

ARTIFICIALLY BIREFRINGENT ALUMINUM GALLIUM
ARSENIDE/ALUMINUM OXIDE-BASED SUBMICRON
WAVEGUIDES AND RESONANT CAVITIES FOR
NONLINEAR OPTICS

A DISSERTATION SUBMITTED TO
THE DEPARTMENT OF APPLIED PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Luigi Scaccabarozzi

June 2006

© Copyright 2006 by Luigi Scaccabarozzi
All Rights Reserved

I certify that I have read this dissertation and that in my opinion it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

James S. Harris, Principal Advisor

I certify that I have read this dissertation and that in my opinion it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Martin M. Fejer

I certify that I have read this dissertation and that in my opinion it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Shanhui Fan

Approved for the University Committee on Graduate Studies:

Abstract

In fiber optics systems there is an increasing need for fast, compact switches. Currently employed opto-electro-optical converters are relatively slow and power demanding. Nonlinear optical devices, such as lithium niobate (LiNbO_3) waveguides, can provide fast, all-optical wavelength conversion. However, they require centimeter-long interaction length and they are not semiconductors, hence optical integration with active devices is extremely challenging.

On the other hand, GaAs has a nonlinear coefficient four times larger than LiNbO_3 and a well-established fabrication technology. Moreover, passive and active devices could be potentially integrated to realize on-chip photonic circuits. Gallium arsenide quasi-phasematched (QPM) waveguides have also been demonstrated and although the required interaction length is smaller (millimeters), they require a complex fabrication process.

In this work we present the design, fabrication and characterization of tightly-confining aluminum gallium arsenide (AlGaAs)/aluminum oxide (commonly denoted with the acronym AlOx) nonlinear waveguide devices. Because of the high index contrast of the AlGaAs/AlOx material system, phasematching can be achieved by artificial birefringence. This approach greatly simplifies the fabrication process,

compared to QPM waveguides. Artificial birefringence in weakly-confining AlGaAs/AlO_x waveguides has been already demonstrated; however, the conversion efficiency remained low, due to poor lateral confinement.

We employed the AlGaAs/AlO_x material system to achieve both birefringent phasematching and sub-micron confinement. The normalized conversion efficiency of our design is larger than 20,000 %/W/cm². This value is one order of magnitude higher than previously reported works. We developed a new fabrication process to realize high aspect-ratio, extremely smooth waveguides, which are characterized by very low propagation loss (~2 dB/mm) at the fundamental wavelength.

Moreover, we showed that nonlinear effects can be further enhanced using a cavity embedded in the waveguide, resonant at the fundamental wavelength. For this purpose a novel dichroic mirror, highly reflective at the fundamental wavelength but with high transmission at the second harmonic, was designed and fabricated. Using these mirrors, we demonstrated for the first time enhanced second harmonic generation in resonant, collinear waveguide devices.

In conclusion, we demonstrated tightly-confining, birefringently phasematched waveguides and resonant cavities with the highest normalized conversion efficiency ever reported. Furthermore, we designed and fabricated the first dichroic mirror for tightly-confining waveguides. Cavities realized with these mirrors achieved large resonant enhancement of second harmonic generation.

Acknowledgements

This work would not have been possible without the help and support of many people. First of all, I would like to thank my advisor, Prof. Harris. In the first place, without him I would probably not even be at Stanford; in fact, he invited me to join his group as a undergraduate visiting student, when I was still in Italy. This first experience convinced me to pursue my Ph.D research in his group. He gave me the freedom to follow my own ideas, although he was available to discuss any issue and to direct me. Among other things, from him I learned how to develop my own ideas and how to work independently.

I'm also deeply grateful to my co-advisor, Prof. Fejer. Since I met him in the first group meetings as a visiting student, he always impressed me with his amazing knowledge of the field (and beyond). He taught me not only part of this knowledge, but also how to do research and how to “think”. Moreover, his suggestions and comments have been invaluable in putting me in the right track to solve a problem and to progress in this project.

I also have to express my gratitude to my other co-advisor, Prof. Fan. I met him for the first time in his class and I was amazed by his clarity and ability to shine a different light on physical concepts (even on familiar ones), such that a simple and

elegant explanation derived, which was source of inspiration to find new ideas. I also have to thank him for the insights and great ideas that led to the success of this project.

I would also like to thank Prof. Brongersma, with whom I had helpful discussions and with whom I interacted several times along the course of my Ph.D., and Prof. Geballe for serving in my orals committee.

Harris group has been a great place to spend my graduate life, and there are many people I wish to thank: Thierry, my first mentor, who introduced me for the first time to nonlinear optics, and Vijit, who taught me about waveguides and how to simulate them; Xiaojun, who has been a great colleague and friend, and to whom I owe all the MBE growths, which have been absolutely essential to this project; Yijie, who has joined the project only the past year, but has learned very quickly and has been able to give significant contributions to the project; Alireza, Rafael and Mike, who have not only helped me with useful discussions and exchange of experimental knowledge, but have also been great friends; Meredith, who has helped me while she was an undergraduate student; Angie, who now has taken over the role of Xiaojun as grower; Hopil, Seth, Homan, Mark, Ofer, Lynford for their help with the MBE growths and for useful discussions. I mentioned here the people that have contributed directly to this work; however, I am also grateful to all the other past and present students who contributed to make Harris group a wonderful place to be.

In Fejer group, I wish to thank: first of all Khrishnan, who spent much time teaching me about optical measurements, always being very patient with me (and my mistakes); Paulina, for her friendship and suggestions, and the other people in the GaAs “sub-group”: Konstantin and Joe; Rosti, David and Xiuping, have been great friends and lab-mates, always ready to answer my questions and help me out.

In Fan group I would like to express my gratitude to Michelle, for the many discussions on many topics and for teaching me about supercomputing simulations, and to Zheng, Wahtung and Fatih, for their help with the early simulations.

Many people in the SNF clean-room have been of indispensable help: a special thanks goes to Eric P., who was always able to answer my questions on the most

disparate lab topics, and gave me priceless hints for developing a successful fabrication process; Jim McVittie, for his precious advices on any dry-etching matter; James Conway, for is support with the e-beam; Uli, Maurice, Manhaz, Ed and all the other SNF staff people for their help and for making the lab a nice and friendly place to work. I also wish to thank Tim Brand, in the Ginzton crystal shop, Tom Carver (Ginzton clean-room), who more than once “saved” my samples with his invaluable practical knowledge; Bob for his help with the SEM; Huatang and Caterine at OEPIC for the cleaving.

I’m truly grateful to Gail, who has done an incredible job in managing Harris group and in solving all the administrative problems; to Paula and Claire, for their warmth and friendliness (and for solving many bureaucratic problems, frequent with international students).

During these years I have been very fortunate to meet many excellent friends, of whom it would be impossible to give an exhaustive list. They contributed to make my out-of-the-lab life full and pleasant and I believe that the friendships developed in these years will last for many years to come.

Without my family, my Stanford experience would not have been possible. I am deeply grateful to my parents for their support in all my decisions: I would not be where I am now and what I am now without them.

These last six years would not have been the same without my wife, Ke. We met shortly after arriving at Stanford and we went through both joys and sufferings of Ph.D. life together. My Stanford experience would not have been as rich and happy without her: she is truly the most precious among everything I have found here.

Table of Contents

ABSTRACT.....	v
ACKNOWLEDGEMENTS	vii
TABLE OF CONTENTS	xi
LIST OF TABLES	xv
LIST OF ILLUSTRATIONS.....	xvii
CHAPTER 1 INTRODUCTION.....	1
1.1 MOTIVATION	1
1.2 NONLINEAR WAVELENGTH CONVERTERS	3
1.3 REVIEW OF PREVIOUS WORK	5
1.3.1 <i>Early works</i>	5
1.3.2 <i>PPLN waveguides</i>	6
1.3.3 <i>AlGaAs as a nonlinear material</i>	6
1.3.4 <i>QPM AlGaAs waveguides</i>	7
1.3.5 <i>Artificial birefringence</i>	8
1.3.6 <i>Resonant enhancement of nonlinear processes</i>	9
1.4 OVERVIEW OF THE THESIS	11

CHAPTER 2	THEORETICAL BACKGROUND	13
2.1	FUNDAMENTALS OF NONLINEAR OPTICS	13
2.1.1	<i>Coupled mode theory</i>	14
2.1.2	<i>Conversion efficiency</i>	16
2.2	NONLINEAR OPTICS IN WAVEGUIDES	17
2.2.1	<i>Guided wave coupled mode theory</i>	18
2.2.2	<i>Effect of loss</i>	20
2.2.3	<i>Effect of disorder</i>	21
2.3	PHASEMATCHING TECHNIQUES	23
2.3.1	<i>Birefringence</i>	23
2.3.2	<i>Quasi-phase-matching</i>	24
2.3.3	<i>Form birefringence</i>	25
2.3.4	<i>Modal phase-matching</i>	27
2.4	RESONANT CAVITIES	27
2.4.1	<i>Principles of resonators</i>	27
2.4.2	<i>Nonlinear resonant cavities</i>	30
CHAPTER 3	DEVICE DESIGN	32
3.1	BACK-OF-THE-ENVELOPE ANALYSIS	33
3.2	MICROCAVITY DESIGN	37
CHAPTER 4	THE FABRICATION PROCESS.....	42
4.1	EPITAXIAL GROWTH	43
4.2	ELECTRON BEAM LITHOGRAPHY	44
4.2.1	<i>E-beam resist</i>	45
4.2.2	<i>Pattern design</i>	46
4.2.3	<i>Exposure</i>	49
4.3	HARD MASK TRANSFER.....	50
4.3.1	<i>SiN mask</i>	51
4.3.2	<i>Nickel mask</i>	52
4.3.3	<i>Chromium mask</i>	56
4.4	ETCHING.....	57
4.5	OXIDATION.....	61
4.6	THINNING AND CLEAVING.....	63

CHAPTER 5	CHARACTERIZATION TECHNIQUES	64
5.1	MEASUREMENT SETUP	65
5.2	WAVEGUIDE CHARACTERIZATION	68
5.2.1	<i>The Fabry-Perot method</i>	68
5.2.2	<i>Dispersion (group index)</i>	70
5.2.3	<i>Propagation loss using the Fabry-Perot method</i>	72
5.2.4	<i>Propagation loss using a variant of the cutback method</i>	73
5.3	LINEAR CHARACTERIZATION OF THE MICROCAVITY	75
5.3.1	<i>Transmission spectra and resonance frequency</i>	75
5.3.2	<i>Cavity-Q</i>	77
5.3.3	<i>Improved geometry cavities</i>	78
5.4	PROBLEMS WITH SECOND HARMONIC GENERATION.....	79
CHAPTER 6	BIREFRINGENT WAVEGUIDES	82
6.1	DESIGN	83
6.2	FABRICATION	90
6.3	OPTICAL CHARACTERIZATION	92
6.3.1	<i>Linear characterization</i>	92
6.3.2	<i>Nonlinear characterization</i>	95
6.3.3	<i>Second harmonic loss: possible causes</i>	101
CHAPTER 7	BIREFRINGENT RESONANT CAVITIES.....	104
7.1	THE CONVENTIONAL MIRROR	105
7.2	DICHROIC MIRROR.....	107
7.3	FABRICATION	113
7.4	LINEAR CHARACTERIZATION.....	116
7.5	ENHANCED SECOND HARMONIC GENERATION.....	123
7.6	OPTIMIZATION OF THE RESONANT CAVITY	128
CHAPTER 8	CONCLUSIONS	136
8.1	SUMMARY OF RESULTS	136
8.2	FUTURE PERSPECTIVES.....	137
8.2.1	<i>The loss problem</i>	138
8.2.2	<i>Predicted performance of low-loss devices</i>	139
8.2.3	<i>Suspended structures</i>	139

8.2.4	<i>Improving the coupling efficiency</i>	141
8.2.5	<i>Integration with active devices and other applications</i>	142
APPENDIX A	FABRICATION RECIPE	144
APPENDIX B	ALIGNMENT TECHNIQUE	148
APPENDIX C	FITTING SOFTWARE	152
REFERENCES	156

List of tables

Table 4.1 Summary of metals etch rate and selectivity.....	53
Table 6.1 Sensitivity and amount of drift of geometrical parameters.....	99
Table 6.2 Contributions of roughness at different interfaces to the loss coefficient. .	102

List of Illustrations

Figure 1.1 Schematic of SHG and DFG.....	4
Figure 2.1 Efficiency vs. length and increase in FWHM vs. $\Delta\alpha$	21
Figure 2.2 Tuning curves for an ideal, a lossy and a disordered waveguide.....	22
Figure 2.3 Birefringent phasematching.....	24
Figure 2.4 Conversion efficiency vs. length for different phasematching conditions.	24
Figure 2.5 Schematic of a artificially birefringent multilayer stack	26
Figure 3.1 Schematic of QPM waveguide and resonant device	33
Figure 3.2 Bandwidth and coherence length vs. mode diameter	35
Figure 3.3 Schematic of the microcavity device.....	37
Figure 3.4 Mode field profile of fundamental and SH modes.	38
Figure 3.5 Overlap efficiency vs. waveguide width, for different core thicknesses ...	38
Figure 3.6 Bandgap size in function of waveguide width and hole radius.	39
Figure 3.7 Q for smaller and shifted hole structures.....	40
Figure 4.1 Schematic of the fabrication process.....	43
Figure 4.2 Schematics of positive and negative mask layouts.	47
Figure 4.3 Positive and negative PMMA mask after developing.....	48
Figure 4.4 SiN mask and sample at the end of the fabrication process.	51
Figure 4.5 Reflectivity measurements for the calculation of metals etch rates.....	53
Figure 4.6 Ni hard mask transferred by wetetch and by liftoff	54

Figure 4.7 Schematic of the Ti/Ni mask structure with the GaAs sacrificial layer.....	55
Figure 4.8 Waveguides fabricated with different masking techniques.....	57
Figure 4.9 SEM pictures of devices after dry-etching.....	59
Figure 4.10 Cross section of devices with different hole diameter.....	60
Figure 4.11 Pictures of devices oxidized in air.....	62
Figure 5.1 Schematic of the measurement setup.....	66
Figure 5.2 Schematic of the input and output coupling stages.....	67
Figure 5.3 Schematic of Fabry-Perot waveguide resonator.....	69
Figure 5.4 Theoretical and experimental group index vs. waveguide width.....	71
Figure 5.5 Experimental loss coefficient vs. waveguide width.....	73
Figure 5.6 Cutback method: schematic and fitting.....	74
Figure 5.7 Theoretical and experimental transmission spectra of microcavities.....	76
Figure 5.8 Resonance wavelength vs. core width for different cavity lengths.....	77
Figure 5.9 Experimental and fitted transmission spectrum of a resonance peak.....	78
Figure 5.10 Experimental Q vs. waveguide core for different cavity geometry.....	79
Figure 6.1 FH TE and SH TM electric field mode profiles of the waveguide.....	83
Figure 6.2 Phasematching wavelength vs. waveguide width.....	85
Figure 6.3 Derivative of the FH and SH index difference vs. PM wavelength.....	86
Figure 6.4 Overlap Efficiency vs. waveguide width.....	87
Figure 6.5 Bandwidth (FWHM) vs. PM wavelength.....	89
Figure 6.6 SEM picture of a birefringent waveguide.....	90
Figure 6.7 Transmission spectrum at the FH wavelength.....	93
Figure 6.8 Illustration of FDTD simulations of facet reflectivity.....	93
Figure 6.9 Facet reflectivity vs. waveguide width at different wavelengths.....	94
Figure 6.10 Propagation loss vs. waveguide width.....	95
Figure 6.11 SH tuning curves, normalized to the square of the FH output power.....	96
Figure 6.12 Experimental and simulated PM curve vs. waveguide width.....	96
Figure 6.13 SH output power normalized by the square of the FH input power.....	97
Figure 6.14 Normalized SH efficiency: experimental curve and fitting.....	98
Figure 6.15 Area and FWHM of the normalized efficiency curve.....	100
Figure 7.1 Conventional DBR mirror and simulated spectra.....	105

Figure 7.2 Band diagram of a plain waveguide and of a periodic structure.....	106
Figure 7.3 Intensity of fundamental and SH modes in the external mirror.....	107
Figure 7.4 Simulated spectra for different external mirror configurations.	108
Figure 7.5 Fundamental transmission and loss spectra for the external mirror.....	109
Figure 7.6 Schematic of the dichroic mirror.....	110
Figure 7.7 Simulated reflection and loss spectra for different tooth depths.....	111
Figure 7.8 Simulated reflection spectrum for different modulation duty cycles.....	112
Figure 7.9 SEM pictures of final dichroic mirror structure.....	114
Figure 7.10 SEM pictures of fabricated mirror	115
Figure 7.11 Cleaved section of the mirror showing the trench depth.	115
Figure 7.12 Mirror structures employing trenches and deeply etched teeth	116
Figure 7.13 Schematic of the cavity device.....	117
Figure 7.14 Cavity experimental and fitted fundamental transmission spectrum	118
Figure 7.15 Plain waveguide fundamental transmission and power spectrum.....	119
Figure 7.16 Experimental and fitted resonant peaks height and finesse	119
Figure 7.17 Fundameantl transmission spectra for different number of mirror pairs.	120
Figure 7.18 Fundamental experimental and extrapolated transmission spectrum.....	121
Figure 7.19 Extrapolated reflectivity vs. length of the main mirror part	122
Figure 7.20 Cavity experimental and simulated SHG spectrum.	124
Figure 7.21 SH Transmission of a cavity device, a reference ad a plain waveguide..	125
Figure 7.22 Experimental and simulated enhancement factor.	127
Figure 7.23 Experimental spectra for cavity devices with different cavity lengths. ..	128
Figure 7.24 Cavity enhancement vs. mirror reflectivity and cavity length.....	129
Figure 7.25 Conversion efficiency vs. mirror reflectivity and cavity length.	130
Figure 7.26 Conversion efficiency divided by L vs. R and L.	131
Figure 7.27 Cavity lifetime vs. mirror reflectivity and cavity length.	132
Figure 7.28 Phasematching bandwidth vs. length.	133
Figure 7.29 Resonant bandwidth vs. mirror reflectivity and cavity length.....	134
Figure 7.30 System bandwidth.	135
Figure 8.1 Schematic of the suspended structure.	140
Figure 8.2 Phasematching curve for different suspended geometries.....	141

Figure 8.3 Coupling tapers: throughput vs. waveguide width 142
Figure C.1 Screenshot of the user interface. 153
Figure C.2 Sample of control file..... 154

Chapter 1

Introduction

This dissertation is primarily an account of contributions in the field of guided-wave nonlinear optical devices based on the aluminum gallium arsenide (AlGaAs)-Aluminum Oxide (AlO_x or Al_xO_y) material system. This work focuses on the design and fabrication of ultra-small nonlinear devices, in which the short waveguide length is compensated by the tight confinement and the resonant enhancement. The goal is not only to realize novel wavelength converters, but also to develop a technological platform, based on selective thermal oxidation and deeply etched structures. Such technology will enable the fabrication of a wide range of passive devices, such as splitters or filters, simply by changing the geometry and not material properties. In our vision, this technological platform will be suitable for the design of complex photonic circuits and could considerably facilitate the integration with active devices. This Chapter begins by explaining the motivations behind our research and its practical application.

1.1 Motivation

Although the growth of the internet traffic has been slowing down the past few years, it is still increasing at a factor of 50-70% per year, and its growth is predicted to

continue. In order to meet this increasing demand for bandwidth, ultra-high speed optical communication systems have been developed¹, such as Dense Wavelength Division Multiplexing (DWDM) systems². This technique is very effective in exploiting the fiber bandwidth by multiplexing the data directly in the wavelength domain. Using such systems, transmission rates greater than 6 Terabits per second (Tb/s) have been achieved³.

However, a faster transmission link alone is not sufficient to significantly improve the performance of an optical network. In fact, optical cross-connects along the line are necessary to route specific information to and from different optical sub-networks. In these routers, it is often required to convert data from one wavelength to another. These converters not only have to be fast, but they also need to possess low latency, low power requirements and large tuning bandwidth. Moreover, transparency (i.e., preservation of both amplitude and phase of the input signal) is highly desirable for the future evolution of the optical network, where information could be encoded also in the phase⁴. In addition, transparency facilitates the implementation of sophisticated security techniques⁵.

Currently, most optical networks employ opto-electro-optical (OEO) converters. Fundamentally, these devices are based on the conversion of the optical signal to an electrical signal, which in turn drives a laser at the new wavelength. However, because of this intermediate electrical conversion, this class of converters is relatively slow (~tens of Gb/s), with relatively high latency (~microseconds) and limited transparency. Moreover, these systems are usually bulky, since each channel requires a different module, and power demanding.

On the other hand, all-optical, nonlinear wavelength converters could solve most of these problems. These converters are based on the nonlinear interaction of input signal with a nonlinear process, which, under certain conditions, can convert one wavelength into another. Because this conversion takes place completely in the optical domain, these processes can be extremely fast ($>10\text{Tb/s}$), with almost no latency and

(relatively) low power requirements. Moreover, this is the only class of converters which offers strict and intrinsic transparency².

Nonlinear optical converters have already been realized in a variety of materials and configurations. However, as discussed in detail in the next Section, waveguide converters, which are typically more efficient than bulk devices, require either a very large interaction length (e.g., several centimeters for periodically-poled lithium niobate waveguides) or a complex fabrication process (e.g., quasi-phaseshifted AlGaAs waveguides). Because of these reasons, integration of nonlinear wavelength converters into photonic circuits is extremely challenging.

The motivation behind our research is then the need for all-optical (nonlinear) frequency converters as efficient as the ones mentioned above, with large bandwidth, but much more compact (i.e., shorter than $\sim 100\mu\text{m}$) so that they could be more easily integrated in future optical circuits. The development of a simple and robust fabrication process will also be a primary guideline in our project.

1.2 Nonlinear wavelength converters

Before reviewing the previous research, it is useful to introduce the basic concepts of nonlinear optical frequency conversion. In a nonlinear material, the polarization generated by an electromagnetic wave cannot be simply approximated with a linear dependency on the electric field. Higher order terms must be taken into account, as follows:

$$P = \epsilon_0 \chi^{(1)} E + \epsilon_0 \chi^{(2)} E^2 + \dots \quad (1.1)$$

where ϵ_0 is the dielectric constant of vacuum and E is the electric field. The second term in the expression is responsible for the second-order nonlinear frequency conversion processes, such as Second Harmonic Generation (SHG). In this process, illustrated in Figure 1.1a, because of the dependency of the nonlinear polarization term on the square of the electric field, a pump beam at a frequency ω generates a beam at its second harmonic (SH), at frequency 2ω . Because of the simplicity of this

process, it will be used in the course of this work to characterize the nonlinear performance of our devices.

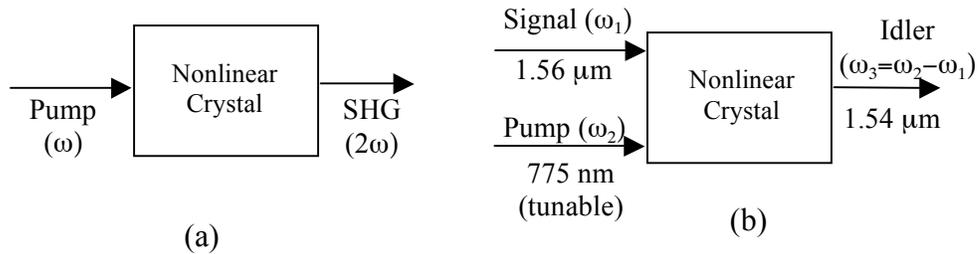


Figure 1.1 Schematic of Second Harmonic Generation (a) and Difference Frequency Generation (b).

A second process of more practical interest is difference frequency generation (DFG). In this case two input beams, called signal and pump, respectively at frequency ω_1 and ω_2 , combine to generate the output, called idler, at frequency $\omega_3 = \omega_2 - \omega_1$ (Figure 1.1b). This process is interesting because it can be used to optically shift data from one channel to another in DWDM applications. Mixing a data channel at $1.55 + \Delta \mu\text{m}$ with a 775 nm pump, would result in an idler (output channel) at $1.55 - \Delta \mu\text{m}$. By tuning the pump, the output channel can then be shifted to a different one. Furthermore, multiple channels could be processed simultaneously, within the 1.5 μm band and between the 1.3 and 1.5 μm bands. It is also possible, with multiple pumps, to perform multiple broadcasting, in which each of the input signals is converted into multiple output wavelengths⁶.

However, in order for these processes to be efficient, the interaction has to be phasematched. In the next Section we will review different phasematching techniques and material systems used in telecommunication applications.

1.3 Review of previous work

Nonlinear optics was born with the first experimental and theoretical works by, respectively, Franken⁷ and Bloembergen⁸ in the early 1960s. It has since then evolved from pure scientific interest into a wide range of commercial applications, such as spectroscopy, metrology, biomedical instrumentation and, of course, optical telecommunications. In particular, second-order nonlinear processes (see Chapter 2 for a detailed discussion of nonlinear optics principles) are extremely attractive for the realization of wavelength conversion. However, the phasematching condition (Section 2.1.2), which represents the momentum conservation law for the interacting photons, must be satisfied in order to obtain appreciable conversion efficiency.

1.3.1 Early works

The first phasematched nonlinear interaction was demonstrated in 1962⁹ using birefringent bulk crystals. Birefringent phasematching, explained in detail in Section 2.3.1, exploits the birefringence of a material to balance its dispersion. Although widely employed at the beginning, this technique has several limitations. For example, the nonlinear coefficient corresponding to the birefringent interaction may not be the highest achievable for that specific crystal (e.g., in LiNbO₃). In other cases, the nonlinear crystal of interest is simply not birefringent (e.g., GaAs).

In a second technique, called quasi-phasematching, or QPM, (Section 2.3.2) the sign of the nonlinear coefficient is periodically alternated to correct for the accumulated phase shift that is caused by dispersion. This technique, which involves complex material engineering, has been successfully employed in different crystals, including LiNbO₃¹⁰ and GaAs¹¹.

Although the invention of QPM increased significantly the nonlinear performance of these bulk devices, the efficiency was fundamentally limited by diffraction. This difficulty was overcome by the use of dielectric waveguides. Guided-wave interactions, first demonstrated by Tien in 1969¹², can be two to three orders of magnitude more efficient than bulk interactions: the efficiency, in fact, scales

quadratically with the interaction length in waveguides, whereas the dependency is only linear for optically focused bulk devices (see Section 2.2). The price to pay is an increased complexity in the fabrication process and during the optical testing, which requires more accurate alignment.

1.3.2 PPLN waveguides

Currently, periodically-poled LiNbO_3 (PPLN) waveguides¹³⁻¹⁵ are the best commercially available devices. Their normalized conversion efficiencies are greater than $150 \text{ \%}/\text{W}/\text{cm}^2$, corresponding to total conversion efficiencies exceeding $3000 \text{ \%}/\text{W}$. The waveguide core is typically defined by diffusion of ions, such as Titanium or zinc, or by exchange of lithium ions with hydrogen ions, resulting in a smooth and graded interface and propagation losses lower than $\sim 0.2 \text{ dB}/\text{cm}$. However, because of the graded profile and of the low index contrast, the achievable confinement is low, such that only weakly confining waveguides can be fabricated. As we will see in Section 2.2.1, the efficiency is inversely proportional to the area of interaction. Thus, in order to achieve high efficiency, a device length of several centimeters is required. Moreover, LiNbO_3 has other disadvantages, such as low photorefractive damage, significant clean-room processing difficulties (e.g., very rough etched surfaces) and inability to integrate with active semiconductor optoelectronic materials. Because of these issues, although LiNbO_3 waveguides can provide excellent discrete devices, enormous challenges exist in its use for integrated optical circuits.

1.3.3 AlGaAs as a nonlinear material

AlGaAs, as a nonlinear material, offers an interesting alternative for the development of novel frequency converters. First of all, its nonlinear coefficient ($d_{14} \sim 90 \text{ pm}/\text{V}$) is 4x larger than the coefficient of LiNbO_3 . It also possesses a high damage threshold ($\sim 1 \text{ J}/\text{cm}^2$), a wide transparency range (from 800 nm up to $12 \mu\text{m}$), which is useful for mid-IR applications, and a large thermal conductivity ($55 \text{ W}/\text{m}/^\circ\text{C}$), essential for high power applications. Moreover, because of the symmetry of the nonlinear tensor, devices could be made polarization insensitive¹⁶.

Another important advantage is the ability to alter the optical properties by varying the semiconductor composition during the material growth and the geometry during the fabrication process. For example, the index of refraction can be controlled by the ratio of aluminum flux to gallium flux during the MBE growth. Since AlAs is lattice-matched to GaAs, two arbitrary compositions can be grown as core and cladding, allowing large design flexibility. This adjusting “knob” is very important for active devices as well, since the absorption edge of the material can be shifted.

Two important properties are the ability to create a stable oxide and the strong dependence of the thermal oxidation rate on the aluminum concentration. By carefully adjusting the composition, it is possible to selectively oxidize an AlGaAs layer (cladding) vs. another (core). The discovery of this property in 1990¹⁷ enabled the realization of current apertures for active device structures, such as vertical cavity surface emitting lasers¹⁸ (VCSELs), the fabrication of strongly modulated distributed bragg reflectors¹⁹ and, most importantly for this work, high index contrast structures^{20, 21} (the index of AlOx at 1.55 μ m is \sim 1.6 vs. the index of AlGaAs \sim 3), which are essential for achieving tight confinement in waveguides.

Finally, a well-developed fabrication technology and the possibility of integrating passive and active devices on the same substrate make the AlGaAs-AlOx material system an excellent choice for nonlinear applications. However, there is a challenge in using this material: the implementation of an effective phasematching technique.

1.3.4 QPM AlGaAs waveguides

Because AlGaAs is not birefringent, Quasi-phasematching has been the dominant phasematching technique. In the most recently developed implementation of this technique, called orientation-patterning²² QPM, the sign of the nonlinear coefficient is inverted by changing the crystal orientation by 90 degrees during the growth. A template with periodic domain inversion is fabricated by Molecular Beam Epitaxy (MBE), optical photolithography and selective wet-etching. A second epitaxial growth is then required to create the waveguide structure on the template. This class

of devices has achieved total efficiencies as high as $\sim 50\%/W$ with only 1-cm-long waveguides²³. The normalized efficiency is 4x larger than PPLN waveguides, due essentially to the larger nonlinear coefficient. Although the most serious problem with these devices (i.e., the high loss due to the template corrugation) has been almost entirely solved²³, the fabrication complexity remains elevated, because of the multiple epitaxial growths. Moreover, the current waveguides are weakly-confining; thus, any attempt to improve the efficiency by increasing the confinement would also increase the loss due to the residual template corrugation.

1.3.5 Artificial birefringence

Another technique to achieve phasematching in isotropic materials, such as AlGaAs, is artificial birefringence. This technique was originally proposed in 1975 by Van der Ziel²⁴, but, because of the lack of material systems with sufficiently high index contrast, an experimental demonstration had to wait until 1998²⁵. The discovery of AlGaAs selective thermal oxidation could provide, in fact, enough index contrast in a suitable nonlinear material system.

This technique consists in inducing birefringent behavior inside an isotropic, non-birefringent material simply by geometry. This effect can be obtained by a multilayer stack of two isotropic materials with different refractive indices, as explained in detail in Section 2.3.3. If the index contrast is large enough, the induced birefringence can balance the material dispersion and the phasematching condition is satisfied.

The first demonstration of this technique in 1998 by Fiore et. al.^{25,26} involved DFG in weakly confining waveguides between $1\mu\text{m}$ and $1.3\mu\text{m}$ radiation to generate an idler at $\sim 5\mu\text{m}$. SHG at telecommunication wavelengths ($1.6\mu\text{m}$)²⁷ was also demonstrated. A conversion efficiency of $\sim 0.12\%/W$ for a 1.7mm-long waveguide (normalized efficiency of $4.5\%/W/\text{cm}^2$) was achieved, limited by the huge propagation loss ($>200\text{dB}/\text{mm}$) at the SH. These losses were attributed to absorption in the AlOx caused by the oxidation process.

As a side, we should note that pure AlAs (instead of, e.g., $\text{Al}_{0.93}\text{Ga}_{0.07}\text{As}$, which we used in this work) was employed as the material to be selectively oxidized; this fact could have significantly contributed to the high SH loss (more on this topic in Sections 6.3.3 and 8.2.1). SHG at $2\mu\text{m}$ ²⁸ was also demonstrated with a normalized conversion efficiency greater than $1000\%/W/\text{cm}^2$. Here the longer wavelength is the reason for both larger efficiency and lower losses. However, in all these cases, weakly-confining waveguides (i.e., with an unoxidized AlGaAs cladding) were employed, resulting in relatively low normalized conversion efficiency, and device lengths greater than $\sim 1\text{mm}$.

1.3.6 Resonant enhancement of nonlinear processes

Since the experiments by Ashkin et al. in 1966²⁹, resonant cavities have been employed to enhance nonlinear optical processes. Originally macroscopic cavities with bulk crystals as nonlinear elements were used for optical parametric oscillators³⁰,³¹ and other frequency conversion processes³². More recently, enhanced nonlinear effects have been studied in a variety of monolithic structures. A vast amount of research has been done in this field. Examples include regular Fabry-Perot resonators, and other more complicated structures, such as fiber ring cavities³³, photonic crystal coupled resonator optical waveguides³⁴ and microring resonators³⁵. Here we will limit our review to two kinds of devices for SHG, which are of interest to this work: 1D singly and doubly resonant Fabry-Perot cavities and resonantly enhanced, surface emitting waveguides. We will finally discuss collinear, resonant devices for SHG and for optical parametric processes.

Since there is no radiation loss in 1D systems, high-finesse cavities can be more easily realized. Simple Fabry-Perot structures resonating either at the fundamental³⁶ or the SH³⁷ wavelength have been realized. However, the enhancement and the efficiency were limited because the cavity was usually short (i.e. shorter than the coherence length, to achieve phasematching) and the interaction was not waveguided. Doubly resonant 1D monolithic cavities have also been studied, both theoretically³⁸,³⁹ and experimentally⁴⁰. Although theoretically much larger

enhancements can be obtained, matching the phase conditions for both fundamental and SH fields, is extremely difficult.

A second area of importance for this work is resonantly enhanced surface emitting waveguides. In these devices, first demonstrated in 1993 by Lodenkamper et al.⁴¹, two counter-propagating modes at the fundamental wavelength generate SH radiating vertically outside of the waveguide. By sandwiching the waveguide between two DBR mirrors, the SH can be resonated. In this work, two-hundred-fold enhancement of the conversion efficiency was demonstrated. Although improvements can be made to this design (i.e., by adding a QPM structure⁴²), two fundamental limitations remain: (i) the SH is generated along the entire length of the waveguide and radiated vertically into free space, so that efficient collection of the generated signal is quite difficult; (ii) the interaction length is limited to the thickness of the waveguide, so that the efficiency is usually much lower than in collinear devices.

Finally, a third category of interesting devices are collinear, resonant, waveguide devices. These devices have been realized in LiNbO_3 and employed multilayer dielectric stacks deposited on the waveguide facets as mirrors. In this case, since the typical waveguide mode size is much larger than the wavelength, the interaction with the mirrors can be considered one-dimensional. Both SHG⁴³ and optical parametric processes have been demonstrated. In optical parametric oscillators (OPO), a pump at the shortest wavelength, with opportune feedback, generates a signal and idler at longer wavelengths. The feedback typically consists of a cavity resonant at either signal or idler (singly resonant OPO, or SRO) or at both (doubly resonant OPO, or DRO). Both SROs^{44, 45} and DROs^{46, 47} have been realized. However, these devices cannot be easily integrated into a photonic circuit, since the nature of the mirrors requires the presence of a flat waveguide facet. Reducing the device length by increasing the lateral confinement would also be impossible for SHG devices: the 1D-mirror approximation would in fact break down, leading to large SH loss, as we will discuss in Chapter 7.

No collinear, resonant, waveguide devices, with mirrors integrated directly in the waveguide, have been demonstrated yet, especially for the case of SHG (i.e., with the cavity resonant at the longer wavelength). The fundamental reason lies in the difficulty to fabricate a dichroic mirror which can resonate the fundamental wave, while being transparent at the SH. In this work we will discuss this problem and demonstrate a novel mirror design suitable to this task, which can be easily integrated in any waveguide section (i.e., without requiring the presence of the facets).

1.4 Overview of the thesis

This research was focused on the design and fabrication of AlGaAs-AlO_x-based, highly-efficient, ultra-short (<100- μ m long), nonlinear wavelength converters. With the goal of increasing the efficiency and reducing the device size, we focused on resonant cavities embedded in waveguides, where the generation is collinear to the propagation direction (i.e., as opposed to surface emitting devices). In both the design and the process development stages, one of the primary objectives was the realization of a device tolerant to fabrication imperfections and to variations of the process parameters.

This dissertation is organized as follows: Chapter 2 describes the theoretical background necessary to understand the rest of the thesis. This includes the fundamentals of nonlinear optics in waveguides and resonators and an overview of the physics behind the different phasematching techniques. Chapter 3 contains back-of-the-envelope calculations and scaling laws of efficiency and bandwidth for ultra-small cavity devices. We also discuss the design of our first generation device: the photonic crystal microcavity. Chapter 4 explains in detail the fabrication process and how it was developed, including particularly instructive “failed” attempts. Chapter 5 describes the optical testing setup and the techniques used for the linear characterization of plain waveguides and microcavities. It also discusses the problems of the microcavity device and the reasons for its extremely low conversion efficiency. Chapter 6 discusses the design of the second generation devices: the tightly-confining,

birefringently-phasematched waveguides. Because the fabrication process is essentially the same as for the microcavities, only minor fabrication issues are addressed here. Finally, the results of both linear and nonlinear characterization of the waveguides are presented. Chapter 7 treats of the design of a novel dichroic mirror employed in our tightly-confining resonant cavities. Fabrication issues and the characterization of mirror and cavities are discussed. The results of the SHG experiments are also reported. Chapter 8 contains a summary of the contributions of this research and future perspectives.

Chapter 2

Theoretical background

In this Chapter we review the background material necessary to understand the work of this thesis. We start with the basic concepts of nonlinear optics and guided waves. We then review the phasematching techniques developed in the past, such as quasiphasematching, and recently discovered ones, such as modal phasematching. Special attention is given to form birefringence, which is the main approach used in this work. Finally, the theory of nonlinear resonant cavities is presented.

2.1 Fundamentals of nonlinear optics

Nonlinear optics is the study of the nonlinear polarization response of a material to an intense electric field. The induced polarization in a medium written in Taylor series expansion, is given by:

$$P_i = \sum_j \epsilon_0 \chi_{ij}^{(1)} E_j + \sum_{j,k} \epsilon_0 \chi_{ijk}^{(2)} E_j E_k + \sum_{j,k,l} \epsilon_0 \chi_{ijkl}^{(3)} E_j E_k E_l \dots, \quad (2.1)$$

where P_i is the i -th component of the induced polarization, $\chi^{(2)}$ and $\chi^{(3)}$ are the 2nd and 3rd order nonlinear optical susceptibility tensors, and E_i is the i -th component of the electric field. In general, E_j , E_k and E_l can have different frequencies, so that the second and third terms of the above expression are responsible for frequency

conversion processes, such as second harmonic generation (SHG), optical parametric amplification (OPA) and third harmonic generation (THG).

The physics of the process can be explained qualitatively by considering the material as a collection of harmonic oscillators⁴⁸. For low applied fields, electrons oscillate at the applied frequency in a potential that can be approximated as a parabola. For stronger fields the parabolic approximation is no longer valid and more terms need to be added to the potential. Thus, nonlinear effects will occur and components at higher harmonics will appear. If more than one frequency is applied, mixing of these frequencies will result in polarization components at sum- and difference-frequencies. This polarization will act as an electromagnetic radiation source at those frequencies. Generating new frequencies or amplifying the starting frequencies is usually the goal of nonlinear optics.

2.1.1 Coupled mode theory

A quantitative analysis of nonlinear processes is provided by the coupled-mode theory. In the following, we will discuss three-wave mixing processes, such as DFG or SHG, as they are of direct interest for this work. Analysis of third-order nonlinear processes can be found in the literature (i.e., in the book of Saleh & Teich⁴⁸). We start from the wave equation in a homogenous and isotropic medium:

$$\nabla^2 \vec{E} - \vec{\nabla}(\vec{\nabla} \cdot \vec{E}) - \frac{1}{c^2} \frac{\partial^2 \vec{E}}{\partial t^2} = \mu_0 \frac{\partial^2 \vec{P}}{\partial t^2}. \quad (2.2)$$

By writing the polarization as $\vec{P} = \epsilon_0 \chi^{(1)} \vec{E} + \vec{P}_{NL}$, where P_{NL} is the nonlinear term of the polarization, and substituting it in the previous equation, we obtain the basic wave equation for the propagation in a second-order nonlinear medium:

$$\nabla^2 \vec{E} - \frac{n^2}{c^2} \frac{\partial^2 \vec{E}}{\partial t^2} = \mu_0 \frac{\partial^2 \vec{P}_{NL}}{\partial t^2}, \quad (2.3)$$

where $n^2 = \epsilon = \epsilon_0(1 + \chi^{(1)})$. Then, we write the field E as a superposition of three plane waves of frequencies ω_1 , ω_2 and ω_3 propagating in the z direction, as follows:

(2.4)

$$E(z, t) = \sum_{i=1,2,3} A_i(z) \exp(j\omega_i t - jk_i z),$$

where A_i is the amplitude and k_i the wavevector. By substituting Equation (2.4) into Equation (2.3), we obtain a set of three Helmholtz equations with sources:

$$(\nabla^2 + k_i^2)E_i = -S_i(\omega_i), \quad (2.5)$$

where S_i is the amplitude of the nonlinear source (right-hand side of Equation (2.3)) at frequency ω_i . For the waves to be coupled by the medium, the frequencies of the plane waves have to be commensurate, otherwise the sources S_i vanish. Assume for example that one frequency is the sum of the other two; this translates into the frequency matching condition:

$$\omega_3 = \omega_1 + \omega_2, \quad (2.6)$$

which simply represents the energy conservation condition. The sources S_i can then be simplified in the following way:

$$\begin{aligned} S_1(\omega_1) &= 2\mu_0 d_{eff} E_3 E_2^* \\ S_2(\omega_2) &= 2\mu_0 d_{eff} E_3 E_1^* \\ S_3(\omega_3) &= 2\mu_0 d_{eff} E_1 E_2 \end{aligned} \quad (2.7)$$

where d_{eff} , which is proportional to $\chi^{(2)}$, is the effective nonlinear coefficient adjusted to take into account the specific crystal orientation. If the interaction is weak (i.e., A_i varies slowly within a distance of a wavelength), we can employ the slowly varying envelope approximation⁴⁸, which consists in neglecting $d^2 A_i / dz^2$ relative to $k_i d A_i / dz$. Using this approximation and Equations (2.6) and (2.7), we can simplify Equations (2.5) as:

$$\begin{aligned}
\frac{dA_1}{dz} &= -i \frac{\omega_1 d_{eff}}{n_1 c} A_3 A_2^* \exp(-i\Delta k z) \\
\frac{dA_2}{dz} &= -i \frac{\omega_1 d_{eff}}{n_2 c} A_3 A_1^* \exp(-i\Delta k z) , \\
\frac{dA_3}{dz} &= -i \frac{\omega_1 d_{eff}}{n_3 c} A_1 A_2 \exp(i\Delta k z)
\end{aligned} \tag{2.8}$$

where $\Delta k = 2k_3 - k_2 + k_1 = 2\pi(n_3 / \lambda_3 - n_2 / \lambda_2 - n_1 / \lambda_1)$ is the phase mismatch vector and n_i and λ_i are the refractive indices and the wavelengths in vacuum at frequencies ω_i . These equations provide a quantitative description of the frequency conversion process. The condition $\Delta k=0$, called phasematching condition, is necessary in order to have appreciable buildup of the amplitudes A_i over a finite length and simply represents the momentum conservation law.

