -10 and -5 dBm of output power, depending on the pump current. The same trend of gain sharply dropping off as power is decreased also presents itself here. By checking the output power against the input power in Figure 5.36, the perturbations in the linear trend are easily visible, indicating at which input power levels each current sees the gain sharply drop off. In this case, the drop appears corrective, as it occurs once the output power has deviated from the linear trend, bringing it back to the expected value. B6-L-04 also demonstrated a sharper drop-off in gain at saturation than B6-L-08.




Figure 5.35: Gain vs. output power of B6-L-04 with 1560 nm input at various current levels.




Figure 5.36: Output power vs. input power of B6-L-04 with 1560 nm input at various current levels. Note how the data deviates from the linear trend, then self-corrects.

The increasing gain at the low-power tails is clearly not real, as having the gain trend to infinity at low input power would indicate a system creating energy from nothing. This observed behaviour is most likely the result of underestimating the loss at the output of the system. Figure 5.37 shows the gain of B6-L-08 against input power with increased output loss. Along with smoothing out the low-power tails, this loss correction also increased the internal gain of the device. While the loss was meticulously characterized before each measurement, there were some factors that affected loss during measurement such as the device slipping on the thermal paste due to the force of the probes. This noticeably caused changes in the coupling point after optimization.




Figure 5.37: Gain vs. input power of B6-L-08 with 1580 nm input at 200 mA. Filter set to 1 nm width, and loss increased by 5 dB over Figure 5.29.

The coupling loss was also noted to change with current, as seen in Figure 5.38 for B6-L-04. The loss decreases as the ASE power of the device increases as seen in Figure 5.34.




Figure 5.38: Facet coupling loss vs. pump current for B6-L-04. Values were derived by subtracting measurements of the ASE through the fiber from measurements directly into the integrating sphere.

Finally, it is noted that device performance degraded over time, and the more they were tested. This often resulted in device failure to an open circuit condition. Suspected causes of this are stress on the devices from the electrical probes, inadequate heat sinking, and prolonged operation at higher voltage levels.

6 Conclusion

The objective of this project was to fabricate a semiconductor optical amplifier (SOA) from III-V material, using the facilities of the McMaster Centre for Emerging Device Technologies (CEDT). The SOA would be based on the design of a multiple quantum well (MQW) laser diode, modified with an anti-reflective (AR) coating applied to the device's waveguide facets. The amplifier was intended to introduce a positive gain into the system to offset optical losses from processes such as modulation of the input signal.

The researcher used an older process manual [25] as a guideline for creating the laser diodes. Over the course of the project, this process was refined and tailored to suit the present capabilities of the CEDT. Multiple avenues of optimization were explored, such as cutting device bars apart with a dicing saw and using a focused ion beam (FIB) to polish the facets. As a result, an updated processing document for fabricating III-V SOA devices was composed, and can be found in Appendix A: III-V Laser Diode/SOA Fabrication Guide.

As a result of the new process, InGaAsP MQW laser diodes were produced with maximum measured optical power in excess of 20 mW, with a peak wavelength around 1580 to 1600 nm. The operating point of such devices was up to 500 mA, with a laser threshold current below 50 mA. These results definitively show that the updated process can successfully produce high quality laser diodes.

After coating the device waveguide facets with a four-layer stack of silicon oxide/nitride thin films, characterization measurements were performed to determine if the AR coating had suppressed lasing behaviour. Optical power vs. current (LI) curves were measured before and after coating. For devices with waveguide widths of 2 μm , the coating induced successful suppression of lasing for 57% of devices. Changes were seen in the LI curves for other widths of waveguides, but nothing suggested that the coating had effectively

suppressed the output power at those widths, even though maximum power was sometimes seen to drop.

The optical spectrum was measured before and after depositing the AR coating. The devices which saw suppressed lasing also experienced changes to their spectrum, in that there was no longer a distinct peak output power. Instead, optical power was distributed over a wide range of wavelengths indicative of amplified spontaneous emission (ASE), with no laser peak emerging at any pump current.

Finally, the gain curve of the coated device was measured against the input and output optical power of the system. The shape of the gain curve resembled those of commercial SOA's, such as one from Thorlabs. The maximum small-signal gain was measured between 0 and 1 dB. However, it was noted that some loss at the output of the system may not have been accounted for, due to the shape of the gain curves. For example, the thermal paste applied under the chip occasionally caused the device to slip under pressure from the electrical probes, upsetting the optimal coupling point. After applying a 5 dB increase to the output power which made the shape of the curve more reasonable, the maximum small-signal gain was found to be 4 dB.

With evidence that the AR coating suppressed lasing behaviour, as well as converting the output spectrum to an ASE spectrum and displaying a gain curve characteristic of a commercial device, it can be stated with confidence that a working SOA was fabricated.

6.1 Future work

The fabricated SOA was not optimal, and further steps can be taken by a future researcher to improve upon its design. Some changes are to the fabrication process while others involve the way in which the devices are measured.

One change would be to grow a material with an emission wavelength closer to 1550 nm. This would allow the SOA to operate in the C-band for optical communications, as well as allowing for interaction with more available equipment such as erbium-doped fiber amplifiers (EDFA) and optical filters. Producing a new growth at this wavelength would better align with the purposes of this project and allow for more accurate characterization measurements.

Another major change would require the future researcher to devise a method of depositing the AR film directly normal to the waveguide facets. This would significantly improve the accuracy of the thickness of each film layer on the facet, leading to better control over the reflectivity. Additionally, the approach described in this thesis of depositing on the top of the chip could be augmented by testing the layer composition of the film stack after deposition. This could help to calculate appropriate deposition rates for such a method.

One comparative study would check the difference in performance of cleaved facets against those polished with a FIB. Additionally, the FIB could be used to cut the facets at an angle, which could significantly reduce the reflectivity of the facet. Device cavity lengths could also be compared to determine the point at which lengthening the cavity becomes detrimental to amplification.

To improve measurements, a vacuum stage could be used to hold the device instead of the current setup with thermal paste on a smooth copper heat sink. The vacuum stage would prevent devices from slipping under pressure from the electrical probes, maintaining the position of the optimal fiber coupling point during measurement cycles. An on-chip thermistor could also be employed to prevent damage and improve performance by providing a more accurate reading of the device temperature.

The distributed Bragg reflector (DBR) gratings side project could also be revisited. Further simulation work may show that the gratings do not need to be cut as deep into the

waveguide as initially thought; if this is the case, the plasma FIB (PFIB) would be effective in creating the desired DBR waveguide structures.

6.2 Skills acquired

Several skills were acquired by the researcher over the course of this project, mostly relating to semiconductor fabrication techniques and equipment. The researcher learned how to use an alpha step profilometer, plasma-enhanced chemical vapour deposition (PECVD), photolithography (PL) equipment, contact PL mask design, wet etching techniques for III-V materials, e-beam evaporation of metals, operating a dicing saw, scanning electron microscopy (SEM), FIB operation, mechanical lapping of III-V material using both manual and automatic approaches, and measurement techniques for characterizing laser diode and SOA devices.

7 References

- [1] Cisco, “Cisco Global Cloud Index: Forecast and Methodology, 2016–2021,” *Cisco*, 2018.
- [2] Bell Canada, “Bell Gigabit Fibe bringing the fastest Internet to Toronto residents with a billion-dollar+ network investment, creation of 2,400 direct jobs,” *BCE*, June 25 2015. <https://www.bce.ca/news-and-media/releases/show/bell-gigabit-fibe-bringing-the-fastest-internet-to-toronto-residents-with-a-billion-dollar-network-investment-creation-of-2-400-direct-jobs-1>.
- [3] M. Isberto, “What is the environmental impact of a data center?” *Colocation America*, May 3 2018. <https://www.colocationamerica.com/blog/data-center-environmental-impacts>.
- [4] I. Burrington, “The Environmental Toll of a Netflix Binge,” *The Atlantic*, December 16 2015. <https://www.theatlantic.com/technology/archive/2015/12/there-are-no-clean-clouds/420744/>.
- [5] W. Bogaerts et al., “Silicon microring resonators,” *Laser & Photonics Review*, **6**, No. 1, 2012, pp. 47–73.
- [6] “Comparison of Different Types of Optical Amplifiers”. *Tutorials of Fiber Optic Products*. Accessed January 2020. <http://www.fiber-optic-tutorial.com/comparison-of-different-optical-amplifiers.html>.
- [7] N.M. Ibrahim and A. Abubaker, “A Comparison of Optical Amplifiers in Optical Communication Systems EDFA, SOA and Raman,” *International Journal of Current Research*, Vol. 6, Iss. 09, September 2014, pp. 8738-8741.
- [8] D. Griffiths, *Introduction to Electrodynamics*, Fourth Edition, Pearson Education, 2013, pp. 210.
- [9] K.M. Gupta and N. Gupta, “P-N Junctions and Their Breakdown Mechanisms”, *Advanced Semiconducting Materials and Devices*, 2016, pp. 193-234.
- [10] A. Kitai, *Principles of Solar Cells, LEDs and Diodes*, John Wiley & Sons, 2011, pp. 75.