Note finally that the equations describing the degenerate case of SHG cannot be obtained directly from Equation (2.8), but must be derived from the initial Equations (2.5) by substituting $E_1=E_2$, resulting in:

$$\begin{aligned}
\frac{dA_1}{dz} &= -i \frac{\omega_1 d_{eff}}{n_1 c} A_3 A_1^* \exp(-i\Delta k z) \\
\frac{dA_3}{dz} &= -\frac{i}{2} \frac{\omega_1 d_{eff}}{n_3 c} A_1^2 \exp(i\Delta k z)
\end{aligned} , \tag{2.9}$$

where there is an extra-factor of 2 dividing the right hand side of the equation for A_3 .

2.1.2 Conversion efficiency

We now consider the case of SHG in the non-depletion regime, which will be the case in our experiments. In this regime the amount of field converted to the second harmonic (A_3) is small compared to the fundamental initial field (A_1). We can then assume A_1 constant, and if the nonlinear medium is homogeneous (n_i constant), the second of Equations (2.9) is easily integrated. By writing the field intensities as $I_i = (\epsilon_0 n c / 2) |A_i|^2$, the SH conversion efficiency can be expressed as:

(2.10)

$$\eta = \frac{I_{2\omega}}{I_{\omega}} = \left(\frac{8\pi^2 d_{eff}^2}{n_{2\omega} n_{\omega}^2 c \epsilon_0 \lambda_{\omega}^2} \right) I_{\omega} L^2 \text{sinc}^2 \left(\frac{\Delta k L}{2} \right).$$

where $\text{sinc}(x) = \sin(x)/x$, L is the interaction length, λ_{ω} is the fundamental wavelength and n_{ω} and $n_{2\omega}$ are the refractive indices at the FH and SH¹, respectively. The first factor in the expression is a function of the wavelength and material properties only. The SH intensity is proportional to the square of the fundamental intensity, as expected from a nonlinear process. This means, as we will see in more detail later, that the higher the fundamental intensity, the more efficient the conversion process. The last factor is the phasematching term. If $\Delta k = 0$, the *sinc* factor is equal to one, and the SH intensity increases quadratically with the length of interaction. If $\Delta k \neq 0$, the SH intensity oscillates with a sine squared behavior. The shortest length at which the efficiency is maximum, is reached when the *sinc* argument is equal to $\pi/2$. This characteristic length, called the “coherence length”, is equal to:

$$L_c = \frac{\lambda_{\omega}}{4(n_{2\omega} - n_{\omega})} \quad (2.11)$$

and depends on the wavelength and on the optical properties of the material. The concepts of phasematching and coherence length will be further discussed in Section 2.3.

2.2 Nonlinear optics in waveguides

In the previous section, we analyzed the ideal case of plane waves. In the real case of a bulk crystal, the overlap between the nonlinear polarization and the generated SH must be taken into account. To obtain high field intensities a tight focus is required. However, the smaller the spot-size, the stronger the effect of diffraction, so that the beam diverges quickly and the intensity decreases. Waveguides are an excellent solution to this problem; in fact, because of total internal reflection, the field is confined in the waveguide and diffraction does not constitute a limitation. This fact

¹ From now on the fundamental and second harmonic will be denoted as FH and SH, respectively.

allows a virtually infinite interaction length with high intensity fields. To estimate the improvement in efficiency, we can calculate the ratio between the waveguide and bulk conversion efficiencies, assuming for the latter a gaussian beam with optimal confocal focusing⁴⁹, as follows:

$$\frac{\eta_{\text{waveguide}}}{\eta_{\text{bulk}}} = \frac{L\lambda_{\omega}}{2n_{\omega}A_{\text{wg}}}, \quad (2.12)$$

where the effective waveguide area A_{wg} , discussed later in this Section, is usually of the order of the cross-sectional mode area. Typically, this ratio is much greater than one, which translates in efficiencies many orders of magnitude larger than bulk efficiencies. This reason justifies the efforts to fabricate waveguide devices, although a more complex design and fabrication are required.

2.2.1 Guided wave coupled mode theory

A description of frequency conversion in waveguides can be derived in a way similar to the one described in Section 2.1.1, where we have to include the effect of spatial distribution of the waveguide modes. A general coupled mode theory for waveguides has been developed and is available in the literature^{50,51}. Here we will just present the results for the case of SHG. We assume the existence of a complete set of orthogonal modes such that an arbitrary transverse field can be expressed in terms of these modes, as follows:

$$E_t = \sum_{\nu} A_{\nu}(z)\bar{E}_{\nu}(x, y)\exp(-ik_{\nu}z + i\omega t) \quad (2.13)$$

where we assumed an axial dependence of the form $\exp(-i\beta_{\nu}z)$ and a time dependence of the form $\exp(i\omega t)$. The modes \bar{E}_{ν} are normalized according to the following relationship:

$$\iint \bar{E}_{\nu}\bar{E}_{\mu}^* dx dy = \delta_{\nu\mu}. \quad (2.14)$$

Substituting Equations (2.13) into Maxwell's equations, and solving for A_{ν} , we obtain the coupled mode equation for SHG in waveguides:

$$\frac{dA_{2\omega}}{dz} = -\frac{i}{2} \frac{\omega d_{eff} J}{n_3 c} A_{\omega}^2 \exp(i\Delta k z), \quad (2.15)$$

$$\text{where } J = \iint \bar{d}(x, y) \bar{E}_{\omega}^2(x, y) \bar{E}_{2\omega}^*(x, y) dx dy \quad (2.16)$$

is the overlap integral. $\bar{d}(x, y)$ takes into account eventual non-uniformities of the nonlinear coefficient over the cross-section and is defined as $d_{eff}(x, y) = d_{eff} \bar{d}(x, y)$.

The quantity $A_{eff} = |J|^{-2}$, called the effective area, is a measure of the waveguide confinement and of how well the FH and SH modes overlap with each other. Typical values for the effective area are 20-40 μm^2 for weakly confining waveguides²³ and as small as 1-2 μm^2 for tightly confining waveguides, as in this work. From these figures we can appreciate the value of tight confinement, as it can potentially increase the efficiency by a factor 10-40, compared to weakly confining waveguides.

Solving the previous equation we obtain the conversion efficiency for a waveguide:

$$\eta' = \frac{P_{2\omega}}{P_{\omega}^2} = \eta_{norm} L^2 \text{sinc}^2\left(\frac{\Delta k L}{2}\right), \quad (2.17)$$

$$\text{where } \eta_{norm} = \frac{8\pi^2 d_{eff}^2}{n_{2\omega} n_{\omega}^2 c \epsilon_0 \lambda_{\omega}^2 A_{eff}}. \quad (2.18)$$

This formula, typically expressed in units of %/W, is similar to Equation (2.10) but has the advantage of being independent of the input power. η_{norm} is called normalized conversion efficiency and is expressed in units of %/W/cm². This quantity is very useful to characterize the intrinsic efficiency of a waveguide structure, since it depends only on material properties and on the cross-sectional geometry (and not on the waveguide length, the propagation losses, or the phasematching condition).

2.2.2 Effect of loss

So far we have considered ideal structures, but in real devices other effects must be taken into account in order to properly describe the nonlinear process. One crucial effect is loss from absorption or scattering due to roughness at the waveguide interfaces or defects in the material. We can describe the combined effect of the different sources of loss using a single, constant power attenuation loss coefficient, α , at each frequency. The modified equation for SHG in the non-depletion regime is then:

$$\frac{dA_{2\omega}}{dz} + \frac{\alpha_{2\omega}}{2} A_{2\omega} = -\frac{i \omega d_{eff} J}{2 n_{2\omega} c} A_{\omega}^2 \exp(i\Delta k z) \quad (2.19)$$

where $A_{\omega} = A_{\omega}(0) \exp(-\alpha_{\omega} z / 2)$. By solving this equation, we obtain the SH amplitude $A_{2\omega}$:

$$A_{2\omega}(z) = i \frac{\omega d_{eff} J}{2 n_{2\omega} c} A_{\omega}^2(0) \frac{\exp(i\Delta k z - \alpha_{\omega} z) - \exp(-\alpha_{2\omega} z / 2)}{i\Delta k - (\alpha_{\omega} - \alpha_{2\omega} / 2)} \quad (2.20)$$

As before, we can write the expression for the conversion efficiency for a waveguide of length L , where we use the combined loss parameter $\Delta\alpha = \alpha_{2\omega} / 2 - \alpha_{\omega}$:

$$\eta' = \frac{P_{2\omega}}{P_{\omega}^2} = \eta_{norm} L^2 \exp(-\alpha_{2\omega} L) \left| \frac{\exp(i\Delta k L + \Delta\alpha L) - 1}{i\Delta k L + \Delta\alpha L} \right|^2 \quad (2.21)$$

Losses have the double effect of decreasing the efficiency and broadening the tuning curve (i.e., the generated SH power vs. wavelength). Figure 2.1a shows the efficiency vs. length for different values of the loss coefficients. We see that for loss at only one wavelength the efficiency saturates, whereas for loss at both wavelengths, an optimal length exists, after which the efficiency decreases.

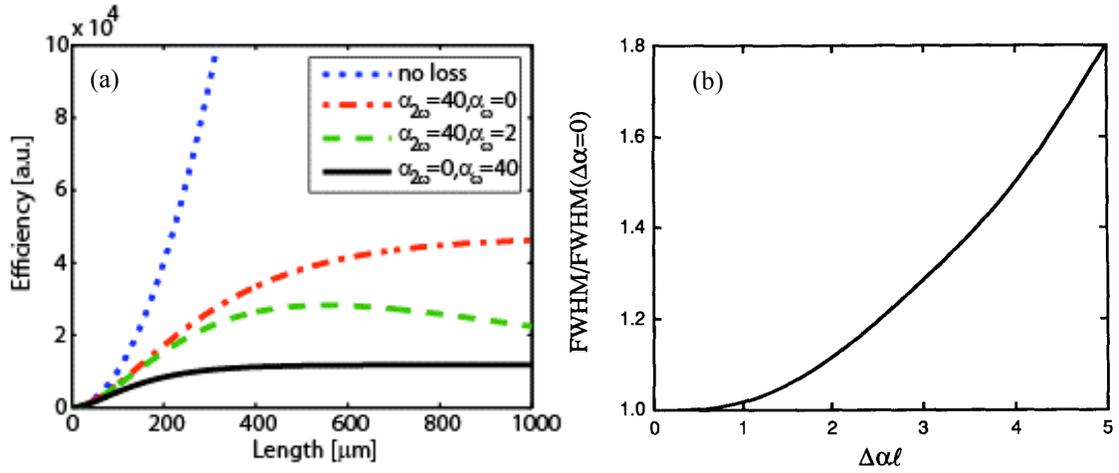


Figure 2.1 (a) Efficiency vs. length for different values of α_{ω} and $\alpha_{2\omega}$ (in dB/mm) and (b) increase in FWHM vs. $\Delta\alpha = \alpha_{2\omega}/2 - \alpha_{\omega}$. This figure was adapted from Bortz et al.⁵².

Figure 2.1b (adapted from Bortz et al.⁵²) shows the increase in Full-Width-Half-Maximum (FWHM) of the bandwidth versus the combined loss parameter $\Delta\alpha$. It is interesting to notice that the FWHM depends on the difference between α_{ω} and $\alpha_{2\omega}$, and that the broadening is significant only if the total loss at one wavelength (typically $\alpha_{2\omega}L$, since the loss due to roughness increases steeply with the frequency) is much larger than the other.

2.2.3 Effect of disorder

A second important issue in real waveguides is disorder. In an ideal birefringently phasematched waveguide (see Chapter 6), designed for SHG at wavelengths λ_{ω} and $\lambda_{2\omega}$, the effective indices n_{ω} and $n_{2\omega}$ are equal and constant. However, in a real device the presence of roughness and fabrication imperfections may alter the geometry or the optical properties of the waveguide along the propagation direction z . In this case, the effective indices and the phase mismatch vector Δk become a function of z . Equation (2.19) must be modified to take into account the phase variations due to disorder as:

$$\frac{dA_{2\omega}}{dz} + \frac{\alpha_{2\omega}}{2} A_{2\omega} = -\frac{i \omega d_{\text{eff}} J}{2 n_{2\omega} c} A_{\omega}^2 \exp(i\Delta\phi(z)), \quad (2.22)$$

$$\text{where } \Delta\phi(z) = \int_0^z \Delta k(z) dz = \int_0^z \frac{2\pi}{\lambda_\omega} (n_{2\omega}(z) - n_\omega(z)) dz .$$

It is important to notice that the effect of uncorrelated, high-spatial-frequency disorder (i.e., of period shorter than the coherence length) on the phase $\Delta\phi$ is relatively small, since the phase variations tend to average out. Thus, disorder due to typical sidewall roughness in dry-etched waveguides (see Section 4.4) can usually be neglected. More problematic is a *drift* of Δk , due for example to non-uniformities of the core layer thickness (see Section 6.3.2). In fact, this drift could result in a severe phase variation which significantly decreases the peak efficiency.

Usually, Equation (2.22) must be solved numerically and the solution depends on the functional form of Δk . Figure 2.2 shows the effects of loss and disorder, where for the disorder, we assume a linear variation of Δk along the entire length of the waveguide. We note also that disorder simply alters the phase of the SH spectrum; thus, although the peak efficiency and the shape of the curve might be severely affected, the total amount of generated SH (i.e., the area under the curve) remains constant. This property of the disorder will be very useful to characterize artificially birefringent waveguides (Section 6.3.2).

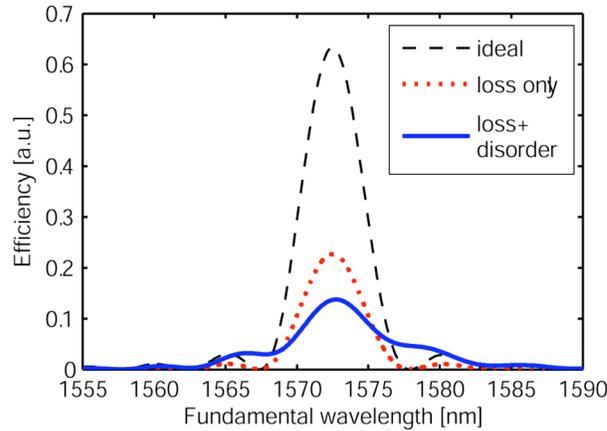


Figure 2.2 Tuning curve for an ideal(a), a lossy (b) and a lossy and disordered waveguide.

2.3 Phasematching techniques

As mentioned in Section 2.1.2, phasematching is the condition necessary to obtain non-negligible conversion efficiency. It can be expressed by the condition $\Delta k=0$, which for SHG corresponds to:

$$\Delta k = k_{2\omega} - 2k_{\omega} = \frac{4\pi}{\lambda_{\omega}}(n_{2\omega} - n_{\omega}) = 0. \quad (2.23)$$

In an ideal, non-dispersive material, energy conservation (Equation (2.6)) implies momentum conservation. The interaction is automatically phasematched and the conversion efficiency increases quadratically with length. However, in a real, dispersive material, the FH and SH waves travel at different velocities and after a certain length, called the coherence length (Equation (2.11)), the SH wave has accumulated a π phase difference with the nonlinear source (i.e., FH wave). The SH generated over the next coherence length is out of phase with that of the previous, reducing the total amount of SH. This results in an oscillatory behavior of the conversion efficiency vs. length, as shown in Figure 2.4b. In this Section we will review different approaches to solve the problem of phasematching in real devices.

2.3.1 Birefringence

The oldest and most straightforward technique for achieving phasematching consists in taking advantage of birefringent materials, in which different polarizations propagate with different indices, as shown in Figure 2.3. By choosing a geometry such that the higher frequency wave sees the lower of the two allowed indices (n_e in the figure), it is sometimes possible to fulfill the ideal phasematching conditions. The problem with this technique, however, is that often the material of interest (such as GaAs) is not birefringent. Alternatively, the nonlinear tensor $\chi^{(2)}$ may not allow for efficient conversion between two fields polarized in different directions (such as in LiNbO₃¹³).

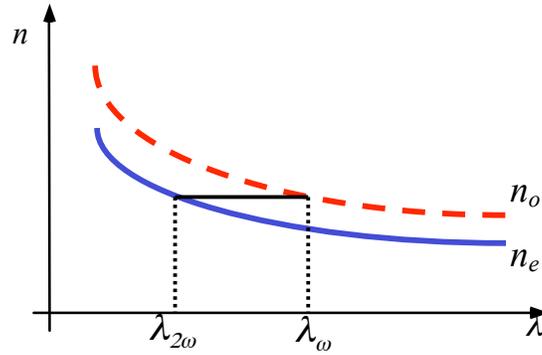


Figure 2.3 Birefringent phase matching in a negative birefringent material.

2.3.2 Quasi-phasematching

Another solution to the phasematching problem is known as quasi-phasematching (QPM)^{8, 53}. If the material can be engineered so that the sign of the nonlinear susceptibility is flipped every coherence length, L_c , an extra π phase shift is artificially added, so that the power continues to flow into the generated frequency in a quasi-quadratic way. Mathematically, this behavior is described by introducing the k-vector K_g (Figure 2.4a). K_g represents the Fourier component due to the grating periodicity, which is generated by the alternating sign of $\chi^{(2)}$. The phasematching condition then becomes:

$$\Delta k = k_{2\omega} - 2k_\omega - K_g = 0. \quad (2.24)$$

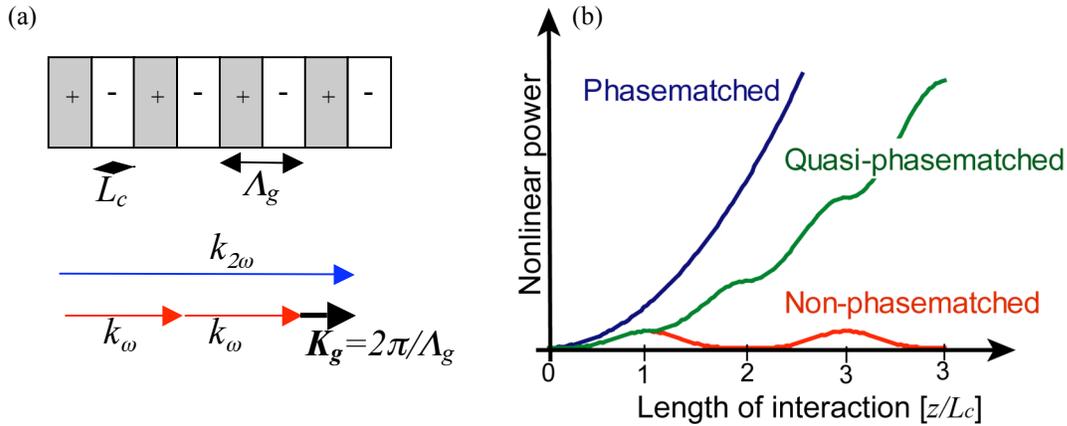


Figure 2.4 (a) Schematic of quasi-phasematching. (b) Conversion efficiency vs. length, in units of coherence lengths L_c , for different phasematching conditions.

Figure 2.4b shows the conversion efficiency vs. length of interaction for the cases of perfectly phasematched, non-phasematched and quasi-phasematched interactions. The coherence length, L_c , strongly depends on the wavelength of operation and the type of nonlinear process. For GaAs, it can range from several hundred microns for generation of terahertz radiation, to $\sim 30\mu\text{m}$ for SHG in the mid-IR region (FH at $4\mu\text{m}$), to $\sim 2\mu\text{m}$ for SHG at $1.55\mu\text{m}$ (generating 775nm). For waveguide devices, the coherence length is also function of the confinement, and can drop from $\sim 2\mu\text{m}$ to below $1\mu\text{m}$ for tightly confining waveguides. As we have seen in Section 2.1.2, high confinement would be desirable for high efficiency, however, the fabrication of very short (e.g., sub-micron) period gratings is very challenging. In the next section, we will discuss a technique that can achieve high confinement, avoiding the complication of fabricating submicron QPM gratings.

2.3.3 Form birefringence

As mentioned earlier, one of the disadvantages of birefringent phasematching is that it cannot be used with isotropic materials, such as GaAs. However, artificial birefringence can be induced in a multilayer structure of two alternating isotropic materials if the propagation direction is *parallel* to the layers (Figure 2.5a). This effect can be intuitively explained by looking at the electric field pattern inside the multilayer stack as illustrated in Figure 2.5b, adapted from Fiore et al.²⁵. The effective index of a propagating mode inside a composite layer increases proportionally with the fraction of field in the high- ϵ material. For a TE polarized wave, the electric field is parallel to the layers and thus continuous at the interfaces. Conversely, for a TM polarized wave, the field is discontinuous at the interfaces, and tends to be more concentrated in the low- ϵ material, thus decreasing the effective index.

Mathematically, for two layers of dielectric constants ϵ_1 and ϵ_2 , and thicknesses d_1 and d_2 , the effective dielectric constants can be expressed, to the lowest order as²⁴:

$$(2.25)$$

$$\varepsilon_{TE} = (\varepsilon_1 d_1 + \varepsilon_2 d_2) / (d_1 + d_2)$$

$$\varepsilon_{TM} = (d_1 + d_2) / (d_1 / \varepsilon_1 + d_2 / \varepsilon_2)$$

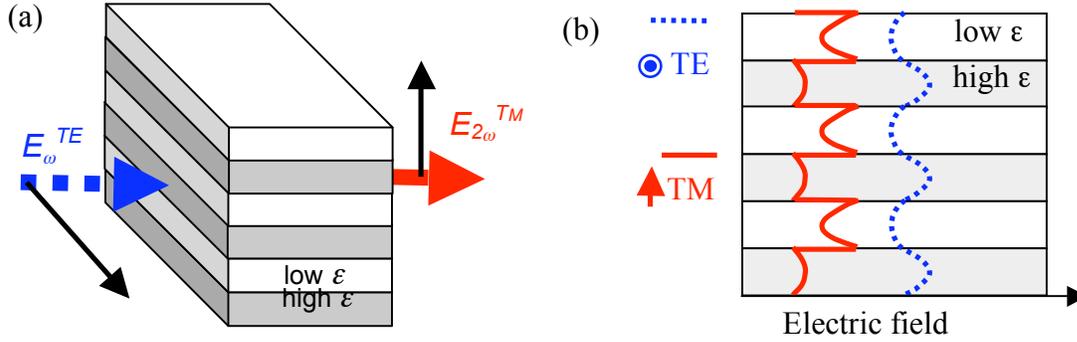


Figure 2.5 (a) Schematic of an artificially birefringent multilayer stack (b) Schematic of electric field distribution inside the stack

From the previous expressions, we can see that the amount of birefringence depends on the index contrast between the two materials and this technique can be used for phasematching only if the amount of artificial (or “form”) birefringence is enough to compensate for the material dispersion. After initial attempts to use GaAs and AlAs as alternating materials, this technique was abandoned since the index contrast was not sufficiently high. Recently, advances in GaAs processing technology, allowed the fabrication of GaAs ($n \sim 3.5$) /aluminum oxide ($n \sim 1.6$) structures, and artificially birefringent phasematching was successfully demonstrated in bulk and weakly confining waveguides (See Section 1.3.5). This approach is promising for GaAs waveguides since the fabrication process does not involve multiple crystal growths and the realization of QPM gratings. Moreover, it has the potential of achieving efficiencies much higher than QPM structures⁵⁴. Implementing this approach in tightly confining waveguides is one the main goals of this work, and will be discussed in detail in Chapter 4.

2.3.4 Modal phasematching

We finally mention, for completeness, the technique of modal phase matching (MPM). This technique has been of limited interest because of the poor spatial overlap between the interacting modes. The expected normalized conversion efficiency in optimized structures is in fact lower, by a factor 20, than in birefringent semiconductor waveguides and comparable with periodically poled LiNbO₃⁵⁵.

The concept is simple: in standard (i.e. GaAs/AlGaAs) waveguides, the 1st even mode at the FH wavelength cannot be phasematched to the 1st even mode at the SH; however, under certain conditions, it can be phasematched to the 2nd even mode. Unfortunately, as mentioned earlier, the serious disadvantage is the poor overlap efficiency of the two modes. However, this technique has recently gained interest because it does not require a thermal oxidation, making the integration with active devices easier⁵⁶.

2.4 Resonant cavities

In this Section we will discuss resonant cavities as a means for enhancing nonlinear optical processes. Cavities can increase the peak efficiency when the maximum length is limited by loss (as explained in Section 2.2.2), or if a short device length is desirable (i.e., for photonic integration). In Chapter 3, the concept of micro-cavities shorter than the coherence length (and thus inherently phasematched) will be discussed, in which the high cavity enhancement could compensate for the short length. The disadvantage of cavities, apart from the higher fabrication complexity, lies in the reduction of the optical bandwidth (Section 7.6).

2.4.1 Principles of resonators

In this Section we will review the basic principles of resonators. Consider a Fabry-Perot cavity of length L formed by two mirrors, identical for simplicity, of field reflectance r and field transmittance t . The cavity medium has refractive index n and power attenuation loss α . If E_i is the incident field of a monochromatic wave of

wavelength λ , E_o the output field and E_c the intra-cavity field after the input mirror, the self-consistency condition can be written as:

$$tE_i + r^2 \exp(-i2kL - i2\phi_m - \alpha L) E_c(0) = E_c(0), \quad (2.26)$$

where $k=2\pi n/\lambda=2\pi n\nu/c$ and ϕ_m is the mirror phase shift. Solving for E_c , and taking the square of the absolute value, we obtain the intensity enhancement of the cavity:

$$\frac{I_c}{I_i} = \frac{|E_c|^2}{|E_i|^2} = \frac{T}{[(1 - R \exp(-2\alpha L))^2 - 4R \exp(-2\alpha L) \sin^2(kL + \phi_m)]^2} \quad (2.27)$$

where R and T are the power reflectance and transmittance, respectively. If we assume n independent of the frequency, the resonant peaks are equally spaced in the frequency domain. The spacing between the peaks, called free spectral Range (FSR) is equal to:

$$\Delta\nu_{FSR} = c / 2nL . \quad (2.28)$$

A useful quantity, which is proportional to the cavity enhancement, is the finesse (F), defined as the ratio of the free spectral range to the full-width-at-half-maximum (FWHM) of the resonance. To obtain an expression for the finesse, we combine the mirrors' reflectances and the propagation loss in a single cavity loss coefficient α_c , as follows:

$$\exp(-2\alpha_c L) = R_1 R_2 \exp(-2\alpha L), \quad (2.29)$$

from which we obtain:

$$\alpha_c = \alpha + \frac{1}{2L} \log\left(\frac{1}{R_1}\right) + \frac{1}{2L} \log\left(\frac{1}{R_2}\right) \equiv \alpha + \frac{1}{2L} \log\left(\frac{1}{R^2}\right). \quad (2.30)$$

Rewriting Equation (2.27) using the previous expression, we can express the FWHM of a resonance peak, as follows:

$$(2.31)$$

$$\Delta\nu_{FWHM} = 2 \frac{\Delta\nu_{FSR}}{\pi} \arcsin\left(\frac{1 - \exp(-\alpha_c L)}{2 \exp(-\alpha_c L/2)}\right),$$

and from here we calculate the finesse:

$$F = \frac{\Delta\nu_{FSR}}{\Delta\nu_{FWHM}} = \pi / \left\{ 2 \arcsin\left(\frac{1 - \exp(-\alpha_c L)}{2 \exp(-\alpha_c L/2)}\right) \right\} \approx \frac{\pi}{\alpha_c L}, \quad (2.32)$$

where the approximation is valid for $F \gg 1$. Since the finesse is a quantity easily measurable, it is very useful for comparing experimental and simulated cavity spectra (Section 7.4).

A second useful quantity, especially for very short devices, where the cavity length is not well defined (Section 5.3.2), is the quality factor Q , defined as the ratio of the resonance frequency to the resonance FWHM:

$$Q = \frac{\nu}{\Delta\nu_{FWHM}}. \quad (2.33)$$

Combining Equations (2.28), (2.32) and (2.33) we obtain the relationship between F and Q :

$$F = \frac{\lambda}{2nL} Q. \quad (2.34)$$

It should be noted that, if the cavity is as short as the wavelength, as in the case of microcavities (Chapter 3), at the first order we can use the approximation $F \sim Q$.

The value of Q can be decomposed in two terms, as follows:

$$\frac{1}{Q_{tot}} = \frac{1}{Q_{rad}} + \frac{1}{Q_{wg}}, \quad (2.35)$$

where Q_{rad} (radiation- Q) includes the effects of loss and scattering and Q_{wg} (waveguide- Q) describes the coupling of the cavity mode with the input and output waveguide sections. If $Q_{rad} > Q_{wg}$ the cavity is overcoupled, and the cavity enhancement is lower than the maximum achievable. If $Q_{rad} < Q_{wg}$, the cavity is undercoupled, and the input power is not efficiently coupled to the cavity. In most

cases of interest, such as this work, the optimum cavity performance is achieved for $Q_{rad}=Q_{wg}$, which corresponds to a value of the mirror transmission equal to the cavity loss.

2.4.2 Nonlinear resonant cavities

Since the field amplitude in a cavity can be much higher than the input amplitude, cavities can be effective in enhancing nonlinear optical effects. We will start our discussion by deriving the enhancement factor of a cavity resonant at the FH, in the case of SHG. We use the expression for the cavity field (Equation (2.27)) and substitute it into Equation (2.19), which results in the differential equation for SHG:

$$\frac{dA_{2\omega}}{dz} + \frac{\alpha_{2\omega}}{2} A_{2\omega} = \frac{i \omega d_{eff} J}{2 n_{2\omega} c} \frac{A_{\omega-in}^2 t_{\omega}^2 \exp(-\alpha z)}{[1 - R_{\omega} \exp(-i2kL - i2\phi_m - \alpha_{\omega} L)]} \exp(-i\Delta k z) \quad (2.36)$$

By solving this equation, we obtain as before, the expression of the SH amplitude *inside* the cavity. The *output* SH amplitude is obtained by multiplying the result by the transmission of the mirror at the SH ($T_{2\omega}$). Computing the efficiency and dividing it by the efficiency of a plane waveguide of the same length, we obtain the enhancement factor Γ_{ω} :

$$\Gamma_{\omega} = \frac{\eta_{cav}}{\eta_{wg}} = \frac{T_{\omega}^2 T_{2\omega}}{|1 - R_{\omega} \exp(-i2kL - i2\phi_m - \alpha_{\omega} L)|^4}. \quad (2.37)$$

The maximum of Γ_{ω} is obtained when the resonant condition $2kL + 2\phi_m = 2\pi q$ (with q integer) is satisfied. The previous expression can be generalized to doubly resonant cavities^{38, 39}, resulting in:

$$\Gamma^d = T_{\omega}^2 T_{2\omega} \left| \frac{1 + R_{\omega} \sqrt{R_{2\omega}} \exp(-i\delta_c - \alpha_{\omega} L - \alpha_{2\omega} L / 2)}{\{1 - R_{\omega} \exp(-i\delta_{\omega} - \alpha_{\omega} L)\}^2 \{1 - R_{2\omega} \exp(-i\delta_{2\omega} - \alpha_{2\omega} L)\}} \right|^2 \quad (2.38)$$

where:

$$\begin{aligned} \delta_{\omega} &= 2\phi_m^{\omega} + 2k_{\omega} L \\ \delta_{2\omega} &= 2\phi_m^{2\omega} + 2k_{2\omega} L \\ \delta_c &= 2\phi_m^{\omega} + \phi_m^{2\omega} + 2k_{\omega} L + k_{2\omega} L \end{aligned}$$

In this case, the maximum enhancement is achieved when δ_ω , $\delta_{2\omega}$ and δ_c are all simultaneously multiples of 2π . From this analysis we see that resonating both wavelengths is not a trivial task, since satisfying three phase relationships implies tight control of fabrication parameters and different degrees of tunability. Because of these complications, we decided to focus our research on cavities resonant at the FH only.

Chapter 3

Device Design

In the previous chapter, we have seen the advantage of tightly vs. weakly confining waveguides, and how critical it is to find a phasematching technique which does not involve prohibitively complex fabrication issues (e.g., sub-wavelength QPM gratings). We have also seen that resonant cavities can be useful to enhance nonlinear effects. In this Chapter we explore the concept of a microcavity waveguide device, where nonlinear-optical phenomena can be enhanced by a 1D-photonic-crystal microcavity in the propagation direction and by a tightly confining waveguide in the two other dimensions. The cavity length is shorter than the coherence length ($\sim 1 \mu\text{m}$), such that the interaction is automatically phasematched.

We will start with a theoretical analysis of the device and compare it to traditional QPM waveguides. We will then describe the quantitative design, the fabrication process and the linear characterization of the cavity. Although we were not able to measure SHG from these microcavities, their study led to the understanding of critical issues in guided wave resonant nonlinear optics. With this knowledge, we were able to design and fabricate a new generation of cavity devices (the second part of this dissertation), which efficiently generated SH. Moreover, the fabrication process, developed for the microcavity device, because of its robustness

and flexibility, has been used with only minor modifications for the fabrication of the second generation of devices. Similarly, the tools and the experience gained in characterizing these devices were applied directly to the characterization of the new devices.

3.1 Back-of-the-envelope analysis

In this section we analyze the enhancement of nonlinear optical phenomena due to the tight confinement and to the presence of a resonant cavity. Scaling rules will be deduced, and we will compare the performance of microcavities with weakly guiding QPM waveguides. For simplicity, we will consider only SHG processes, although our analysis can be easily extended to other cases of interest, such as DFG.

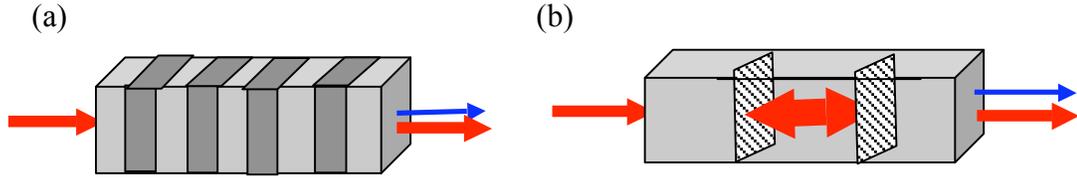


Figure 3.1 Schematic of QPM waveguide (a) and resonant device (b)

We start rewriting the SHG efficiency, expressed by Equation (2.17), for a traveling wave (i.e, non-resonant) QPM waveguide (Figure 3.1a) in the following way:

$$\eta_{\text{wg}} = \left(\frac{8\pi^2 d_{\text{eff}}^2}{n_{2\omega} n_{\omega}^2 c \epsilon_0 \lambda_{\omega}^2} \right) \left(\frac{L_{\text{wg}}^2}{A_{\text{wg}}} \right) \text{sinc}^2 \left(\frac{\Delta k L_{\text{wg}}}{2} \right) = \eta'_{\text{norm}} \left(\frac{L_{\text{wg}}^2}{A_{\text{wg}}} \right) \text{sinc}^2 \left(\frac{\Delta k L_{\text{wg}}}{2} \right) \quad (3.1)$$

where L_{wg} is the interaction length, A_{wg} is the effective area and $\Delta k = k_{2\omega} - 2k_{\omega} - K_g$ is the phase-mismatch vector, as defined in Section 2.2.1. We assume that η'_{norm} is

dependent only on the material properties and on the wavelength, neglecting, to first order, the dependence of n_ω on the geometry. This assumption is justified since in most cases n_ω varies at most by a factor two, whereas in our discussion we are interested in order-of-magnitude variations. For the same reason, at this stage we also neglect the effects of loss and disorder.

For our purposes, we will consider the following figures of merit: the normalized maximum conversion efficiency and the bandwidth, defined respectively as:

$$\eta'_{wg} = \frac{\eta_{wg(\max)}}{\eta'_{norm}} = \frac{L_{wg}^2}{A_{wg}} \quad \text{and} \quad (3.2)$$

$$BW_{wg} = \frac{\Delta\lambda_{FWHM}}{\lambda} = \frac{\pi}{\lambda \frac{\partial\Delta k}{\partial\lambda} L_{wg}}. \quad (3.3)$$

η'_{wg} represents the maximum attainable conversion efficiency in the case of perfect phasematching, and is a function of only the geometrical parameters. BW_{wg} describes the wavelength tunability range within which the conversion efficiency remains greater than $\sim 50\%$ of its maximum value (η'_{wg}). This expression has been derived by approximating Δk as:

$$\Delta k \sim \frac{\partial\Delta k}{\partial\lambda} \Delta\lambda. \quad (3.4)$$

Given the dependence of η'_{wg} , ideally we would like large L_{wg} and small A_{wg} ; however, decreasing the mode area will also sharply decrease the bandwidth, due to increased modal dispersion. This behavior is shown in Figure 3.2a, where we simulate a square waveguide with $\text{Al}_{0.5}\text{Ga}_{0.5}\text{As}$ core and air cladding, and varied the core size.

Now imagine removing the QPM structure and building instead a cavity of finesse F inside the waveguide, resonant at the FH frequency (Figure 3.1b). If losses are sufficiently low and we assume perfect mirror transmission at the SH, the resonant

enhancement described in Section 1.3.6 can be approximated as $\sim (F/\pi)^2$. We can then write the conversion efficiency for a resonant cavity as:

$$\eta'_{res} = \frac{L_{res}^2 F^2}{\pi^2 A_{res}}. \quad (3.5)$$

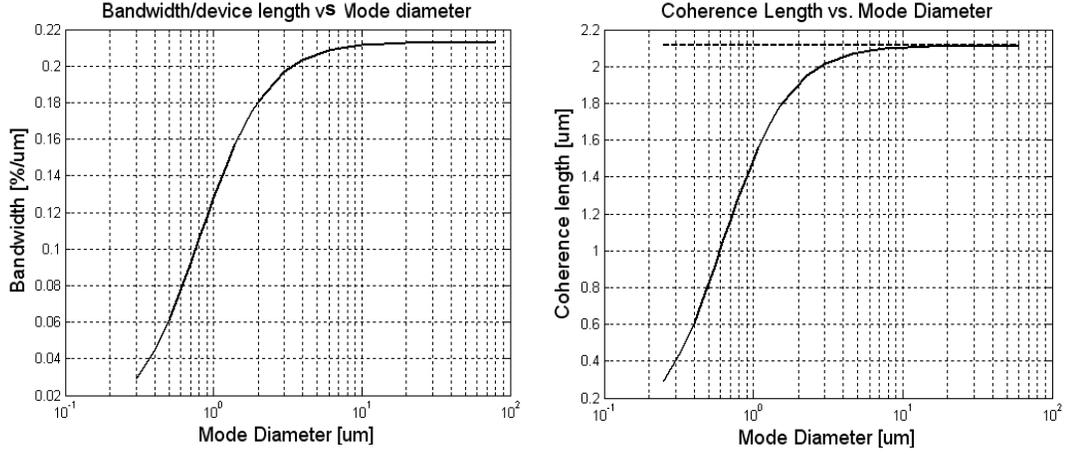


Figure 3.2 Bandwidth (a) and coherence length (b) vs. mode diameter

Finally, if we want to avoid using any phasematching technique, the interaction length, L_{res} , must be smaller than the coherence length, defined in Equation (2.11). It is instructive to look at how the coherence length varies as a function of the confinement (Figure 3.2b). Also in this case, we simulated a square $\text{Al}_{0.5}\text{Ga}_{0.5}\text{As}_{0.5}$ waveguide. We see that the coherence length decreases sharply into the sub-micron regime as we increase the confinement. To design a resonant cavity with efficiency as high as the QPM waveguide, we need a finesse equal to:

$$F_{res} = \pi \frac{L_{wg}}{L_{res}} \sqrt{\frac{A_{wg}}{A_{res}}}. \quad (3.6)$$

If $L_{res} \ll L_{wg}$, the bandwidth is no longer determined by the phasematching bandwidth (BW_{wg}), but is rather limited by the square of the resonance bandwidth (hence the factor 2 in the formula), given by:

$$BW_{res} = \frac{\Delta\lambda_{FWHM}}{\lambda} \sim \frac{1}{2Q} \sim \frac{\lambda_{\omega}}{4n_{\omega} L_{res} F_{res}} \quad (3.7)$$

We consider now realistic values for the waveguide parameters, where we approximate, to the first order, the effective area with the mode size. A mode diameter of $0.5 \mu\text{m}$ can be attained using the AlGaAs/ Al_xO_y system in a tightly confining waveguide, compared to the $10\text{-}\mu\text{m}$ mode diameter of the QPM waveguide. The coherence length for this structure is $\sim 1 \mu\text{m}$, so the cavity length must be of the same order. From the previous considerations, we deduce that we need a cavity with a finesse of ~ 700 in order to have the same “effective length” as a 1 mm long QPM waveguide, which we take as a reference. Such a cavity can be built using a Photonic Bandgap Crystal structure as described in the next Section.

	Device length L_D [μm]	Mode diam. [μm]	Coherence length L_c [μm]	Normalized conv. eff. η'_{norm} []	Bandwidth BW [%]	$\eta'_{norm} * BW / L_D$ [μm^{-1}]
QPM	1000	10	2.2	1e5	~ 0.07	~ 0.07
Resonant	5	0.5	~ 1	1e5	~ 0.028	~ 5.6

Table 3.1. Comparison between AlGaAs QPM waveguides²³ and resonant microcavity device. The resonant device length of $5 \mu\text{m}$ includes the length of the mirrors.