- [11] A. Kitai, *Principles of Solar Cells, LEDs and Diodes*, John Wiley & Sons, 2011, pp. 195.
- [12] R.G. Hunsperger, *Integrated Optics: Theory and Technology*, Sixth Edition, Springer Media, 2009, pp. 228.
- [13] R.G. Hunsperger, *Integrated Optics: Theory and Technology*, Sixth Edition, Springer Media, 2009, pp. 227.
- [14] R.G. Hunsperger, *Integrated Optics: Theory and Technology*, Sixth Edition, Springer Media, 2009, pp. 230.
- [15] R.G. Hunsperger, *Integrated Optics: Theory and Technology*, Sixth Edition, Springer Media, 2009, pp. 233.
- [16] P.W. Milonni, J.H. Eberly, *Lasers*, John Wiley & Sons, 1988, pp. 1.
- [17] R.G. Hunsperger, *Integrated Optics: Theory and Technology*, Sixth Edition, Springer Media, 2009, pp. 377.
- [18] W.T. Tsang, “Extremely low threshold (AlGa)As modified multiquantum well heterostructure lasers grown by molecular-beam epitaxy,” *Applied Physics Letters*, **39**, 1981, pp. 786.
- [19] C. Weisbuch and J. Nagle, “The Physics of the Quantum Well Laser,” *Physica Scripta*, Vol. T19A, January 1987, pp. 209-214.
- [20] “SOA1117S – C-Band Semiconductor Optical Amplifier, CWL=1550 nm (Typ.), Butterfly Pkg, SMF, FC/APC Spec Sheet”, *Thorlabs*. Accessed January 2020.
<https://www.thorlabs.com/thorproduct.cfm?partnumber=SOA1117S>
- [21] R.G. Hunsperger, *Integrated Optics: Theory and Technology*, Sixth Edition, Springer Media, 2009, pp. 270.
- [22] J.E. Spencer and P. Young, “Contrasting the Photodigm DBR Laser Diode Architecture with Competing DFB Designs,” *Photodigm Inc.* Accessed January 2020.
<http://www.photodigm.com/difference-between-dbr-and-dfb-lasers>
- [23] “S145C – Integrating Sphere Photodiode Power Sensor, InGaAs, 800 – 1700 nm, 3 W Spec Sheet,” *Thorlabs*. Accessed January 2020.
<https://www.thorlabs.com/thorproduct.cfm?partnumber=S145C>

- [24] R.G. Hunsperger, *Integrated Optics: Theory and Technology*, Sixth Edition, Springer Media, 2009, pp. 385.
- [25] J. Czaban, “Fabrication of a Laser Diode for Fiber Optic Telecommunications,” *CPFR Photonics Fabrication: Graduate Short Courses*, August 2007.
- [26] “Reflectance Calculator,” *Filmetrics*. Accessed December 2019.
<https://www.filmetrics.com/reflectance-calculator>
- [27] J. Pankove, “Temperature dependence of emission efficiency and lasing threshold in laser diodes,” *IEEE Journal of Quantum Electronics*, Vol. 4, Iss. 4, April 1968, pp. 119-122.

Appendix A: III-V Laser Diode/SOA Fabrication Guide

Introduction

There are a few orders of business to attend to before launching into the actual device processing. A few tips and tricks, and an outline of the traps. Never be afraid to seek help! Even though almost nobody physically handles III-V material at Mac anymore, something they say might help you out.

Tips

It should go without saying, and it will go without saying once you break your first wafer, that III-V material is incredibly brittle. In fact, it will remain brittle until, funny enough, you thin it down and anneal the metal contacts into the device. I dropped a whole case of annealed devices off of a desk and the only one that exploded was the one I hadn't annealed...I've had heart palpitations ever since. So, rule number one when handling your devices is don't drop them. Rule number two is DON'T DROP THEM. Rule number three is don't talk about Beer Club. Again, you will break some of this material and you will learn, so don't freak out when it happens the first time!

From my personal experience, other researchers who are used to handling silicon chips are ill-equipped to handle your precious. Silicon can take much more abuse than III-V, so they're used to being able to apply just that little bit more force that makes the difference between one laser diode and fifteen laser diodes. If you require a tool that a silicon user is the expert on, the best I can suggest is that you politely ask to tag along, and guard your sample like a mama grizzly. Request to handle your sample personally, and clear your schedule to be able to stick around if there's any potential need to handle the sample.

Another personal preference of mine was to use regular metal tweezers for III-V material instead of Teflon-tipped tools. A lot of people swear by the Teflon as it makes it more difficult to scratch chips, and the springiness of the tweezers is a bit different. That being said, I always found it more difficult to feel the chip through the Teflon, and so had more control with the bare metal tweezers. The 2A flatheaded tweezers from SPI do a very nice job of this, and were all-stars throughout my degree.

Make sure to verify your etch and growth rates on test substrates before trying it on your real material, and re-verify these rates whenever any significant changes are made to your fabrication equipment. Such things could be cleaning the CVD growth chamber or re-tuning the mask aligner. Spending the little bit of time to verify your process will save you massive headaches further down the line! Also make sure to keep in contact with the CEDT technicians so that you know when these maintenance operations happen.

To verify acid etch rates, perform your etch on a test film or substrate (there are plenty of trash samples in TAB if you ask nicely), then use an alpha stepper to ensure that the proper

etch depth is achieved. You can also slap that baby in an SEM to get a real good look, but that's probably overkill; you do you though.

To verify growth rates, you'll be using the variable angle spectroscopic ellipsometer (VASE) to check film thickness and index. You can also use a thin film colour guide: for example, a 100 nm film of SiO_x should appear royal blue when properly grown.

Pitfalls

I tried two different methods of separating laser bars: dicing and cleaving. The thinning process required for cleaving is quite long, and requires a certain amount of practice and motor skill to get right; as such, I attempted to dice first, before eventually switching to cleaving. I suggest you put the time into learning the thinning and cleaving approach, as the dicing will only give you headaches. I have two masks designed, one for dicing and one for cleaving (refer to my thesis for the other differences between the two). The dicing mask is labelled as "III-V_laser_mask_submission", and the cleaving mask is labelled as "III-V_laser_cleave_w_gratings".

From my experience with Dr. Xu's dicing saw in TAB, even equipped with III-V blades from Disco the process is too traumatic for the III-V material. It could have to do with this particular tool, as III-V should be able to be diced with the proper equipment. In any case, with this tool the chip fractures while dicing, sending pieces flying; if a better dicing saw is acquired (or if you're able to secure time on the ETB saw), it may be possible to revisit the dicing process. The dicing process can be learned from a current operator, ensuring that you use the Disco blades of model ZHZZ-SD4800-H1-50-A1726 or model ZHZZ-SD4500-H1-50-A1726 which should be floating around the Knights lab somewhere.

In addition to the trauma inflicted on the chip (and on you), the facets cannot be used directly off the saw, but must be FIB'd to smooth them out. The FIB process was performed with Dr. Luxat's tool in TAB under the eye of Dr. Zhilin Peng; you will need radiation safety training to access it due to its location. My personal process was to make an initial cut across the facet at 3 nA beam current to remove the bulk of unwanted material. If the facet didn't seem to be clean enough, I would polish at 1 nA, then 100 pA if still necessary. The facet does not have to be perfect, but you should be looking for reduced curtaining and no debris or large defects in the surface or waveguide. I also found it helpful to only do the FIB once I had bonded my sample to a copper backing plate already (see the Bonding section of this document). This provided stability to the sample and saved me a lot of stress. Just make sure that the copper plate is as flat as possible!

In an effort to simplify the thinning process, I went to the Materials department and got permission to use one of their auto-polishers, which are typically used with SiC sheets to polish steel samples. There is certainly potential in this technique, especially as the system allows for up to 6 samples to be thinned simultaneously; however, I would not suggest trying it outside of a fume hood or without a proper fume extraction setup in place.