Table 4.1 compares the two structures in terms of efficiency and bandwidth. For the QPM waveguide bandwidth, a realistic value of $\partial\Delta k / \partial\lambda \sim 0.1 \mu\text{m}^{-2}$ obtained from simulations is assumed. We see that, with a finesse of ~ 700 , the resonant microcavity has the same peak efficiency as a 1-mm -long QPM AlGaAs waveguide with a bandwidth only 50% smaller. Consider now the combined figure of merit:

$$\eta_{comb} = \frac{\eta'_{norm} BW}{L_D} \quad (3.8)$$

which takes into account efficiency, bandwidth and total device length, L_D . This last parameter is important in the case of dense device integration. Using this figure of merit, we see that a resonant, tightly confining device would perform theoretically $\sim 80\text{x}$ better than a 1 mm QPM waveguide. It is also interesting to note that, using the scaling laws previously derived, η_{comb} for QPM waveguides becomes independent of length, so that longer waveguides would have the same “combined” performance.

These considerations motivated us to focus our attention on ultra-short, high- Q , resonant devices.

We must note finally that this back-of-the-envelope analysis assumes perfect transmission of the mirror at the SH wavelength. This assumption will turn out not to be true, and will constitute the main reason for abandoning this design. However, the principles of this analysis remain valid if a mirror with sufficiently high transmission, such as the one presented in Chapter 7, can be invented for ultra-short cavities.

3.2 Microcavity design

A schematic of the device, similar to the one originally realized by Ripin et al.²¹ for linear optical filtering applications, is shown in Figure 3.3. It consists of an AlGaAs core supported by an AlO_x cladding. The waveguide is used to couple light into a cavity realized by introducing a defect in a 1D array of holes (photonic bandgap crystal), so that a defect mode is created. The cavity is resonant at the wavelength of 1.55 μm (input fundamental wave) and designed to have high Q .

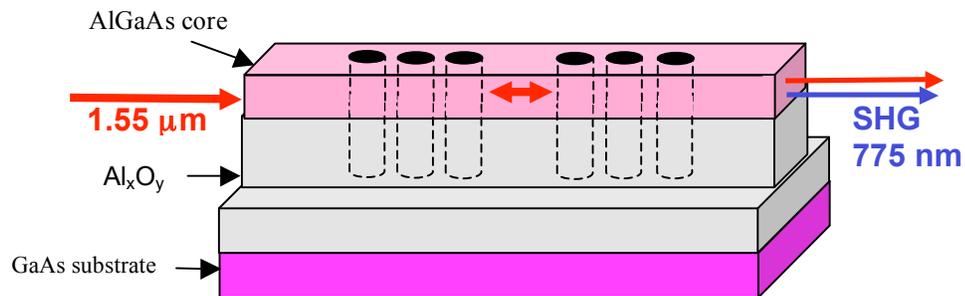


Figure 3.3 Schematic of the microcavity device.

The first element to design, in order to narrow down the parameter space, is the geometry of the waveguide. This task was accomplished using a 2D frequency domain finite difference algorithm⁵⁷. We designed a single mode waveguide at the FH wavelength and we simulated both FH (1550 nm) and SH (775 nm) mode profiles (Figure 3.4).

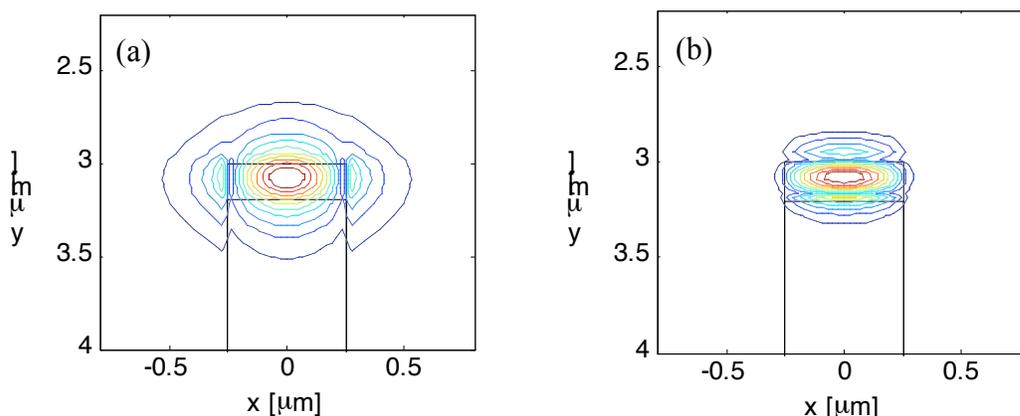


Figure 3.4 Mode field profile of FH (a) and SH (b) modes.

From the overlap integral between these modes, we estimated the theoretical conversion efficiency, assuming perfect phasematching. Figure 3.5 shows the conversion efficiency for different geometries of the waveguide core. The maximum achievable normalized efficiency is $\sim 4.2 \cdot 10^5 \text{ \%}/\text{W}/\text{cm}^2$ for a core thickness of ~ 180 - 200 nm and a width of ~ 480 - 500 nm . This value is more than three orders of magnitude larger than that of commercial LiNbO_3 waveguide devices.

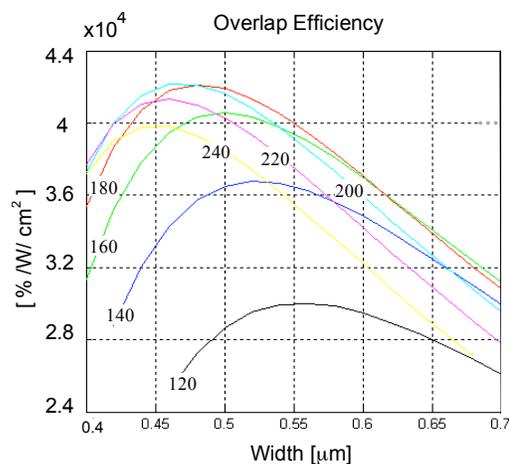


Figure 3.5 Overlap efficiency vs. waveguide width, for different core thicknesses (in nm)

The next step consists of the design of the mirrors. We considered a periodic structure consisting of a waveguide with an array of vertical holes. For particular

values of pitch and diameter of the holes, the structure does not possess any guided mode over a certain range of wavelengths, called the optical bandgap. The bandgap center wavelength is primarily determined by the hole pitch, whereas its width is largely a function of waveguide width and hole radius. Figure 3.6 shows the variation of the bandgap size, expressed as fraction of the center wavelength, as function of the geometrical parameters. The importance of the optical bandgap size lies in the fact that it determines the attenuation of the cavity mode in the mirror and ultimately the cavity mode size.

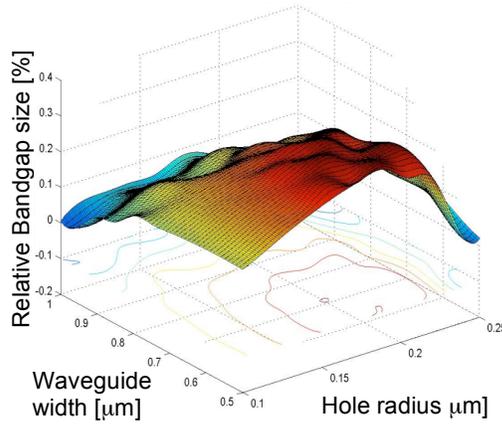


Figure 3.6 Optical bandgap size, expressed as fraction of center wavelength, in function of waveguide width and hole radius.

Once we have established the range of the geometrical parameters of the mirrors, we can move to the cavity design. A defect in the hole array (simply obtained by making the spacing between two holes larger) introduces a defect mode in the optical bandgap, thus creating a cavity. We employed 3D FDTD simulations to adjust the resonant wavelength of the cavity (which depends on center spacing) and to optimize the radiation Q of the cavity (Section 2.4.1). Figure 3.7a shows the radiation Q vs. number of holes for a symmetric, ideal, air-cladding structure. If we use a periodic row of holes, we see that the radiation Q saturates at a value of ~ 3000 . However, this value can be increased using more complex cavity designs.

A first cavity optimization method⁵⁸ consists in reducing the size of the inner holes (i.e., the holes closest to the cavity). If the size is carefully adjusted, radiation multipoles cancel each other and the radiation loss decreases. This fact can significantly increase the *radiation Q* (up to 2.5x), as shown in the explorative simulations of Figure 3.7a. Although this approach is theoretically very effective, fabricating holes of small diameter is extremely challenging.

A second technique consists in shifting the inner holes, instead of reducing their size. Figure 3.7b shows simulations of the *total Q* for the real AlGaAs/AlO_x structure. According to theoretical predictions, shifting the inner holes outwards with respect to their center lattice position by 25 nm, would increase the *Q* by 40%. Although a more modest increase is expected from this technique, this structure is more viable from the experimental point of view, and will ultimately result in a higher experimental total *Q*.

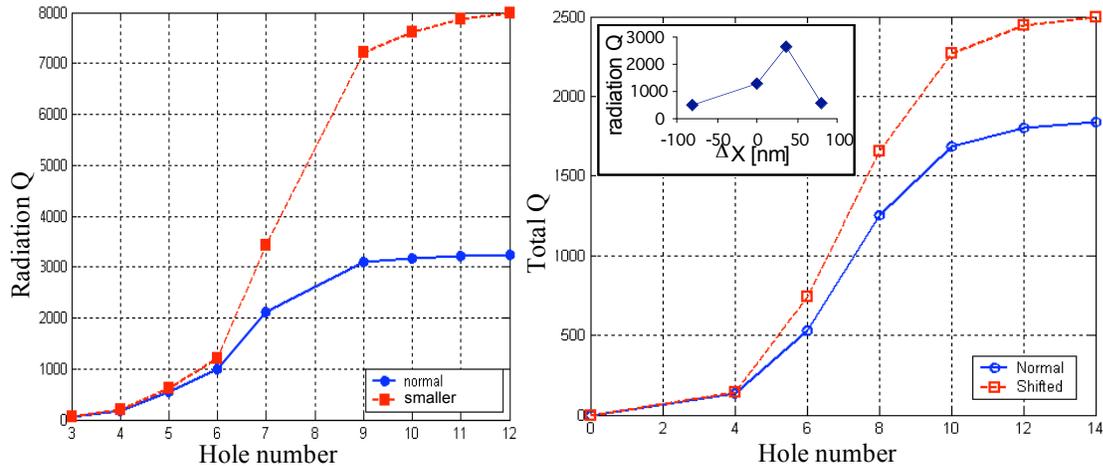


Figure 3.7 (a) Radiation *Q* vs. hole number (on each side) for an ideal, symmetric air cladding case. Normal case (solid line) and case with inner hole diameter halved (dashed line). For the real case (Al_xO_y cladding) the absolute value of *Q* is lower, although the behavior is similar. (b) Total *Q* vs. number of holes for the real device geometry for the normal case (solid line) and for the case with the inner holes shifted by 25 nm (dashed line). The inset shows the radiation *Q* vs. the amount of shift.

Chapter 4

The Fabrication Process

In this chapter we will discuss in detail the fabrication process of the microcavity device. The goal is to develop a flexible and robust process, which can be employed as a technological platform for the realization of a broad range of devices based on the AlAs/AlOx material system. A proof of such flexibility will be given in Chapter 6 and Chapter 7, where the same process will be used to fabricate the 2nd generation waveguide and cavity devices based on a completely different working principle and design.

A second intent of this Chapter is to provide a guide for future researchers in this field. With this purpose in mind we will discuss also some intermediate stages of the process development. These will include the problems encountered and the solutions or workarounds that we discovered. In particular, we included a study of different hard masks that can be used for GaAs dry-etching and their performance. A summary of process, with the technical details useful to reproduce it, is contained in Appendix A.

The fabrication process, summarized in Figure 4.1, starts with the growth of the waveguide structure by Molecular Beam Epitaxy (MBE). $\text{Al}_{0.93}\text{Ga}_{0.07}\text{As}$ is used for the cladding layers, and will be thermally oxidized to AlOx. Then, the wafer is

patterned by e-beam lithography on PMMA and the pattern is transferred to a metal hard mask by lift-off. The sample is then dry-etched in chlorine-based plasma and the residual hard mask is removed. Finally, after selective thermal oxidation of the cladding layer, the wafer is thinned and cleaved for testing. In the following sections we will describe in details each process step.

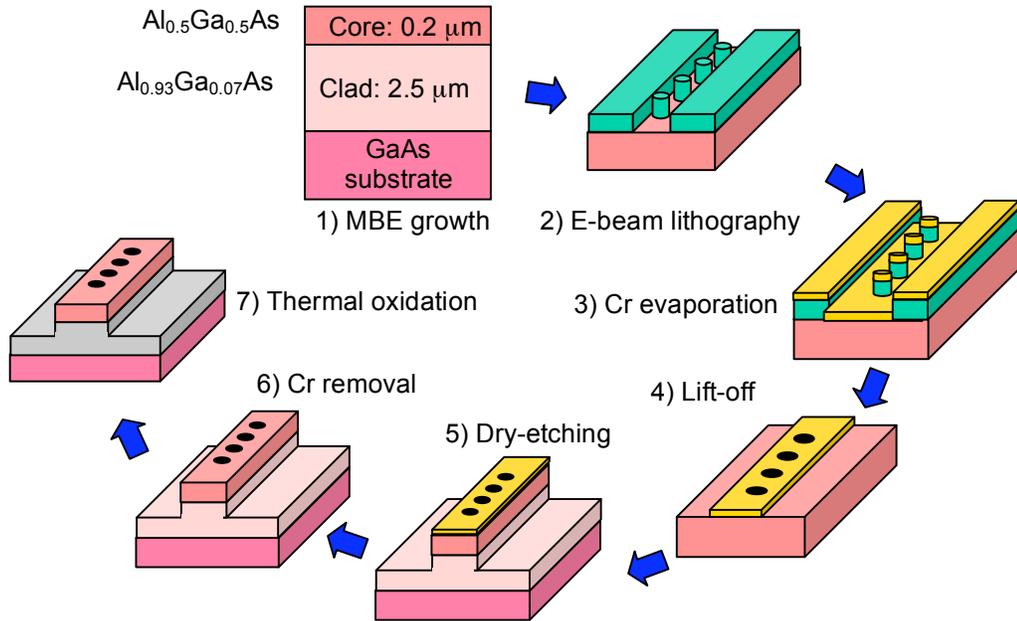


Figure 4.1 Schematic of the fabrication process of the photonic crystal microcavity.

4.1 Epitaxial growth

Molecular Beam Epitaxy (MBE) is a technique for epitaxial growth, in which the constituent materials are provided to the substrate in the form of molecular beams. This characteristic, which is due to the extremely low growth pressure ($\sim 10^{-10}$ Torr), enables the growth of layer thicknesses with theoretically atomic precision⁵⁹. After heating the substrate at a temperature of $\sim 700^\circ\text{C}$, the molecular beams (of Ga, As and Al in our case) are directed to the surface of the substrate, where they combine to form new crystalline material. The substrate is continuously rotated to improve the growth uniformity. Because of the ultra-high vacuum, the mean free path of the molecules in the beams is much longer than the distance between the sources and the substrate.

Thus, the molecules react only after reaching the substrate. Shutters placed in front of the effusion cells are used to switch on and off the different beams, enabling the growth of AlGaAs layers with different Al concentration.

Since GaAs and AlAs are (almost) lattice-matched, it is theoretically possible to grow any AlGaAs composition. Practically, the quality of the material decreases with the Al concentration and the thickness of the layer. The first challenge was then the growth of a thick layer of material with high Al content. Different key-elements contributed to the development of a successful recipe: (i) growing on a GaAs wafer misoriented by 4° towards the $[111]A$ plane, where A indicates Ga rich planes, gives a smoother surface because of the decreased reactivity of the Al atoms with surface contaminants; (ii) initial AlAs/GaAs buffer layers reduce the surface roughness due to the misoriented substrate and help burying surface contaminants; (iii) a slow growth rate is necessary to minimize defect formation.

A N^+ or P^+ -doped substrate is used to avoid charging during e-beam lithography and to facilitate the optical testing (by absorbing uncoupled stray light). However, both core and cladding layers have to be undoped, to avoid loss from free-carrier absorption. For the core, we choose $Al_{0.5}Ga_{0.5}As$ instead of pure GaAs to avoid absorption at the SH wavelength (~ 770 nm). For the same reason, we do not grow any GaAs cap on top. Although a GaAs cap could be helpful to protect the wafer against oxidation, we estimated from simulations that even a few nanometers of GaAs could produce significant loss at the SH. To minimize the formation of a thick native oxide layer, the wafer could be kept in a dry box or capped with a low index insulator (i.e., SiO_2). Alternatively, the oxide could be removed before the fabrication.

4.2 Electron beam lithography

Because of the sub-micron size and the precision required in the positioning of the features (i.e., hole pitch), electron beam (e-beam) lithography had to be used. The patterning of our devices was entirely done using a Raith150 e-beam lithography system, available at the Stanford Nanofabrication Facility. This system is designed for

experimental applications (i.e., not for production). Consequently, compared to typical commercial systems, it offers extreme flexibility in terms of adjustable parameters, at the expense of a steep learning curve.

The main advantages of e-beam lithography vs. optical photolithography are (i) nanometer-accuracy in feature-positioning; (ii) ultra-small attainable feature size (~20nm); (iii) ultra high resolution (~4nm). Another significant advantage is that no physical mask is needed. Thus, the pattern can be easily modified from run to run, which is a highly desirable feature in the context of process development. The big downside, however, is the throughput, which can be orders of magnitude lower than the throughput of photolithography. For this reason only small areas of a chip (millimeters square) can typically be written in a single (several-hours-long) session.

In the following sections we will discuss three issues of capital importance for e-beam lithography: the choice of resist, the design of patterns and the choice of exposure parameters.

4.2.1 E-beam resist

The resist choice affects both key parameters of the e-beam lithography itself, and the processing steps following lithography. For our device, the most important criteria for choosing the e-beam resist are resolution, dry-etching resistance and undercut. These properties are prevalently intrinsic to the resist, although they can be modified by resist thickness, type of substrate, beam energy and processing conditions.

In our process, we utilize PMMA 950k, although we experimented with several resists. We will briefly describe here the problems encountered with the other resists and the advantages of using PMMA 950k. Since we need to deeply etch our structures, in our first experiments we employed chemically amplified resists, such as Shipley UVN5 and UVN30. These resists offer excellent dry-etch resistance, but we were forced to abandon them because of their poor resolution. Zeon ZEP 520 has very good dry-etching resistance and superior resolution compared to UVN resists. However, its extreme sensitivity ($< 10 \mu\text{C}/\text{cm}^2$, increased by the significant GaAs

backscattering) resulted in poor control of the feature size and poor repeatability. Moreover, the accentuated undercut at low beam energies makes this resist unsuitable for deep dry-etching.

Due to these drawbacks, we switched to Poly-Methyl-Methacrylate (PMMA) resists. Since PMMAs have very poor dry-etching resistance, a hard mask is required. Hard masks add complexity to the fabrication process, but eliminate any etching selectivity concern. We have extensively experimented with both PMMA 495k (5% in Chlorobenzene) and PMMA 950k (2% in Chlorobenzene or Anisole). PMMA 495k was used in conjunction with a “positive” mask (see Section 4.2.2) and a SiN hard mask (see Section 4.3.1). The smallest features obtained with a 3500 Å-thick film were 300nm holes in 600nm-wide waveguides. Beside the insufficient resolution, other disadvantages of this resist are the slight undercut and poor dry-etch resistance, which are source of extra-roughness during the mask transfer step.

We found that the best choice for the resist was the PMMA 950k (in Anisole), combined with a pattern transfer by lift-off to a chromium hard-mask. Due to the superb resolution of this resist, enhanced by the ultra-thin layer (80 nm), we were able to fabricate waveguides as narrow as 100 nm and holes as small as 90 nm in diameter using a “negative” mask pattern (see Section 4.2.2). The slight undercut, obtained with 10 keV beam energy, favored lift-off, so that a double layer process was not required. The only disadvantage is the extremely poor dry-etching resistance, but the pattern transfer by lift-off completely solves the problem.

4.2.2 Pattern design

Once the resist and the type of mask transfer (dry-etching or lift-off) have been selected, we can design the pattern to be written (to which we will refer as “mask”, analogously to UV lithography). As shown in Figure 4.2, the transfer process determines the tone of the mask to be written. Since PMMA is a positive tone resist, a *dry-etching* transfer requires the exposure of the regions that will be etched, such as the holes and the regions outside the waveguide (“positive” mask), whereas a *lift-off* transfer involves writing the waveguide and leaving the holes regions unexposed

(“negative” mask). In the first case, the mask layout consists simply of rectangles and arrays of circles (Figure 4.2a), so that any modification of the cavity geometry is straightforward. The positive mask layout was used mainly with a SiN hard mask (see Section 4.3.1), which is non-conducting. Consequently the substrate charging and the proximity effect (discussed in next Section) were the main issues. These problems were reduced respectively by decreasing the beam current and by adjusting the feature size (Figure 4.3a).

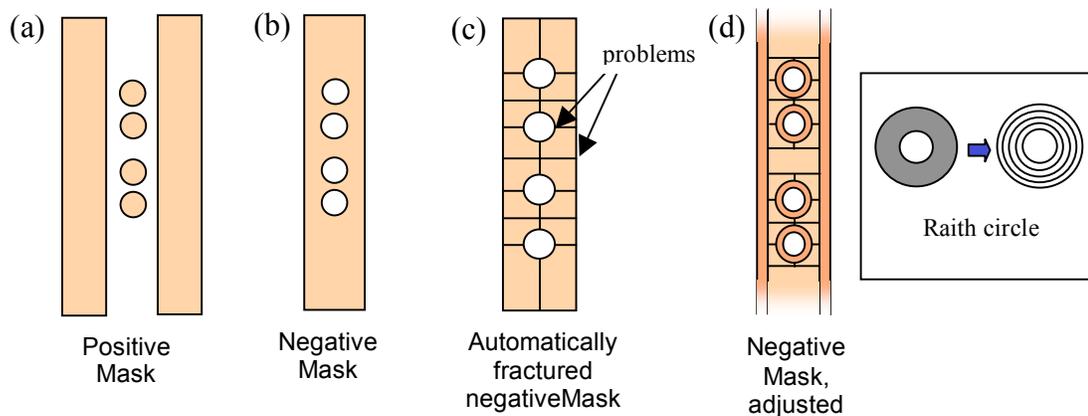


Figure 4.2 Positive mask, transferred by dry-etching (a) and negative mask, transferred by lift-off. The shaded areas are exposed. (c) Negative mask illustrating fractured polygons. (d) Negative mask with improved layout. The inset illustrates the concept of Raith circle.

On the other hand, the layout of the negative mask, employed in combination with metal hard masks, is considerably more complex. The main problem is the practical implementation of non-connected, curved polygons, such as donuts. Most layout editing programs, such as L-Edit or the Raith150 software, do not allow the definition of non-connected polygons. Any polygon of this type, must be obtained by a boolean operation between two layers (i.e. a square and a circle) and is automatically fractured into elementary connected polygons (Figure 4.2c). These boundaries introduce asymmetries at the time of exposure, causing imperfections, notches, or other artifacts. The pattern in the SEM picture of Figure 4.3b exhibits an especially clear case of this effect. The problem was solved by fine adjustment of the exposure parameters and by introducing additional polygons (Figure 4.2d), which completely

eliminate the presence of junctures at the edges of the pattern (Figure 4.3c). In particular, for the holes we employed a special feature of the Raith150 editor: the so-called Raith-circle. A Raith-circle (or ring) is not fractured into simple polygons, but is written as concentric single-pixel lines, as shown in Figure 4.2d. This simple, but powerful, feature retains the circular symmetry at the time of exposure and produces extremely smooth, round features.

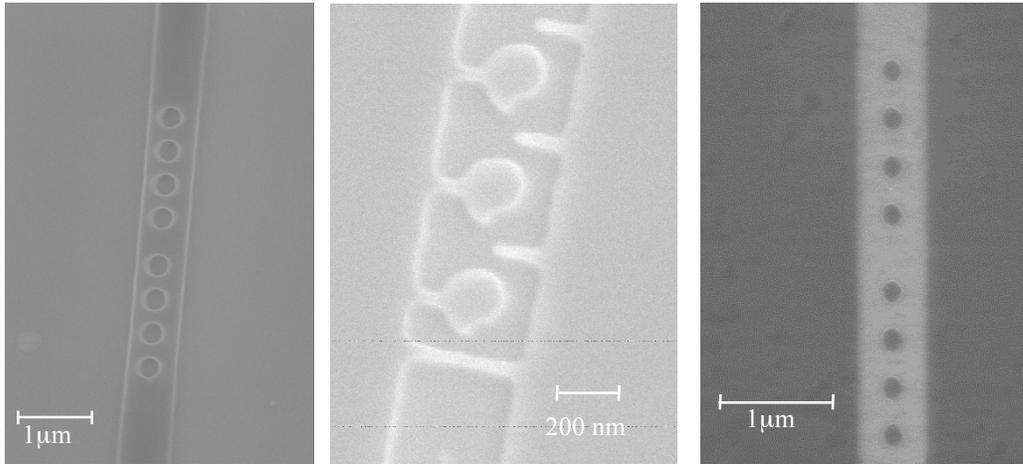


Figure 4.3 PMMA after exposure and developing: (a) Positive mask (b) Negative mask with fractured polygons (c) Negative mask with improved layout.

Manually modifying the cavity geometry in negative masks requires significantly more effort than in positive masks. To expedite this tedious task we wrote a Matlab script that, given the geometrical parameters, generates automatically ASCII files describing the pattern according to a Raith protocol. These files can be directly imported into the Raith150 editor.

We also mention that the exposure time of a negative mask is generally $\sim 10x$ shorter than for the same positive mask, since the area of the region outside of the waveguide is much larger (greater than $3 \mu\text{m}$ wide on each side) than the waveguide itself ($\sim 0.5 \mu\text{m}$ wide).

4.2.3 Exposure

In this Section, we will briefly discuss the most important exposure parameters. The *beam energy* determines the penetration depth of the electrons into the resist (and the substrate) and, consequently, the dose. This parameter was maintained fixed at 10 keV in all our experiments. Based on the experience of previous users, the choice of this relatively low value reduces both the required dose and the proximity effect in the presence of a highly scattering substrate, such as GaAs. Furthermore, the low energy beam produces a slight undercut that is desirable for liftoff processes.

Next, the *dose* (equal to $130 \mu\text{C}/\text{cm}^2$ for our patterns) is determined by a test exposure array. Finally, the *beam current* (which is determined by the column aperture) and the *area step size* must be selected in order to maintain the *beam speed* low enough to avoid stitching and placement errors. The area step size (in units of nanometers) is the beam incremental step used during an area exposure, and represents the resolution. These two quantities are related to the dose by the following equation:

$$Dose = \frac{(Beam\ current)(Dwell\ time)}{(Area\ step\ size)^2} \quad (4.1)$$

where the *dwell time* is the time of transit of the beam on a location. Using the previous equation, the beam speed can then be calculated as:

$$Beam\ speed = \frac{Area\ step\ size}{Dwell\ time} = \frac{Beam\ current}{(Dose)(Area\ step\ size)} \quad (4.2)$$

Given a required dose and maximum beam speed, a trade-off between speed (higher current) and resolution (lower area step size) must be made. The theoretical maximum value of the beam speed is 10 mm/s, but, in our case, we write with a beam speed below 4 mm/s (10 μm aperture, 20 pA of current and 4 nm area step size). With a low beam speed, we reduce charging problems, increase accuracy and avoid stitching errors. The drawback of writing with such a low current is the increased exposure time.

After exposure the sample is developed in Methyl-Isobuthyl-Ketone (MIBK): Isopropanol (IPA) = 1:3 for 30 seconds and rinsed in IPA for another 30 seconds. In the case of lift-off processes, no further treatment of the PMMA mask is performed. Soft-bake (90 °C) and/or UV hardening may be used to harden the resist, but these treatments did not produce promising results in our case.

4.3 Hard mask transfer

After developing, the pattern is ready to be transferred into a hard masking material. In this section, we will discuss some important characteristics that a “hard mask” should possess and the various materials that we employed in our experiments. The material of choice for our final process is chromium; however, we will also describe other materials (SiN, Ni) that we employed in early samples and what we learned from their use.

In general, a hard mask is required for deep etching, or if the resist, as in the case of PMMA, has very poor etching resistance. The main properties of a mask are the selectivity versus the material etched, the smoothness of the profile and the verticality of the *mask* sidewall. Although the *selectivity* can be slightly varied by modifying the process conditions, it is mostly an intrinsic property of the material. The more selective the mask, the thinner is its required thickness. The benefits of a thin mask will be described later in this Chapter. The *roughness* of the mask edge translates directly into sidewall roughness, so an initial smooth mask profile is essential to obtain smooth etched sidewalls. The mask edges should also be as *vertical* as possible, since erosion during the etching results in sloped etched-sidewalls. A thin, robust mask is usually of great help in minimizing this problem.

In our early experiments, we employed simple resists, such as ZEP 520 and PMMA 495k for dry-etching. Although these materials were able to stand the dry-etching step, the edges were eroded during the process, adding roughness and altering the original feature size. For these and for previously mentioned reasons, we decided to employ hard masks.

4.3.1 SiN mask

Silicon Nitride (SiN) was used in our early samples for the “positive” mask process, in combination with the PMMA 495k resist. A 2000 Å-thick layer of SiN was deposited directly on the AlGaAs surface by Plasma Enhanced Chemical Vapor Deposition (PECVD) and 3500 Å of PMMA 495k was spun on top. After exposure, the PMMA pattern was transferred into SiN by Reactive Ion Etching (RIE) in SF₆ plasma.

Using high bias, relatively straight sidewalls were obtained, as can be seen in Figure 4.4a. However, the transfer etching introduced significant roughness at the mask edges. Moreover, the erosion of the resist altered the features sizes, i.e. enlarged holes. The waveguides and cavities were then dry-etched and the remaining SiN mask was removed by RIE (SF₆ plasma). The sample at the end of the fabrication process is shown in Figure 4.4b.

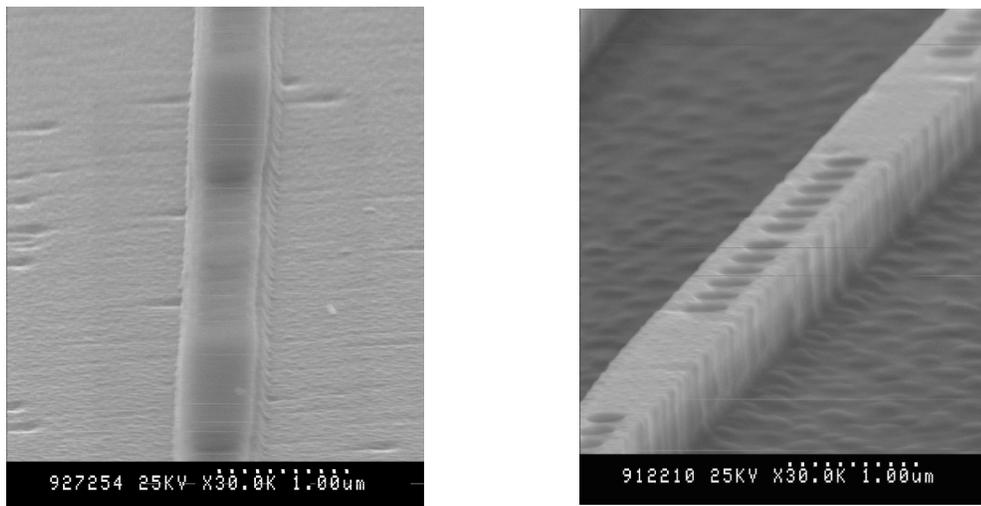


Figure 4.4(a) SiN mask after mask transfer, with PMMA still in place. (b) Sample at the end of the fabrication process.

The advantage of this process is the ease of deposition and removal of the SiN layer. However, there are several drawbacks: (i) the insulating SiN layer increases the *charging effect* during exposure, which decreases the resolution. The minimum hole size that could be obtained using this process was ~300 nm, unsatisfactory for our devices; (ii) as mentioned above, the dry-etching transfer method introduces *extra-*

roughness into the mask and consequently into the waveguide sidewalls; (iii) the thick mask (2000 Å) limits the maximum hole depth (see Section 4.4). These problems convinced us to explore other types of hard masks, such as metal masks, which can produce superior results.

4.3.2 Nickel mask

Metal hard masks are not as stable as dielectrics and they may require careful adjustments of the process conditions. However, if all the experimental problems are resolved, they can produce remarkably better results compared to dielectric masks. In fact, they generally have very high selectivity; therefore, only very thin layers are required for deep etching. This characteristic helps reducing the sidewall roughness, since the mask undercut, one of the primary causes of roughness, is proportional to the thickness. Moreover, metals can be easily e-beam evaporated; this fact makes them suitable for lift-off processes, which are preferable to etching transfer processes.

Based on the chemistry employed for dry-etching, one metal might be preferable to another. For etching AlGaAs, we employed a chlorine-based chemistry (see Section 4.4). According to Scherer et al.⁶⁰, nickel is one of the best masks in a chlorine atmosphere; therefore we started experimenting with this metal. We experimentally measured the etch rate of thin (200-300Å) metal layers on bare GaAs and silicon by in-situ reflectivity measurements (Figure 4.5). If the metal layer is thin, the reflectivity decreases during the etching until when the metal is thin enough that interference fringes appear. Shortly after, when the GaAs substrate is reached, the reflectivity becomes constant. We were able to verify that nickel has excellent selectivity vs. AlGaAs (~1:100, see Table 4.1). However, because this metal is such a good hard-mask, it is also very difficult to transfer patterns into it and to remove it afterwards.

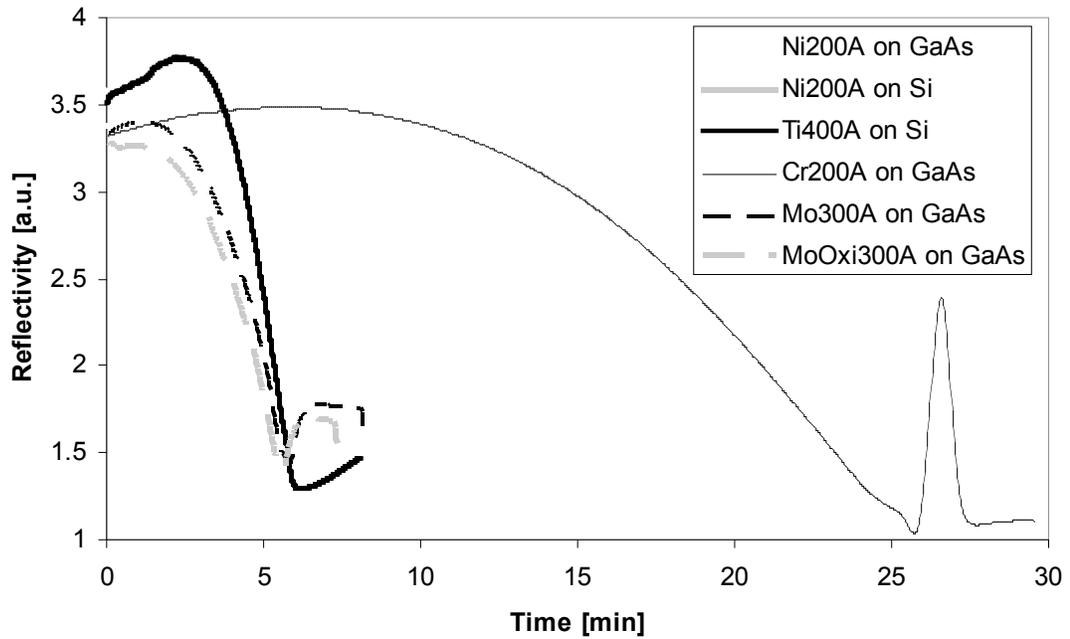


Figure 4.5 In-situ reflectivity measurements for the calculation of metals etch rates.

metal	etch rate [Å/min]	selectivity vs AlGaAs
(SiN)	320	4
Ti	67	19
Mo	50	25
Ni	12	104
Cr	7	178

Table 4.1 Summary of etch rate and selectivity of different metals vs. AlGaAs. SiN is added for comparison.

In our first attempts, we deposited a thin nickel layer ($\sim 200\text{-}300 \text{ \AA}$) directly on the AlGaAs surface; then, we did e-beam lithography on PMMA. Next, we *wet-etched* the exposed nickel with Transene Ni-etchants. The results in terms of sidewall quality were remarkably better than with dielectric masks. However, as can be seen in Figure 4.6a, the mask edges were still rough and the feature sizes difficult to precisely control. *Argon-ion sputtering etch* may also be a viable option for mask transfer, but preliminary tests did not show promising results. Therefore, because of this lack of success and because of the inherent roughness of the dry-etching transfer, we did not further investigate this method.

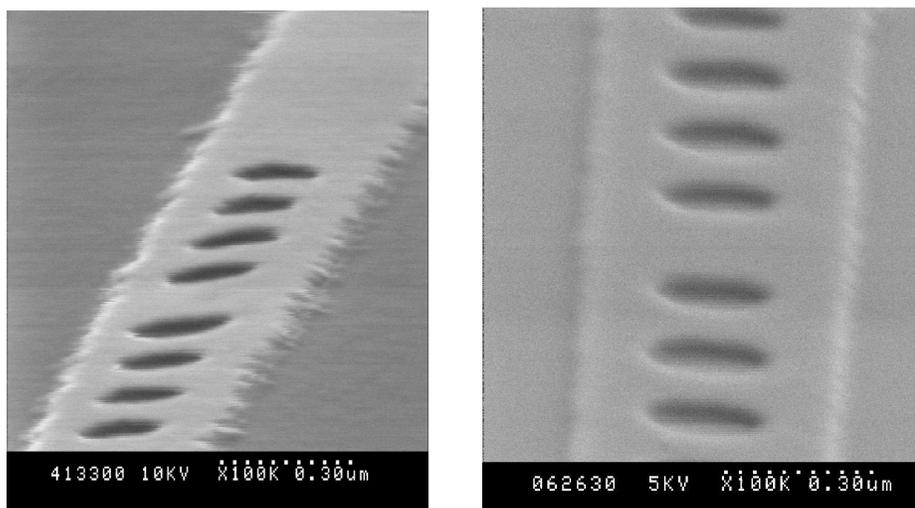


Figure 4.6 Ni hard mask transferred by wet-etch (a) and by lift-off (b)

We solved the problem of mask transfer using a *lift-off* process: we first generated the pattern in PMMA on AlGaAs and then evaporated a thin nickel layer. Because of the slight undercut in the PMMA, even an 80 nm-single layer of resist was thick enough to allow a “clean” lift-off of the 20-30 nm nickel film. We point out the importance of the lift-off process for the control of the feature size, which is now determined only by e-beam lithography (and no longer by dry/wet etching parameters). In this way, the robustness and repeatability of the process are significantly increased. The mask edge profile is improved as well, since no extra roughness is introduced by the lift-off process. The only disadvantage, as we mentioned earlier, is the complexity of using a negative mask (see Section 4.2.2).

The next problem we had to face was the mask removal. Common nickel etchants cannot be used, since they attack the now-exposed high-aluminum content AlGaAs cladding. This problem can be solved with a *sacrificial GaAs layer*. This approach requires the addition of a thick (~ 500 Å) GaAs layer on top of the AlGaAs core during the MBE growth. After dry-etching, the GaAs layer can be selectively etched using a solution of citric acid:hydrogen peroxide = 4:1⁶¹⁻⁶³, thus lifting-off the residual nickel mask. Unfortunately, the thin residual nickel film tends to adhere to

the GaAs surface; therefore, the lift-off requires five to ten minutes to complete, causing the solution to attack the waveguide structure as well. We resolved this complication by using a *double metal mask*, consisting of 150 Å of nickel on top of 150 Å of titanium. The titanium induces tensile stress in the metal film, thus favoring the lift-off. The final structure is summarized in Figure 4.7.

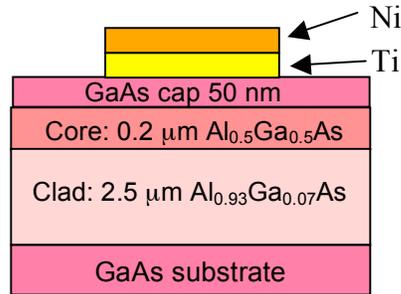


Figure 4.7 Schematic of the Ti/Ni mask structure with the GaAs sacrificial layer.

Although the resulting sidewall roughness is quite small (~10 nm peak-to-peak) and the feature size can be well controlled, this solution still presents some *inconveniences*: although the acid solution does not etch the cladding, it oxidizes its outer layer. As we will see in Section 4.5, the Al_{0.93}Ga_{0.07}As oxide is stable only if it is oxidized under specific conditions. Oxidation caused by a strong oxidizer, such as hydrogen peroxide, leads to an expansion in volume and the oxide tends to crack. Notches at the interface between core and cladding were also observed, as shown in Figure 4.9a. One might consider removing the mask after the thermal oxidation; unfortunately, the GaAs layer forms a thin oxide layer at the surface during the oxidation step, preventing the citric:peroxide solution from etching it.

Another problem is the increased effective thickness of the mask (30 nm Ti/Ni + 50 nm GaAs), which, for a fixed dry-etching time, results in a decreased final hole depth (see Section 4.4). Another drawback of this approach is the limited maximum feature size: since the GaAs layer is etched from the mask edges, larger features require longer etching times, which degrade the surface quality. These problems can be elegantly solved using a chromium mask, as described in the next Section.

4.3.3 Chromium mask

As we have seen, the most serious problem with the nickel mask is its removal. Therefore, we looked for other metals with comparable selectivity, which could be more easily removed. In our tests (Table 4.1), we discovered that chromium has an even better selectivity than nickel ($\sim 180:1$ vs. AlGaAs), although it presents the same removal challenges. Moreover, probably because of electrochemical inhibition, most chromium wet-etchants, such as Ciantek Cr-100, do not etch chromium on GaAs.

However, it was accidentally discovered that chromium can easily be removed in an oxygen plasma at high temperature ($> 300\text{ }^\circ\text{C}$)⁶⁴. One possible explanation for this etching reaction is that, in the plasma, chromium forms oxides that are volatile at low temperatures, such as CrO_3 , whose boiling point is at 250°C . This fortunate circumstance, together with the fact that AlGaAs can be oxidized by chemical reaction only in the presence of water molecules, allows us to safely remove the residual chromium mask before the thermal oxidation. The oxygen plasma oxidizes only the surface of AlGaAs, without damaging the structure. This thin layer of oxide is not only safe, but also helps to protect the cladding surface from oxidation due to water vapor in the air between the mask removal and thermal oxidation steps.

The mask removal is performed in a STS PECVD tool (the same as that used for depositing silicon nitride), which is the only one in our fabrication facility capable of both generating an oxygen plasma and heating the sample to the required temperature. After this treatment, the surface appears smooth and clean. We also performed X-ray photoelectron spectroscopy (XPS) measurements to verify that the chromium was indeed etched and not simply oxidized. These measurements showed that only minute traces of chromium were left on the surface.

We noticed that this etching occurs only if the chromium is deposited at a high deposition rate ($>2\text{ \AA/s}$) and the evaporation is performed at a very low pressure ($<10^{-7}$ Torr). We believe that these conditions help to decrease the number of impurities incorporated into the chromium film during the deposition, which could prevent the volatile chromium oxide from forming. Another important requirement is the

“cleanliness” of the PMMA pattern: PMMA residues, if covered by the evaporated chromium, prevent the chromium film from being etched, presumably because of the formation of non-volatile organo-metallic compounds. Finally, we notice that the chromium mask does not require any GaAs sacrificial layer, thus also solving the problem of the decreased hole depth encountered with nickel masks.

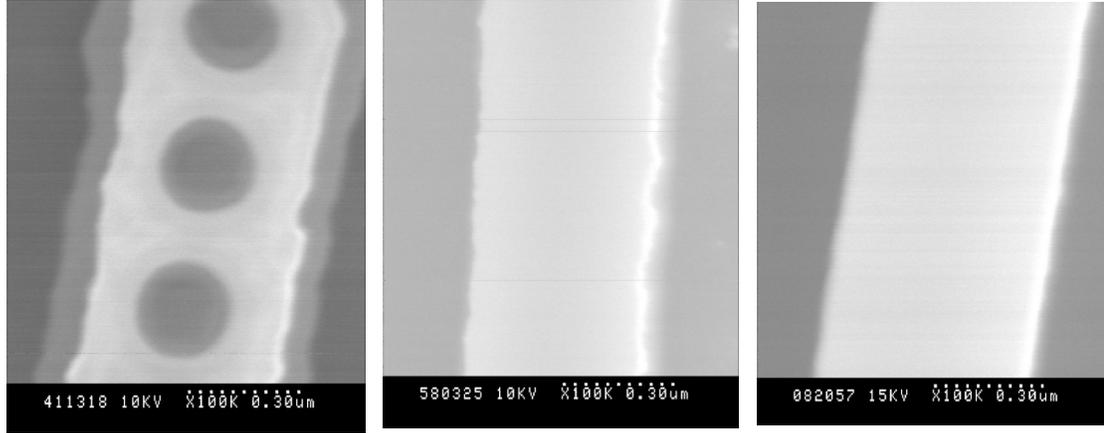


Figure 4.8 Top view of waveguide after dry-etching, using dry-etched SiN mask (a), wet-etched Ni mask (b) and lifted-off Cr mask (c). Pictures are taken with the mask still in place.

In summary, in this Section we have described several types of hard masks and pattern transfer methods. Dielectric masks are easier to utilize, but the results in terms of roughness and control of the feature size are not satisfactory for our purposes. Metal masks, on the other hand, offer superior results. We conclude by showing SEM pictures (Figure 4.8) of waveguide profiles (top view) etched using different masks. The improvement in edge roughness proceeding from dielectric to chromium mask is clearly visible.

4.4 Etching

In this Section, we describe the dry-etching of the waveguide and cavity structure. Dry-etching may appear to be the most critical step in the fabrication process, since it defines the sidewall surface. However, once a recipe has been established, the sidewall roughness is primarily determined by the lithography and mask

characteristics; the etching step simply transfers the mask features, including the roughness, into the sidewall.

The definition of the etching recipe is a challenging task per se, but, once defined, the same recipe can be used for several different types of patterns and masks, producing consistent results. This is usually because, unlike other steps, the dry-etching is entirely performed by a machine, and, therefore, is less sensitive to processing conditions and human error.

For our process, we use a Plasmaquest Electron Cyclotron Resonance (ECR)-enhanced Reactive Ion Etching (RIE) system. RIE is one of the most popular etching techniques; it combines the chemical activity of reactive species generated in the plasma by RF radiation, with physical effects caused by ion bombardment. In an ECR-RIE system, an additional microwave generator is used to provide a high density plasma through the ECR effect, whereas the RIE generator determines the energy of the ions impinging on the surface. The combination of high ECR power with relatively low RF power results in a high-density, low-energy plasma. Using this type of plasma, we can obtain high etch rate as well as low surface damage, without sacrificing anisotropy.

The characterization of this machine and the development of an efficient, highly anisotropic recipe for etching AlGaAs structures has already been done in one of our past works⁶⁵. Here we will describe the results of this characterization and the recipe that we used in our process. To limit the number of parameters in developing an optimal recipe, we set the gas flows to typical values (Ar:Cl₂:BCl₃=15:10:1.5 sccm) obtained from literature⁶⁶⁻⁶⁸ and from previous users, while we use the ECR and RF powers and the pressure as variable parameters. These three parameters determine the ion mean free path and, consequently, the degree of anisotropy (i.e, the ion energy). High *ECR power* increases the plasma density, thus increasing the etch-rate and decreasing the mean free path, whereas the *RF power* primarily affects the ion energy. The *chamber pressure* also affects the plasma density and the degree of anisotropy.

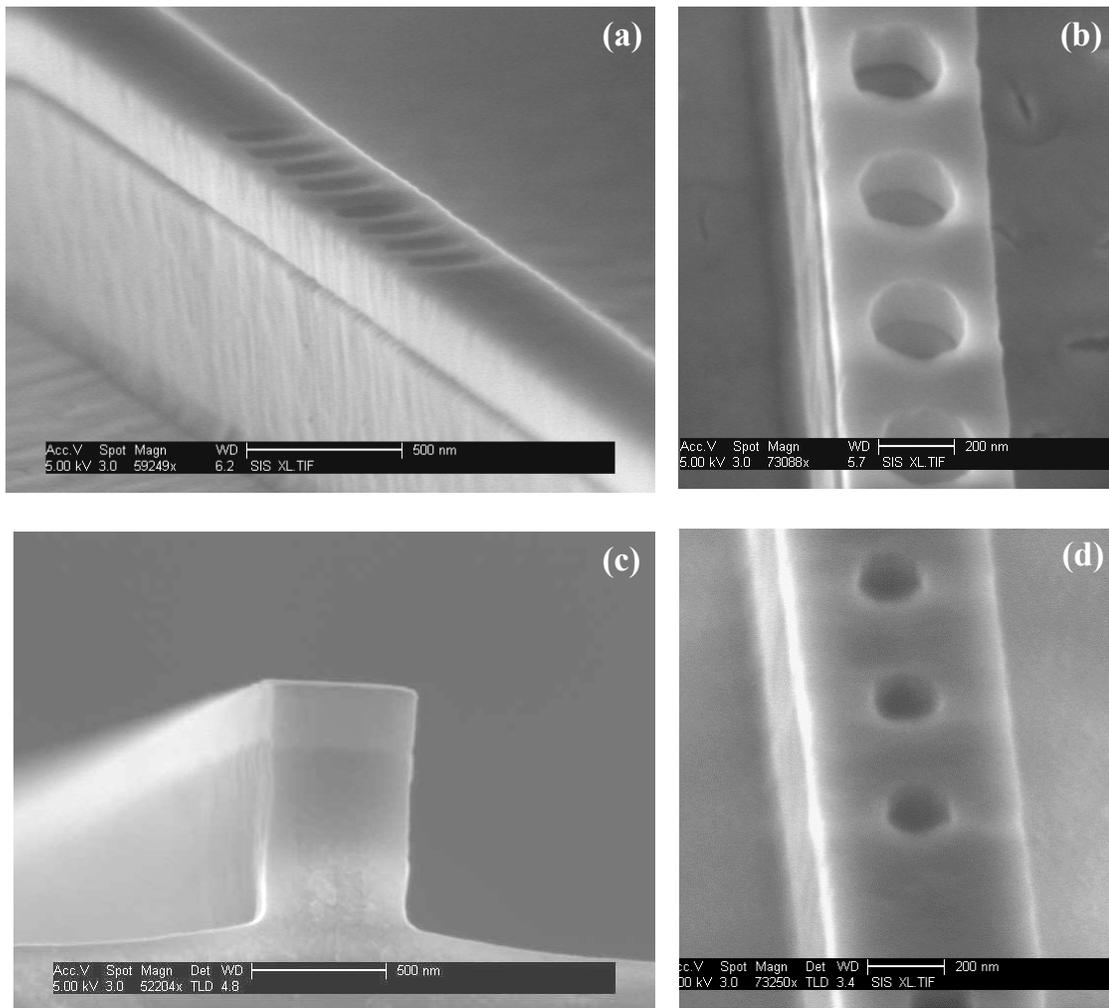


Figure 4.9 (a and b) SEM pictures of cavity devices after the fabrication process using a Ti/Ni mask with GaAs liftoff layer. The notch at the interface between core and cladding is due to the mask removal step. (c and d) Cavities processed using a chromium mask. Note the smaller holes and the absence of the notch on the sidewalls.

These three parameters must be balanced to obtain a mean free path sufficiently long to achieve anisotropy (lower values of the parameters), while keeping the etch rate sufficiently high (higher values of the parameters). Our recipe utilizes a pressure of 2 mTorr, 40-60W RF and 200W ECR power, resulting in an ion energy of 50-60eV (50-60V DC bias).

One final important parameter is the *chuck temperature*, which can be adjusted by both a helium-based chilling system and a heater. High temperatures tend to favor

chemical etching versus physical etching, thus increasing the etch-rate and decreasing the anisotropy. Although the chuck temperature can be adjusted, the heating resulting from the etching itself increases the wafer temperature, altering the chemical etch-rate (i.e., as opposed to the physical etch-rate, due only to ion-bombardment). Our relatively low-power recipe reduces the substrate heating; therefore, we can maintain the chuck temperature stable at 13°C for the entire etching period. This recipe produces very smooth and straight sidewalls and, because of its low power, the mask/AlGaAs selectivity is enhanced. Typical etched profiles are illustrated in Figure 4.9.

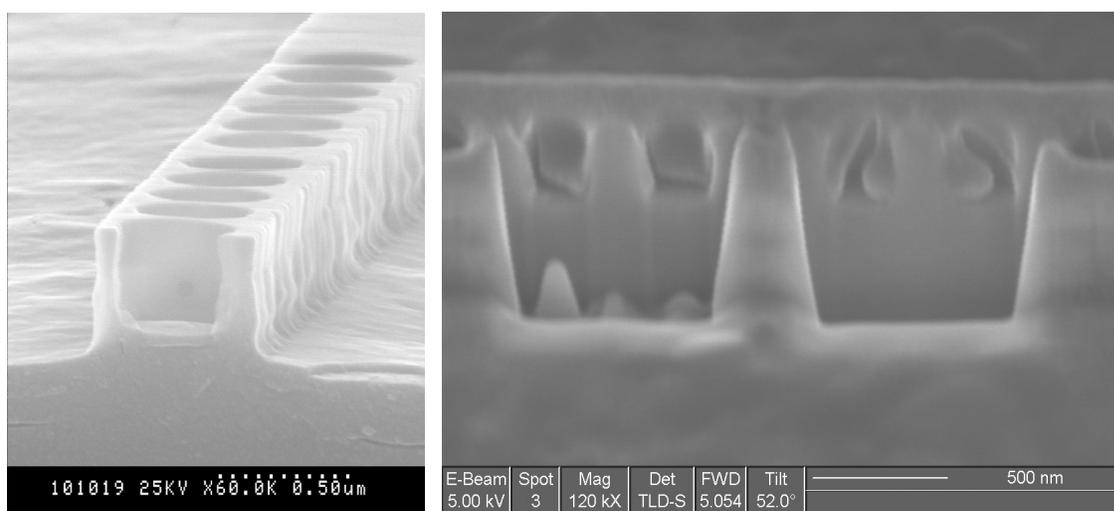


Figure 4.10 Cross section of devices with different hole diameter. For large (360 nm) holes (a) the hole depth (~800nm) is close to the ridge depth. For smaller (200 nm) holes (b) the etch depth(~400nm) is only half the ridge depth. The cross section in (b) was obtained by Focused Ion Beam.

This recipe, although very effective, is not problem-free. This set of parameters was originally developed to etch ridges (waveguides) rather than photonic crystal microcavities (holes). When the aspect ratio of a hole (= depth/diameter) becomes greater than one, the hole etch-rate decreases significantly compared to the ridge etch-rate. This well-known phenomenon⁶⁹ is primarily due to transport problems of neutral species at the bottom of the hole. Practically, this issue limits the maximum hole depth: in fact, the longer the etching time, the taller the ridge

becomes; consequently, the structure becomes too fragile, with high risk of breaking during the subsequent processing steps. Figure 4.10 shows cross sections of holes of different diameters and correspondingly reduced etch-depth. In our specific case, for a hole diameter of 200nm, the maximum etch depth is ~400nm. Although this value slightly decreases the cavity Q, the behavior of the device is not significantly altered.

4.5 Oxidation

After the dry-etching and the mask removal steps, the sample is thermally oxidized in a furnace at 420°C in a water-vapor-saturated nitrogen atmosphere. At this temperature, the water molecules diffuse into the AlGaAs cladding, which is converted to aluminum gallium oxide ($\text{Al}_x\text{Ga}_z\text{O}_y$), commonly called aluminum oxide (AlOx). The arsenic is (almost) entirely removed in the form of arsine (AsH_3)⁷⁰. With this step, the refractive index of high-aluminum content AlGaAs is decreased from ~3 to ~1.6, allowing the fabrication of tightly confining structures. A very important characteristic of this step is the selectivity of the reaction; the oxidation rate increases exponentially with the aluminum content, dropping almost to zero for aluminum concentrations lower than 80%^{71,72}.

The oxidation system consists of a quartz tube maintained at a desired, adjustable temperature. Nitrogen (100 sccm) flows through a DI water bath (called a bubbler) kept at 95°C, and then injected into the quartz tube. The sample is placed on a quartz boat, manually introduced into the furnace and left inside until the cladding is completely oxidized. The typical oxidation rate for $\text{Al}_{0.93}\text{Ga}_{0.07}\text{As}$ at 420°C is ~100 nm/min in the vertical direction and ~50% faster in the horizontal direction⁷¹. The oxidation step typically requires a total time of 15 to 25 minutes, depending on the etch depth and on the width of the structure.

As a matter of fact, this conceptually simple step is experimentally extremely critical, for several reasons: (i) the oxidation of the cladding must be complete. Unoxidized $\text{Al}_{0.93}\text{Ga}_{0.07}\text{As}$, if exposed to air (for example after cleaving), would oxidize at room temperature, expanding and cracking (as opposed to the wet-thermal

oxidation process, during which the oxide volume shrinks⁷³). (ii) The effects of room-temperature oxidation are illustrated in Figure 4.11. Incomplete oxidation leads to thermal fragility of the oxide/substrate interface, such that any further processing involving thermal cycles, such as baking on a hotplate, would make the oxide crack. (iii) Prolonged over-oxidation is harmful as well. If the sample is left in the furnace too long, the GaAs substrate will start to oxidize, forming gallium oxide, which is considerably more fragile than aluminum oxide. Therefore, the oxide/substrate interface becomes extremely fragile⁷⁴; in fact, delamination of the oxide layer from the substrate has been observed during the scribing or the cleaving steps.

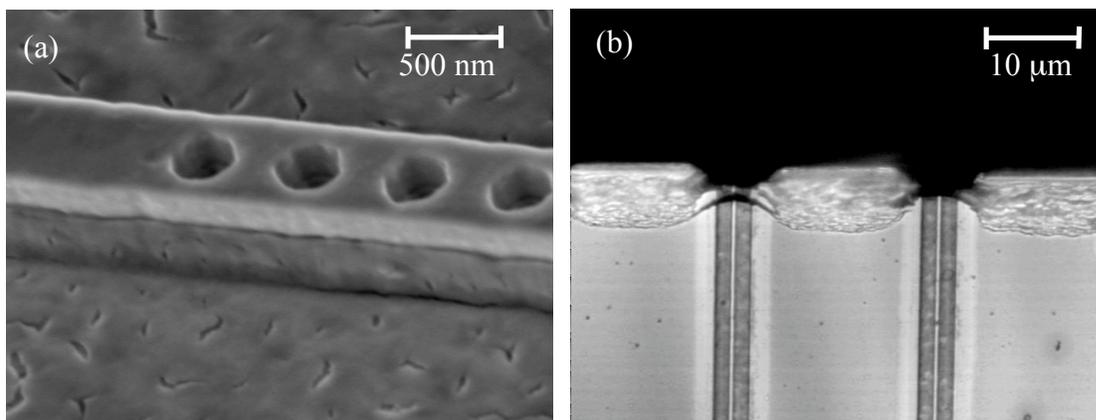


Figure 4.11 (a) SEM picture of device oxidized in air. (b) Picture of broken waveguide facets taken at the optical microscope. After cleaving, the unoxidized portion of cladding exposed to air expands and induces stress in the surrounding region, cracking and destroying the entire facet.

As a result, we found that the oxidation time is a very critical parameter and must be carefully calibrated. A partial solution to this problem involves growing a thin (50-100 Å) $\text{Al}_{0.5}\text{Ga}_{0.5}\text{As}$ layer between the GaAs/AlAs buffer layers and the $\text{Al}_{0.93}\text{Ga}_{0.07}\text{As}$ cladding. The $\text{Al}_{0.5}\text{Ga}_{0.5}\text{As}$ oxide is less fragile than pure-GaAs oxide and can provide better mechanical stability. Consequently the oxidation-time window is widened.

4.6 Thinning and cleaving

After oxidation the sample is spin coated with a protective thick (3-6 μm) resist layer to protect the waveguide structure while the substrate is thinned down to 100-150 μm . This step is necessary for cleaving sub-millimeter long samples. As a rule of thumb, in order to obtain a clean cleavage for GaAs, the ratio of thickness:length should be at least $\sim 1:5$. The width of the sample also plays a role and should be no larger than $\sim 100\times$ the thickness. Another constraint is given by the minimum thickness of $\sim 100\mu\text{m}$, dictated by the fragility of the material. Thinner samples would easily crumble during scribing, and could easily simply break during handling.

Samples longer than 1mm can usually be scribed and cleaved manually: the wafer (facing upward) is notched on the edge, and is flipped and placed on a concave surface. A small force is then applied with a sharp blade on the back side of the wafer (now facing upwards) in correspondence of the notch.

Although in our early experiments we were able to manually cleave sub-millimeter samples, we later decided to employ a professional machine (Dynatex Dry Process Dicing machine) in a private company, which provided excellent precision of the cleavage location (essential for our 2nd generation cavity devices) and repeatability.

Chapter 5

Characterization techniques

In this Chapter we discuss the techniques used to optically characterize the waveguides and the microcavity devices. In the first section we describe the experimental apparatus. The testing setup, as well as our experimental skills, have been undergoing continuous improvements for the entire length of the project. In particular, an effective alignment technique specific for tightly confining waveguides was developed. This technique, with adequate training, can be performed in minutes and requires only the aid of an IR-sensitive card. The details of this technique are described in Appendix B. Finally, the samples' quality has improved as well, so that optical noise and other forms of interference were minimized in the later measurements. This improvement is obvious, for example, when comparing a waveguide Fabry-Perot spectrum of early devices (i.e., extremely noisy and difficult to interpret) with the spectrum of the last birefringent devices, where the Fabry-Perot fringes are clear and usable without the need of any signal processing.

In Section 5.2, we review the theory of the Fabry-Perot method for the determination of the group index and the propagation loss. Finally we report the loss measurements using the Fabry-Perot method and compare them with results obtained using a variation of the well-known cut-back method. In Section 5.3 we discuss the linear characterization of the microcavity, including the transmission spectra and the

determination of the Q. We finally discuss the causes of the extremely low conversion efficiency and the reasons that forced us to abandon this design.

5.1 Measurement setup

In this Section we describe the optical setup, illustrated in Figure 5.1. A tunable laser with a large tuning range is necessary to characterize the microcavities. Depending on the availability of the source and on the type of measurement, in the course of our experiments we employed different types of lasers, including a New Focus Vidia (with a tuning range of 1520-1570 nm), a HP 8168F (1440-1590 nm), an Agilent 81680A (1456-1590 nm) and an Agilent 81640B (1500-1640 nm).

The typical output power of these sources is limited to a few milliwatts; therefore, for alignment purposes, or when higher power is needed, the laser is connected to an Erbium-Doped Fiber Amplifier (EDFA). Although the EDFA has an output power of up to 2 W, to avoid damage to the waveguide input facet, it is typically not driven at power levels higher than ~ 300 mW. We also noticed that, because of dust particles, any fiber connector between the EDFA and the free-space coupler is quickly damaged if the power is increased above ~ 200 -300 mW. To solve the problem, we removed all the connectors by splicing the fiber terminals.

Since in our experiments the input beam must be TE-polarized, a polarization controller is placed downstream from the laser. The fiber-guided light is then converted to a free-space beam with a collimating lens. The distance from the lens to the fiber can be adjusted to obtain the best spot size for the input coupling lens. Since the polarization may change while tuning the wavelength, a polarizer ensures that the input beam is TE-polarized. After the polarizer, a beam sampler monitors the input power.

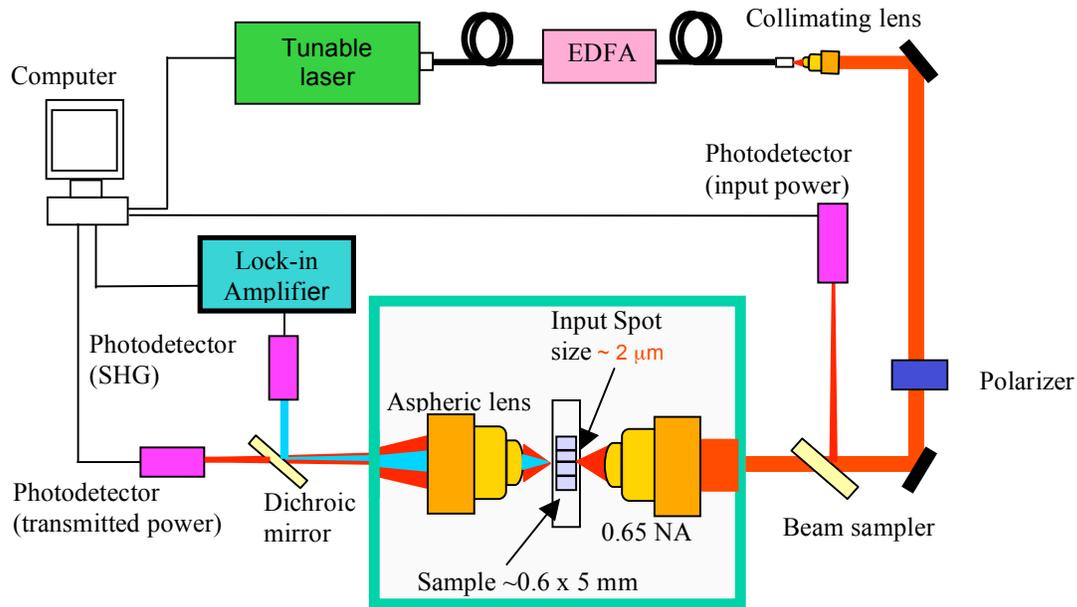


Figure 5.1 Schematic of the measurement setup. The frame represents an enlarged view of the sample with the coupling lenses.

High numerical aperture ($NA = 0.65$ or 0.68), anti-reflection (AR) coated, aspheric lenses couple light in and out of the waveguide. The lenses and the sample are mounted on 3-axis and 6-axis translation stages, respectively. Electrostrictive actuators are employed only for the alignment in the most sensitive directions: the translation of the input lens in the z direction (the beam propagation direction, as shown in Figure 5.2), and the translation of the sample in the x and y directions. The output lens images the light from the waveguides onto either a photodetector for power and transmission spectrum measurements, or onto an infrared camera for alignment and sample navigation purposes.

A dichroic mirror placed after the output lens is used to pick off the SH signal, which is recorded by a separate silicon photodetector. Since the focal length of the output lens is different for FH and SH light, an additional 3-cm-focal-length, AR-coated lens is necessary to focus the SH on the detector.

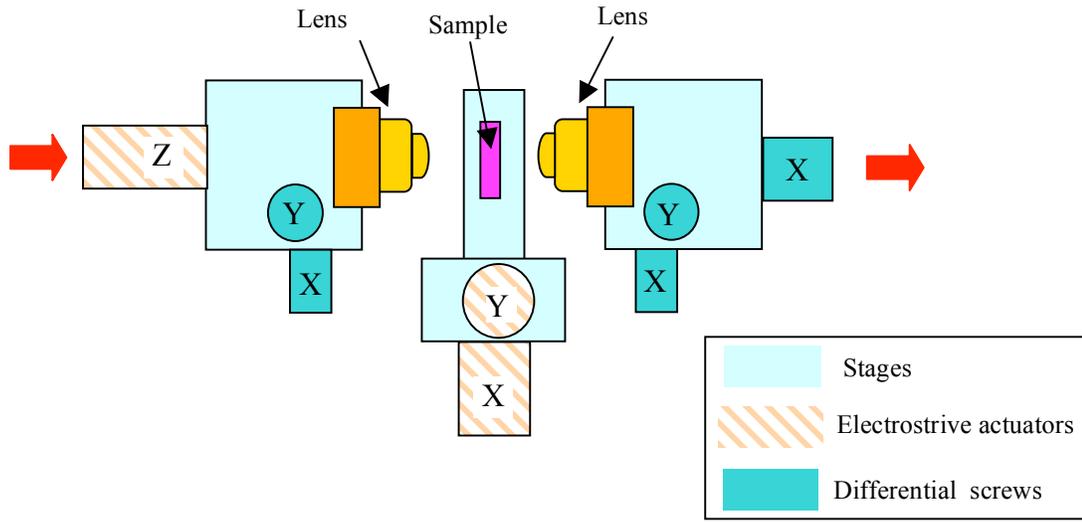


Figure 5.2 Schematic of the input and output coupling stages. Only three electrostrictive actuators are necessary.

For SHG experiments, a lock-in amplifier is employed. The input signal modulation is provided by a chopper or directly by internally modulating the laser. The laser, the photodetectors and the lock-in amplifier are controlled via a GPIB interface using a Labview program. The software records the input FH, the output FH and the output SH power simultaneously, ensuring accurate normalization of the signal. The advantage of using the GPIB interface for recording the signal is the simplicity and precision of the measurement. The slow acquisition speed is not a limiting factor for SH measurements, since the small signal necessitates long integration times. However, if it is not required to record the SH power, such as in the case of linear characterization of the cavities, recording using a Data Acquisition card (DAQ) is a far more convenient and faster method. This alternative system enables the acquisition of high-resolution scans (0.05 nm or less) of the entire spectrum (150 nm) in a few seconds (compared to tens of minutes for the SH measurements).

5.2 Waveguide characterization

In this section we discuss the characterization of plain waveguides (i.e., without cavities). In our case, the method of choice is the Fabry-Perot technique. The advantage of this technique lies in the fact that the measurements of the group index and of the waveguide loss do not depend on the coupling conditions. In fact, these conditions, in the case of tightly confining waveguides, are neither easily determinable nor repeatable. An alternative technique, called the cutback method, exists, but because of the coupling conditions problem, it is very difficult to apply. However, we will see that a variant of this method can be employed and such variant will be useful to verify the results obtained by the Fabry-Perot measurements.

5.2.1 The Fabry-Perot method

This technique is based on the measurement of the waveguide transmission vs. wavelength. The waveguide behaves as a Fabry-Perot etalon and from the interference fringes in the transmission spectrum, it is possible to calculate the waveguide effective index and propagation loss. This method has been described in detail by Regener et al.⁷⁵; therefore, here we only briefly review the theory and report the results needed for our characterization.

The transmitted intensity, I_T , of a symmetrical single-mode waveguide resonator (Figure 5.3a), derived in Section 2.4.1, can be rewritten in this slightly different form:

$$I_T = \frac{T^2 \exp(-\alpha L)}{(1 - R \exp(-\alpha L))^2 + 4R \exp(-\alpha L) \sin^2(\phi/2)} I_0 \eta_c \quad (5.1)$$

where I_0 is the input intensity; η_c is the coupling efficiency to the (lowest) waveguide mode; R and T are the mode reflectivity and transmissivity of the facets, respectively; α is the intensity attenuation coefficient; L is the sample length; $\phi = 2\pi n_{eff}/\lambda L$ is the round-trip phase, with λ being the wavelength and n_{eff} the effective index of the mode, respectively. As illustrated in Figure 5.3, as the wavelength is tuned, the transmitted intensity varies periodically.

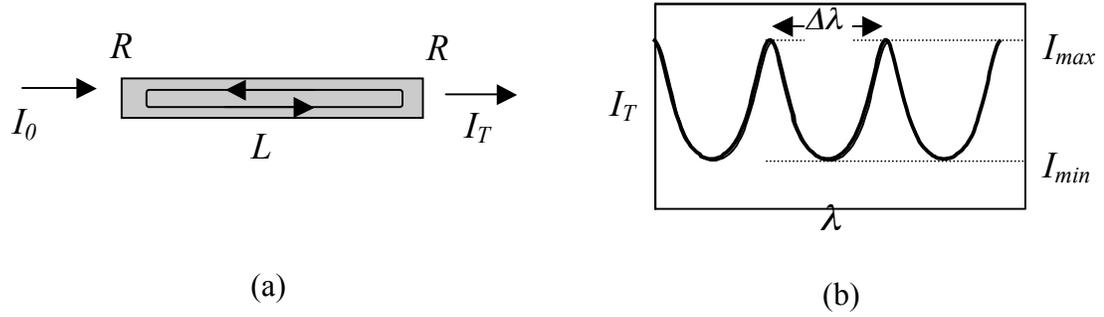


Figure 5.3 (a) Schematic of Fabry-Perot waveguide resonator. (b) Transmitted intensity as a function of the wavelength.

By measuring the period of the fringes ($\Delta\lambda$) and the sample length (L) and using the expression for the round-trip phase, the group index can be calculated as:

$$n_{group} = \frac{\lambda^2}{2L\Delta\lambda}. \quad (5.2)$$

To calculate the propagation loss, we start by defining the contrast, (K), of the Fabry-Perot resonance as:

$$K = \frac{I_{max} - I_{min}}{I_{max} + I_{min}} \quad (5.3)$$

where I_{max} and I_{min} are the maximum and minimum of the fringes, as shown in Figure 5.3b. Using Equation (5.1), we rewrite the expression for the contrast as follows:

$$K = \frac{2R \exp(-\alpha L)}{1 + [R \exp(-\alpha L)]^2}. \quad (5.4)$$

Notice that this quantity is independent of I_0 and η and is only a function of R and α . Solving for α , the previous expression can be rewritten as:

$$\alpha = -\frac{4.34}{L} \ln \left[\frac{1}{RK} \left(-\sqrt{1 - K^2} \right) \right], \text{ in units of dB/mm.} \quad (5.5)$$

By measuring the contrast and the sample length and estimating the facet reflectivity by numerical simulations, we can then calculate the waveguide attenuation coefficient without requiring knowledge of the coupling efficiency.

This discussion is strictly valid only for single-mode waveguides. For multimode waveguides, the total signal is a superposition of the signals of different modes. However, if the effective index of the lowest mode is well separated from the indices of the higher-order modes (as is usually the case for tightly confining waveguides), they do not easily couple to each other. Moreover, the input coupling to odd symmetry modes is usually not efficient, so that only the higher order even modes could generate spurious interference. The Fabry-Perot technique is then still usable provided that, in the Fourier transform of the spectrum, the peak corresponding to the mode of interest and its higher harmonics are sufficiently separated from the peaks of other modes.

Although simple and powerful, this technique presents a few limitations: very lossy waveguides may have low fringe contrast, which could be difficult to measure. A low facet reflectivity would decrease the amplitude of the fringes as well, thus decreasing the precision of the measurement. In the opposite case of extremely low-loss waveguides, the contribution of the reflectivity to the contrast would be much larger than the contribution of the propagation loss. Thus, the uncertainty in the estimated loss value would also be large. This problem could be solved by fabricating longer waveguides, thus increasing the total round-trip loss.

In the case of our waveguides, both the reflectivity and the loss are sufficiently high so that these issues are of no concern. However, in our case, a different set of problems has to be solved.

5.2.2 Dispersion (group index)

We started by characterizing the geometrical dispersion of our waveguides. We measured the transmission spectra of waveguides of different width and, using Equation (5.2), we calculated the group index. The experimental value of the group

index was then compared to the expected theoretical value, according to the following expression:

$$n_{group}(w, \lambda_0) = n_{eff}(w, \lambda_0) - \lambda_0 \left. \frac{d n_{eff}(w, \lambda)}{d \lambda} \right|_{\lambda_0}, \quad (5.6)$$

where n_{eff} is the effective index of the propagating mode, λ_0 is the wavelength (maintained fixed at 1.55 μm), and w is the waveguide width. The second term, called the “dispersion term”, results from the chromatic dispersion of the waveguide. The theoretical values of the effective index as a function of the width and of the wavelength were obtained by 2D frequency-domain simulations, employing the same Matlab code used in Chapter 3. The experimental results are in excellent agreement with the theory (Figure 5.4).

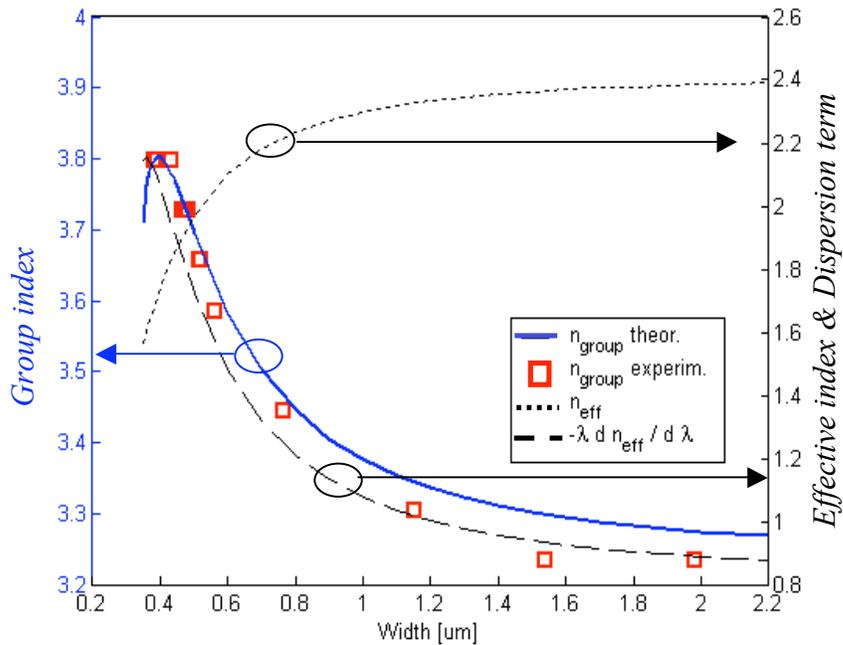


Figure 5.4 Theoretical and experimental group index vs. waveguide width.

Looking at the behavior of each term of the group index, we notice that, because of the extremely tight confinement, as the waveguide width is decreased, the

mode is pushed out of the core, and consequently, the effective index (dotted line) decreases significantly. However, the dispersion term (dashed line), which is typically negligible in weakly confining waveguides, is in this case as large as the effective index itself, resulting in an unusually high group index (solid line). The maximum value of the group index, obtained for a width of 400 nm, is ~ 3.8 , much higher than the core refractive index (~ 3.1). We observe that a high value of the group index is useful for enhancing nonlinear processes, since the light interacts for a “longer” time with the material. On the other hand, a high group index is also responsible for higher propagation losses. The loss coefficient is in fact inversely proportional to the group velocity⁷⁶ (see Section 6.3.3). We finally mention that this observation is consistent with the lower loss coefficient of birefringent waveguides, described in Chapter 6, which exhibit a much lower group index (~ 3).

5.2.3 Propagation loss using the Fabry-Perot method

Using the Fabry-Perot technique described in Section 5.2.1, we characterized the propagation loss as a function of the waveguide width. The contrast and the sample length were measured, whereas the reflectivity of the facets was estimated by 3D finite-difference time-domain simulations; the value was found to be close to the Fresnel reflectivity between air and the core material ($\sim 28\%$), although a significant fraction of the mode was in air. The results are shown in Figure 5.5.

Because of the sub-micron size of the waveguides, the coupling efficiency was very poor, and consequently, the throughput very small. This fact is at the origin of the spurious frequencies in the measured signal spectrum; however, the physical cause of this optical noise is not clear, since it does not depend on the coupling conditions. Although the raw signal was filtered by a Fourier transform, the contrast ratio fluctuates from wavelength to wavelength. Statistics over these fluctuations resulted in the average loss value and the error bars displayed in the figure.

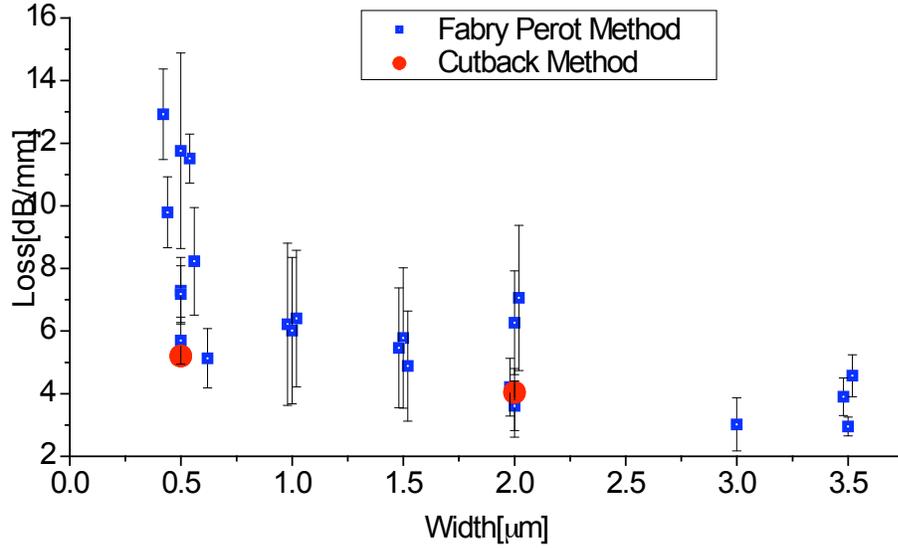


Figure 5.5 Experimental loss coefficient vs. waveguide width measured with both the Fabry-perot method and the cutback method.

As expected, the loss decreases steeply with increasing waveguide width until a width of ~ 500 nm (our nominal design), where the loss is ~ 5 -6 dB/mm; it then decreases slowly to 3-4 dB/mm for larger multimode waveguides. For wide waveguides, the loss is no longer dominated by sidewall scattering, but, we believe, by the scattering at the top and bottom surfaces of the core. These loss values may seem large; however, we point out that the total length of our waveguides is only a fraction of millimeter (typically $600 \mu\text{m}$), so that the total loss is only ~ 2.5 -3 dB. Furthermore, the length of the microcavity device is $< 10 \mu\text{m}$, so that in this case the scattering due to roughness is absolutely negligible.

5.2.4 Propagation loss using a variant of the cutback method

We also measured the losses using an alternative technique, to verify the experimental results. This technique is simply a variant of the cutback method, which consists in measuring the throughput of a waveguide vs. length. The throughput (P_{out}/P_{in}) of a waveguide can be written as:

$$\frac{P_{out}}{P_{in}} = A \exp(-\alpha L) \quad \rightarrow \quad \ln\left(\frac{P_{out}}{P_{in}}\right) = -\alpha L + \ln(A), \quad (5.7)$$

where L is the length and A includes all the losses independent of length, such as coupling loss or reflections at the facets. Plotting the logarithm of (P_{out}/P_{in}) vs. L and calculating the slope of the curve, we can calculate the propagation loss α . For weakly confining waveguides, this is typically accomplished by measuring the loss, then cleaving the sample and re-measuring the shorter samples. This technique is effective only if A is constant from measurement to measurement, which implies that the coupling conditions must be repeatable. Unfortunately in our case, because of the small size of the waveguides and the critical alignment, it is highly improbable to obtain exactly the same coupling conditions for two samples. Moreover, given the initial length of our sample, it is prohibitively difficult to cleave it in smaller pieces.

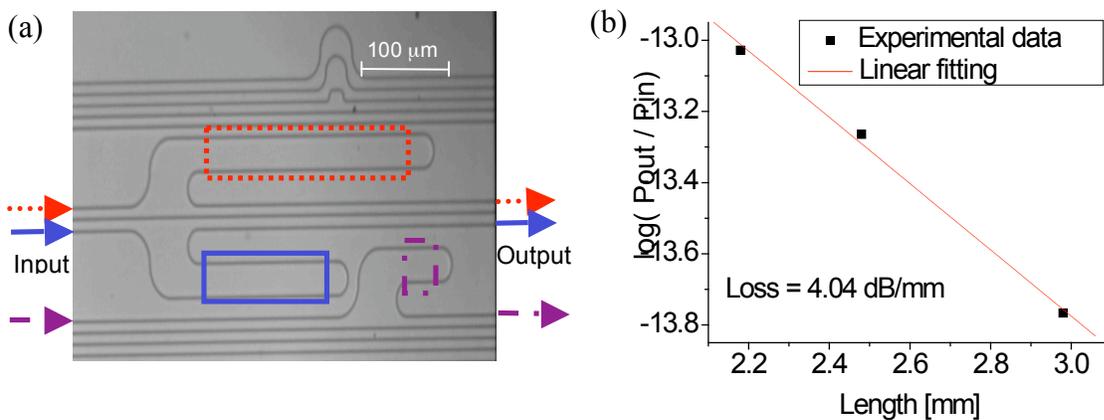


Figure 5.6 (a) Top view (optical microscope) of the cutback method sample. (b) Transmission vs. sample length and fitting curve for one set of three waveguides.

A way to overcome this problem is to use the design shown in Figure 5.6a. This layout allows us to design a set of waveguides of different lengths on the same sample. It also has the advantage, compared to the similar structure used by Lee et al.⁷⁷, of collinear input and output, which makes the alignment considerably simpler. In the set of Figure 5.6a, the three waveguides have the same number of turns and differ only in their length. The length difference is known and highlighted in the figure.

In this case the coupling conditions are very similar since we only need to translate the stage from one waveguide to the other. The only other assumption is that all the waveguides are identical (except for the length) and have identical loss. This is a strong assumption, but reasonable in our case since our waveguides are e-beam-defined and the growth and fabrication conditions are very homogeneous within the area of interest. Only macroscopic defects would significantly change the transmission through the waveguide, invalidating the measurement. However, such a sample could be easily recognized and discarded, since it would have a low correlation with the other waveguides of the same set.

We have performed measurements on sets of three waveguides for two different waveguide widths (500 nm and 2 μm). Figure 5.6b shows the experimental data and the fitting function of one of the sets. The (averaged) results are in very good agreement with the Fabry-Perot method, as shown in Figure 5.5.