Now with all that out of the way, onto the part you probably skipped over all of this for: the fabrication procedure I used!

Fabrication Procedure

Take a quarter of a 2” wafer or a comparably sized piece of a 3” wafer as your sample (at least 1.5 cm x 1.5cm). Make sure you know which side is the top side and which direction the major flat lies. Your waveguides must be aligned perpendicular to the major flat for vertical sidewalls, as the InP etch is anisotropic. Aligning parallel to the major flat will result in sloped sidewalls, leading to optical losses from the waveguide. I also suggest making the sample as square as possible, as it will make spincoating easier – just ensure that you remember where the major flat is! This procedure will also assume that you are using the material procured from Landmark Optoelectronics, the structure of which is as follows:

Composition	Doping (cm ⁻³)	Purpose	Layer Thickness
InP	Undoped	Capping	10 nm
In _{0.53} Ga _{0.47} As	P ⁺ = 1x10 ¹⁹	Contact Layer	200 nm
InP	p = 5x10 ¹⁷	Ridge	1.5 μm
In _{0.7322} Ga _{0.2678} As _{0.581} P _{0.419}	Undoped	Confinement	105 nm
In _{0.53} Ga _{0.47} As	Undoped	Quantum Well	8 nm
In _{0.7322} Ga _{0.2678} As _{0.581} P _{0.419}	Undoped	Barrier	10 nm
In _{0.53} Ga _{0.47} As	Undoped	Quantum Well	8 nm
In _{0.7322} Ga _{0.2678} As _{0.581} P _{0.419}	Undoped	Barrier	10 nm
In _{0.53} Ga _{0.47} As	Undoped	Quantum Well	8 nm
In _{0.7322} Ga _{0.2678} As _{0.581} P _{0.419}	Undoped	Barrier	10 nm
In _{0.53} Ga _{0.47} As	Undoped	Quantum Well	8 nm
In _{0.7322} Ga _{0.2678} As _{0.581} P _{0.419}	Undoped	Barrier	10 nm
In _{0.53} Ga _{0.47} As	Undoped	Quantum Well	8 nm
In _{0.7322} Ga _{0.2678} As _{0.581} P _{0.419}	Undoped	Barrier	10 nm
In _{0.53} Ga _{0.47} As	Undoped	Quantum Well	8 nm
In _{0.7322} Ga _{0.2678} As _{0.581} P _{0.419}	Undoped	Confinement	105 nm
InP	n = 5x10 ¹⁷	Buffer	0.5 μm
n ⁺ -InP	n ⁺ = 2x10 ¹⁸	Substrate	500 μm

Remove any dust or debris from the quarter using a nitrogen gun.

TOOL: UV Ozone

The purpose of this step is to clean off any organic substances which may be stuck to the surface of the sample. **Do not use a piranha etch for this step, or any other method involving sulphuric acid.** Sulphuric acid is an etchant of InGaAs, which is one of your device layers and especially the one which facilitates a better ohmic contact with the waveguide.

1. Open the door of the UV Ozone tool and check that there are no samples on the platform. If empty, close it up, set the timer on it for **5 minutes**, then turn the power on. This is a pre-heating stage.
2. Once the 5 minutes have elapsed, turn the power off and let the system cool for another **5 minutes**.
3. Load your sample in the chamber for **20 minutes**, maybe take some notes or check your phone.
4. Turn the power off and let the chamber cool for another **5 minutes**. In the meantime, make sure the DI water is ready to go.
5. Transfer your sample from the UV Ozone to a running DI rinse in the laminar flow bench. Let the sample sit under the rinse for **5 minutes**.
6. Remove the sample and blow it dry with the nitrogen gun. Transfer to a case to prepare for CVD.

Due to the structure of the material, we must remove the InP capping layer. This will be done via a **flash etch** (fancy way of saying dip it in acid and pull it out quickly).

1. **Create a solution of 1:3 HCl:H₃PO₄ in a 250 mL beaker.** Measure the HCl first in a 25 mL graduated cylinder through a glass funnel, then multiply that quantity by 3 to get your H₃PO₄ volume. The H₃PO₄ is very viscous, making it easy to precisely pour.
2. Place the sample in a wafer basket, and submerge in the acid solution for **3 seconds**, pulling it out as the timer expires and dunking into a large beaker of DI water to quench the reaction.
3. Transfer the sample to a running DI rinse for 5 minutes.
4. Dry the sample with a nitrogen gun. **Make sure that you blow as parallel to the surface of your chip as possible!**

Next up is the deposition of the mask oxide. This is a ~100nm layer of SiO_x used as a hard mask to etch the mesas and ridge waveguides out of the InP. The thickness and quality of this oxide isn't incredibly important, as it will be removed and replaced later on. However, if the oxide is too thick then the isotropic SiO_x etch will erase some of your finer features, such as 2 um waveguides, leading to sadness. For this step, I used the Technics Silicon Oxynitride CVD in the clean room, although any CVD capable of SiO_x deposition will do.

TOOL: Clean Room Technics PECVD

Follow the instructions paired with the tool. A few things to note:

1. Before loading the sample, make sure to use the nitrogen gun to remove any detritus in the chamber. As a precaution, I would slowly close the lid on my sample, wait a second, then lift the lid and check for debris again, to make sure nothing was flaking off the lid (can be an issue if the chamber is not cleaned often enough).
2. Once the sample is loaded into the chamber and the lid is closed, set the pressure controller switch left to “Closed”. Toggle the VAC switch (about 1 Hz) until you hear a click for the higher-speed vacuum engaging. Doing this will mitigate the movement of your sample in the chamber due to the sudden suction in the chamber.
3. Wait until the chamber pressure is as low as you think it will go before turning on the stage heater. The heater will affect the chamber evacuation rate slightly, and from my experience it seemed to get the chamber a few mTorr lower by doing this.
4. **Operational parameters are: 90 sccm SiH₄ / 70 sccm N₂O / 50 W / 650 mTorr / 300 C / 3 minutes**
5. Wait until the chamber gets to about 297 C before turning on the gas flows. The gas flow will drop the chamber temperature a little bit, so no point in wasting process gases. If you want to wait longer to get to 300 C then it’s your call, but the last few degrees always take a long time to achieve, and you’ll be waiting 15+ minutes by this point already.
6. Start your 3 minute timer as soon as you see the plasma ignite, and **quickly** crank the power up to 50 W.

Follow the shutdown instructions paired with the tool, then come back in 3-4 hours once the chamber has cooled off to retrieve your sample. I preferred having a pair of wafer tweezers and a pair of Teflon tweezers to remove the sample, as it’s difficult to get underneath it with only one tool. Look for a royal blue colour indicating a 100 nm film, then take your sample to the VASE to check it out and validate your film. Don’t forget to come back and write down what you grew in the log book!

TOOL: Ellipsometer (VASE)

The ellipsometer use is as follows, in case you have difficulty tracking down someone to show you the ropes; keep in mind that this is a very rudimentary method of operating the system, and any advanced analysis using the VASE should be informed by somebody other than me.

1. Remove the plastic dropsheet over the VASE. This is there to protect the system from excessive dust.
2. Turn on the main power (green button on right of top module)
3. Turn on lamp power (green button at bottom of bottom module)
4. Fire lamp (green button at top of bottom module)
5. Place your sample on the stage, with your area of interest approximately aligned with the horizontal groove. This is the plane on which the VASE will sweep to measure film

- thickness. For example, the grown film is typically a bit different across the surface of the sample, so you want to try to line up the middle of the sample with this plane.
6. Log into QReserve on the computer, then open CompleteEase if it isn't already open. Log into CompleteEase as user Jeremy Miller and hit enter (there is no password for this account)
 7. Wait for the system to finish initial calibrations, then set your measurement type as dielectric film on si and click "Measure"
 8. Name your file, etc.
 9. You will be asked to align the system. Simply adjust the tuning knobs on the sides of the stage until the red cross lines up with the centre of the screen and hit Enter or "Cancel Alignment". Your measurement will begin!
 10. Once the measurement is finished, switch to the Analysis tab
 11. Change the film type to a Cauchy film if it isn't already, and change your substrate material from Si-JAW to the proper material (unless you've grown this on silicon you silly III-V researcher you). The correct substrate material property file will normally be found under Recent, but if not then hunt around the directories for a bit and you should stumble on it eventually.
 12. When you have changed the parameters to properly match your device, click "Fit" on the left side of the screen and take a look at the film thickness and index that pop up. If the error in either is significant (you judge), then there may be something in the path of the measurement, and all you have to do is shift your sample a little bit and re-measure. Overall, use your common sense: if the numbers coming up for film thickness and index don't make sense, then keep searching for potential errors in the measurement before assuming the fabrication tool is broken. If you're truly stuck, never be afraid to seek help!
 13. Once you've finished all the measurements, the file automatically saves.
 14. Click on the Hardware tab, then Angle. Change the angle to 45 degrees to move the source and detector arms up.
 15. Log out of CompleteEase, log out of QReserve
 16. Shut off the lamp power first (green button on the bottom of the bottom module)
 17. Finally, shut off the main power (green button on the top module)

Next up, we'll be patterning the MESA mask onto the sample. This will allow us to etch out the ridges for the waveguides. There are two wet etches involved in this step, so set aside a good chunk of time for this.