5.3 Linear characterization of the microcavity

The next step is the characterization of the microcavity. The experimental technique is the same as the one used for a plain waveguide (measurement of the transmission spectrum). However, we now obtain a single resonance peak, modulated by the spurious Fabry-Perot cavities formed by the facets of the waveguide and the cavity mirrors. Although the spectrum is complicated by these fringes, useful information, such as the resonance wavelength and cavity Q, can nevertheless be extracted using filtering methods. Finally, we will discuss the enhancement of the Q derived from the more complex cavity designs described in Section 3.2.

5.3.1 Transmission spectra and resonance frequency

Since the bandgap spans several hundred nanometers (Figure 5.7a), the spectral region of interest for the characterization of the microcavity exceeds the tuning range of the EDFA, so that only the laser must be used. Moreover, the signal in the bandgap is very weak; thus, transmission spectra were often recorded with the aid of the lock-in

amplifier, even at the fundamental wavelength. The laser source used in the experiments has a tuning range of 150 nm (1440-1590 nm), which covers only a part of the spectrum of the gap. In order to study the band edges and the tunability of the resonance frequency, we varied the geometrical parameters of the cavity so that the different spectral features were shifted within the laser tuning range. For example, in Figure 5.7a, we can see the right edge of the band gap and the resonance peak obtained by varying the waveguide width. The experimental spectra, which have been filtered to remove the spurious Fabry-Perot oscillations, are in excellent agreement with the simulations, as shown in Figure 5.7b. The peak width appears larger in the simulations because of the coarse spectral resolution used, necessary for the numerical simulation of a broad spectral region in a reasonable computation time.

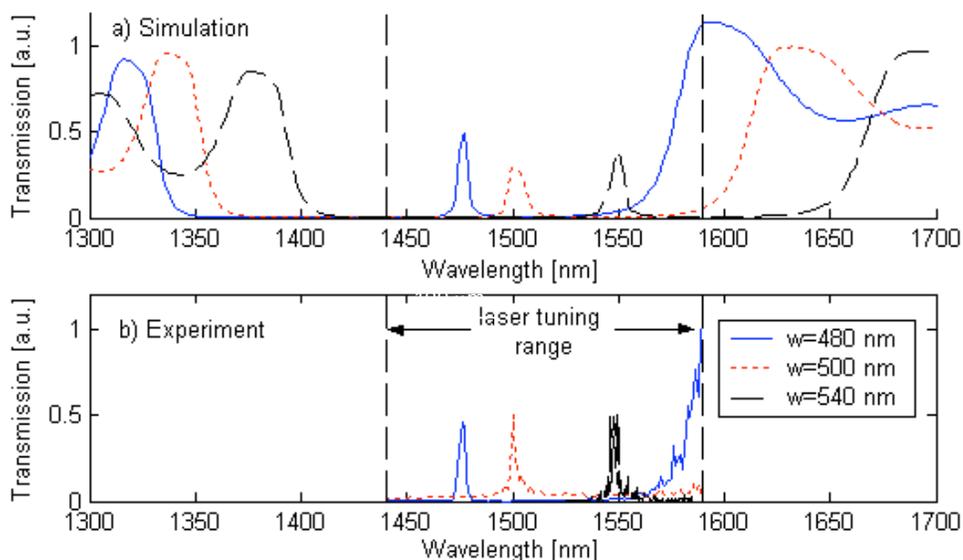


Figure 5.7 Theoretical (a) and experimental (b) transmission spectra for cavities of different core width (w).

The resonance wavelength can be finely tuned simply by changing the geometry. This feature is demonstrated by the results of Figure 5.8. In this experiment we fabricated waveguides of different width and cavity length and recorded the corresponding resonance wavelength. Note the linearity of the curves for

narrow widths and the roll-over for wider widths, due to the decreased dispersion of the waveguide.

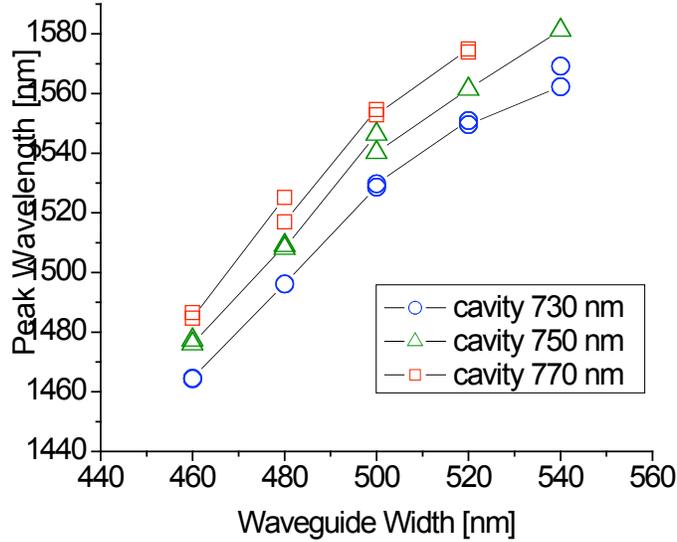


Figure 5.8 Resonance wavelength vs. core width for different cavity lengths.

5.3.2 Cavity-Q

Next, we measured the Q value of the resonance peaks. The experimental data appears “noisy” (solid line in Figure 5.9a), but, as mentioned earlier, this is fixed-pattern optical “noise” due to the presence of two spurious Fabry-Perot cavities (between the cavity mirrors and the waveguide facets). Consequently, it cannot be removed by improving the signal-to-noise ratio. However, if the period of the fringes is short compared to the resonance peak width, these oscillations can be easily filtered out. The filtering technique that we adopted consists in taking the Fourier transform of the experimental data and of a Lorentzian curve, for which we guess the Q-value. We then fit (using a fitting algorithm) the low frequency section of the Fourier spectrum (i.e., from DC to the first spurious frequency-components). This procedure allows one to easily recover the Lorentzian profile of the cavity mode (red dashed line in Figure 5.9a), from which we extract the Q-value.

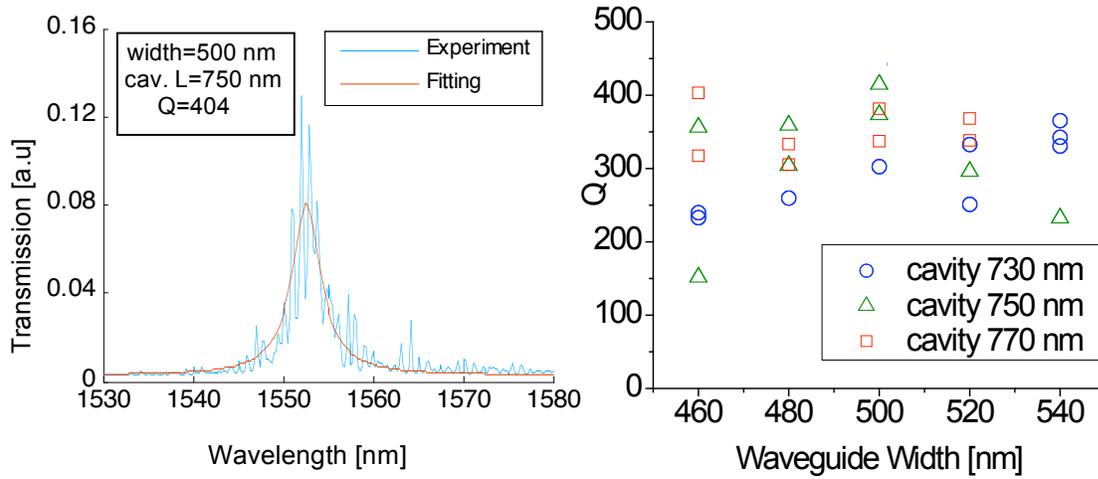


Figure 5.9 (a) Experimental (solid line) and fitted (dashed line) transmission spectrum of a resonance peak vs. wavelength. (b) Distribution of Q vs. width and cavity length

The Q value of a set of cavities of different width is plotted in Figure 5.9b. A clear trend in the dependence of the Q on the waveguide width or cavity length is difficult to trace for this set of cavities. This is mainly due to two reasons: (i) the dependence of Q on these two parameters is weak; (ii) the fabrication imperfections have a very strong effect on Q , thus contributing to the statistical variation of its value. A maximum Q of 404 was measured for cavities with six holes for each mirror. This value is in good agreement with an expected Q of 500 (Figure 3.7). The discrepancy can be explained by sidewall roughness and other imperfections.

5.3.3 Improved geometry cavities

In this Section we discuss the measurements of the cavities with improved design, as described in Section 3.2. In all the waveguides, the center defect spacing (or cavity length) and the number of holes are kept constant at 750 nm and 6 holes, respectively, while the size or position of the holes are changed. Figure 5.10 shows the Q value for these cavities. Data points of the same type with the same abscissa represent two waveguides with nominally identical geometry. These measures, compared to the one shown in the previous section, show well defined trends, which are evidence of improved uniformity of fabrication conditions. Moreover, the curves corresponding to

a modified cavity geometry (i.e., smaller or shifted holes) exhibit a higher Q than the cavities with a simple geometry. The maximum measured Q of ~ 700 corresponds to a cavity with a hole diameter of 200 nm, a waveguide width of 520 nm and the inner holes shifted outwards by 25 nm. This experimental result is in very good agreement with the expected value of ~ 750 (simulations of Figure 3.7).

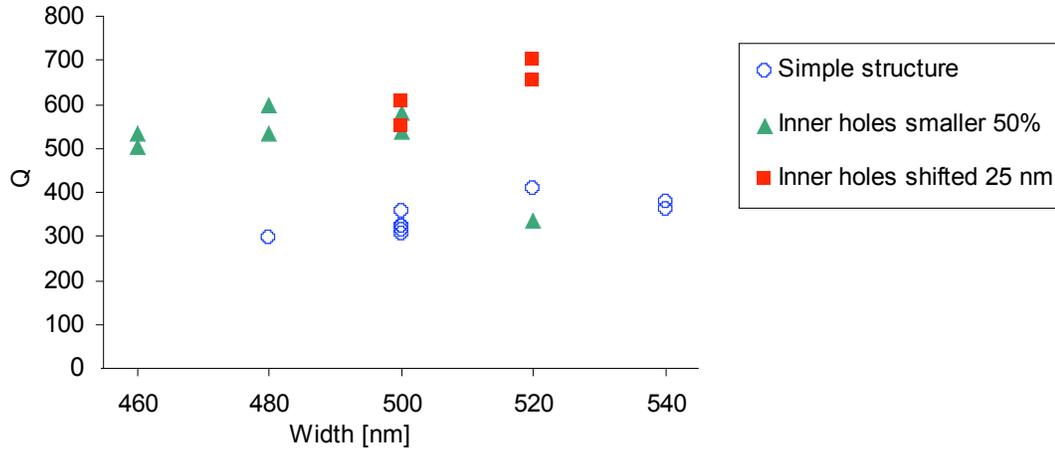


Figure 5.10 Experimental Q vs. waveguide core for different cavity geometry.

5.4 Problems with Second Harmonic Generation

If we summarize the results achieved so far, we can see that our experiments at the FH wavelength are in very good agreement with the theoretical expectations for both plain waveguide and cavity properties. However, although the required Q to significantly enhance nonlinear optical processes was attained, no SH was detected. The causes of this behavior were discovered when the numerical tools to fully simulate (in three dimensions) the nonlinear behavior of the device were developed and accurate simulations were performed. Two critical design problems emerged concerning the confinement ability of the cavity and the transmission properties of the mirror at the SH. The first problem was the effective coherence length being shorter than originally calculated using 2D simulations. Moreover, the length of the cavity mode was longer than expected. These combined effects resulted in the cavity mode being longer than

one coherence length (i.e., including the portions in the mirrors). The additional effect of strong mode distortion inside the mirrors further affected the efficiency.

The second problem concerned the mirror transmission at the SH. Although a 1D quarter-wave multilayer stack provides high reflectivity at the design wavelength and perfect transmission at the SH, this assumption is not valid in three dimensions. On the contrary, the SH transmission through a (tightly-confining) waveguide mirror is more and more suppressed as the FH reflectivity is increased. An explanation of this behavior will be given in Chapter 7, together with a solution to this problem.

With a study of generation efficiency vs. cavity length, we discovered that the first problem could have been considerably alleviated with a more confining design, involving a shorter cavity. Such structure could have improved the generation efficiency by more than 50x. However, this new design would not have fixed the fundamental and much more serious transmission problem. Furthermore, since essentially all the degrees of freedom in the cavity design were used up to create the necessary confinement, no room was left to design a solution for the second problem. Because of these reasons we decided to abandon this design. The two aforementioned problems were faced and solved with new designs of both waveguide and cavity, which are the topics of Chapters 6 and 7, respectively.

We conclude this chapter with one final remark. Although these two problems would have considerably decreased the generation efficiency, one could think that with enough power (i.e., obtained by using an EDFA), it might have been possible to detect at least a weak trace of SH. Unfortunately, a third complication of practical origin, that became clear only by characterizing the waveguides of Chapter 6, would have ruled out this possibility. It turns out, in fact, that the SH propagation loss in these waveguides is much higher than the FH loss. Ignoring this fact, the microcavity devices were placed in the center of the waveguide, typically at least 300 μm from the output facet, so that an additional propagation loss of several dBs would have most probably lowered the signal into the noise level.

CHAPTER 5. CHARACTERIZATION

Chapter 6

Birefringent waveguides

In this Chapter we will deal with the issue of increasing the coherence length of a tightly confining waveguide. We will see that this problem can be simply and elegantly solved by using artificial birefringence (Section 2.3.3). Employing this technique in regular GaAs/AlGaAs waveguides can be complicated. The required selective thermal oxidation may in fact be incompatible with other processing steps. However, it can be easily implemented in our devices since AlO_x is already employed as a cladding layer, so that no additional fabrication steps are required.

Moreover, not only the fabrication process is essentially the same as for the microcavity, but also the optical setup and the characterization techniques are essentially unchanged. Because of these reasons we were able in a very short time to fabricate and test these new birefringently-phaseshifted structures. Although in our case the fabrication was straightforward, the real challenge was the design of a structure that is tolerant to fabrication imperfections and deviations from the nominal design.

6.1 Design

In this Section we discuss how we can achieve phasematching in a tightly confining waveguide using a multilayer core. Consider the waveguide structure shown in Figure 6.1. This structure consists of a $\sim 300\text{nm}$ -thick $\text{Al}_{0.5}\text{Ga}_{0.5}\text{As}/\text{AlOx}/\text{Al}_{0.5}\text{Ga}_{0.5}\text{As}$ core on top of $2.5\mu\text{m}$ AlOx cladding. Air provides the lateral and top confinement. Notice that this design is similar to the one described in Chapter 3, except for the thin AlOx layer in the center of the AlGaAs core. As before, the aluminum concentration of the AlGaAs layers has been chosen to obtain transparency at the SH. The thin AlOx layer does not significantly affect the TE mode profile at the fundamental wavelength (Figure 6.1a), but generates large discontinuities in the TM mode at the second harmonic (Figure 6.1b). These discontinuities have the effect of pushing a significant fraction of the field in the low-index (AlOx) layer. The effective index can then be lowered enough to be matched with the FH index, thus achieving phasematching.

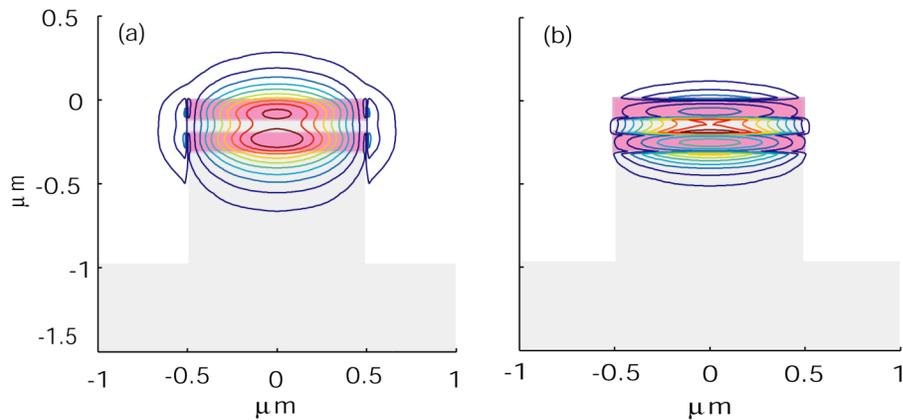


Figure 6.1 FH TE (a) and SH TM (b) electric field mode profiles of the waveguide.

However, any roughness or deviation from the ideal geometry will result in phase variation (Section 2.2.3) and, ultimately, in a finite coherence length. So, in designing an optimum structure, we have to find a balance between high tolerance to fabrication imperfections, which would suggest a large, less confining structure, and high conversion efficiency, which, conversely, requires tight confinement.

We simulated the waveguide modes using the same 2D semivectorial, finite-difference, frequency-domain algorithm⁵⁷ used in Chapter 3. This approach solves a scalar equation for quasi-TE (or quasi-TM) modes, with the assumption that the longitudinal field components are small compared to the dominant transverse component. This approximation may seem too rough for tightly confining waveguides; however, this algorithm produced results in good agreement with exact 3D FDTD simulations and, more importantly, with the experiments (e.g., the group index measurements in Section 5.2.2 and dispersion curve measurements in Section 6.3.1). The advantage of using this algorithm is the computational speed: given a specific geometry, the mode profiles and the effective indices of FH and SH modes can be typically calculated within seconds on a desktop PC.

In order to find the optimal design, we had to explore a multidimensional parameter space. Of course, the brute-force approach of simulating every possible structure would have been impractical, even with a very fast code; thus, we adopted the following strategy. We fixed the parameters of the vertical structure (indices and thicknesses of the layers) and simulated the waveguide modes for different wavelengths (distributed over the optical testing windows) and different widths. We chose wavelength and width as sub-parameter space, since these parameters can be more easily varied experimentally (by, respectively, tuning the laser wavelength and changing the e-beam lithography pattern). The layer parameters are in fact set at the time of the epitaxial growth. Changing them would require another time (and money)-consuming wafer growth. Moreover, the incremental step between each simulation can be fairly coarse, since the modes profiles and the effective indices, vary slowly with the wavelength and the width.

Then, for each set of simulations, we find the locus in the wavelength-width plane for which the effective indices of FH and SH are equal. This locus, that we called the “phasematching (or PM) curve”, describes the relationship between phasematching (PM) wavelength and waveguide width. Because of the computational speed of our algorithm, one PM curve can typically be calculated in a few hours, using a 1GHz, PentiumIII processor.

We then calculate the PM curve for different vertical structures. Initially, we reduced the number of parameters to two: the thicknesses of the AlGaAs layers, assumed equal, and the thickness of the AlOx layer. For convenience, from now on, we will refer to the AlGaAs layers as “Core” and to the AlOx layer as “Space”. For the Core, we used the index of $\text{Al}_{0.5}\text{Ga}_{0.5}\text{As}$, calculated using the expression in the paper by Deri et al.⁷⁸; for the Space, we employed an index of 1.58, measured by ellipsometry on test samples, and in agreement with the reported values^{79, 80}. The effects of the AlOx index variations will be discussed in detail in Section 6.3.2.

The dependence of the PM curve on Core and Space thickness is illustrated in Figure 6.1a and Figure 6.1b, respectively. A first observation is the large tunability of the structure with a relatively small variation of the waveguide width. This characteristic provides a simple tunability “knob”: complex structures with waveguide sections generating SH at different wavelengths could in fact be realized simply by changing the width in each section.

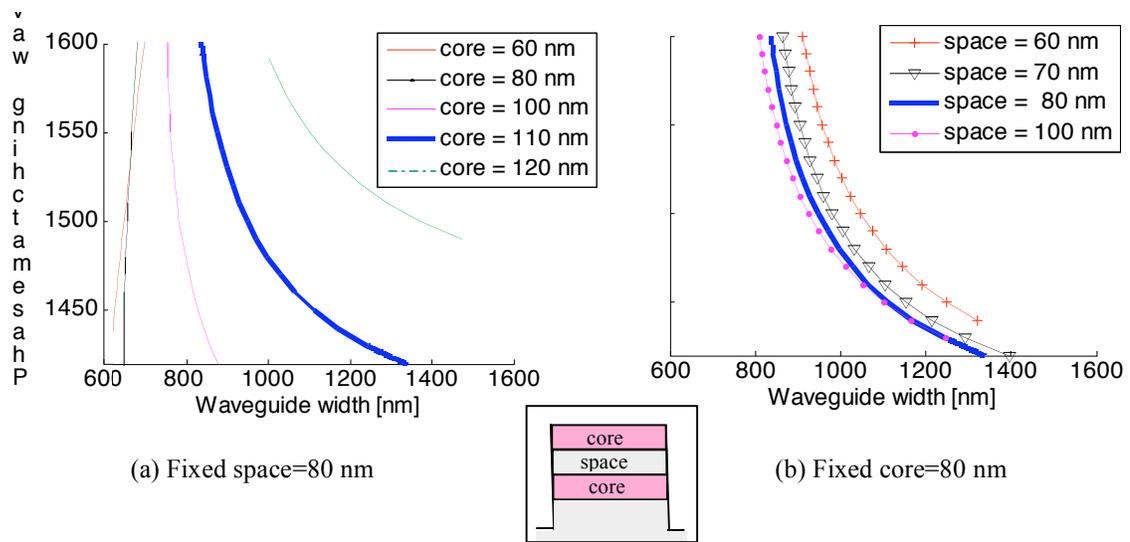


Figure 6.2 Phase-matching wavelength vs. waveguide width for different thicknesses of Core (AlGaAs layer) (a) and Space (AlOx layer) (b).

However, this large sensitivity to geometrical parameters becomes a robustness issue: the PM curve, in fact, has a strong dependence on the Core thickness. In particular, for a Core thinner than $\sim 100\text{nm}$, the curve is practically vertical. This

means that width variations of only a few nanometers ($\sim 0.1\%$ of the width) can shift the PM wavelength by several tens of nanometers.

This sensitivity problem can be better understood looking at Figure 6.3. This graph shows the derivative of the FH and SH index difference with respect to width, as a function of the phasematching wavelength. This quantity gives a measure of the tolerance of the structure: the smaller this value, the more tolerant will be the structure to variations in waveguide width. The curves in this figure confirm that the sensitivity to width variations grows with the wavelength and is inversely proportional to the Core thickness (while it does not depend significantly on the Space thickness).

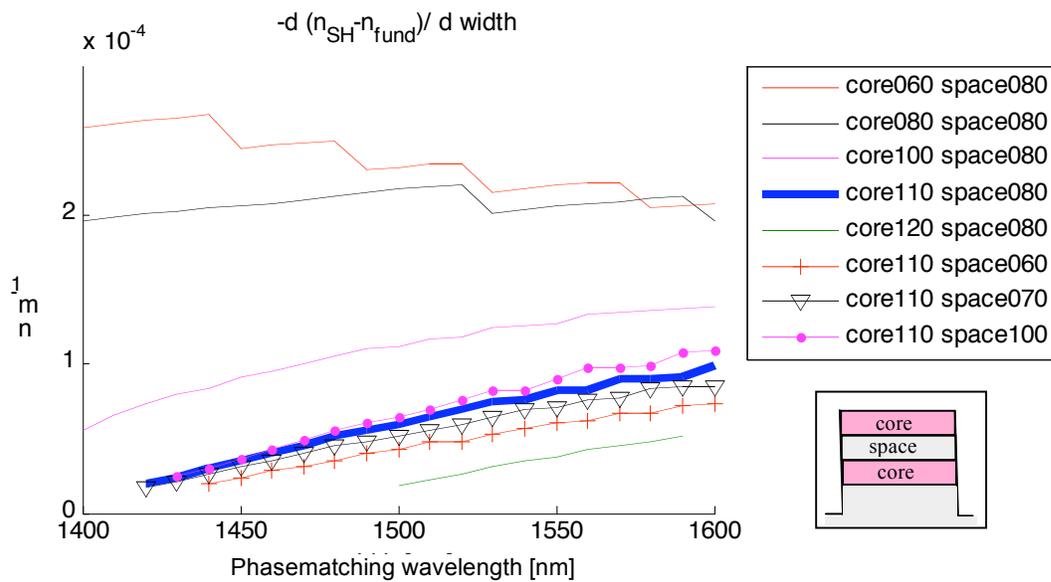


Figure 6.3 Derivative of the FH and SH index difference in respect to the waveguide width vs. PM wavelength for different thicknesses of Core (AlGaAs layer) and Space (AlOx layer)

This discussion brings us to the definition of our *first design criterion*: the sensitivity to width variations must be sufficiently low that width fluctuations due to sidewall roughness do not significantly affect the PM wavelength. Practically, this

results in a lower bound for the Core thickness of $\sim 110\text{nm}$, which in turns imposes a lower limit of $\sim 900\text{nm}$ on the width.

Next, we will consider the normalized conversion efficiency (or overlap efficiency) of the structure. Its expression can be easily calculated as an overlap integral of the electric field distributions of the FH and SH modes profiles, as defined by Equations (2.16) and (2.18). Figure 6.4 shows the overlap efficiency vs. waveguide width, where the efficiency is calculated at the corresponding PM wavelength. Also in this case, we see a much stronger dependence on the Core thickness than on the Space thickness. The *second criterion*, corresponding to the maximization of the efficiency, translates to a Core thickness of $\sim 100\text{-}110\text{nm}$.

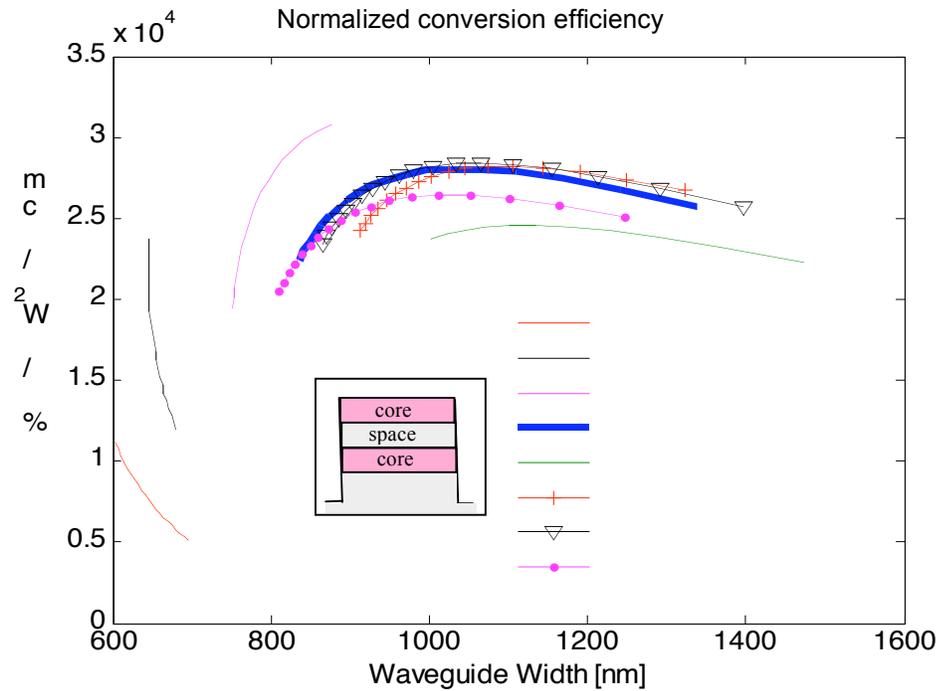


Figure 6.4 Overlap Efficiency vs. waveguide width for different thicknesses of core (AlGaAs layer) and space (Al_xO_y layer).

The absolute value of efficiency is, for most structures, relatively flat (vs. width) and greater than $20,000 \text{ \%}/\text{W}/\text{cm}^2$. To our knowledge, this value is the highest ever reported: for comparison, it is 20x higher than the efficiency of previously

reported for AlGaAs birefringently phasematched waveguides²⁸ and more than 200x higher than commercial periodically poled LiNbO₃ waveguides⁸¹.

We notice, finally, that this value is $\sim 20x$ smaller than that of the microcavity device (Section 3.2). This is expected, since this waveguide cross section is 4x larger and a considerable fraction of the fields is inside the AlOx layer (i.e., does not contribute to the integral of the effective area). However, this lower efficiency is largely compensated by a much longer device length.

A *third criterion* is given by the fundamental 3rd mode cutoff width. To achieve high conversion efficiency, the power injected in the waveguide should be entirely coupled to the 1st even mode (i.e., the mode shown in Figure 6.1a), since it is the only one which is phasematched. A strictly single mode waveguide would then be our first choice. However, this condition would impose a waveguide width narrower than $\sim 800\text{nm}$, which is incompatible with our first criterion (see Figure 6.2a). This condition can be relaxed allowing a second guided mode. The second mode is in fact of odd symmetry, and thus has a theoretical zero coupling-efficiency with a symmetric, input gaussian beam. This condition can be practically obtained if the input beam is perfectly perpendicular to the waveguide. In fact, any tilting would break the symmetry and increase the coupling with odd-symmetry modes. Experimentally, this condition is well satisfied; thus, our new design constraint becomes the 3rd mode cutoff width, which is $\sim 1150\text{nm}$, compatible with our first criterion.

From the criteria imposed so far, the optimal design structure corresponds to a Core and Space thicknesses of, respectively, 110 and 80-90 nm (i.e., the thick, solid, blue curve in the previous figures). Although the tolerance on the Space thickness is large ($>10\%$), the tolerance on the Core thickness is much tighter ($<5\%$): this problem will result in a major fabrication challenge.

Other geometrical parameters, such as etch depth, Core and Space refractive indices, cladding thickness and index, have been studied in a similar way. However, the PM curve proved to be sufficiently robust to variation of these parameters. One

exception (discussed in the next Section) is the AlOx refractive index, whose variation with experimental and environmental conditions can be surprisingly high.

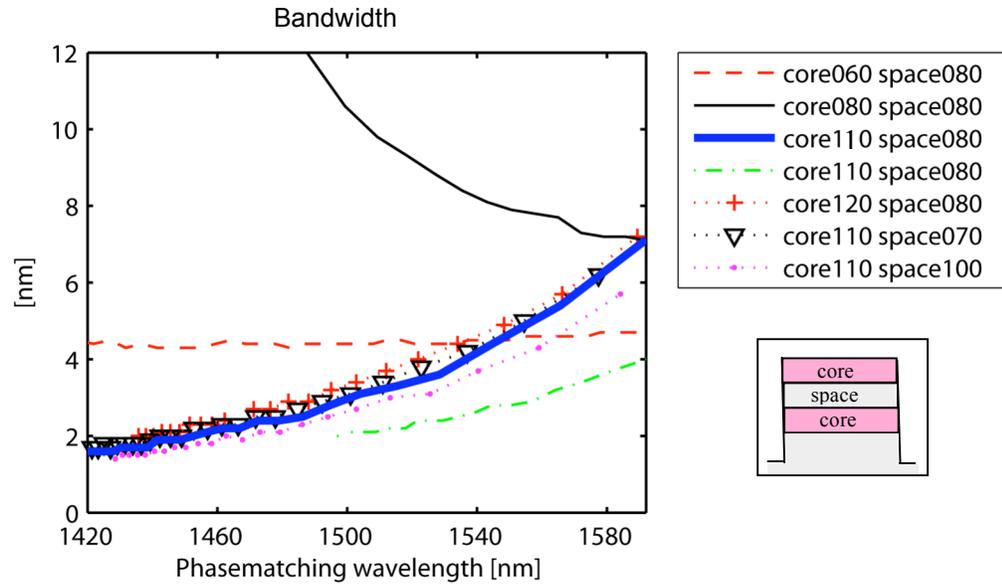


Figure 6.5 Bandwidth (FWHM) vs. PM wavelength for different thicknesses of Core (AlGaAs layer) and Space (AlOx layer), calculated assuming a device length of $600\mu\text{m}$.

One last property that is worth discussing is the bandwidth of the tuning curve for a typical device length of $600\mu\text{m}$. As we can see in Figure 6.5, the bandwidth (FWHM) for an ideal, lossless device strongly depends on the Core thickness. For the nominal design, the bandwidth increases with the PM wavelength and decreases with width; thus, once more, we are faced with the trade-off between larger bandwidth and tighter tolerances. As the slope of the PM curve becomes steeper, the bandwidth increases, and diverges as the PM curve approaches the vertical (i.e., for Core=80nm and Space=80nm). Such a structure would possess ultra-large bandwidth; however, the fabrication tolerances would become correspondingly tighter, making such a design (currently) impossible to realize.

6.2 Fabrication

The fabrication process is essentially the same as the one described in Chapter 4, except for minor differences that will be discussed in this section. One first obvious difference is the presence of the multilayer core in the epitaxially grown structure. The three core layers are visible in the SEM picture of the fabricated and cleaved waveguide shown in Figure 6.6a. Although the growth recipe is essentially the same, particular attention must now be paid to thickness tolerances. The typical thickness uncertainty of MBE-grown layers is $\sim 5\%$, which is theoretically sufficient for our purposes. However, because of the initial thick cladding layer, the calibration of the growth rate may shift during this time; thus, the thicknesses of the core layers may deviate by more than 5%. This problem may be solved with a growth interruption and recalibration; however, this adds to the fabrication complexity. This and other (apparent) “details” may result in a final structure sufficiently different from the ideal one to radically alter the PM curve.

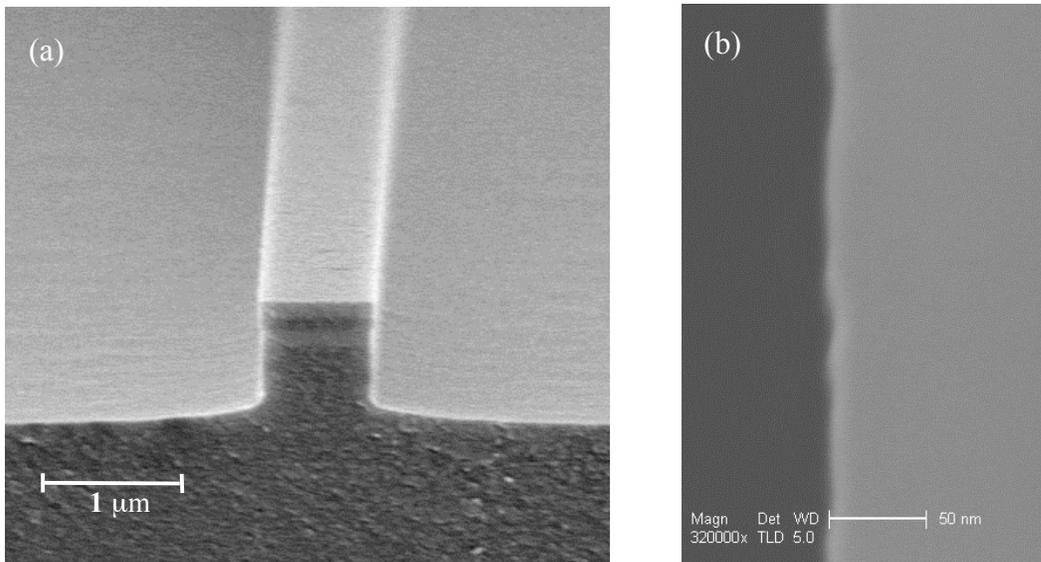


Figure 6.6 (a) SEM picture of a birefringent waveguide. The roughness on the front surface is due to sputtered gold particles, used for improving the SEM imaging. (b) Details of the sidewall roughness (top view).

A second issue is the thickness and composition non-uniformity *across* the wafer. Since the wafer is continually rotated during the growth, the non-uniformity is essentially in the radial direction. We can minimize this effect, by writing the waveguide patterns perpendicularly to the radial direction. Another source of non-uniformities is the dry-etching step: the etch rate is in fact higher in proximity (typically within 1-2mm) of the edges of the sample (edge-effect), thus limiting the usable area on the chip.

A third important fabrication issue is the roughness of the sidewalls and especially the drift of the waveguide width along the device length. In fact, as explained in the previous Section, any deviation from the ideal geometry would result in phase variation and ultimately, in decreased efficiency. Figure 6.6b shows the detail (top view) of a typical sidewall. Using image processing algorithms we extracted the roughness profile. For our best devices we were able to measure roughness as small as 4nm RMS, over a distance of 2 μ m. This value is, to our knowledge, the state-of-the-art for tightly confining AlGaAs waveguides. By taking multiple pictures along the entire length of the waveguide and stitching them together, we analyzed the width variations along the length. A width drift smaller than 20 nm over a distance of 500 μ m was measured.

This experimental width distribution was used as the functional form for Δn in Equation (2.22): the impact of this roughness and drift on the PM condition was calculated and the SH tuning curve was found to be only marginally affected. Finally, we mention that, although we could not perform a similar calculation for the deviation of the layer thicknesses (because of the difficulty of the measurement itself), we expect that these deviations should be sufficiently small.

Another issue is the partial oxidation of the AlGaAs core layers before and during the thermal oxidation. Since we do not use a GaAs cap (Section 4.1), if the wafer is left in air for a long time, a thin layer (typically tens of angstroms thick) of the top AlGaAs surface is oxidized. Although the oxidation rate is very low for

$\text{Al}_{0.5}\text{Ga}_{0.5}\text{As}$, the AlGaAs layers become thinner also during the thermal oxidation. Although small, the thickness variations deriving from these effects are important and they have to be compensated for by growing slightly thicker layers.

One last problem is due to the variation of the AlOx refractive index due to thermal oxidation parameters and the environmental conditions. It is well known that the refractive index of the thermal oxide can vary up to $\sim 5\%$ (between ~ 1.54 to ~ 1.62) depending on many factors, such as aluminum content, temperature and gas flow⁸². We found, for example, that the oxidation rate and the oxide index depend on the amount of water in the bubbler and on the conditioning time of the furnace tube. Although small variations of the index can be neglected, variations of the order of a few percent can strongly affect the slope and the position of the PM curve. It is then extremely important, for consistency of results, to maintain the experimental conditions as constant as possible.

AlOx has also the problem of being highly hygroscopic. Its index will drift over time until a saturation value that can be up to 10% different from the original value⁸². The saturation time greatly depends on the humidity conditions, going from a few days to several months. A way to mitigate this problem has not been investigated, but it could involve rapid saturation by exposing the sample to a high-humidity and (possibly) high temperature environment.

6.3 Optical characterization

After fabrication, the waveguides were tested using the optical setup described in Chapter 5. In this Section we first describe the linear characterization of the devices, using the Fabry-Perot technique (Section 5.2.1), and we then move to the analysis of SHG spectra.

6.3.1 Linear characterization

Figure 6.7 shows a typical Fabry-Perot spectrum: the regularity of the fringes across the entire spectrum indicates high sample quality and good coupling conditions.

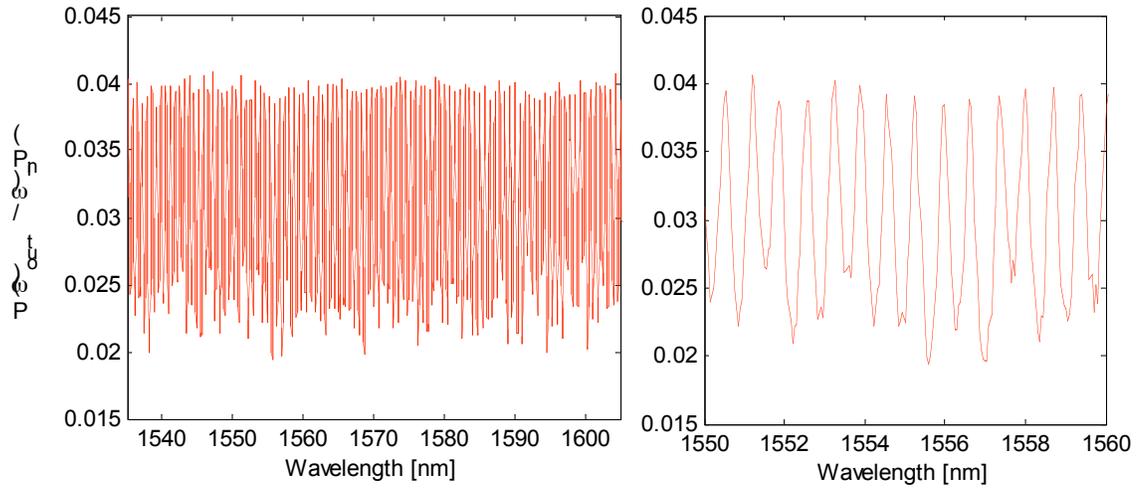


Figure 6.7 (a) Transmission spectrum at the FH wavelength. (b) Detail of the same spectrum.

To extract a value of the loss, the reflectivity of the facet was estimated using 3D FDTD simulations. To simulate the facet, the waveguide was truncated along the propagation direction. A plane source of the size of the waveguide core, with a gaussian profile in the temporal domain, excites the guided mode of the waveguide. The intensities of the forward propagating and of the reflected pulses are recorded.

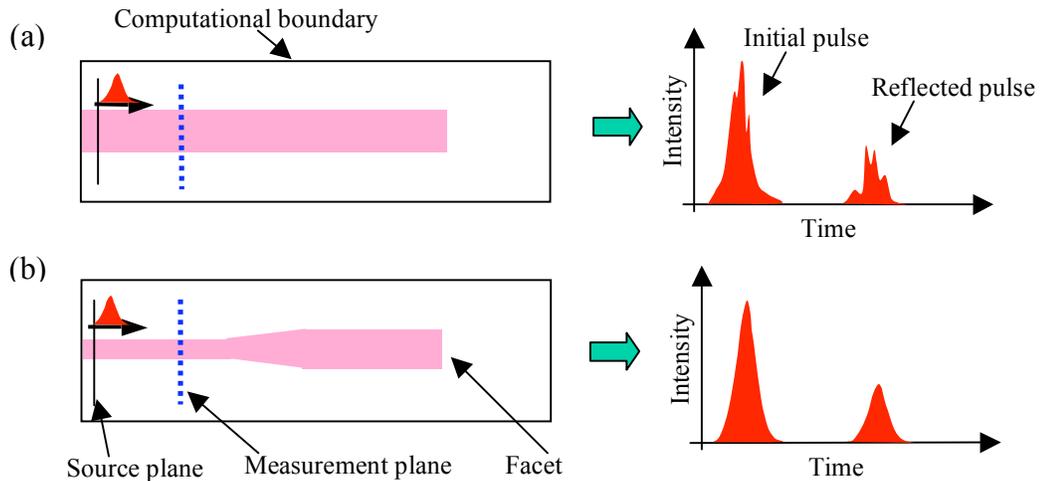


Figure 6.8 Illustration of FDTD simulations of facet reflectivity without (a) and with (b) mode filter.