TOOL: Spinner

The spinner is used to distribute primer and photoresist evenly across the surface of your chip. The apparatus pulls a vacuum through the center spindle to hold the sample in place; once a vacuum is confirmed here, the spindle will get up to speed. For our purposes, we'll be using **MPHP Primer and S1808 Photoresist** with the spinner set to **4000 RPM for 30s**. Before starting, it is important to check the hydrophilicity of the sample surface. Take the sample, dip it in a shallow dish of DI water, then lift it out and check whether water clings to the SiO_x layer or not. If the water forms a uniform layer across the surface, you

have hydrophilic behaviour, should blow-dry the chip, and proceed. **If the water runs off of the surface or otherwise forms beads**, stick the sample into the UV Ozone for a run as before. You most likely have organic contamination on the surface of your chip, and that will cause delamination of your photoresist, ruining the film. Doing this quick check will literally save you an hour of tedious work.

1. Find a chuck that is just smaller than your sample, and attach it to the spindle, pushing down to make sure the seal is good.
2. Load the calibration piece of silicon (should be in a dish next to the spinner) and turn the spinner on. Time this, then adjust the timer and spindle speed as necessary. Remove the calibration piece and put it back in the dish.
3. If there isn't one already there, get a large watch glass and blow it off with the nitrogen gun to remove any dust. Set it atop the spinning well **concave side down**. It should just cover the well. This will act as a dust cover when you spin. Take this time to also blow-clean the inside of the well to get rid of any dust in there. Any dust particle that lands on your chip as you spin will result in streaks in the photoresist, and you will likely have to go through the process of cleaning the resist off and reapplying it.
4. Lay a sheet down on the flow bench top, and grab **two 10 or 20 mL glass beakers, two plastic syringes and two syringe filters**. Ensure that the glassware is clean, then blow into each beaker with the nitrogen gun to remove any excess dust, and set upside down on the sheet.
5. Pour some primer into the first beaker. You don't need a lot, and you'll likely slop some over the side of the bottle but that's ok, you'll get used to pouring this stuff real soon. Once the primer is in the beaker, quickly pull it into a syringe. Remove air bubbles from the syringe and apply the filter to the tip, again removing bubbles. Rest it on the lip of the primer beaker, you'll still need it. Repeat with the photoresist.
6. Shift the watch glass off of the well and place your sample on the chuck. Turn the spinner on to ensure contact, then quickly take the nitrogen gun and blow **directly normal** to the surface of your sample to remove any dust on it while it's spinning on the spindle. Take this opportunity to double check the spindle speed, as it will change with chips of different sizes. Quickly replace the watch glass and stop the spinner.
7. Take your syringe of primer and let a few drops fall into the beaker to clean any dust out of the filter tip. Keeping the syringe inverted, let **2 or 3 drops** fall onto the surface of your sample. Cover up with the watch glass and wait **30 s** for the primer to spread on its own. Once time has elapsed, turn the spinner on and wait for it to finish.
8. Take your syringe of photoresist and let a few drops fall into the beaker to clean any dust out of the filter tip. Keeping the syringe inverted, let **12 to 15 drops** fall onto the surface of your sample, however much is necessary to get a good spread of resist across the chip surface without flooding it. Cover up with the watch glass and again wait **30 s** for the resist to spread on its own. Once time has elapsed, turn the spinner on and wait for it to finish.
9. Check the quality of your film. It should be mostly uniform. If you see areas where the photoresist hasn't adhered, or there are numerous "comet tail" streaks from particles on the surface, see the next section on fixing a bad film. If your chip has corners (which it


should), you may observe a slightly different thickness of film at the corners than the rest of the chip. This is called edge blende and is unavoidable with radially asymmetrical samples; it will not have an impact on your fabrication. If you have any doubts on the film quality, take it to the microscope with the yellow filter turned on and inspect. If your film looks good, skip ahead to the section on soft baking!

Repairing a bad film

So you had a bad film eh? Happens to the best of us. There are a couple of common types of bad film, some more easily fixed than others.

First, if you have “comet tails” in your film, check under a microscope to see the extent of the film disruption. If there is no interruption of the resist (i.e. the streak is there but it’s just a difference in thickness as opposed to fully missing resist), then you’re fine to continue with the lithography, you may just have one bad device because of this. **Don’t forget that you’re making almost 400 devices here.** If you lose one or two, it isn’t even significant, it happens in industry too!

Delamination can be much more of a pain to deal with. Under the microscope, it may look something like this:



You can try to spin another layer of photoresist on over this to see if you can cover up the most egregious flaws in the film. If that doesn’t work, you’ll have to remove the film and re-spin. For this process, grab a large glass dish with sidewalls and a couple of beakers.

1. Place the sample in a wafer dipper and put it in the dish. Put **acetone** in a beaker. Spray the surface of your sample with acetone from the squirtbottle to remove the majority of the resist and primer, then submerge the sample in the beaker of acetone for **5 minutes**.
2. Put methanol in another beaker. Once the acetone soak has elapsed, dunk the sample into a beaker of DI water, then bring it back to the dish. Spray the surface with **ethanol**

- first, then **methanol**. Soak in methanol for another **5 minutes**. When finished, dunk in a new beaker of DI (keep flushing the DI beaker to remove trace organics).
- Put isopropyl alcohol (IPA) in a beaker. Bring the sample back to the dish, then spray with **ethanol** again, followed by **IPA**. Soak in IPA for **5 minutes**. Once finished, move to dish and spray with ethanol again, then dunk in DI and transfer to a running DI rinse for another 5 minutes.
 - When you remove the sample from the rinse, check the hydrophilicity. You are now back to square one.

TOOL (sort of): Hot plate, Soft bake

Hopefully you have a good film by this point, so this should be pretty straightforward. Set a hot plate to 110 C and let it get up to temperature. Once at temperature, transfer your sample to the surface and partially cover with a glass dish to allow off-gassing but protect from dust. **Soft bake for 2 minutes**. Remove the sample onto a paper sheet on the flow bench and cover with the glass dish until cool.

TOOL: Mask Aligner

The aligner is a temperamental tool. The spark gap has oxidation issues, meaning that occasionally the lamp will refuse to fire and you'll be set back. It's a pretty easy fix but you should contact the CEDT technician to help. To avoid causing this problem for yourself, **always let the tool cool down until the heat sink fins on top of the lamp are cool to the touch**. Also, I highly recommend checking if the aligner will work **before** starting the rest of your work in the clean room, as it will save you some frustration. Operate the aligner as per the instructions on the sheet next to it.

Once the lamp fires and heats up, do a test fire to check the aligner power. Some people swear by external gauges, but I find that the power on the lamp module itself is fairly accurate. Set the display to show mW of power, then bring the contact lever up (without your sample in it) and hit Expose. Check the maximum power that appears (in my case, typically around 9.8 mW). Run a quick calculation: you want to expose for **30 J** of energy at this stage, so set the timer to 3.06 s in our case. **Then, knock off another 0.1 s**, making it 2.96 s in our case. This is an adjustment I made based on observed over-exposure with the properly calculated time. The 0.1 s adjustment applies universally to all energy exposures you will do, so keep it in mind.

This first alignment is fairly straightforward. Try to center the pattern on your sample as much as possible, and ensure that your MESA ridges are aligned **perpendicular to the major flat**. The only other thing to watch out for is that the edges of your pattern are parallel to the edges of your sample. Expose your sample and let's keep on truckin'!

Chemistry time!

Now that the pattern has been laid into the photoresist, it's time to develop it. This process can be quite subjective, and requires vigilance on your part, so don't turn away.

1. Mix up a **1:5 solution of 351 Developer:DI Water**. As before, measure out the 1 part (351 in this case) into a 250 mL beaker, then add DI water as necessary. Swirl to combine. **Have a beaker of DI ready for quenching**.
2. Load the sample onto a wafer dipper and submerge in the solution. You should notice red wisps of photoresist begin to curl off the sample, and the pattern reveal itself. After **about 40 s**, when you don't see any more significant wisps of red coming off the chip, pull it from the developer and dunk in the DI quench. Record the time spent in solution.
3. Transfer the sample to a running DI rinse for **5 minutes**, then blow-dry with nitrogen.

TOOL: Back to the hot plate

Turn the hot plate up to 130 C. Once at heat, place your sample on the plate and partially cover as before. **Hard bake for 2 minutes**. Remove the sample from the plate onto a paper sheet in the flow bench, then cover with the glass dish until cool.

More chemistry!

This next step will use **hydrofluoric acid (HF)** to etch through the mask oxide exposed by the photoresist mask. This will transfer the pattern from the “soft” photoresist mask to the “hard” oxide mask. As the SiO_x etch is isotropic, care must be taken with the timing of this etch to avoid accidentally etching away finer features. The waveguides and other mask features have been designed to accommodate for this isotropic etching.

1. Safety first! Slap on a pair of nitrile gloves and the plastic faceshield before handling HF. Take all precautions to ensure you don't spill or splatter any HF while working with it. This stuff can kill you, so be careful: don't fear it, but respect what it can do.
2. In a **400 mL Teflon beaker**, pour enough HF that it will cover your sample in a wafer dipper.
3. The etch rate is approximately **1 minute per 100 nm** of oxide, so plan your etch time according to your film thickness (again, if your film is more than 100 nm thick be cautious of the effects of the isotropic etching). **Verify the etch rates ahead of time**.
4. Load your sample onto the dipper, then submerge for the time you calculated. Once the timer is done, **immediately transfer to a DI quench to stop the reaction**. Then transfer to a running DI rinse for **5 minutes**.
5. Blow-dry the chip when the rinse is done. Inspect under a microscope: you should be able to see distinct differences in colour where the oxide has been etched away.

We will now remove the soft mask, leaving behind the hard oxide mask.

1. Over a glass dish with sidewalls, place the sample in a wafer dipper and spray with acetone to remove most of the resist mask. Transfer to a beaker of acetone and soak for **2 minutes**.
2. Transfer the sample to a quench beaker, then a running DI rinse for **5 minutes**.
3. Inspect the sample under the microscope. The resist should all be gone, leaving you with the oxide mask alone. Note any features which seem to have over-etched; if the

damage is significant, you may have to remove the rest of the oxide with a **2 minute** HF etch, and redeposit the mask oxide.

4. It is prudent at this point to use the alpha step to measure the thickness of the oxide to confirm your ellipsometer reading, ensure that you actually got all of the oxide, and for comparison during the next step.

The III-V etch

Alright, this is the point of no return. Up until now, as illustrated in the last point, it's possible to fully reverse (almost) everything and return to the start point. Once you begin with the III-V etch, though, you can't turn back, as the changes on the chip will be permanent. First up will be the InGaAs contact layer.

1. In a 600 mL beaker (big one), prepare a **1:8:80 solution of H₂SO₄:H₂O₂:DI**. Again, measure out the 1 part first, this time **sulfuric acid** in a 10 mL graduated cylinder. **Make sure you wear latex gloves for this and not nitrile**, as sulfuric acid will eat through nitrile and then your hands. Due to the poor design of the acid bottle, you will likely slop some down the side of the sulfuric acid bottle while trying to pour it. Don't panic, but re-cap the bottle and transfer it to the bench sink. Run city water (no need for DI here) over the cap and down the sides of the bottle for a good while, and carry on with the rest of the prep. Pour the acid from the 10 mL cylinder into the bottom of the beaker. You want to aim for **4 mL or less** of the acid, as you will be heavily diluting it.
2. Next, measure out the **80 parts DI water** and add that to the beaker.
3. This etch will take **30 s per 100 nm** to eat through the material, so for our 200 nm layer we'll set up a 1 minute timer.
4. Measure out the hydrogen peroxide and keep it to the side until you're ready to dunk the sample into solution. It is important to add the H₂O₂ last, as it is an activating agent for the solution. Once you have your sample on a wafer dipper and a timer ready, add the peroxide to the beaker and swirl to combine. Then quickly dunk the sample in and start the timer.
5. Once the timer has elapsed, quickly transfer to a quench beaker, then a running DI rinse for **5 minutes**.

Next up will be the InP etch which fully defines the ridge waveguides.

1. In a 250 mL beaker, prepare a **1:3 solution of HCl:H₃PO₄**, the same composition as the flash etch at the beginning of the process. Swirl to combine.
2. Dip your sample into the solution for **3 minutes 30 s (240 s)** for a ridge depth of 1.5 um. When finished, transfer to a quench beaker then a running DI rinse for **5 minutes**.
3. Verify the depth of the etches using the alpha step and your knowledge of the mask oxide thickness. If the etch depth is good, proceed. If the alpha step says nothing has happened, ignore it and proceed anyways.
4. Inspect under a microscope. You should be able to see some discrepancies in the pattern where the anisotropy of the III-V etch left material behind at the edges of the waveguides or mesas.

5. Remove the mask oxide with a **2 minute** HF dip, transferring to a quench beaker and running DI rinse for **5 minutes**.
6. Inspect your sample under the microscope. The MESA pattern should have transferred to the III-V material now.


You're now done etching your III-V material!

TOOL: CVD Again

The next part of the process is to grow the barrier oxide on the sample. You can do this in any CVD that will produce a reliable oxide, since the thickness is a little more important this time around, due to the vias. If the oxide is too thick now, you will ruin your vias while trying to define your vias. Deposit **100 nm of SiO_x**. If the chip has been sitting around for a while, it may be prudent to redo the initial clean step involving the UV Ozone to remove any organics which may have stuck themselves on in the meantime, **before** depositing the barrier oxide. Use the VASE to check your film again.

VIA Photolithography and Definition

The next steps will be to cut the vias into the waveguides. The vias are the openings on the tops of the waveguides which allow for electrical contact. Make sure you check the hydrophilicity of the barrier oxide again. As long as you properly use the alignment marks, you should be fine. Properly aligned marks for these masks look like this:



Repeat the spincoating process as for the MESA mask, including the **soft bake**. **The mask aligner power process will be the same as for the MESA mask too**, except that you will be aligning the VIA mask instead, and ensuring that the alignment marks are properly lined up, including parallax considerations. I found it useful to check opposite corners of the mask to check for rotation, and slowly work my way to the proper alignment.

Develop the VIA pattern in the same way as the MESA mask, although it will be a little more difficult to see the wisps coming off the chip this time. Rinse and **hard bake** as before too.


Once the sample has been developed and hard baked, the vias will be cut with HF.

1. In a **Teflon** beaker, pour enough HF that it will cover the sample in the dipper.
2. Use the **60 s per 100 nm** etch rate again to determine your etch time.
3. Dip your sample into the HF. Once the time has elapsed, remove it to a DI quench and running DI rinse for **5 minutes**. Blow-dry the sample once finished.
4. Check under the microscope to see if the vias have been cleared of oxide. If not, drop back into the HF again for small increments of time until the vias are clear.
5. Strip the photoresist with acetone **as if the film was bad**. You are about to apply another layer of resist on top of the chip, so do as much as possible to remove excess organics. Check under the microscope to confirm the via clearance again.
6. Check the hydrophilicity of the chip. If it's good, you're clear to proceed! If you have water running off the surface, you will need to slap it back in the UV Ozone to burn off organics again, **as before**.

METAL Photolithography and Deposition

The MET mask protects the majority of the chip and defines contact pads and labels, so that when the metal is deposited the devices don't short-circuit. There are a few key differences in the lithography steps here, so make sure to pay attention!