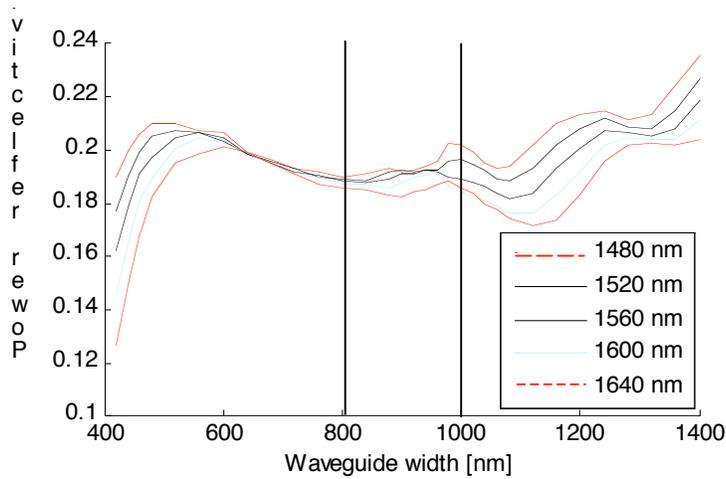


Figure 6.9 Facet reflectivity vs. waveguide width at different wavelengths. The width range of interest is highlighted.

However, because of the larger core, (i.e., compared to the microcavity case), the coupling to leaky modes by the source or by the reflecting facet is not negligible. This coupling results in noisy and irregular reflection spectra. The problem can be mitigated by introducing a mode filter in the source section. This filter consists of a narrow, single-mode section of waveguide slowly tapered to the desired width, as illustrated in Figure 6.8. Since we are interested in the reflection of the 1st guided mode back *into* the same mode, the taper is used to filter both the forward and the reflected wave. The only inconvenience is the numerical roughness added by the taper; however, this problem can be solved with a proper normalization.

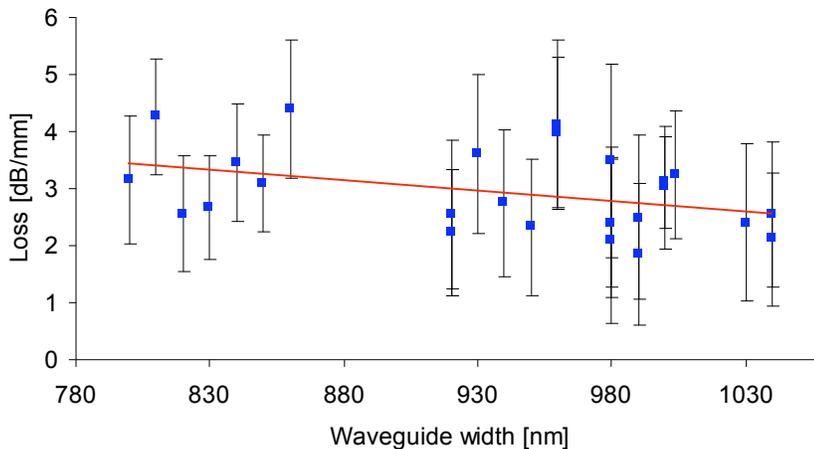


Figure 6.10 Propagation loss vs. waveguide width. A trend line s also plotted

The calculated reflectivity vs. waveguide width, shown in Figure 6.9, has fluctuations smaller than 1%, which are sufficiently low for our purpose. The reflectivity of the waveguides at 1550 nm, that we assume constant for the range of widths considered (800-1000nm), is $\sim 19 \pm 0.5\%$.

Using the previous value for the reflectivity, we can extract the propagation loss from the Fabry-Perot data. A plot of the propagation loss for waveguides of different width appears in Figure 6.10. The loss dependence on width is very weak; this indicates that the contribution to the loss from sidewall roughness is small. The lowest measured loss is 1.84 dB/mm, with typical values of 2-3 dB/mm. To our knowledge, this loss constitutes the state-of-the-art for tightly confining AlGaAs waveguides.

6.3.2 Nonlinear characterization

The SH transmission spectra were recorded simultaneously with the FH spectra, which were used for the normalization. Normalizing the SH power to the *input* fundamental power squared would be an intuitive normalization; however, we will see later in this Section that normalizing to the *output* fundamental power is more convenient for a number of reasons.

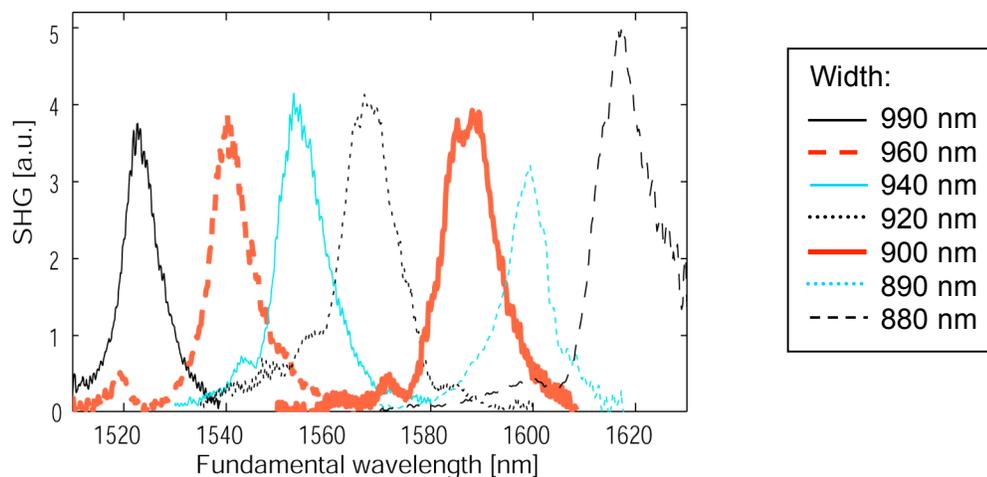


Figure 6.11 (a) SH tuning curves, normalized to the square of the FH output power, for waveguides of different width.

We measured transmission spectra for waveguides of different width and normalized them using the output FH power. The resulting tuning curves are plotted in Figure 6.11. Notice that the entire tuning range of our laser (>120 nm) can be covered by varying the waveguide width by less than 100nm. By measuring the peak wavelength vs. waveguide width (Figure 6.12), we plot the experimental PM curve, which is in excellent agreement with the theoretical predictions (Section 6.1).

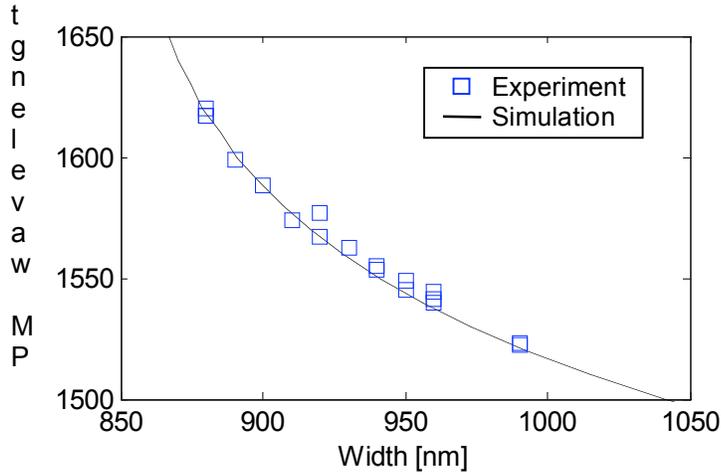


Figure 6.12 Experimental (squares) and simulated (line) PM curve vs. waveguide width. The Core and Space of the theoretical structure are respectively 110nm and 90nm-thick.

Once we have verified that our PM curve model is valid, we can proceed to the analysis of the single tuning curves. Such analysis can be theoretically carried out using either one of the two aforementioned normalizations.

In Figure 6.13 we plot the SH power normalized by the *input* fundamental power. This curve represents the total external conversion efficiency and is expressed by the following equation:

$$\frac{P_{2\omega}(out)}{P_{\omega}(in)^2} = \eta' \left(C_{2\omega} C_{\omega in}^2 T_{2\omega} \right) \left(\frac{T_{\omega}}{|1 - R_{\omega} \exp(-ik_{\omega} L - \alpha_u L)|^2} \right)^2, \quad (6.1)$$

where $P_{2\omega}(out)$ and $P_{\omega}(in)$ are, respectively, the SH output and FH input powers measured at the detectors; η' is the single-pass conversion efficiency of the waveguide of length L (i.e., without considering reflections at the facets), as defined in Equation(2.21); $C_{2\omega}$ and $C_{\omega in}$ are, respectively the SH output and FH input coupling coefficients; T_{ω} and $T_{2\omega}$ are the facet transmissions at the fundamental and SH wavelengths; R_{ω} , α_{ω} are the facet reflectivity and the propagation loss at the FH wavelength. The last term of this expression describes the modulation due to the Fabry-Perot etalon effect.

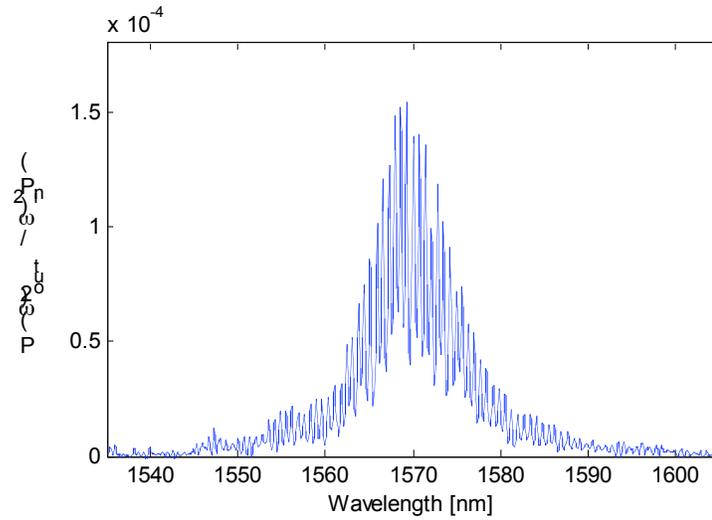


Figure 6.13 (a) SH output power normalized by the square of the FH input power.

Although Equation (6.1) gives the total amount of generated SH, it is not suitable for the characterization of the intrinsic properties of the device. In fact, it depends on the input coupling coefficient ($C_{\omega in}$), which is usually small ($\sim 10\%$), not repeatable from waveguide to waveguide and difficult to estimate. Moreover, because of the Fabry-Perot fringes, it is not convenient for the fitting of the experimental data. For this purpose, a better normalized curve is obtained using the *output* FH power. If we define the FH throughput as:

$$(6.2)$$

$$\frac{P_{\omega}(out)}{P_{\omega}(in)} = C_{\omega in} C_{\omega out} \frac{T_{\omega}^2 e^{-\alpha_1 L}}{|1 - R_{\omega} e^{-ik_{\omega} L - \alpha_1 L}|^2},$$

where $C_{\omega out}$ is the FH output coupling coefficient, the new “output”-normalized efficiency is obtained by substituting Equation (6.2) into Equation (6.1):

$$\frac{P_{2\omega}(out)}{P_{\omega}(out)^2} = \eta' e^{2\alpha_1 L} \frac{C_{2\omega} T_{2\omega}}{T_{\omega}^2 C_{\omega out}^2}. \quad (6.3)$$

The previous expression is independent not only of the Fabry-Perot term, but also of the FH input coupling efficiency. The required coefficients are now only the output coefficients, which are much larger (40-60%) and less dependent on the alignment conditions. The internal conversion efficiency, η' , can then be easily calculated from the experimental, output-normalized tuning curve using Equation (6.3). As expected, in the resulting curve, plotted in Figure 6.14, the Fabry-Perot fringes normalize out, leaving a smooth profile.

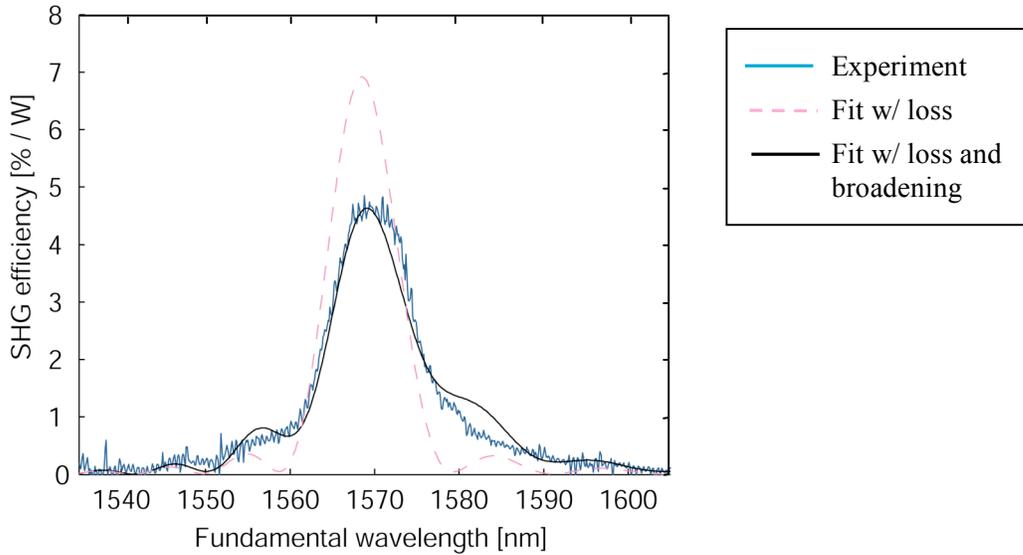


Figure 6.14 Normalized SH efficiency: experimental curve (solid light blue), fitting including loss (dashed pink) and fitting including loss and broadening effect (solid black).

From this experimental curve, we can now extract the value of the loss at the SH. As discussed in Section 2.2, the area under the curve depends only on the loss and

not on the disorder. Then, using the values of (i) the measured FH loss (2-2.5 dB/mm); (ii) the theoretical normalized conversion efficiency (Section 6.1) and (iii) the theoretical dispersion of the waveguide (verified with the PM curve, Figure 6.12), we fit the area using the SH loss as fitting parameter. The resulting SH loss coefficient is estimated to be 30-40 dB/mm. The causes of such a large loss coefficient are not yet clear and possible explanations will be discussed in the next section.

Although the area of the fitted curve (pink-dashed line in Figure 6.14) is equal to the area of the experimental curve, their shapes do not match well. This observation suggests that we need to include disorder in our model. As discussed in Section 2.2.3, disorder broadens the tuning curve without affecting its area. Since we do not know its functional form, we modeled the disorder as a linear drift of a geometrical parameter along the entire length of waveguide. By fitting the experimental shape using the amount of drift per unit length as a fitting parameter, we were able to closely match the shape of the experimental and theoretical (black line in Figure 6.14) curves. The value of the extracted drift coefficient corresponds to a variation of the SH and FH index difference of $\sim 3.3 \cdot 10^{-5} \mu\text{m}^{-1}$. The corresponding variation of the geometrical parameters has been estimated from 2D frequency domain simulations.

Parameter	Sensitivity in $n_{2\omega} - n_{\omega}$ with respect to parameter	Total drift required over 600 μm
Core thickness [nm]	$3.8\text{e-}3 \text{ nm}^{-1}$	5 nm=4.5%
Oxide index	0.5	0.04 = 2.5%
Width [nm]	$2.6\text{e-}4 \text{ nm}^{-1}$	70 nm=7.3%

Table 6.1 Sensitivity and amount of drift of geometrical parameters required to fit the shape of the efficiency curve.

The results for three parameters are summarized in Table 6.1. The first column represents the change in index difference corresponding to a change in the parameter, and is a measure of the sensitivity of the PM curve to that parameter. The second column shows the total amount of required drift, assuming that the other parameters

do not drift. Obviously we can rule out the possibility of the width to be responsible of the curve broadening *by itself*. However, the amount of required drift of core thickness or oxide index (or a combination of both) is not unreasonable and could explain the broadening.

Finally, Figure 6.5 summarizes the results obtained from waveguides of several different widths and two lengths (300 and 600 μm). Different SH-generating lengths can be realized on the same sample by varying the waveguide width, so that only the portion of waveguide close to the output facet is phasematched. Both the area and the FWHM of the tuning curves are plotted in function of the PM wavelength. Corresponding theoretical curves that take into account loss and disorder are also plotted. Although the data variability is quite high, the model can reproduce the experimental trend, especially for the FWHM. It is also interesting to notice that the bandwidths for both lengths are similar. This fact is due to the combined effect of loss and roughness, which impact more heavily longer waveguides

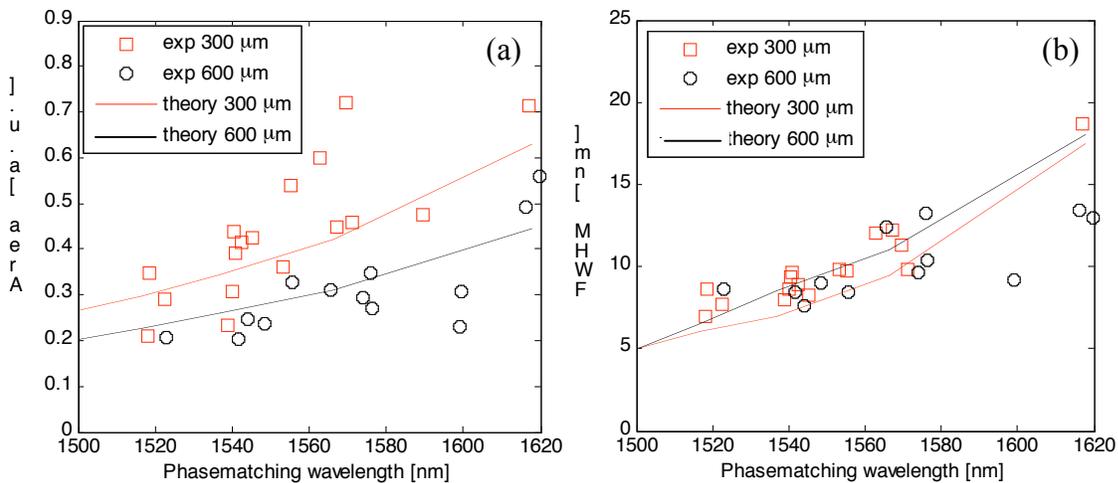


Figure 6.15 Area of the normalized efficiency curve (a) and FWHM of the tuning curve (b) vs. phasematching wavelength for two different waveguide widths. Both experimental data and theoretical curves are showed.

6.3.3 Second harmonic loss: possible causes

In this Section we discuss the possible causes of the large SH losses. Because of technical problems, we were not able to measure the SH loss by direct methods, such as cutback or Fabry-Perot, but only through extrapolation from SHG tuning curves. This fact considerably limited our ability to characterize the device at the SH wavelength and ongoing efforts are made to solve this problem. Nevertheless, using the previously described measurements, other published works and 2D numerical simulations, we can at least delineate the possible origins of such high loss.

We start by considering scattering due to roughness at the interfaces. Currently, an analytical model to describe this type of loss in a complex 3D multilayer structure is not available; moreover, 3D FDTD simulations would be computationally prohibitive. However, for a rough estimation, we can use a semi-analytical model, in which the loss coefficient is proportional to the magnitude of field at the interface. Using the volume-current method^{76, 83} and approximating the roughness as uncorrelated and using the free-space Green's function (i.e., neglecting the presence of the waveguide itself), the contribution to the loss coefficient from roughness at a given interface can be written as:

$$\alpha_I(\omega) \propto \oint_I \left(\Delta\varepsilon E_{\parallel} - \varepsilon \Delta \left(\frac{1}{\varepsilon} \right) D_{\perp} \right)^2 b^2 \frac{\omega^4}{V_g} ds, \quad (6.4)$$

where $\Delta\varepsilon$ and ε are, respectively, the difference and the highest of the dielectric constants of the materials on the two sides of the interface; E_{\parallel} and D_{\perp} are, respectively, the electric field component parallel to the interface and the electric displacement component perpendicular to the interface; b is the amplitude of the roughness; ω is the frequency; V_g is the group velocity, defined as the ratio of the speed of light in vacuum to the group index of the waveguide mode; the integral is taken along the interface, I . For typical waveguide modes, such as the ones shown in Figure 6.1, the contributions from the different interfaces, for FH and SH, are shown in Table 6.2. Before continuing our discussion, we have to point out that this is a very crude approximation and not valid for the interfaces inside the waveguide. However,

if we just consider the order-of-magnitude behavior, we can still make some interesting observations.

	FH TE [a.u.]	SH TM[a.u.]
Sidewalls	6.9	6.5
Interface 1	33	2299
Interface 2	61.3	2383.5
Interface 3	94.2	6147.1
Interface 4	113.2	6196.1

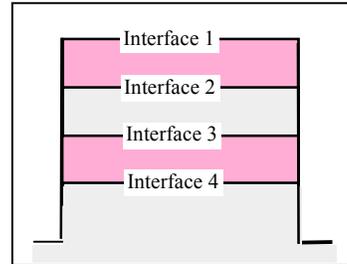


Table 6.2 Contributions of the roughness at different interfaces to the loss coefficient of FH TE and SH TM modes.

Three facts emerge from these simulations: (1) the horizontal surfaces are 10-20x more sensitive than sidewalls to roughness, (2) the sidewall roughness affects FH TE and SH TM modes similarly and (3) horizontal surfaces are $\sim 600x$ more sensitive to roughness at the SH than at the FH wavelength. Several interesting implications follow. Assuming that the losses are due entirely to scattering, from the mild dependence of the FH loss on the waveguide width (Figure 6.10), we deduce that a significant fraction of the loss must come from roughness at the *horizontal* interfaces. This assertion would require, from fact (1) and from the measured sidewall roughness of $\sim 4\text{nm}$ RMS (Section 6.2), a horizontal roughness of $\sim 1\text{nm}$, which is not unreasonable for AlGaAs/AlOx interfaces^{84, 85}.

From this observation and from fact (2), we can also infer that the SH loss is not strongly affected by sidewall roughness. Finally, if these assumptions are correct, from fact (3) we deduce that the huge losses at the SH could be explained simply by roughness at the horizontal interfaces even as small as $\sim 1\text{nm}$. Of course, this could be true only if the horizontal interfaces, are effectively as rough as expected. Attempts to analyze a waveguide section using high-resolution Transmission Electron Microscopy (HR-TEM) were made, but failed because of technical problems. We believe that

such an analysis would be of extreme interest and would help to confirm (or rule out) the possibility of loss entirely due to scattering.

A second possibility would be absorption loss due to residual arsenic from the oxidation process, either in the AlOx layer or at the interfaces. It is well known that a small fraction of the arsenic does not leave as arsine (AsH₃), but is trapped in the AlOx. Free arsenic could create levels within the bandgap and allow two-step absorption processes. Very few works have been published on this topic^{70, 86, 87} and the physics and chemistry of the oxidation process are still not completely clear. However it has been suggested⁸⁸ that irradiation using low-energy hydrogen ions could be useful in removing the excess arsenic and reducing the loss. Such treatment would be worth trying on future samples.

In conclusion, although the causes of the SH loss are presently unidentified, we believe that either scattering from roughness at the horizontal interfaces or absorption mechanisms due to the oxidation process (or most probably a combination of both) are the two most probable sources.

Chapter 7

Birefringent resonant cavities

In the previous chapter we solved the problem of the tight constraints imposed by a short coherence length using artificial birefringence. Although the device length can now be made theoretically infinite, the large SH loss limits the maximum length to 300-400 μm . A way to increase the conversion efficiency without increasing the device length is to use resonant cavities, as explained in Section 2.4. With a cavity resonant at the FH, the conversion efficiency can be increased by a factor proportional to the *square* of the cavity finesse. Such solution could reduce the required length, thus mitigating the problem of high SH loss and reducing the device size. However, in order to use cavities, we first have to solve the second major problem encountered with photonic crystal microcavity devices (Chapter 3): the poor transmission of the mirror at the SH. This problem, and how to solve it, are the topics of the first two sections of this chapter. A discussion of the fabrication issues and the characterization of the new cavity device will follow. Finally, we will discuss the trade-offs between efficiency and other figures of merit, such as bandwidth and tuning range.

7.1 The conventional mirror

In this Section we will discuss the reasons for the poor SH mirror transmission. Consider a conventional waveguide Distributed Bragg Reflector (DBR). This mirror consists of a periodic modulation of a geometrical parameter, such as the waveguide width, along the propagation direction (Figure 7.1a). 3D FDTD simulations of this structure show that, although high FH reflectivity is easily achievable, the SH transmission is almost completely suppressed. Moreover, this behavior does not depend on the specific type of modulation (i.e., holes or trenches). We also point out that increasing the mirror reflectivity to improve the efficiency would result in further suppression of the SH transmission.

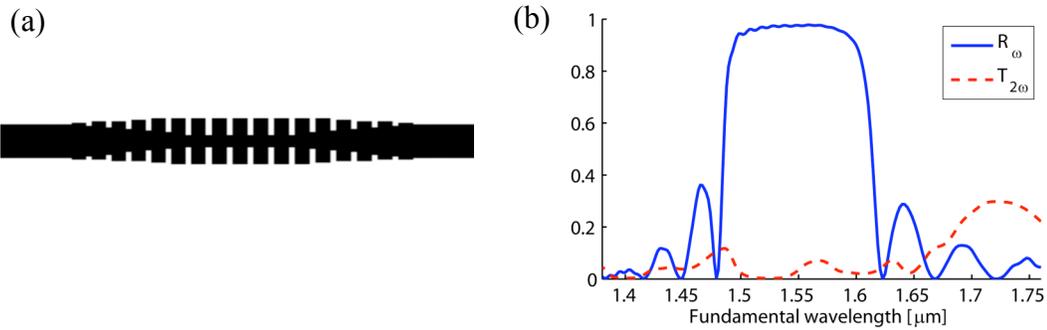


Figure 7.1 Top view of a conventional DBR mirror and simulated FH transmission and SH reflection spectra.

In order to understand this behavior we start by analyzing the band diagram of a plain waveguide, represented in the schematic of Figure 7.2a. The first guided TE and TM modes of an asymmetric, tightly confining waveguide are shown. The abscissas of the intersection points between the two curves and the FH and SH frequency lines represent the k-vectors of the two modes. The region above the light-line, corresponding to the AlOx index, is called the “radiation modes region”. In this area, no guided modes exist. If now we introduce a modulation of period $a \sim \lambda_\omega / 2n$, where λ_ω is the fundamental wavelength and n is the effective index of the waveguide

mode, the band diagram becomes periodic (with period $2\pi/a$) and symmetric around π/a ⁸⁹ (Figure 7.2b). The TE mode curve splits, creating a bandgap, which is at the origin of the FH high reflectivity. However, the folding of the band diagram also results in the SH frequency lying entirely in the radiation mode area. This means that any attempt to propagate the SH through such a structure results in severe radiation loss. From this analysis, we can see that the high SH loss originates from the presence of the modulation itself.

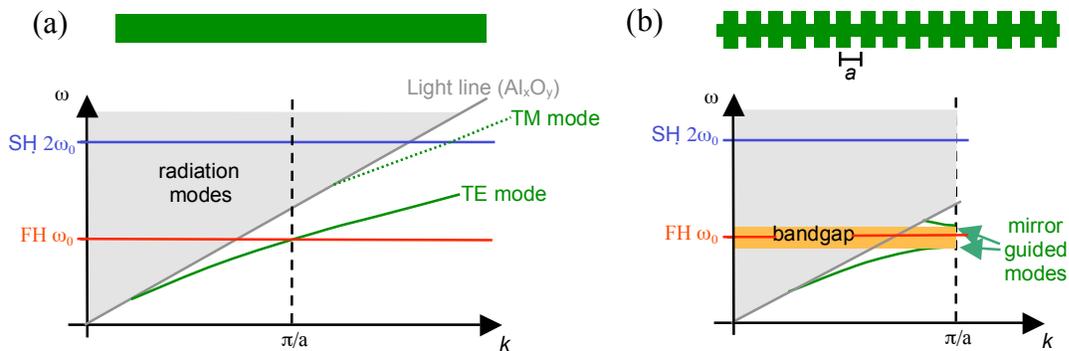


Figure 7.2 Schematic of band diagram of a plain waveguide (a) and of a periodic structure (b).

A direct solution to this problems would be to design a structure in which the SH frequency line would have a portion falling below the light-line and a TM guided mode in that region. This solution may involve the use of suspended structures, in which the light-line (now corresponding to air instead of AlOx) has a steeper slope, and of a larger waveguide core, which could raise the effective index. However, suspended structures are considerably more complicated to fabricate and less stable. Moreover, a large core would result in a weakly confining, multimode waveguide. A second possible solution consists in using unoxidized AlGaAs as the cladding. Such a waveguide could be single mode, but, again, at the price of losing the tight confinement. In the next Section we will present an indirect solution, which does not suffer from any of these drawbacks.

7.2 Dichroic mirror

In this Section we discuss the design of a mirror for tightly confining waveguides with dichroic behavior: highly reflective at the FH and highly transmissive at the SH. The basic idea is to fabricate a structure periodic for the FH wavelength, in which the SH could propagate without “seeing” such periodicity. A hint at designing such a structure is given by the difference in the TE FH and TM SH mode profiles in the waveguide (Figure 6.1): because of the longer wavelength and the discontinuity at the sidewalls, the FH mode extends horizontally much more than the SH mode. This difference can be further accentuated by decreasing the width of the waveguide. The fundamental mode is thus squeezed out of the waveguide along the horizontal direction.

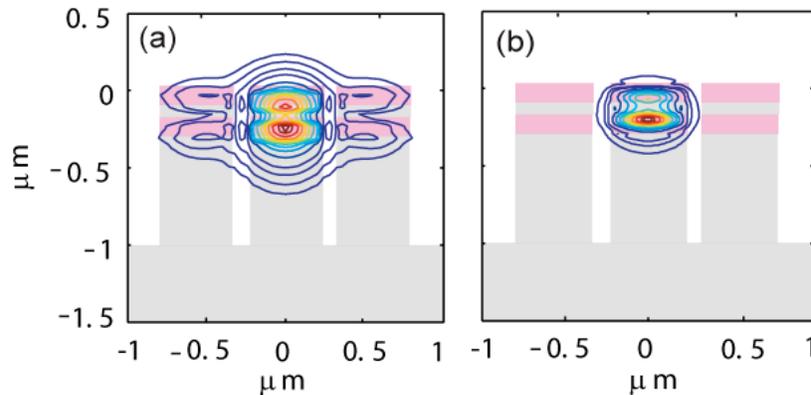


Figure 7.3 Intensity distribution of the FH TE (a) and SH TM (b) modes in the section of the external mirror (with connector, as shown in Figure 7.4b). These profiles are obtained by 3D FDTD simulations. The (inner) waveguide section is 460nm wide.

We can take advantage of this fact by placing the modulation outside of the waveguide core. In this way, while the FH can still interact with the mirror, the SH can be transmitted unperturbed, as shown in Figure 7.3. The FH reflection and SH transmission spectra (Figure 7.4a) confirm these expectations, yielding high transmission at the SH. However because now the modulation is far away from the core, the reflection of the FH is considerably reduced.

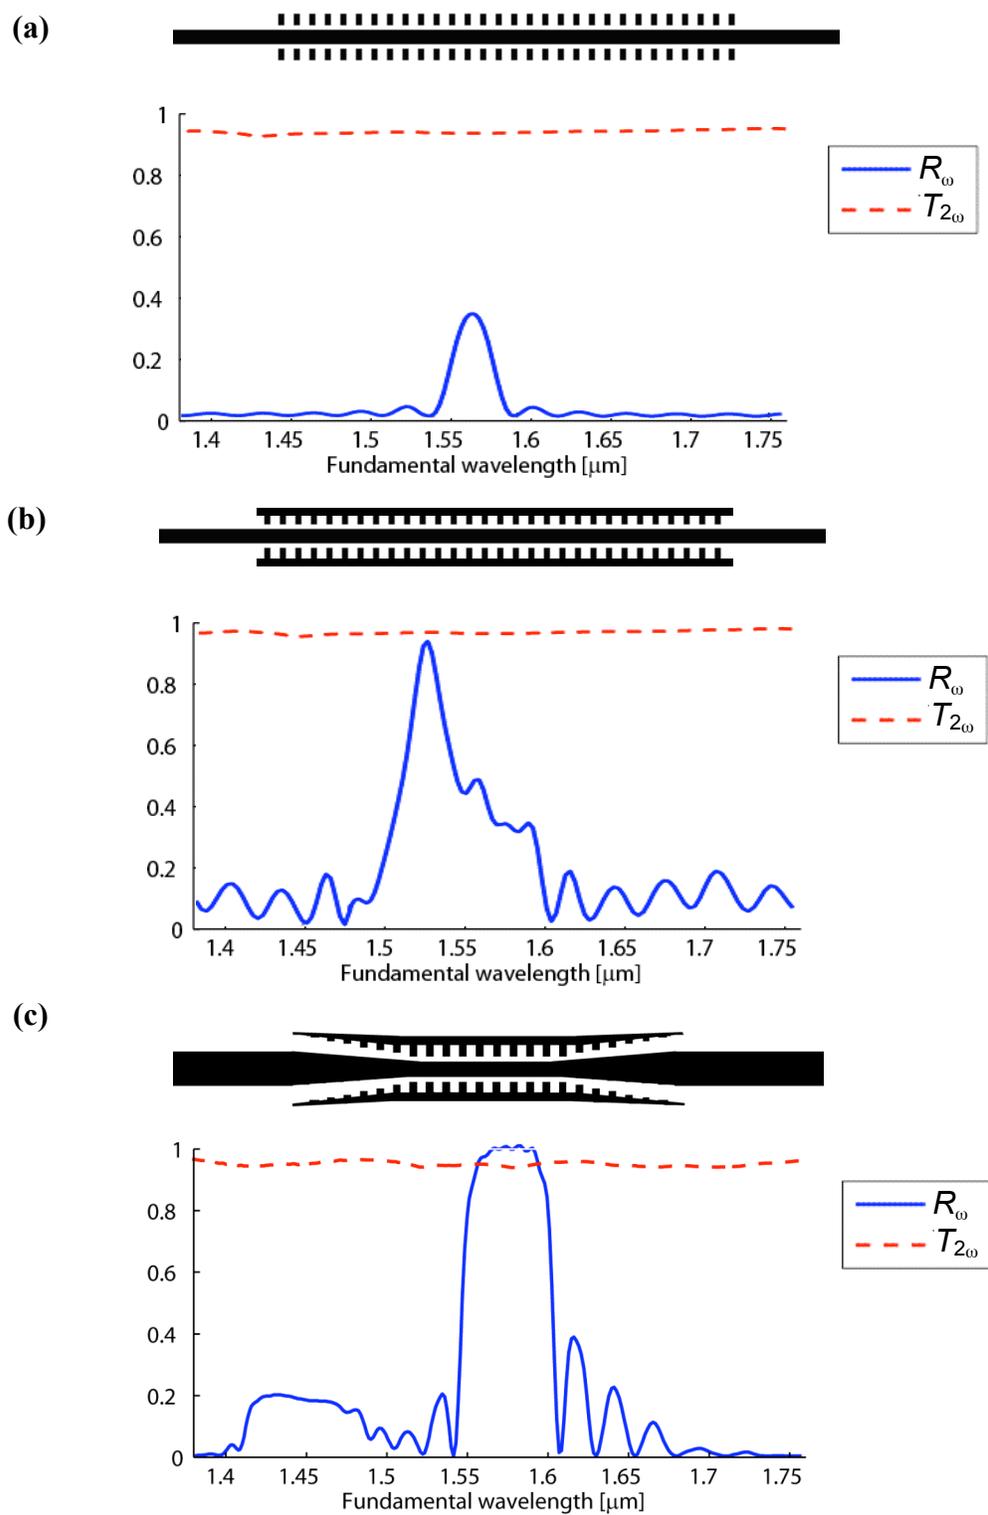


Figure 7.4 Simulated FH reflection and SH transmission spectra for an external mirror without (a) and with (b) an added connector, and final design with tapered ends (c).

This issue can be solved by adding a section connecting the posts, which we call a “connector”, as illustrated in Figure 7.4b. This idea is borrowed from photonic crystal theory: in fact, for in-plane, TE modes a larger bandgap is achieved using a hole array (instead of a post array), in which the dielectric regions are connected. The connector has a similar function in connecting the electric field lines between the posts. However, although the peak FH reflection is very high, its bandwidth is still considerably narrow. This is due to the abrupt interface between waveguide and mirror; this interface is the source of large radiation losses because of mode mismatch. This problem can be solved by tapering the input and output ends of the mirror, as shown in Figure 7.4c. Notice that both the mirror part and the waveguide part have to be tapered: in fact, the waveguide width must match the phasematching width, typically much larger ($>800\text{nm}$) than the width inside the mirror ($440\text{-}480\text{nm}$).

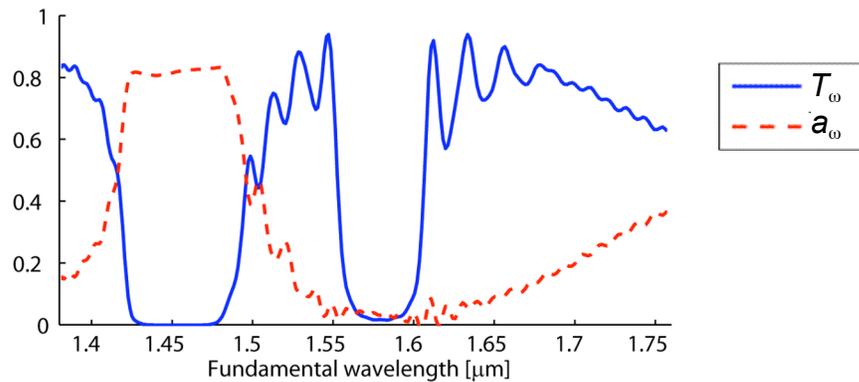


Figure 7.5 Simulated transmission and loss spectra at the FH wavelength for the tapered mirror of Figure 7.4c.

Before describing the optimization of the mirror structure, it is interesting to look at the FH transmission and loss spectra, displayed in Figure 7.5. Notice the dip in the transmission spectrum in the wavelength range of $1.45\ \mu\text{m}$, corresponding to high radiation losses. We believe that this behavior, caused by the presence of the connector, is due to the efficient coupling to a leaky mode (i.e., a mode above the light-line) of the mirror. However, we emphasize that such a dip does not affect the behavior of the mirror in the region of interest at $1.55\ \mu\text{m}$. The increasing loss at

longer wavelengths can be explained by the cutoff behavior of the narrow inner waveguide.

The optimal mirror structure was the result of an optimization process of several geometrical parameters, indicated in Figure 7.6. The rest of this section will be dedicated to a description of the most critical parameters and their interplay in determining the final mirror characteristics. The *inner waveguide width* determines the horizontal extension of the field outside of the waveguide, thus affecting the mirror strength. The narrower the waveguide, the larger the bandgap at the FH, but more important is also the scattering loss due to sidewall roughness. The width must be optimized together with the *waveguide-tooth gap*. This gap should be as small as possible to increase the mirror strength, but without significantly perturbing the SH mode. The ranges of these parameters are 460-500nm for the width and 100-140nm for the gap. These values take into account the effect of losses due to fabrication issues and will be further discussed in the following section.

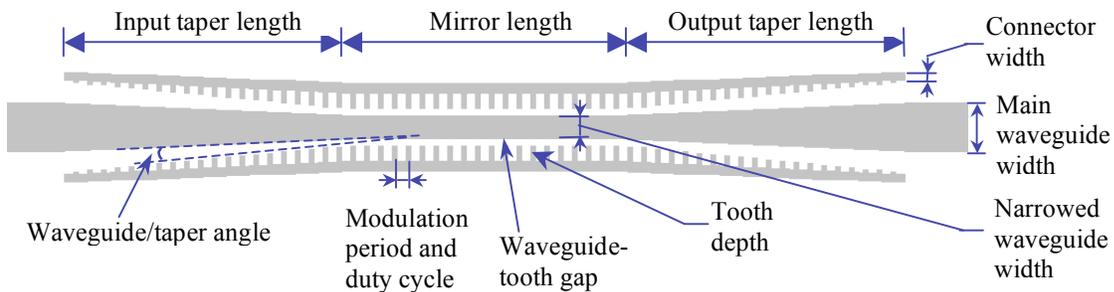


Figure 7.6 Schematic of the dichroic mirror. The geometrical parameters involved in the optimization process are indicated.

The *tooth depth* and *connector width* not only affect the mirror strength, which is an increasing function of these parameters, but also the position of the second “transmission dip”. Figure 7.7 shows the reflectivity and loss spectra for mirrors with different tooth depth (varying the connector width produces a similar behavior). We can see that for increasing tooth depth, the bandgap size and the peak reflectivity increase. However, because the coupling with leaky modes is affected as well, the corresponding “transmission dip” red-shifts until merging with the bandgap. This

results in increased loss in the bandgap area, which ultimately limits the maximum reflectivity and the mirror bandwidth. Also in this case an optimum exists, given by a $\sim 280\text{-}320\text{nm}$ tooth depth and $\sim 200\text{-}240\text{nm}$ connector width. With these conditions, the mirror bandwidth is maximized without spoiling the reflectivity.

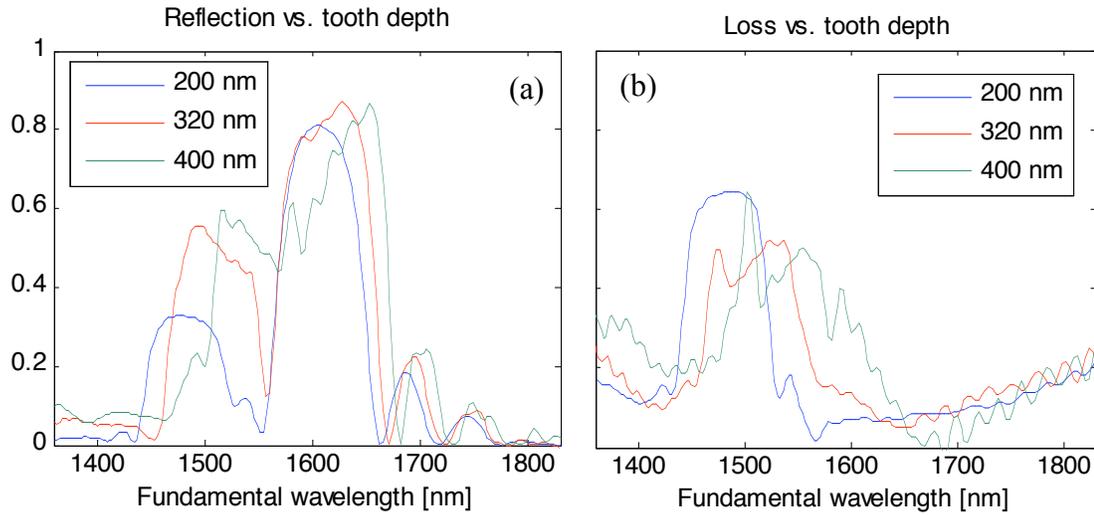


Figure 7.7 Simulated reflection (a) and loss (b) spectra at the FH wavelength for different tooth depths.

The modulation *duty-cycle*, defined as the ratio of the tooth width to the modulation period, also has a similar effect on the mirror reflectivity, as shown in Figure 7.8. For low duty cycle we obtain a clean, symmetric reflection spectrum, corresponding to a weakly modulated mirror. However, by increasing the duty-cycle we increase both the mirror strength and the coupling to the leaky modes, which eventually spoils the reflectivity. The optimal duty cycle has a value of $\sim 40\%$.

The *modulation period*, or pitch, is adjusted to position the center wavelength of the bandgap around $\sim 1550\text{nm}$ and the *number of mirror pairs* is varied to obtain the desired reflectivity. We point out that the mirror loss is almost entirely due to the mode mismatch at the mirror ends and is independent of the number of pairs (at least in the simulations). This is why in our optimization, the direct goal is the minimization of loss (i.e., not maximization of reflectivity).

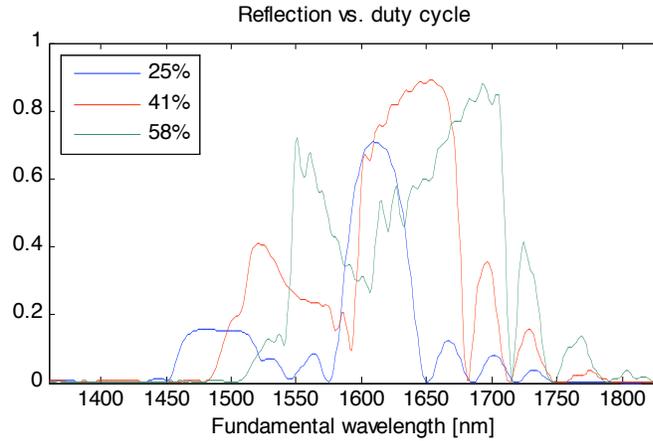


Figure 7.8 Simulated reflection spectrum at the FH wavelength for different modulation duty cycles (=tooth width /modulation period).

The optimization of the *taper geometry* is extremely important, as it determines the mode-mismatch loss. The length and angle of the taper are gradually increased until the point where the losses stabilize at a minimum value (3-4%). Such optimization is required since a taper longer than necessary would be a source of higher scattering loss in a real device because of sidewall roughness. The minimum taper length consists of ~ 15 pairs (we conservatively used 20 in our experiments) and the optimum angle is obtained when the distance between the beginning of the taper and the waveguide is ~ 360 -400 nm. To simplify the optimization and the fabrication, the modulation period in the taper is constant and equal to the period in the mirror part. We mention also that, although the connector width could be tapered as well, its effect on the loss is not significant. Therefore, we decided to keep it constant through the entire structure in order to improve the mechanical stability. Finally, we observe that the losses could be probably further reduced by adding other degrees of freedom such as tapering of the modulation period and the duty-cycle, or a nonlinear taper shape. However, we decided that, at this stage, the added complexity of such optimization was not justified. The result of this complex numerical optimization process was a mirror with reflectivity at the FH as high as 95-96%, limited by the residual radiation loss (~ 3 -4%) due to mode-mismatch. The mirror bandwidth (FWHM) is larger than 40 nm and the SH transmission is greater than 90% (Figure 7.4c). This dichroic mirror will be the key-element in the realization of tightly confining resonant cavities.

7.3 Fabrication

The fabrication process is exactly the same as the one used for the birefringent plain waveguides (Section 6.2). In fact, because of the robustness and flexibility of this process, there is no added complexity in the fabrication of cavities, compared to plain waveguides: only the e-beam mask has to be changed to generate the mirror pattern. For this purpose a Matlab script similar to the one used for the microcavity device (Section 4.2.2) was employed.

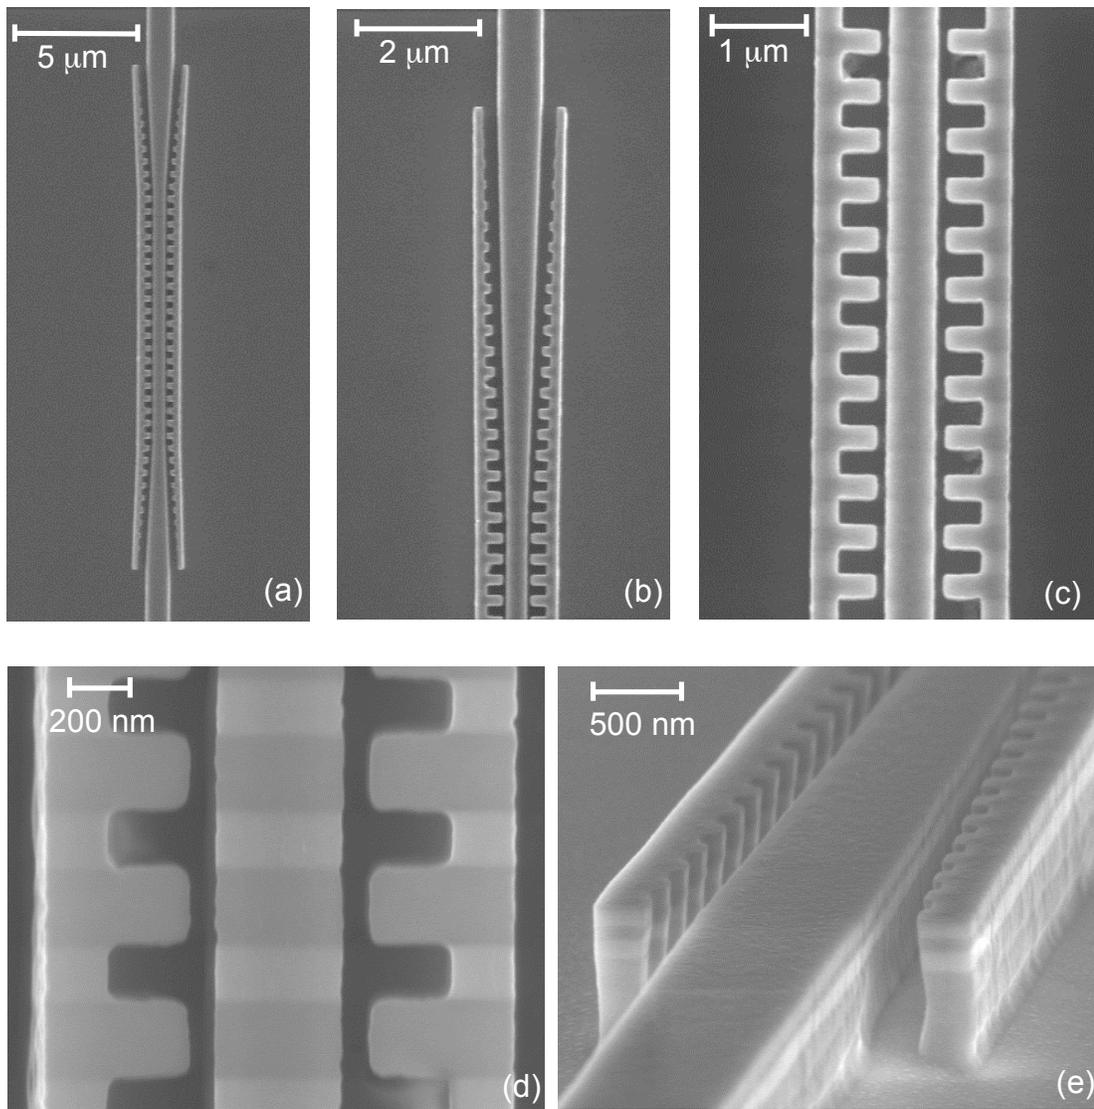


Figure 7.9 Top view (a) and details (b,c,d) of the final dichroic mirror structure, taken at different magnifications. (e) View at an angle of the mirror taper.

Figure 7.9 shows the fabricated mirror structure with different degrees of detail. The features are well defined and the roughness is only slightly higher than in plain waveguides (~ 6 nm RMS). The connector width is maintained constant through the entire structure to improve the mechanical stability and the yield. In fact, very small features (<100 nm) are less tolerant to dose variations and more prone to produce defects. The residual top surface roughness was measured to be similar to the unprocessed wafer roughness, and is due to the MBE growth.

As we just mentioned, features smaller than ~ 100 nm are neither desirable, nor absolutely necessary for good performance of the device (as discussed in the previous section). However, we fabricated mirrors with a tapered connector width in order to assess the capabilities of the process. Such a mirror, shown in Figure 7.10, demonstrates that features as narrow as 40nm (which is the smallest width investigated), with an aspect ratio of $\sim 1:25$, can be successfully fabricated with our process. Notice that the irregular profile of the taper (Figure 7.10b) is intentional and not due to roughness.

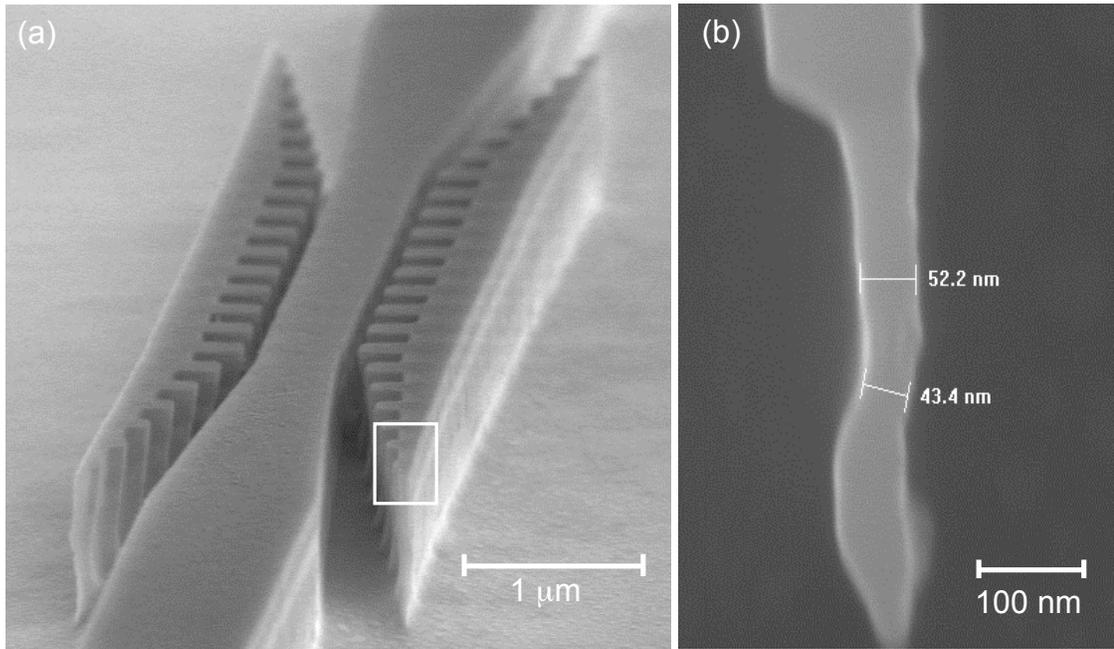


Figure 7.10 (a) Fabricated mirror with tapered connector and (b) detail (top view) of the highlighted section. The irregular shape is intentional and not the result of roughness.

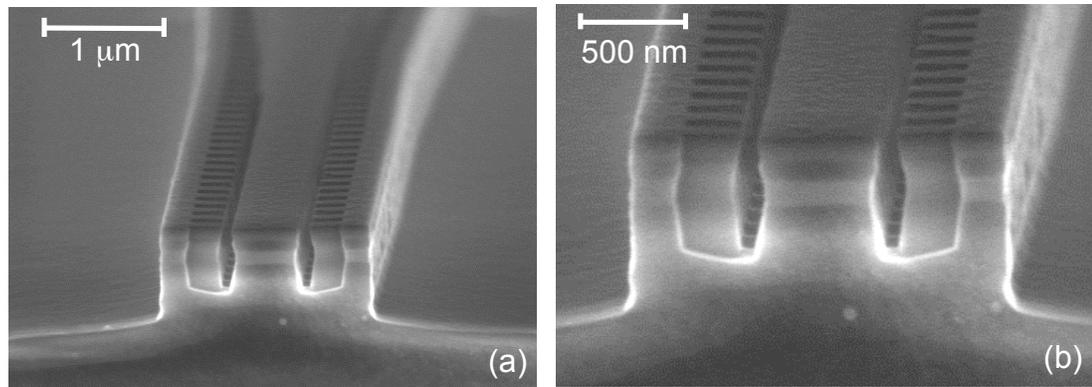


Figure 7.11 Cleaved section of the mirror (a) and detail (b), showing the trench depth.

We also investigated the minimum trench width: the narrowest measured trench (i.e., the distance between waveguide and tooth) was $\sim 60\text{nm}$ wide. Although for such width, the trench depth could not be measured, we believe that it is much shallower than the ridge depth. Figure 7.11 shows the cross section, obtained by cleavage, of a mirror with a minimum trench width of $\sim 100\text{nm}$. As discussed in Section 4.4 for the microcavity's holes, we observe also in this case a similarly

reduced depth of the trenches. However, compared to holes, the dependence of etch rate vs. aspect ratio is not as strong; the trenches, helped by the tooth modulation, reach a depth of $\sim 600\text{nm}$ (aspect ratio 6:1 vs. $\sim 2:1$ for holes). These trenches turn out to be sufficiently deep to avoid significant spoiling of the mirror reflectivity.

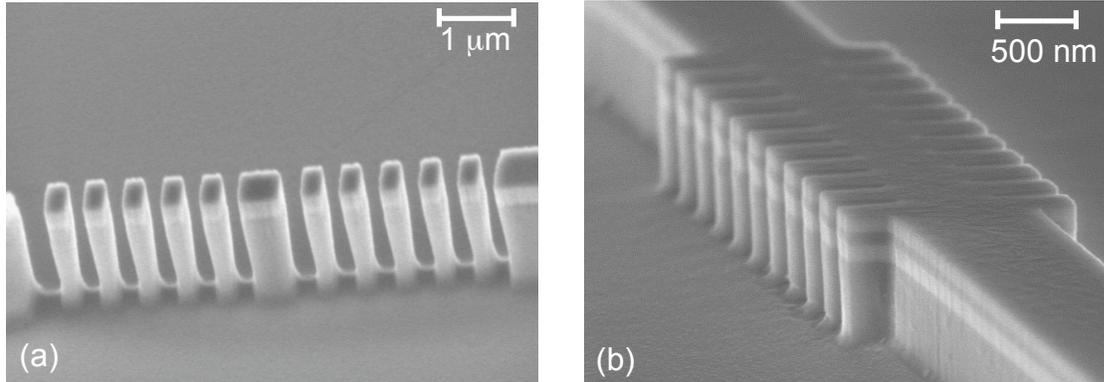


Figure 7.12 Mirror structures employing trenches (a) and deeply etched teeth (b), mainly used to test the process capabilities.

We finally mention that this process has been successfully used to produce other kind of mirrors, such as “post mirrors”, potentially useful for TM modes, or strongly modulated conventional DBR mirrors, as shown in Figure 7.12. These structures, considered in the early stages of this project, are not suitable to be used as dichroic mirrors. However, they could find other applications in linear optics.

7.4 Linear characterization

In this Section we discuss the linear characterization of cavity devices fabricated using the dichroic mirror. An illustration of the device is shown in Figure 7.13. Since now the cavity length is no longer limited any longer by the coherence length, cavities longer than $\sim 100\text{-}200\ \mu\text{m}$ can be fabricated. We also place the cavity as close as possible ($\sim 20\ \mu\text{m}$) to the output end of the waveguide to minimize the effect of the large SH attenuation loss.

There is also a second important consideration of this configuration: for long cavities (i.e., high cavity Q) the FWHM of the resonant peak is considerably narrowed

(~nanometer). It then becomes comparable to the period of the spurious Fabry-Perot cavities (A and C in Figure 7.13) formed by the dichroic mirrors and end facets. Thus, in order to facilitate the extraction of useful information from the spectrum, the lengths of the spurious cavities should be as different as possible from the length of the nominal cavity (cavity B). This is accomplished using the configuration in Figure 7.13, where the ratio of the cavity lengths is roughly A:B:C=4:1:0.25.

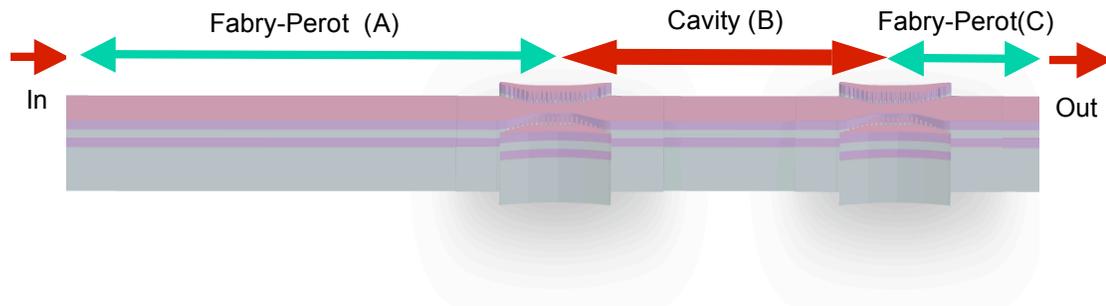


Figure 7.13 Illustration of the cavity device. The nominal cavity (B) and the spurious Fabry-Perot cavities (A and C) are visible.

The FH transmitted power was normalized to the input power. Variations in coupling efficiency due to the chromatic dispersion of the coupling lenses were also taken into account. The resulting transmission spectrum of a typical device is plotted in Figure 7.14 (dark line). The bandgap edges and the resonant peaks are visible. The enlargement of the spectrum (Figure 7.14b) shows more clearly the high-frequency ripples (due to cavity A) and the peak-height variations (mainly due to the low-frequency modulation of cavity C).

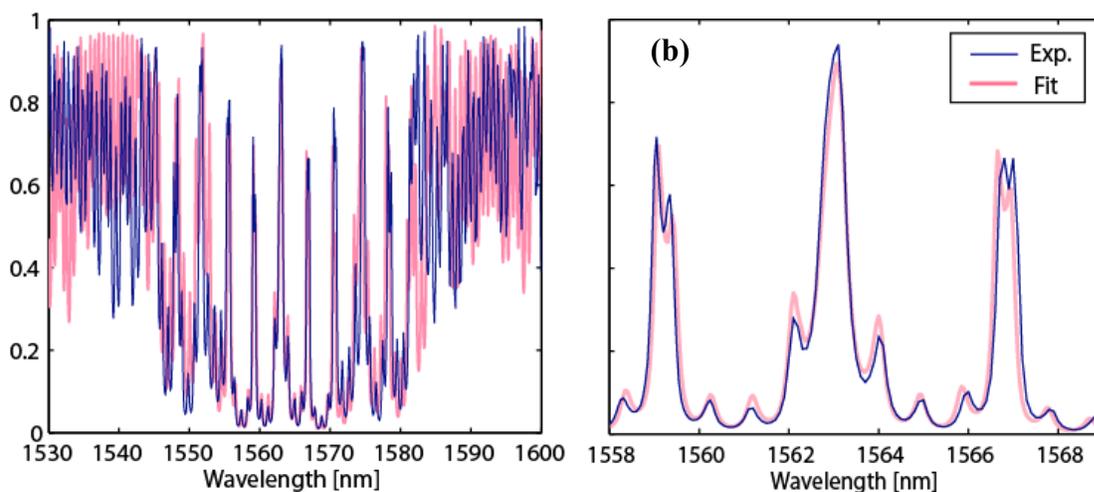


Figure 7.14 (a) Experimental (dark line) and fitted (light line) transmission spectrum of the cavity at the FH wavelength. (b) Detail of the central portion of the spectrum.

Although the cavity lengths are quite different, it is still very difficult to extract useful information directly from this spectrum. Moreover, since we have a system of three coupled cavities, the spurious Fabry-Perot effects cannot be easily modeled analytically. What we did, instead, was to model the entire device (including the end facets) using a 1D nonlinear transfer function algorithm⁹⁰. This algorithm and the Matlab-GUI software that we created to fit the model to the experimental data are explained in Appendix C. Before fitting the complex cavity spectrum, we have to take a step back and first analyze a plain waveguide. By fitting the spectrum of plain waveguide over a large wavelength range, we can determine its dispersion. Figure 7.15 shows the *detail* of such fitting and the power spectrum, calculated using the *entire data set*. The first two harmonics of the Fourier series are visible in Figure 7.15b. The flat-top peaks are due to the waveguide dispersion, which has been modeled with the same 2D frequency-domain code described in Section 6.1. This approach can be very accurate over a spectral region of ~ 20 -30 nm, which is wide enough for our purposes.

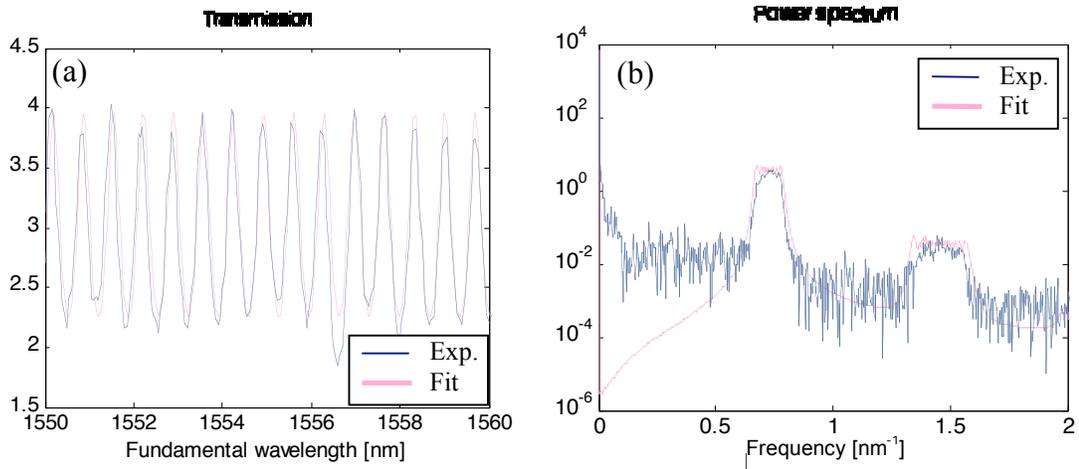


Figure 7.15 (a) Experimental (dark thin line) and fitted (light thick line) **detail** of the FH transmission of a plain waveguide. (b) Power spectrum of the **entire** transmission spectrum.

Once the waveguide dispersion has been determined, we can model the entire device cavity. For the mirrors, we employ simple multilayer stacks, where we use additional degrees of freedom to better simulate the 3D behavior of the structure (e.g., the propagation refractive index is allowed to be different from the index used for the Fresnel reflection at the interfaces). Scattering loss (nonexistent in a 1D system) is modeled with a non-zero absorption coefficient. The fitted spectrum is shown overlaid to the experimental spectrum in Figure 7.14.

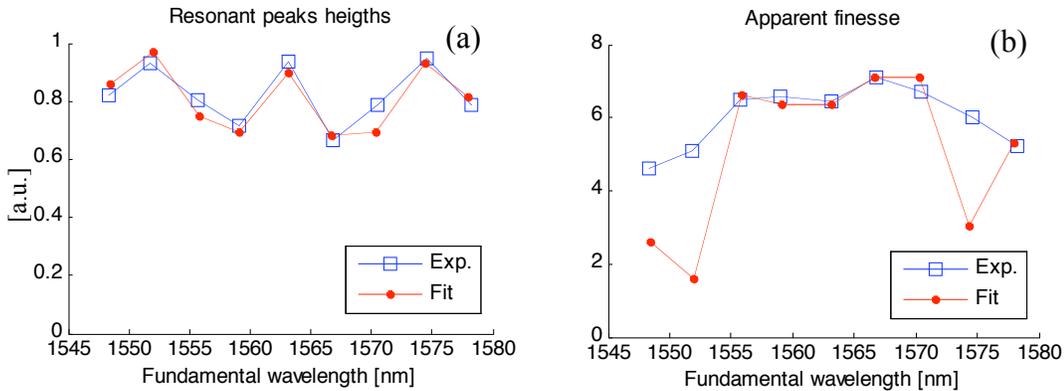


Figure 7.16 Experimental and fitted resonant peaks height (a) and finesse (b) extracted from the spectra of Figure 7.14

In order to verify that the fitting is correct, we compare (Figure 7.16) the experimental and fitted resonant peak heights and “apparent” finesse (i.e., finesse calculated from the Fabry-Perot-modulated resonant peaks). The value of the fitted

finesse is not accurate close to the bandgap edges; this is due to the deviation of the model from the real waveguide dispersion and to the complex behavior of a 3D mirror in proximity of the band-edge. However, in the region of interest (i.e., the central part of the bandgap) the fitting is excellent.

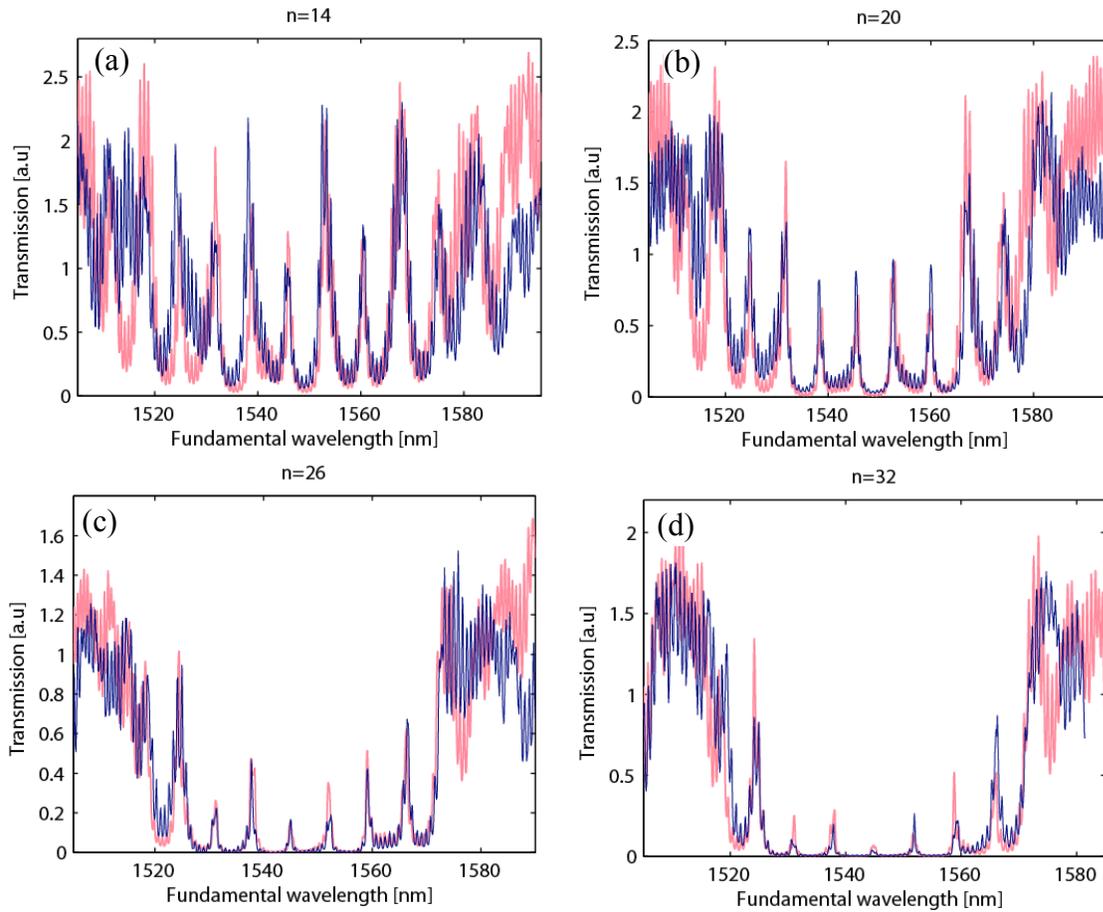


Figure 7.17 Experimental (dark line) and fitted (light line) transmission spectra at the FH wavelength for structures with a different number of mirror pairs.

In order to test the efficacy and flexibility of our software, we fabricated and tested cavities with different geometrical parameters. Figure 7.17 shows an example of such test, in which *only* the number of mirror pairs is varied. We start by fitting one of them, and use the same model to fit the other spectra, by varying in the software *only* the number of mirror pairs. As we can see, the software can track very well the

variations of the transmission spectra. In this way we demonstrate that our model is not only a numerical fitting model, but it has also physical significance. Thus, it can be used to make predictions and as a design tool.

Once we have obtained an accurate fitting of the experimental data, in order to extract useful information, we can artificially “turn off” the reflectivity of the end-facets and recover the transmission of cavity B alone. Figure 7.18 shows this new spectrum, where the resonant peaks have also been artificially suppressed for clarity. From this curve we can easily extract the mirror reflectivity, the cavity finesse and the expected enhancement as a function of the wavelength.

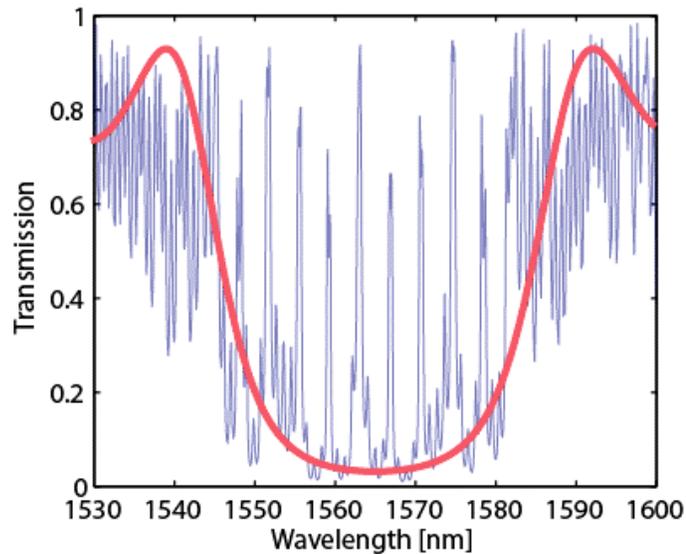


Figure 7.18 (a) Experimental transmission spectrum at the FH wavelength. (light thin line) and extrapolated transmission of the cavity where the resonant peaks and the facet reflectivity have been artificially suppressed (dark thick line).

Once we have calculated the information relative to the mirror, we can easily infer the behavior of the reflectivity as a function of the number of mirror pairs. This last information will be used to design the optimal cavity structure (Section 7.6).

The extrapolated reflectivity curve vs. length of the main mirror part (excluding the tapers) is shown in Figure 7.19. The curve does not start from zero because the tapers provide a reflectivity equivalent to ~ 4 -5 pairs. For this curve, the

maximum achievable reflectivity is $\sim 96\%$ and the mirror bandwidth is $\sim 40\text{nm}$. The highlighted point corresponds to the specific mirror used as an example in this and in the next section. Its structure has the following parameters: an inner width of 480 nm , a waveguide-tooth gap of 140 nm , 28 mirror pairs and 20 taper pairs. The modulation period is 488 nm , with a duty-cycle of 44% .

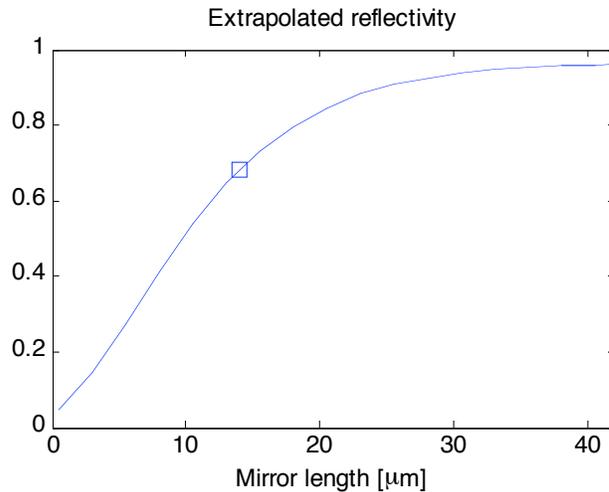


Figure 7.19 Extrapolated reflectivity vs. length of the main mirror part (excluding the tapers). The highlighted point corresponds to the mirror used as example in this and the next section.

We conclude this section with a discussion of the mirror loss, which is only 3-4%. A large contribution to this loss comes from radiation at the mirror entrances. However, by measuring different structures, we noticed that the mirror loss steeply increases as the inner waveguide width is decreased (while is only weakly dependent on other geometrical parameters). This fact indicates that the sidewall roughness becomes dominant for waveguides narrower than $\sim 460\text{-}480\text{ nm}$. This fact ultimately limits the maximum modulation and, consequently, the minimum length of the mirror.

Despite this limitation, more strongly modulating mirrors have been fabricated by slightly decreasing the waveguide-tooth space: a performance similar to the previously mentioned one was obtained with a mirror $\sim 20\text{-}30\%$ shorter.

7.5 Enhanced Second Harmonic Generation

In this section we discuss the nonlinear characterization of the resonant cavities. The main problem in fabricating these devices is the alignment of the phasematching bandwidth with the mirror bandwidth. This difficulty is further increased by the different wafer growth used for these experiments: although this wafer has a better quality compared to the one used for plain waveguide measurements, the AlGaAs core layers were grown 10% thinner than expected (i.e., 100 nm instead of 110nm). Unfortunately, for technical reasons it was not possible to obtain another wafer with the correct layer thicknesses; so, all the results presented in this chapter refers to the thinner layer structure.

Such structure results in several negative consequences: (i) the phasematching width is shifted to narrower waveguides ($\sim 800\text{nm}$), which implies higher scattering loss from roughness; (ii) the sensitivity to width variation is dramatically increased; now a width drift of a few nanometers produces a PM wavelength shift of tens of nanometers. This translates in a significant difficulty in positioning the PM peak within the testing window; (iii) the effect of disorder is also amplified, resulting in further broadening of the tuning curve; (iv) the sensitivity to other parameters, such as the oxide index, is also increased, resulting in decreased repeatability. However, a thinner core has one advantage: the phasematching bandwidth is greatly increased; thus, *if* the phasematching peak is within the optical testing windows, then its bandwidth will span the entire spectrum.

Because of these issues, a comprehensive study of different phasematching conditions and cavity geometries was not possible; instead, we had to limit the nonlinear characterization to the (very few) phasematched cavities in the sample. The particular cavity analyzed in the previous section is one of these few devices. We point out that the results of this section do not represent the best performance of these cavities which could be achieved with the correct layer structure. The potential performance of these devices will be discussed in Section 7.6. Finally, because of these limitations, the absolute value of the efficiency is not an especially meaningful

figure of merit. Instead, more interesting parameters are the cavity enhancement and the mirror transmission at the SH wavelength. The rest of this section is dedicated to the explanation of the techniques used to estimate these parameters and the results obtained.

For the same device of Figure 7.14, we recorded the corresponding SHG spectrum, shown in Figure 7.20. The same figure also shows the SH spectrum simulated using the 1D nonlinear transfer function algorithm. The FH field distribution resulting from the fitted FH spectrum was used as the source.

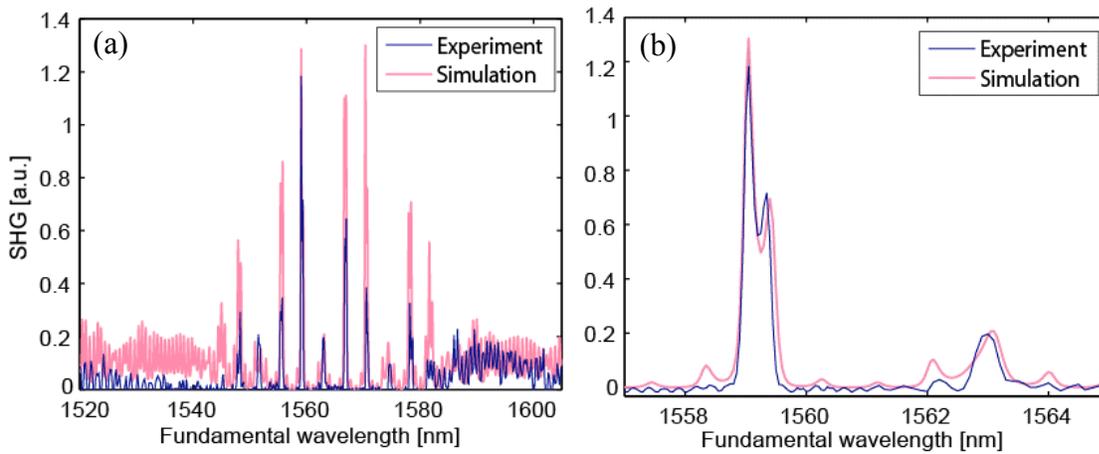


Figure 7.20 (a) Experimental (dark line) and simulated (light line) SHG spectrum. (b) Detail of the central portion of the spectrum.

The model reproduces the qualitative behavior of the SH quite well, including the suppression of certain resonant peaks, due to the interference with the spurious Fabry-Perot cavities. However, because of the uncertainty in the single-pass SHG due to the thinner core, it is very difficult to extract quantitative information from this analysis. Instead, we can compare the SH from the cavity to: (1) the best achievable plain waveguides given the high SH loss (Section 6.3.2), which are $\sim 400 \mu\text{m}$ long; (2) the SH from a plain waveguide of the same length of the cavity ($\sim 80 \mu\text{m}$).

A schematic diagram of these test structures and the corresponding spectra are shown in Figure 7.21. For the cavity device, we use two different widths for the cavity and the other waveguide sections, so that only the cavity is phasematched. Notice that the PM tuning curve of the plain waveguide is much broader than the peaks of Section 6.3.2; also, the PM peak of the reference device, which is much shorter, is essentially flat. The comparison with the best plain waveguides shows that we can indeed enhance the peak efficiency of a plain device. In this case, a 80- μm -long cavity is $\sim 50\%$ more efficient than the best plain waveguide, which is $\sim 4\text{x}$ longer. Moreover, here we do not consider the extra attenuation of the cavity device along the last section of waveguide (L_C). This is, to our knowledge, the first demonstration of enhanced SHG in collinear, tightly-confining resonant cavities.

For the comparison with the waveguide of the same length, the reference device consists of a plain waveguide with sections of different widths and lengths. These sections are connected using simple tapers instead of mirrors, so as to match the widths and lengths to those of the cavity devices.

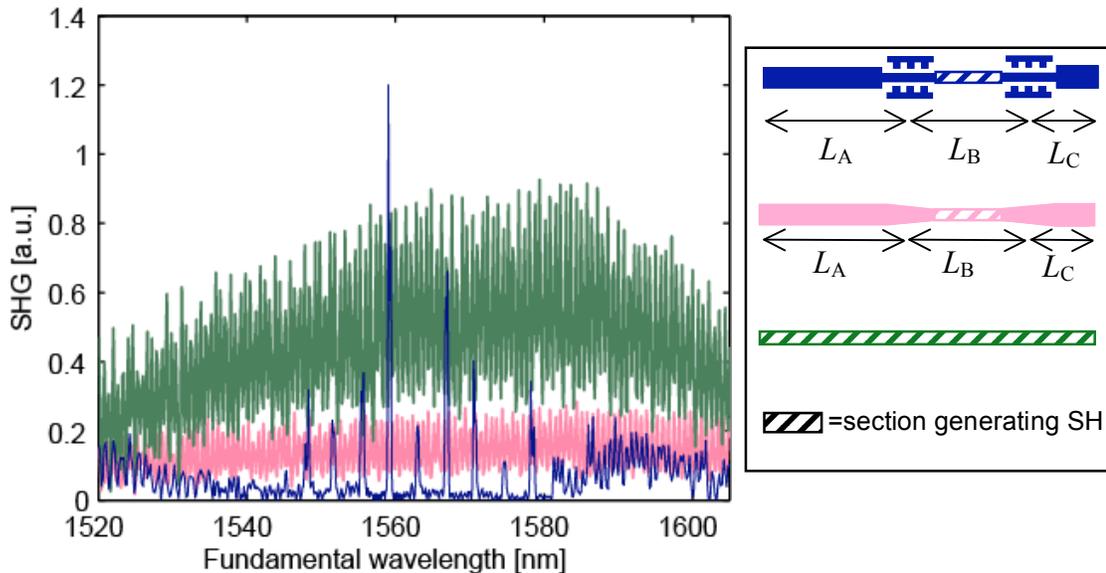


Figure 7.21 Experimental SHG spectrum of a cavity device 80 μm -long, a reference plain waveguide generating over the same length, and a plain waveguide $\sim 400\mu\text{m}$ -long.

The normalized conversion efficiency for the cavity device can be expressed as:

$$\left. \frac{P_{2\omega-out}}{P_{\omega-in}^2} \right|_{cav} = P_{SP} \Gamma_{\omega} T_{2\omega} \exp(-\alpha_{2\omega} L_C), \quad (7.1)$$

where P_{SP} is the single-pass SH power generated in the cavity length, L_B , as expressed by Equation (2.21); Γ_{ω} is the SH enhancement due to the cavity, assuming perfect transparency of the mirrors at the SH; $T_{2\omega}$ is the SH mirror transmission; $\alpha_{2\omega}$ is the SH propagation loss and L_C is the length of the last section of different width (i.e., non-phasedmatched). Using the same notation, the efficiency of the reference device is then written simply as:

$$\left. \frac{P_{2\omega-out}}{P_{\omega-in}^2} \right|_{ref} = P_{SP} \exp(-\alpha_{2\omega} L_C). \quad (7.2)$$

The ratio of the efficiency of the cavity to the efficiency of the reference (practically obtained by dividing the two respective spectra of Figure 7.21) depends only on Γ_{ω} and $T_{2\omega}$, and not on P_{SP} or $\alpha_{2\omega}$, which are difficult to measure with the current samples.

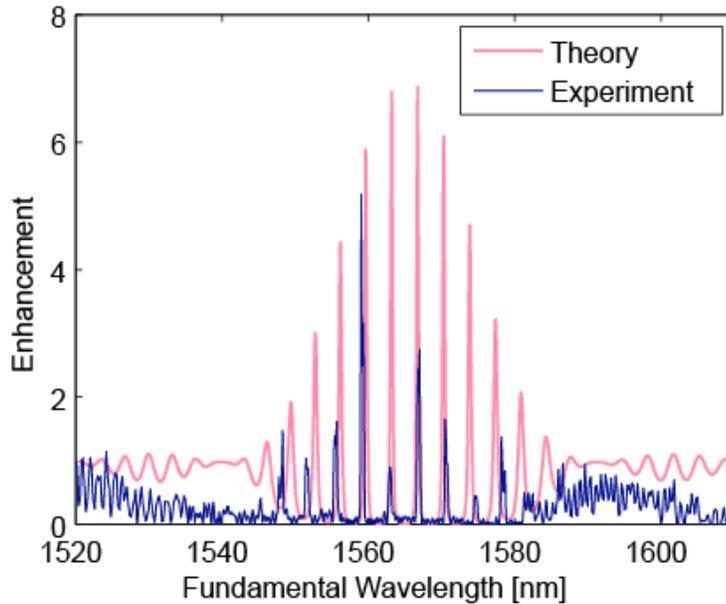


Figure 7.22 Simulated (light thick line) enhancement factor (Γ_{ω}) and experimental data (dark thin line).