1. **As before**, spin on photoresist and primer.
2. Soft bake the newly spun resist, **this time for 90 s at 90 C**. Remove to a sheet to cover and cool.
3. Set the timer on the mask aligner such that you will be exposing for **50 J** instead of 30 J, and again subtract **0.1 s** from the time. For instance, at 9.8 mW, expose for 5 s.
4. Examine the sample under the microscope after exposure. You may see something like this:



5. If you see that, it means you still have too much organic contamination on your chip surface. You know what to do.

6. Prepare a **dry** 250 mL beaker with enough **toluene** to cover the sample in a dipper. Submerge for **7 minutes**. **DO NOT TRANSFER TO A QUENCH OR DI RINSE**. This will cause the photoresist to crack, ruining your film. Instead, dry the toluene off directly using a nitrogen gun. The toluene hardens the top layer of resist, creating a shelf of resist once developed which will help to remove deposited metal.
7. Hard bake, **this time for 90 s at 90 C**. Remove to a sheet to cover and cool.
8. Mix a developer solution, **as before**. Submerge the sample, and again wait until the red wisps subside. As the top layer of resist is much more...resistant to the developer, this step will take about 2 times longer than for the MESA or VIA masks. Be patient, and remove the sample to a quench and rinse once the red wisps subside.
9. Check your film under the microscope to make sure you have good contact pad definition.

TOOL: e-Beam evaporator

The metal contacts are deposited by e-beam evaporation in the RIE side of the clean room. This is the first of two metallization steps in the fabrication process, and involves using the **adjustable angle deposition stage**. Not performing a tilted deposition will result in a contact that does not adequately cover the sidewalls of the waveguide, making it almost impossible to electrically pump without destroying the device (due to a combination of heat sinking problems and mechanical stress on the waveguide itself). Ensure that this step is performed correctly the first time, because nobody wants to spend hours removing a gold-based contact layer. This is a long process (~2-3 hours depending on level of practice) and requires a decent level of motor skill, so be prepared for sweat and frustration, and make sure to monitor deposition rates to ensure one of the crucibles doesn't run out mid-process. Bring something to do while you time the deposition to make sure you don't overshoot by accident (it's not the end of the world if you do by a bit).

Follow the instructions provided alongside the tool, with these addendums, and you'll be fine! You will be depositing **250 Å of Ti, 500 Å of Pt and 2400 Å of Au: 1200 at 30° and the other 1200 at -30°**.

1. Only use gentle pressure to tighten the clips! The chip is still very brittle, and too much pressure will snap it.
2. Make sure the waveguides are aligned in the proper orientation relative to how the sample holder will tilt. The waveguides should be parallel to the axis about which the stage tilts.
3. Leave the stage in its normal configuration for the Ti and Pt contacts, as these must be deposited straight on.
4. Use the manipulator to tilt to 30 degrees as indicated on the holder. Take your time with this, and be careful not to miss/snap back the manipulator and have it make contact with your sample.
5. Once the stage is at the proper angle, deposit the first 1200 Å of Au.
6. Use the manipulator to tilt the stage to -30 degrees (the other 30 degrees on the holder), then deposit the second 1200 Å of Au.

7. Remove the chip from the sample holder, inspect it to make sure everything looks intact, and place it in an acetone beaker. Place this beaker into the ultrasonic bath for 30 minutes to remove excess metal. Repeat as necessary until you're satisfied with the amount of metal removed. An overnight acetone soak is likely necessary to loosen up the toluene-hardened photoresist enough to achieve liftoff.
8. Still having problems getting the metal off? Get a needle, hook it onto a syringe and blast off the metal forcefully, being careful to not scratch the chip with the needle tip.

Thinning

Make sure you have a lab coat for this because it can get messy, and nobody wants to spend time getting InP juice out of their clothes. The devices are thinned in the CCEM prep lab, using a **South Bay thinning chuck**. Make sure that the sample puck you use is properly flat, because this tool takes a lot of abuse. We should have a flat puck squirreled away in the Knights lab. This is another time-consuming process and I hope to god that you haven't been skipping arm day at the gym because it's a bit of a workout. If you haven't already, you will need to clip the curved corners of your sample to make sure it fits on the puck.

1. Place the sample under the micrometer on a microscope slide. Measure the sample thickness in a nine-point grid, noting any gradient in the thickness. **Be careful not to tap your chip too hard with the probe**. You will cause tiny fractures which will propagate when the chip is heated.
2. Measure the puck thickness – if the gradient of the puck is large (i.e. 20 um difference between one side and the other), consider having a layer lathed off of the top to re-level the puck.
3. Put the metal puck on the hot plate and place a few small pieces of crystal bond in the middle of the puck. **Make sure the plate is cold when you do this**. This will allow you to adjust the placement of the crystal bond nuggets.
4. Crank up the plate temperature to a bit past what they say is the max temperature you should use (marked by white lines on the temperature dial).
5. Wait until the crystal bond liquifies to the point that it resembles a water droplet, then place your sample (**etched side DOWN**) onto the puck, applying gentle pressure to the top with a flat object to ensure the sample is parallel to the puck surface; move the chip around a bit in small circular motions using the flat bottom of a pair of wafer tweezers to ensure even coverage of the crystal bond across the chip surface. Make sure that crystal bond is not creeping over the edges of the chip, as that will complicate the thinning process.
6. Once satisfied with the evenness of the crystal bond layer, remove the entire puck from the hot plate with a pair of crucible tongs and place it on one of the larger metal discs next to the hot plate to cool down. This may take 10-15 minutes.
7. Once cooled, clean up any excess crystal bond with a kim wipe soaked with a bit of acetone, in gentle brushing motions so as not to scratch or damage your device.
8. Measure the thickness of the sample now that it is bonded to the puck. This will let you know how thick the crystal bond is, and your target thickness while the chip is attached to the puck.

9. Mount the puck in the South Bay chuck and turn the axial screw to fasten the puck in. Release the set screw on the side so that the middle column moves freely up and down (but does not rotate), and move the micrometer dial to the top. The puck should now move freely along the axis of the chuck, and protrude significantly below the level of the chuck feet, while being able to retract past the level of the feet.
10. With the middle column still free to move, gently place the entire chuck, chip side down, onto a clean and flat surface. This will make your chip level with the feet of the chuck.
11. Tighten the set screw so that vertical movement of the puck is prevented, and spin the micrometer dial down so that it just barely tightens. Note the marking on the dial here.
12. Loosen the dial to set how much material you want to remove. The dial is marked in increments of 25 μm . For instance, if you start at 500 μm thick and are aiming for 175 μm (removing 325 μm), try to remove 200 μm of material first to see how that goes, then adjust your dial settings from there. If you attempt to remove too much material in one go, you risk snagging the edge of your chip on the grinding surface or “doming” the surface, where the middle of the chip will be much thicker than the edges.
13. Once you have set the dial, pick up the chuck and loosen the set screw. Apply pressure to the axial screw with one finger while simultaneously re-tightening the set screw. This should lock in your exposed thickness of chip as whatever you set it. Test that you have exposed your chip by placing the chuck back down on the flat surface, and gently rocking it around in a circular motion. You should be able to feel the outline of your chip as the edge contacts the flat surface. Tighten the screw above the micrometer dial to lock in your measurement.
14. Start with a 600 grit (Standard ANSI) SiC lapping disk/paper strip to remove large amounts of material. In the fume hood, place a glass plate in a tub such that one edge of the plate rests on the lip of the tub, and the other edge rests in the tub itself, forming a glass ramp.
15. Fill a squirtbottle with DI water, and place your lapping paper on the glass ramp. If it's a paper strip, wrap both edges around the near and far sides of the glass, pinning them against the tub to secure the paper. If using a disk, the surface area of the disk should suffice to secure it.
16. Wet the paper with some water from the bottle, then place your sample down on it. Ensure that you apply even, light pressure, and that your sample stays relatively parallel with the plane of the glass ramp. **Let the weight of the chuck do most of the work.** Grind in a figure eight pattern (to ensure you don't create consistent defects in one direction of the chip) until you no longer feel like the chip is grinding on the paper, or no longer see black residue coming off on the paper. Periodically rotate the chuck by 90° too. The feel for when you've stopped grinding your chip takes practice, so don't worry if you miss it the first few times. If you overgrind, all you do is dull the grinding surface, as the feet of the chuck are also made of silicon carbide.
17. Once you've finished a round of thinning, rest your arm and measure the thickness of your sample on the puck to check your progress. Make sure to avoid getting InP residue all over the place (rinse/wipe as necessary). If you haven't removed enough material, return the puck to the chuck and continue polishing again. Ideally you want to aim to

- be within 200 ± 5 μm of a chip thickness (excluding crystal bond). For whatever reason, I always found that the measurements from the CCEM micrometer became inconsistent, and that 200 μm was a good target. If you are within a neighbourhood of 200 ± 5 μm , move on to the next thinning step. Otherwise, continue with the 600 grit. 200 μm is also a good thickness to cleave at.
18. Now, grab some 1200 ANSI grit discs (at least 2). Apply them to the glass plate in the same way as before.
 19. Set your micrometer dial to remove approximately 100 μm . Trust me on this one.
 20. Grind away! You'll feel the chip stop properly grinding quite early on, but keep going at it for a few minutes. Check periodically to see how the finish is developing on the back side of the chip, rinsing as necessary, and target areas that don't appear to be properly thinning.
 21. Once you feel like you're no long making progress with your grinding, check the thickness of the sample and repeat until you have a finish on the back which allows you to see your reflection. It should not be necessary to go beyond 1200 grit. If one of the 1200 discs loses its abrasiveness, switch to the other.
 22. Once you feel like your arm is going to fall off, thin a little bit more, then call it a session. Dispose of the waste water in the tub in the appropriate waste bucket, rinse everything off, etc.
 23. Measure the thickness of your sample, then carefully wrap it in some kim wipe (puck and all) and take it up to the clean room.

Removing your sample

Now that you've thinned your chip down, as you can imagine, it's quite a bit more fragile than before. This part has the potential to go sideways real quick, so keep an eye on everything and make sure your hands are stabilized.

1. To remove the sample from the puck, take it to the clean room and place the puck sample side up on a hot plate.
2. Ramp the hot plate up slowly, as the difference in coefficients of thermal expansion between the three materials (metal, crystal bond and III-V) can cause your freshly thinned chip to fracture and you to weep uncontrollably for a day. Start by ramping to around 100 C, then tick up by increments of 10 C, waiting a few minutes each time to allow the entire system to reach the temperature.
3. While waiting for the temperature to stabilize, prepare a beaker of acetone with a tall wafer basket in it, with enough acetone to cover the width of the chip.
4. Once you reach about 150-160 C, the crystal bond should be loose enough that you can begin to gently slide your ship off of the puck using a pair of tweezers so that about 2/3 of the chip is hanging over the edge of the heated puck.
5. Shift the puck towards the edge of the hot plate using tongs and, using 2A tweezers, not wafer tweezers, pull the chip along the plane of the puck surface until it releases, **making sure to apply force parallel to the plane of the puck**. Then, quickly place the wafer upright in the waiting acetone wafer basket, which should release the chip from your tweezers by dissolving the crystal bond on contact. If the chip doesn't

- release, panic for a second or two, then gently apply pressure to the chip against the basket to pry it off.
6. If you fracture the chip at any point here, don't worry, as long as there are salvageable waveguides on it you'll be fine. The fractures normally occur perpendicular or parallel to your waveguides.
 7. Soak in acetone for **5 minutes**, then transfer to a quench DI beaker and transfer again to a new acetone bath. Repeat this process **six more times** to ensure that all of the crystal bond has been removed. When finished the series of acetone baths, transfer to a running DI rinse for **5 minutes**.
 8. Turn the hot plate off and remove the puck, allowing it to cool on a paper sheet. While it's still warm (not hot), wipe off any crystal bond you can see on the surface with a wipe soaked in acetone. Once it has cooled, wipe it down again with more pressure applied to remove any more excess crystal bond.
 9. Place your chip (or perhaps chips at this point) into a sectioned sample carrier, and take the day off because you earned it!

Metallization Part 2

Wow this is gonna be so easy compared to that last part. In fact, everything after this is easy compared to the last part. You don't even need the adjustable angle holder this time, just the normal one! Again, the samples are even more fragile now that they're thinned, so take extra caution. Now grab a margarita and hit the e-beam evaporator again. The back contact is **250 Å of Ni, 500 Å of Ge and 1200 Å of Au**.

1. Take a scrap piece of silicon that is small enough to clip onto the e-beam holder. A 2" wafer is a good bet, and there should be plenty lying about if you ask around. This will make it so that you don't have to risk directly clipping your thinned devices.
2. **Load up a syringe with 1827 photoresist**. Using the syringe like a brush, slowly squeeze out photoresist and smear it onto the wafer to form a thin layer. Place your device(s), waveguide side down, into this photoresist, making sure none creeps over the edge and onto your chip surface. Place all the chips you want metallized.
3. Ensure that you squeeze as much of the photoresist out from underneath the chips as possible, as the low pressure of the chamber will cause any air bubbles to expand and lift your chip off the carrier wafer.
4. Soft bake for **2 minutes at 110 C**, then **hard bake for 7 minutes at 130 C**. Bake the resist until the pieces do not move, then bake a little more.
5. Clip the silicon into the sample holder stage and deposit the Ni/Ge/Au contact, using the typical instructions. Let the load lock pressure come down as usual, then as you move the sample from the load lock to the deposition chamber, check the sample for any bubbles of photoresist that may have formed and expanded. If the bubbles are large, remove the sample from the deposition chamber, pull it out of the load, lock, remove the bubbles with tweezers, then load the sample back in.
6. Once finished, place the carrier wafer in a glass dish with sidewalls and spray acetone around the edge of your device to remove photoresist. Once you have removed all

visible resist, soak the carrier and device in acetone for **15 minutes**, or until the devices separate from the carrier.

7. Transfer the device to a clean beaker of acetone and let soak for another **5 minutes**, then transfer to a DI quench and running rinse for **5 minutes**.
8. Blow-dry the sample very carefully, or just move it back and forth on a piece of dry paper for a while if you need to, this would be a bad time to break it into pieces.
9. Breathe.

TOOL: Rapid Thermal Annealer

Now that both metal contacts have been deposited, we need to drive them into the semiconductor to create a strong contact. Annealing the metal into the chip will also grant it a surprising amount of structural stability. The tool to be used is not one of the RTA's in TAB, but the one in Dr. LaPierre's lab, as it has superior temperature control over the range that we require. Follow the SOP that goes with the tool and you should be good to go.

The recipe to use is Matt.rcp. This ramps the chamber temperature up to 415 C over 60 s, then holds steady at 415 C for 30 s, then cools down to room temperature over 120 s.

TOOL: Cleaving setup

The devices will be separated by the cleaving tool in TAB. A CEDT technician should be able to train you on its operation. As with a lot of this procedure, this tool requires quite a lot of practice, so I recommend using garbage III-V samples to make sure you get the motor skills down before trying it on your real device.

1. Scribe the chip about 1 mm from the edge, making sure the scribe tip force is set to InP.
2. With the vacuum still on, smooth the plastic tape backing over the chip on all sides, sealing it in a sort of envelope. Try to have some more plastic on the right side than on the left.
3. Move the chip to the microscope slide fixed to the edge of the cleaving tool. Use the angle of the lighting to find the scribe mark, and line it up about half a millimeter over the edge of the slide.
4. Hold the left side of the plastic backing with one hand. Pinch the backing and tape together on the right and pull parallel to the face of the device to create tension on both layers.
5. Snap quickly, applying pressure nearer to the edge of the chip that has been scribed and slightly twisting so as to start the cleave on that end instead of the near end. Do your best to push downwards along a plane normal to your devices (i.e. parallel to the side of the tool). If all goes well you should have a clean snap as the bar comes off.
6. Balance the laser bar on its edge under the microscope and examine the facet. If you've done it properly, there should be a mostly mirror finish on the facet, with some stress striations along the back edge of the chip. We don't care about those because they do not affect device performance, although fewer striations is always better.

Copper bonding

Now that the bar is separated, we need to bond the device to a support that will act both as a handling aid and as a rear contact. For this, we'll use a piece of copper sheet metal and silver paint.

1. Clip off a piece of copper sheet just larger than your device bar, making sure that there will be enough space to land a contact probe once you've mounted the sample.
2. Flatten the copper against the anvil section of a vice, using a hammer. Get this as flat as possible to ensure better rear contact with the devices.
3. In a fume hood, shake up a bottle of SPI silver paint to re-emulsify it. A brush is contained attached to the cap of the paint bottle.
4. Place your copper down, and have your device at the ready to place. Using the paint brush, wipe off most of the excess on the inner lip of the bottle, then lay down a thin layer of paint on the copper.
5. Quickly transfer your device to the paint before it has a chance to dry. Position it how you'd like, and try to avoid pushing it around on the paint too much, as this will normally lead to paint creeping up the sides of the device and ruining the facets/shorting out your device. It may also be possible to apply very light pressure to the top of the device bar to ensure better adhesion.
6. Let the bar sit for 5-10 minutes as the paint cures and dries.

You're done! You're a Master! Seriously, give yourself a pat on the back and a few pitchers of beer because this is a tough process to get right. Now all you have to do is test it to make sure it actually works...

BONUS ROUND: SOA Anti-reflection coating

Now that you've tested your devices and found out they're lasing, it's time stop that lasing and start amplifying instead! This process is far from perfect, so take it with a grain of salt, but it worked for me and it should work for you too. You will need to work out what compositions of SiO_x and Si₃N₄ to use to properly attenuate reflections for your device, based on the output spectrum. This can be easily handled by changing the layer thicknesses in the Filmetrics calculator, <https://www.filmetrics.com/reflectance-calculator>. The films can then be deposited by a more precise CVD tool than the Technics one, namely the ECR-PECVD in TAB. I used silicon nitride layers that were 53 nm thick, and silicon oxide layers that were 73 nm thick, arranged as a four layer stack.

1. Construct a physical mask for the devices using a piece of solid aluminum wire that, when flattened with a hammer, is about 2/3 the width of one of your device cavities.
2. Bend both sides of the flattened section of wire such that the mask lies on the device as below:



- This ensures firm contact with the top of the device with even pressure, without directly clipping onto the device.
3. Clip the copper carrier onto the deposition stage using the tungsten clips, then clip the mask over top.
 4. The operator of the tool should be able to help you from here, and get your AR film deposited!
 5. Once you remove your device from the deposition stage you've got yourself (hopefully) an SOA!


Appendix B: Device Fabrication Matrix

Device		Mask	Piranha etch	Thinned	Diced or Cleaved	FIB'd	Gratings Attempted	AR Coating	Lasing	Amplifying	Operable
Growth	No.										
A	1	0	N	Y*	Diced	Y	N	N	N	N	Y
A	2	0	Y	Y*	Diced	Y	N	N	Y**	N	Y
A	3	1	Y	N	Diced	Y	N	Y	Y	N	Y
A	4	1	N	Y	Cleaved	N	Y	N	N	N	N
B	9	1	N	N	N/A	N	Y	N	N	N	N
B	6	2	N	Y	Cleaved	N	Y	Y	Y	Y	Y
B	3	2	N	N	N/A	N	N	N	N	N	N
B	8	N/A	N	N	N/A	N	N	N	N	N	N
B	2	2	N	N	N/A	N	N	N	N	N	N


*Thinned devices were not tested for operation.

**Low power lasing only, but characteristic LI curve is present.

Growth A (2")



Growth B (3")



Appendix C: Amplification Sweep Matlab Code

```

%Initial author Zhao Wang,
%Adapted by Matthew Vukovic, Feb 2020

close all;
clear all;
instrreset;

device = 'EDFA-L04-260mA'; %device name and conditions, cap at 260mA
file_name = ['C:\Users\Matt\Documents\00_Matlab\' device]; %file name for csv

%static program variables
address = 20; %GPIB address of Agilent 8164A
sigpow_init = 0;%-10; %initial laser source power (dBm)
sigpow_fin = 0;%5; %final laser source power (dBm)
sigpow_step = 2; %power step (dBm)
sigpowvec = [sigpow_init:sigpow_step:sigpow_fin]; %input power vector
atten = 0; %signal attenuation at laser output
attinit = 0;
attfin = 40;
attstep = 2;
attvec = [attinit:attstep:attfin]; %vector of attenuation

% initial variables
lambda_start=1560;%1510; %start wavelength
lambda_stop=1560;%1640; %stop wavelength
lambda_step=2; %wavelength step
lambda=[lambda_start:lambda_step:lambda_stop]'; %wavelength vector
points=length(lambda);
scan_speed=5; % nm/s only 0.5 5 40 allowed
avg_time=2E-4;% photodiode average time

%for JDS Uniphase HA1 Attenuator
att = gpib('ni',0,3);
set(att,'InputBufferSize',100000);
att.EOSMode = 'read&write'; %allow queries
fopen(att); %open instrument
fprintf(att, 'RESET') %restore to default 0 dB attenuation and 1550 nm
fprintf(att, 'DISP 0') %display attenuation

%GPIB setting
obj= gpib('ni',0,address); %create GPIB object
set(obj,'InputBufferSize',100000);
fopen(obj); %open instrument
fprintf(obj,'*CLS'); %instrument setting reset
obj.EOSMode = 'read&write'; %allow query() to pass to instrument
flushinput(obj);flushoutput(obj); %flush the data that was stored in the buffer

%trigger configuration
fprintf(obj,'trig:conf LOOP'); % an output trigger automatically works as input
trigger
fprintf(obj,'TRIG1:OUTP DIS'); % PD output trigger is disabled
fprintf(obj,'TRIG1:INP SME'); % PD will finish a function when input trigger is
enabled
fprintf(obj,'TRIG0:OUTP STF');% TLS will send output trigger when sweep starts
(input trigger generated)
fprintf(obj,'TRIG0:INP IGN'); % (TLS input trigger is ignored)

```



```

% sensor setting
fprintf(obj, 'init1:cont 1'); %continuous detection mode
fprintf(obj, ['sens1:pow:atim ' num2str(avg_time)]); %set the averagetime to lms
for sensor 2
fprintf(obj, 'sens1:pow:rang:auto 1'); %set auto ranging on
fprintf(obj, 'sens1:pow:rang 10DBM'); %set max power expected at output
fprintf(obj, 'sens1:pow:unit 0'); %set the unit of power: 0[dBm],1[W]
fprintf(obj, 'sens1:pow:wav 1560nm'); %set sensor wavelength centered at 1550 nm
fprintf(obj, 'TRIG1:INP IGN'); %ignore input triggers and freely measure optical
power

% tunable laser setting
fprintf(obj, 'outp0:path high');% choose which path of tunable laser. output1 [low
power high sens] output2 [high power]
fprintf(obj, 'sour0:pow:unit 0'); %set source power unit to dbm
fprintf(obj, ['sour0:pow 0dbm']);%set laser power for alignment and polarization
check
fprintf(obj, 'sour0:wav 1560nm'); %set wavelength for calibration
fprintf(obj, 'sour0:AM:stat OFF');
fprintf(obj, 'outp0 on'); %turn laser source on

%log execution time
t=0;
while str2num(query(obj, 'wav:swe?'))==1
    t=t+1;
end

% stepped wavelength sweep

%generate plot body
% figure()
% xlabel('Wavelength (nm)')
% ylabel('Power (dBm)')
% title([device ' Wavelength Sweep'])
leglab = {}; %empty cell to hold legend strings

for leg = 1:length(attvec)
    name = [num2str(sigpowvec(1) - attvec(leg)) ' dBm'];
    leglab{leg} = name;
end

if length(sigpowvec) > 1
for leg = 1:length(sigpowvec)
    name = [num2str(sigpowvec(leg)) ' dBm'];
    leglab{leg+length(attvec)} = name;
end
leglab{leg} = name;
end
end

hold all; %accumulate all on one plot

fprintf(att, 'D 0') %open beam block output

dataset = [];
fprintf(att, ['WVL ' num2str(lambda(1)) 'e-9']) %set attenuation
wavelength

for m = 1:length(attvec)
    if m == 1

```

```

        pause() %pause to align
        fprintf(obj, ['sour0:pow ' num2str(sigpowvec(1)) 'dbm']); %set initial
power
    end
    fprintf(att, ['ATT ' num2str(attvec(m)) ' dB']) %set attenuation

    powin = zeros(1,points); %initialize
    powout = zeros(1,points); %initialize
    realatt = zeros(1,points); %initialize

    for c = 1:points
        fprintf(obj, ['sour0:wav ' num2str(lambda(c)) 'NM']) %set source
wavelength
        fprintf(obj, ['sens1:pow:wav ' num2str(lambda(c)) 'nm']) %move sensor
wavelength with input
        fprintf(att, ['WVL ' num2str(lambda(c)) 'e-9']) %set attenuation
wavelength

        if c == 1 %only on first run through loop
            fprintf(obj, 'outp0 on'); %turn source on, outp:chan state
        end

        powin(c) = str2double(query(obj, 'sour0:pow?')); %record input power
        powout(c) = str2double(query(obj, 'read1:pow?')); %'read1:chan1:pow?');
%record output power
        realatt(c) = str2double(query(att, 'ATT?'));
    end

    dataset = [dataset; [powin' realatt' powout]]; %concatenate data

end

csvwrite([file_name '.csv'], dataset); %write data to named file
fprintf(obj, 'outp0 off'); %turn source off

fclose(obj);
delete(obj);
clear obj

fprintf(att, 'D 1') %shut off attenuator output
fclose(att);
delete(att);
clear att

the_variable_that_goes_BING; %audible indicator of program complete since this
takes a while to run

```