By looking at the ratio of these two experimental spectra, plotted in Figure 7.22, we see immediately that the cavity device can enhance the SHG by a factor ~ 5 , compared to a plain section of waveguide of the same length, corresponding to a total internal conversion efficiency of 5-7%/W. The *experimental* ratio gives the effective enhancement of the cavity, which is the product of the *theoretical* cavity enhancement (Γ_{ω}) and the SH mirror transmission ($T_{2\omega}$). Γ_{ω} , also plotted in Figure 7.22, can be calculated from the fitted model of the FH spectrum (previous Section). By comparing the experimental and the theoretical enhancement factor, we can then extract $T_{2\omega}$. The value of $T_{2\omega}$ is in this case, 80-85%, in good agreement with the expected value of 90%. This analysis, together with the measurement of high FH reflectivity, demonstrates that our mirror indeed possesses a dichroic behavior.

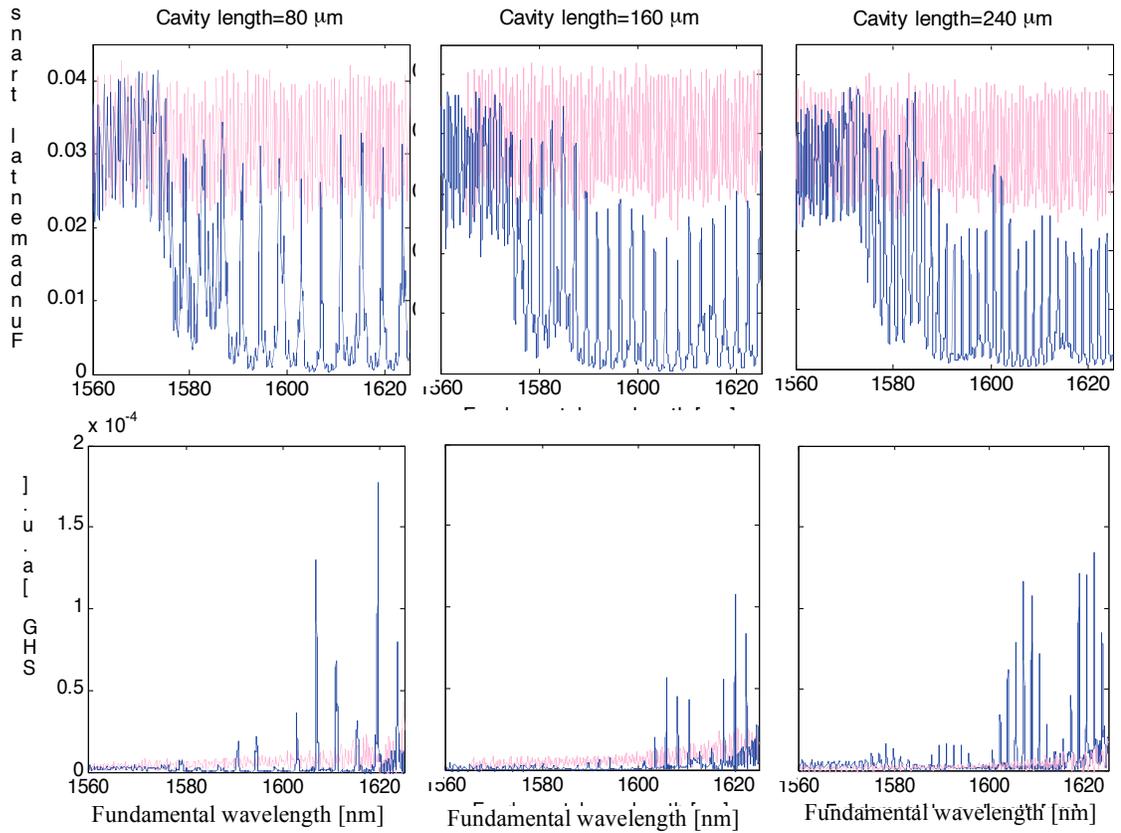


Figure 7.23 Experimental FH (top) and SH (bottom) spectra for cavity devices with different cavity lengths (dark line) and corresponding reference devices (light line).

Cavities with different geometries were also characterized, but the phasematching peaks were often not centered in the testing window; consequently, the SH signal for the reference sample was too weak for any meaningful estimation of the enhancement ratio. However, it is still interesting to look at their FH and SH spectra. Figure 7.23 shows these spectra for cavity devices with different cavity lengths, and the corresponding reference devices. Notice that, as the cavity length increases, the free spectral range decreases and the apparently random distribution of peak heights (at both FH and SH) now displays a precise pattern: the low frequency modulation of the peak heights turns out to be regulated by the length of the last section of the device (L_C).

By fitting the FH transmission spectra, Γ_ω of $\sim 10\times$ have been measured. From this figure and from the best single-pass efficiency of the plain waveguides of Chapter 6, we can predict that a conversion efficiency of $\sim 14\%/W$ (for a $100\mu\text{m}$ -long cavity) should be attainable with the correct layer structure. Furthermore, using our fitting software as design tool and realistic values of loss and mirror parameters, we can calculate the best performance achievable with an optimized geometry. We estimated that efficiencies of $\sim 30\%/W$ should be achievable with samples only $\sim 150\mu\text{m}$ long. This figure is comparable to the best reported efficiency ($=48\%/W$) for QPM AlGaAs waveguides, which are $70\times$ longer²³.

However, the total conversion efficiency is not the only important parameter. In the next Section we will discuss more systematically, different figures of merit and the advantages and trade-offs of different cavity configurations.

7.6 Optimization of the resonant cavity

In this section we discuss the optimal design of the cavity device, according to different figures of merit. In our discussion, the waveguide sections are considered to

have the same characteristics of the waveguides described in Chapter 6. We will then assume a propagation loss of $\sim 2\text{dB/mm}$ at the FH and $\sim 40\text{dB/mm}$ at the SH. For the mirrors, the loss is assumed fixed at $\sim 4\%$ (as inferred from the measurements of the previous Section). The mirror reflectivity (R) and the cavity length (L) will be used as free parameters as they can be easily modified during e-beam lithography.

The first figure of merit that we consider is the total internal conversion efficiency. This figure is given by the product of the total cavity enhancement ($\Gamma_{\omega_{tot}}$) and the single-pass efficiency (P_{SP}), expressed, respectively by Equations (2.37) and (2.21). Figure 7.24 shows $\Gamma_{\omega_{tot}}$ vs. R and L . Because of a non-zero FH loss, $\Gamma_{\omega_{tot}}$ increases for shorter devices, as expected, and peaks for a value of the mirror transmission equal to the cavity loss a_{tot} (i.e., $R=1-2a_{tot}$). In this and in the following figures, the circle represents the device characterized in Section 7.5. For this device, $\Gamma_{\omega_{tot}}$ is ~ 5 , but it could be doubled by increasing the number mirror pairs, at constant L .

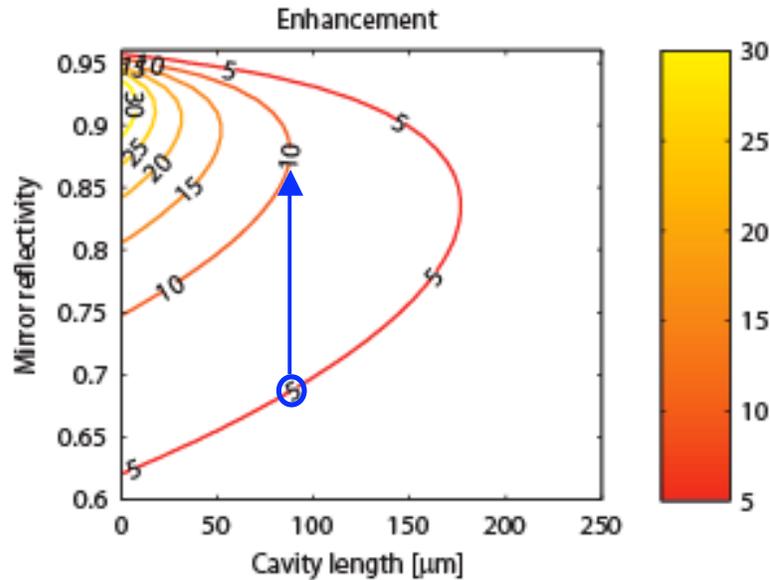


Figure 7.24 Cavity enhancement vs. mirror reflectivity and cavity length. The blue circle represents the previously characterized device.

Figure 7.25 shows the total conversion efficiency. This efficiency has a broad optimum region for L between ~ 150 and $300 \mu\text{m}$ and R between 70% and 80% . We see that, compared to the current device, efficiencies 6x higher should be possible

($\sim 30\%/W$). The large area of the optimal window is due mainly to the high SH loss and would sharply decrease if the device quality is improved.

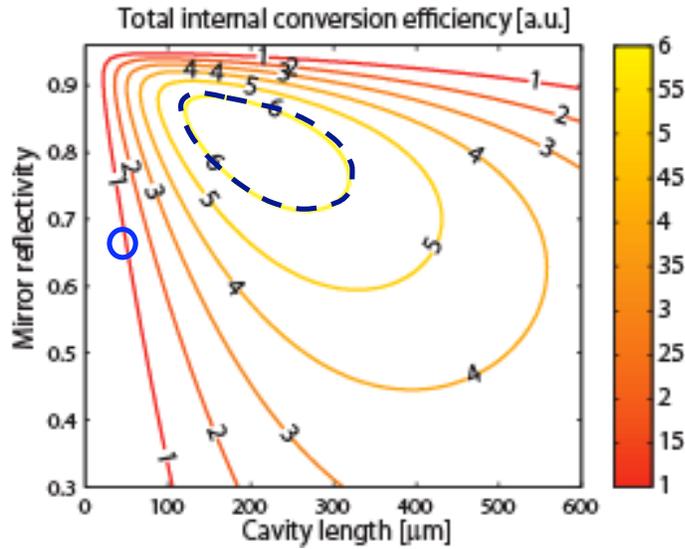


Figure 7.25 Total internal conversion efficiency vs. mirror reflectivity and cavity length. The blue circle represents the previously characterized device.

However, the efficiency is not the only figure of merit. Future integration of these devices will require them to be very compact. We can then plot the combined figure of merit “efficiency divided by length” (Figure 7.26). Taking the intersection of the optimal regions of Figure 7.25 and Figure 7.26 (shaded area in Figure 7.26), the best device would be $\sim 150\mu\text{m}$ -long, with 85% mirror reflectivity.

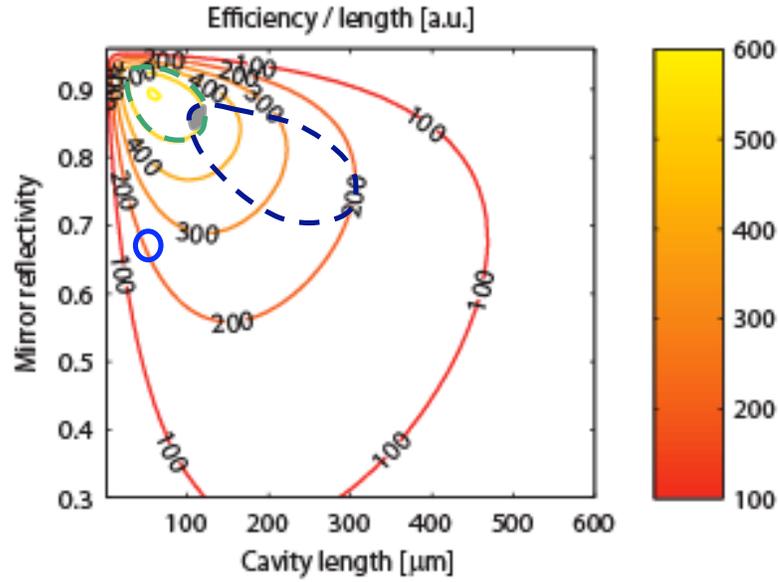


Figure 7.26 Conversion efficiency divided by L vs. R and L . The circle represents the previously characterized device. The dashed area is the region of maximum efficiency, and the shaded area is the overlap between best efficiency-length and maximum efficiency.

Another figure of merit is the temporal bandwidth of the device, characterized by the cavity lifetime, defined as:

$$\tau_{cav} = \frac{\omega}{Q} = \frac{\pi c}{n L F} \quad (7.3)$$

where Q is the quality factor, ω the frequency, n the group index, L the length and F the finesse. The larger the cavity lifetime (plotted in Figure 7.27), the longer the time necessary to switch the signal on and off; consequently, the device becomes slower. For the current device, τ_{cav} is only ~ 2 ps, which translates into several hundred Gigahertz bandwidth. However, the optimized device described earlier would be 5x slower, as shown in Figure 7.27. It is clear then that the actual optimal configuration will depend on the required speed, since a trade-off between efficiency and bandwidth is necessary.

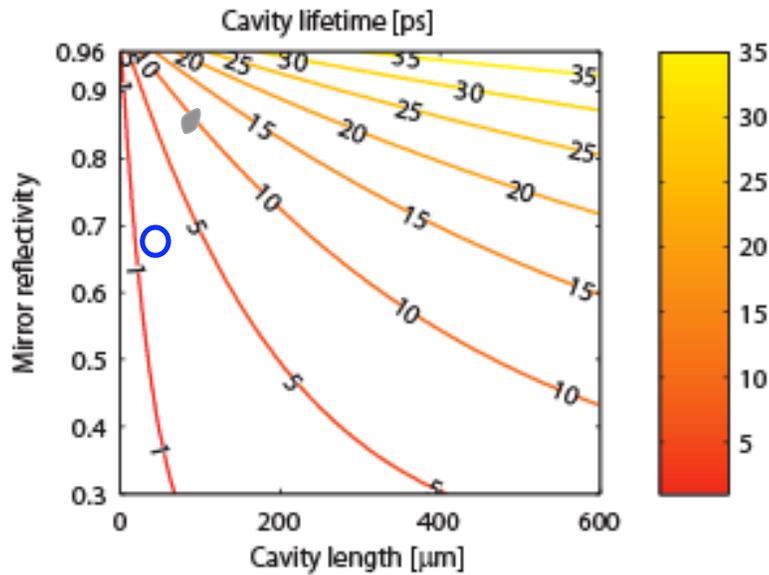


Figure 7.27 Cavity lifetime vs. mirror reflectivity and cavity length. The blue circle represents the previously characterized device. The shaded area corresponds to the optimality region, as defined in Figure 7.26

One last parameter of interest is the tuning range of the device. This is a complex quantity, which is determined by different factors. The first one is the phasematching bandwidth, which is inversely proportional to the length (Equation (3.3)) and independent of the reflectivity. As shown in Figure 7.28, loss and disorder have negligible effects on the bandwidth for device lengths shorter than $\sim 300\mu\text{m}$. For the current device, this bandwidth is smaller than 30nm, but can be increased by varying the layer thicknesses (Section 6.1). Ideally, we would like a phasematching bandwidth at least as large as the mirror bandwidth (35-40nm). In fact, in this way the bandwidth offered by the cavity configuration can be fully exploited.

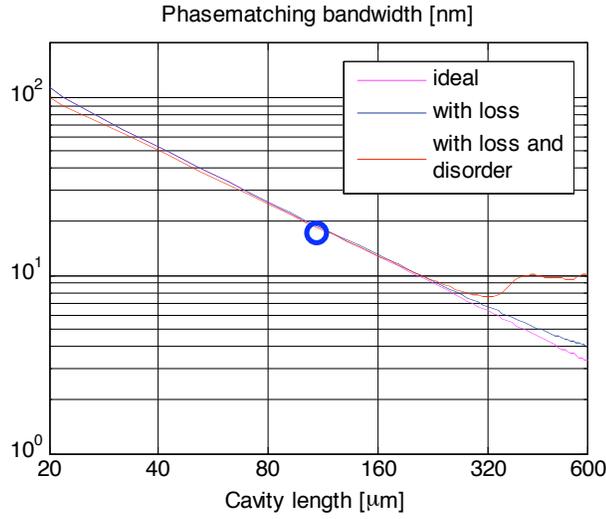


Figure 7.28 Phasematching bandwidth vs. length for ideal (dotted line), lossy (solid line) and lossy and disordered (dashed line) waveguides.

Another factor is the resonance bandwidth, defined as (Section 2.4.1):

$$BW_{RES} = \Delta\nu_{FWHM} = \frac{\Delta\nu_{FSR}}{F} = \frac{c}{2nLF}, \quad (7.4)$$

where $\Delta\nu_{FWHM}$ is the FWHM of the peak, $\Delta\nu_{FSR}$ is the free spectral range, c the speed of light, n the group index, L the length and F the finesse. Notice that this bandwidth is inversely proportional to the cavity lifetime. From Figure 7.29a, we can see that this is the limiting bandwidth for any device length larger than a few microns and any reflectivity of practical interest.

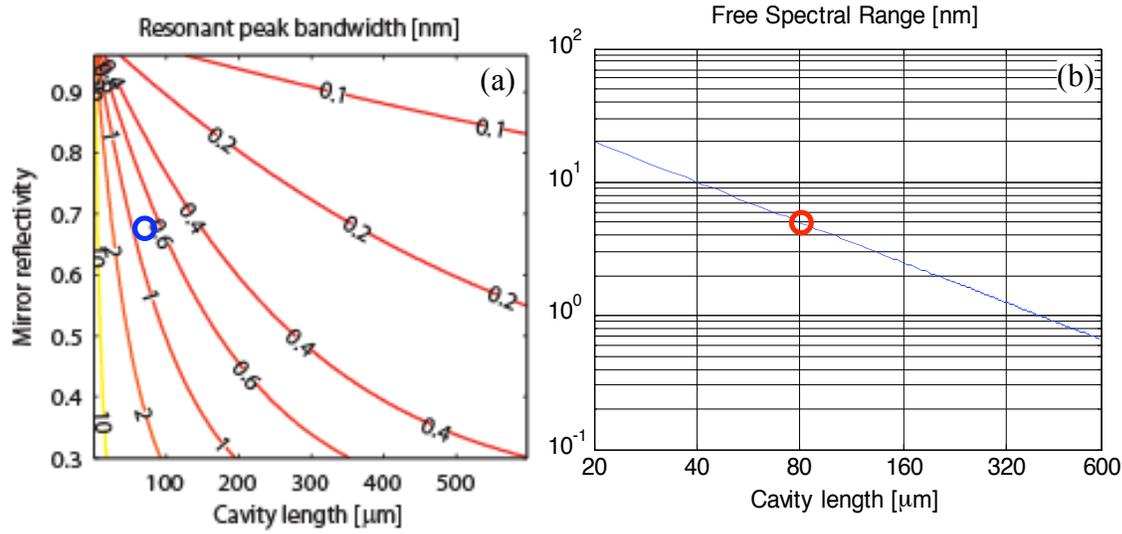


Figure 7.29 Resonant bandwidth vs. mirror reflectivity and cavity length (a) and free spectral range vs. length (b).

However, if the cavity is longer than $\sim 20\mu\text{m}$, the Free Spectral Range (Figure 7.29b) is smaller than the mirror bandwidth, such that multiple peaks are present within the bandgap. Each of these peaks could then be a channel of our hypothetical optical network. Thus, the new limiting bandwidth is the product of the single-resonance bandwidth times the number of useful channels:

$$BW''_{RES} = \Delta\nu_{FWHM} \frac{BW_{PM}}{\Delta\nu_{FSR}} = \frac{BW_{PM}}{F} \quad \text{or} \quad (7.5)$$

$$BW'_{RES} = \Delta\nu_{FWHM} \frac{BW_{MIR}}{\Delta\nu_{FSR}} = \frac{BW_{MIR}}{F}, \quad (7.6)$$

depending on which is smaller, between the phasematching (BW_{PM}) and the mirror (BW_{MIR}) bandwidths. In the first case (Figure 7.30a), the bandwidth depends on both length and reflectivity and, for the actual device, is about 8x larger than the single peak bandwidth (i.e., 8 channels). Of course, increasing the length or the reflectivity (or, in a word, the cavity-Q) decreases this bandwidth, so that, again, a trade-off must be found between bandwidth and efficiency.

However, in the second case (Figure 7.30b), the total bandwidth is independent of the length, so that more freedom in the design can be attained. We could then choose to have a short device ($<100\mu\text{m}$ -long), with fewer, wider channels, useful for a coarse WDM system, or a longer device ($\sim 300\mu\text{m}$ -long) with many, narrower channels, spaced by less than 1nm, more suitable for DWDM applications.

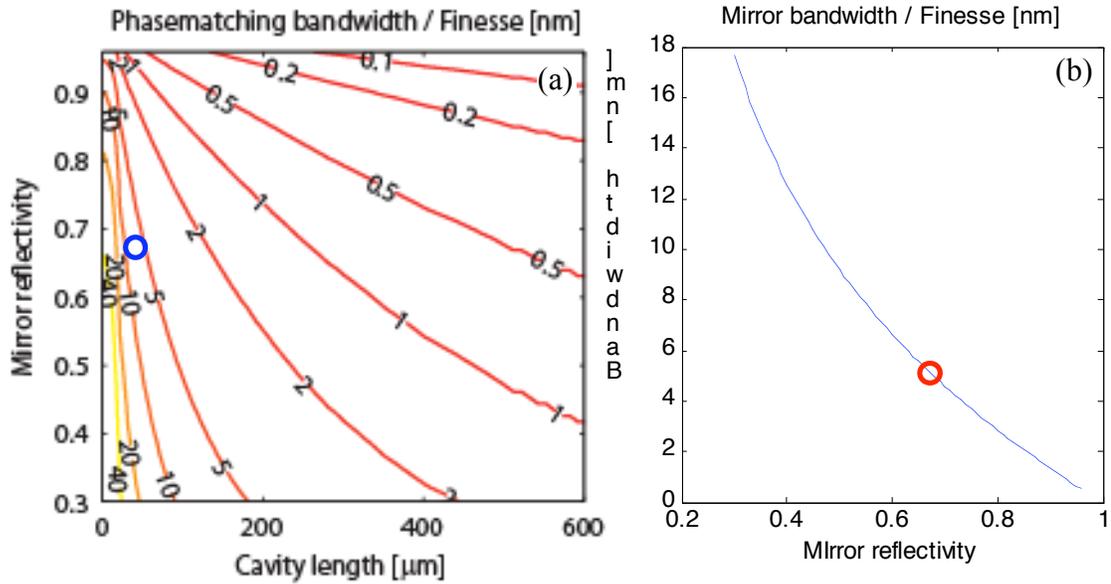


Figure 7.30 System bandwidth in the case of limiting phasematching bandwidth (a) and in the case of limiting mirror bandwidth (b). For the second case, the bandwidth is independent of the device length.

In conclusion, an absolutely optimal design cannot be defined. It depends, in fact, on the particular requisites of the system. A converter for high-speed transmission, for example, will require high speed ($\sim 100\text{s}$ of GHz) and thus low cavity-Q. This configuration limits the efficiency, so that higher input power is required. On the other hand, low-speed ($<100\text{GHz}$) modules could be highly resonant and highly efficient, reducing the size and the power consumption.

Chapter 8

Conclusions

We conclude this dissertation with a review of the results achieved in the previous chapters, followed by suggestions for future research.

8.1 Summary of results

In the first part of this work (Chapter 3 to Chapter 5) we designed and fabricated a high-Q photonic crystal microcavity device, embedded in a tightly confining waveguide. By shrinking or shifting the holes closer to the cavity, the radiation loss could be decreased and Q values as high as 700 were achieved. Although the conversion efficiency of these devices was extremely low, they could still be useful for linear-optical applications, such as sensing⁹¹ or filtering⁹².

A robust and flexible fabrication process was developed (Chapter 4); this process can produce deeply etched, AlGaAs/AlOx ridge structures as narrow as 40nm and with aspect ratios (i.e., width to etch depth) as high as 1:25. Sidewall roughness as small as 4nm RMS and drift of the ridge width smaller than 20nm over 500 μ m were measured. The resulting waveguide propagation loss was as low as 1.84 dB/mm for waveguides ~900nm-wide. To our knowledge, this is the state-of-the-art value of the loss coefficient for tightly confining waveguides in AlGaAs/AlOx.

To solve the problems of the microcavity device, a second generation device, phasematched by artificial birefringence, was developed (Chapter 6). Because of the tight confinement, this waveguide has a normalized conversion efficiency higher than 20,000 %/W/cm². This value is more than 20x larger than the efficiency of previously reported birefringent waveguides, and more than 200x larger compared to commercially available state-of-the-art LiNbO₃ nonlinear waveguides. The total internal conversion efficiency is ~5%/W with a device length of ~300μm. The total efficiency is currently limited only by the high loss at the SH.

The conversion efficiency of these devices was enhanced using a cavity embedded in the waveguide, resonant at the FH wavelength. For this purpose we designed a novel dichroic mirror, highly reflective at the FH wavelength and transmissive at the SH. This is the first demonstration of a dichroic mirror for tightly confining waveguide devices. Cavities with enhancements of ~5-6x were characterized, while enhancements of up to 10x were inferred by indirect measurements. This is, to our knowledge, the first demonstration of resonant SHG in a tightly confining waveguide, where the generation is collinear to the propagation direction (i.e., as opposed to the less efficient surface emitting devices) and with the mirror fully integrated in the waveguide (i.e., without requiring dielectric coatings on the waveguide facets).

8.2 Future perspectives

In this Section we discuss the current limitations of our devices (e.g., the SH loss) and possible ways to overcome them. In that case, the predicted performance will be estimated. Future directions of this research, including integration with active devices and how the coupling efficiency could be improved, will also be outlined. Finally, we will briefly discuss other possible applications of our devices.

8.2.1 The loss problem

As mentioned above, the high loss at the SH ($\sim 30\text{-}40$ dB/mm, compared to the FH loss ~ 2 dB/mm) is currently the major limitation of our devices. Although experiments seem to indicate that this SH loss is due to the oxidation process²⁷, its physical origin is presently unknown.

Since with our current optical setup, it is not possible to measure directly the SH loss (i.e., by coupling SH light and measure the transmission), its origin is very difficult to diagnose. In fact, indirect SH measurements (i.e., using only generated SH), have several limitations, such as the fixed polarization, and the impossibility to test non-phasematched structures (e.g., with no AlOx layer inside the the core).

We believe that, in order to understand and solve the loss problem, it will be essential to perform direct measurements of SH loss. Different structures, not necessarily phasematched, could then be tested. By varying the number and the thicknesses of the AlOx layers, and the polarization of the input light, it should be possible to understand if the SH loss originates from the AlGaAs/AlOx interfaces or from the AlOx bulk. Since similarly high SH loss was also observed in other works²⁷ and lacked an exhaustive explanation, we believe that such study would benefit not only this project, but also the scientific community.

We think that theoretical tools to predict the radiation loss due to roughness should also be developed in parallel. These tools would be essential to confirm (or invalidate) our preliminary estimations (Section 6.3.3). Although loss from absorption can be easily included in our 2D model, loss from roughness requires a 3D model. Although 3D FDTD simulations could theoretically provide an answer, such a job would be computationally very intensive; thus, it would probably not be suitable for a systematic study of the entire parameter space. An analytical or semi-analytical method, similar to the one developed by Barwicz et al.⁹³, would be much more challenging to develop, but could also provide more physical insights.

Once the loss mechanism has been identified and (possibly) understood, three main approaches could be taken to solve it. The first one would involve modifying the

epitaxial growth or the fabrication process. For example, the aluminum concentration of the cladding layer could play a role: compare in fact our work ($\text{Al}_{0.93}\text{Ga}_{0.07}\text{As}$, $\alpha_{2\omega} \sim 35$ dB/mm) with that of Fiore et al.²⁷ (AlAs , $\alpha_{2\omega} \sim 200$ dB/mm). The second approach would be to design a structure that is less sensitive to loss. The third one would involve a design that eliminates this problem by completely avoiding the thermal oxidation. This approach will be discussed later in this section.

8.2.2 Predicted performance of low-loss devices

We will now briefly discuss the performance of our devices in the case that their loss could be significantly decreased. Assuming that the major source of loss is roughness at the interfaces, if the roughness could be decreased by a factor 3, the loss at the SH (and FH) would be theoretically reduced by a factor ~ 9 (i.e. $\alpha_{\omega} \sim 0.2$ dB/mm, $\alpha_{2\omega} \sim 4$ dB/mm). If we consider plain waveguides, the optimal device length would be ~ 3 mm and the total efficiency ~ 150 %/W, 3x higher than the best QPM, 1-cm-long, AlGaAs waveguides.

Let us now consider the cavity device: if the mirror loss could be lowered from 3-4% to 1% (by both reducing the scattering loss and by careful optimization of the taper geometry), a maximum efficiency of 1600%/W could be achieved with a device length of 650 μm and a mirror reflectivity of 95% (cavity enhancement $\Gamma_{\omega} \sim 37$). This value of the efficiency is comparable to state-of-the-art, 5-cm long PPLN devices⁸¹. The best performance in terms of efficiency-size would be given by a device only ~ 140 μm long with 97%-reflective mirrors ($\Gamma_{\omega} \sim 200$). The total efficiency would be ~ 650 %/W. We point out that for such a short device, the limiting factors become the FH propagation and mirror losses, rather than the SH loss. However, as we discussed in Section 7.6, the efficiency and the device length must be ultimately traded off with the bandwidth and the cavity lifetime.

8.2.3 Suspended structures

One solution to the problem of high SH loss consists in designing a structure that can be phasematched without requiring thermal oxidation. Since high contrast is needed,

the only two candidates as low index materials are AlOx and air. According to preliminary simulations, a suspended structure consisting of AlGaAs and air similar to the one shown in Figure 8.1 could theoretically be phasematched. Since no high-Al-content AlGaAs can be exposed to air after fabrication (Section 4.5), pure GaAs could be used as cladding, instead. Holes etched sufficiently far from the waveguide could provide access to the underneath GaAs layers. The GaAs layer could then be removed by a selective wet-etchant, such as the one used to lift-off the Nickel mask (Section 4.3.2).

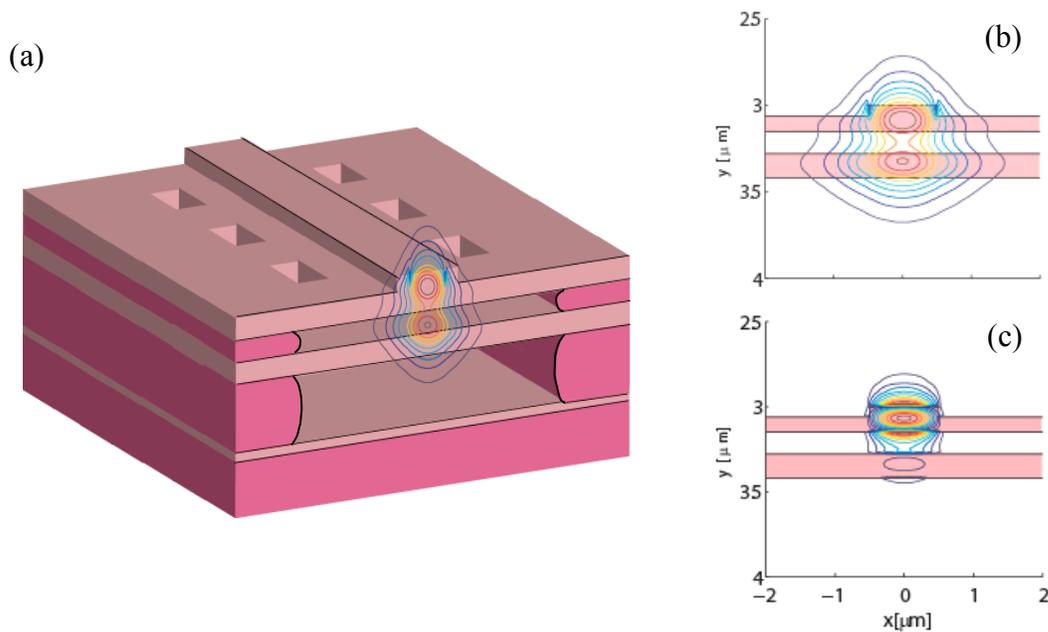


Figure 8.1 (a) Schematic of the suspended structure. The darker parts are GaAs, the lighter parts are AlGaAs. (b) FH TE mode electric field profile (c) SH TM mode electric field profile.

Although intriguing, this design presents several challenges. First of all, undercutting such a structure could not be trivial, because of problems related to the transport of etchant species. The wet-etching step may then need to be carefully developed. Also, simulated PM curves (Figure 8.2) show that the sensitivity to width and etch depth is very small while the sensitivity to the layer thicknesses is much greater. This fact would shift tolerance issues from the waveguide patterning step to the MBE growth (thus, increasing the complexity). Moreover, any deviation in layer

thickness could not be compensated by a different waveguide width. Therefore, changing the PM wavelength would involve growing a new wafer.

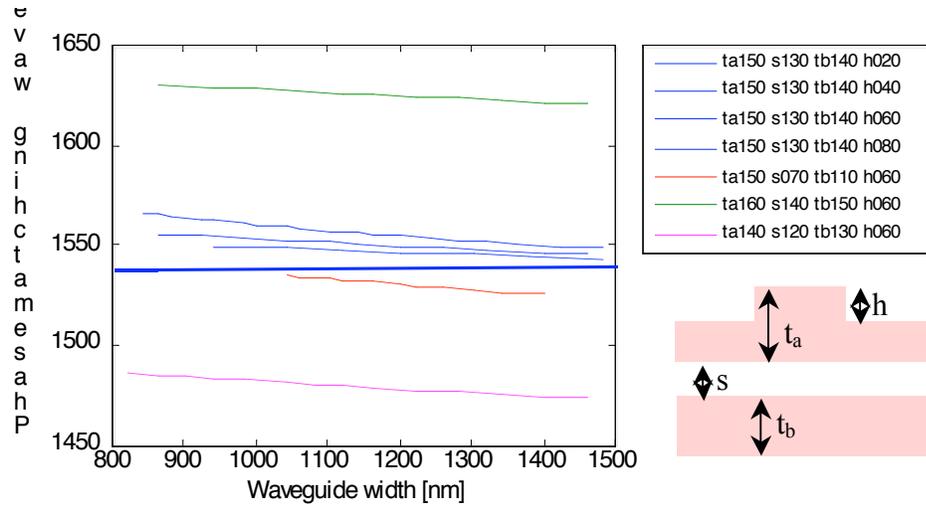


Figure 8.2 Phase-matching wavelength vs. waveguide width for different suspended geometries.

However, if these fabrication issues can be overcome, this structure should outperform our current design. In fact, simulations show that the normalized conversion efficiency is comparable to that of AlGaAs/AlOx waveguides. Moreover, no oxidation is required and the ridge is only shallow-etched: consequently all major sources of roughness or absorption would be (almost) completely eliminated.

8.2.4 Improving the coupling efficiency

Another limitation of tightly confining waveguides is the coupling efficiency with a free space beam or a fiber mode. Because the latter are much larger, the coupling efficiency is typically lower than $\sim 20\%$. Although our devices are ultimately meant to be employed in an integrated optical circuit, it would still be necessary to couple the signal in and out of the chip. For this purpose at least two architectures could be used.

The first one consists in using inverse tapers to increase the waveguide mode size⁹⁴. The challenge of this approach is the realization of a symmetric cladding. In our case, this involves the deposition of a top-cladding material with index similar to AlOx. Preliminary results using a SU-8 resist cladding on single-core waveguides

(Figure 8.3), showed a tenfold improvement in the waveguide throughput (from 4% to 40%). The extrapolated value of the coupling efficiency reached $\sim 90\%$. Both plain waveguides and tapered waveguides (i.e., from the input-width indicated in figure to $\sim 500\text{nm}$ and back to same input width) were analyzed. As we can see from Figure Figure 8.3a, the presence of the taper did not affect significantly the throughput. The next step would be combining this technology with the current waveguide and cavity devices.

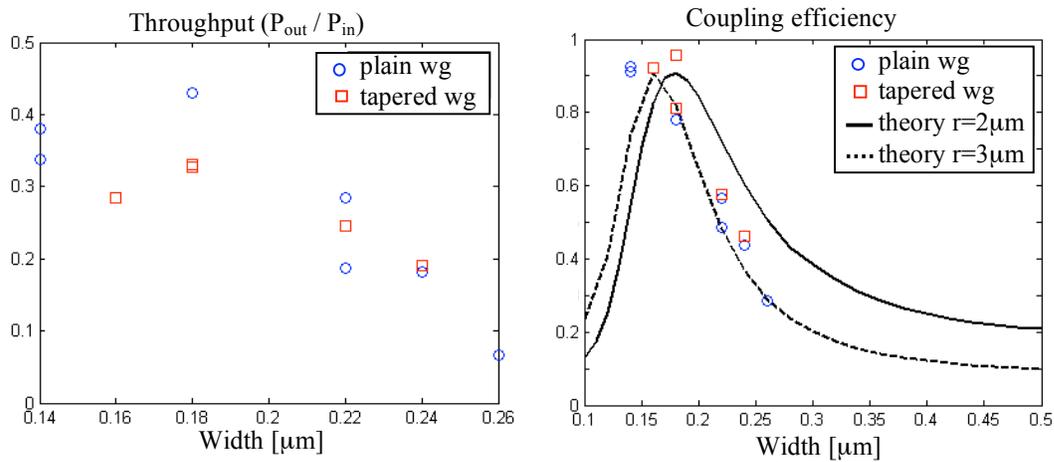


Figure 8.3 (a) Throughput vs. waveguide width for plain waveguides and waveguides tapered from the indicated width to a width of $0.5\mu\text{m}$. The vertical core structure is the same as the waveguides of Chapter 3. (b) Extrapolated coupling efficiency for the same structures. Theoretical predictions for an input gaussian beam of radius equal 2 and $3\mu\text{m}$ are also shown.

A second more ambitious design involves grating-assisted vertical coupling, as described by Taillaert et al.⁹⁵. In such a scheme, a regular single-mode fiber could be vertically butt-coupled directly to a waveguide with the aid of a grating and a spot size converter. Coupling efficiencies as high as 74% have been demonstrated in the work mentioned above. Our devices would benefit from such a scheme, since they could be tested in 2D arrays, directly after fabrication, without requiring thinning and cleaving.

8.2.5 Integration with active devices and other applications

We conclude with a discussion of one possible scheme for integration with active devices and a brief overview of other possible applications of the technology

developed in this work. One great advantage of using GaAs as the base material, is the possibility of growing active structures. However, active layers, such as quantum wells, are typically strong absorbers when used in a passive device. Hence, multiple growths are often necessary when active and passive devices must be integrated on the same chip.

In AlGaAs/AlOx devices, one could think to use the selective oxidation in the passive devices to suppress the absorption in the quantum wells (i.e by partially oxidizing them and shift their absorption wavelength). The oxidation of the active sections of the device could be prevented by depositing a protective layer, such as Silicon Nitride, which could be removed afterward. In this way, if the phasematching structure and the in-plane laser structure were properly designed, only one growth would be required.

The technology developed in this work could also be used for other applications. Our plane waveguides are extremely lossy at the SH, but the FH loss is quite low. We could then think to employ them for nonlinear optical applications at longer wavelengths. In this case, scattering and eventual absorption are expected to be significantly lower: in fact, the roughness amplitude would be the same, but the radiation loss coefficient is inversely proportional to the fourth power of the wavelength (Section 6.3.3). Mid-IR Difference Frequency Generation for spectroscopy could be one of these possible applications^{16,96,97}. Waveguide terahertz generation (by DFG of two frequencies in the near-IR⁹⁸) is also a promising application. Phasematching could be realized using artificial birefringence, and the terahertz radiation could be guided by metallic structures as in the case of quantum cascade lasers⁹⁹. Finally, the dichroic mirror technology could potentially be used to design filters or polarization-sensitive mirrors for linear applications.

Appendix A

Fabrication recipe

In this Appendix we summarize the optimized version of the process employed to fabricate the birefringent waveguide and the dichroic mirror. Technical details useful to reproduce the process are included.

1. MBE growth.

Growth on wafer n+ or p+, <100> misoriented by 4° towards <111>A. The target design structure is:

- 1100 Å Al_{0.5}Ga_{0.5}As (no GaAs cap!)
- 900 Å Al_{0.93}Ga_{0.07}As
- 1100 Å Al_{0.5}Ga_{0.5}As
- ~2 μm Al_{0.93}Ga_{0.07}As
- 100 Å Al_{0.5}Ga_{0.5}As
- GaAs/AlAs buffer layer
- GaAs substrate

The growth parameters should be set as described in Section 4.1.

2. Resist coat

Clean the sample with acetone and isopropanol and prebake on hotplate at 200°C for ~5 minutes. Spin PMMA 950k, 2% (in anisole) for 40 seconds at 2000 RPM

(resulting in ~80nm-thick layer). Clean the backside of the sample with a Q-tip and acetone (very important!). Postbake on hotplate at 200°C for 2 minutes.

2. E-beam lithography

Typical parameters used with the Raith150 system: working distance, 5mm; acceleration voltage, 10 keV; aperture, 10 μ m (corresponding to a current of 20pA); dose, 130-140 μ C/cm²; area step size, 4nm (corresponding to a beam speed of ~4mm/s); magnification, 600x; writefield, 100 μ m x 100 μ m.

After exposure, develop sample in methyl-isobutyl-ketone (MIBK):isopropanol=1:3 for 30 seconds and rinse in isopropanol for 30 seconds.

3. Metal deposition & lift-off

E-beam evaporation of ~240 Å of chromium (in Innotec). Deposition should be done at low pressure (< 9·10⁻⁸ Torr) and high deposition rate (> 2 Å /s).

Lift-off: leave sample in acetone for ~ 1 hour, then splash with acetone bottle to remove most of the metal. Then put the sample in a beaker with fresh acetone and sonicate for ~15 minutes. Finally rinse in isopropanol.

4. Dry-etching & mask removal

Typical parameters for dry-etching in PlasmaQuest: gas flows, Ar:BCl₃:Cl₂ = 15:10:1.5 sccm; RF power, 44W; ECR power 200W; He pressure, 8-10 Torr; chiller temperature, 0°C; heater temperature, 13°C.

Condition the chamber with oxygen plasma (~30 min) and with the above recipe (~1.5 hours). Then etch the sample ~8 minutes for ~1 μ m ech-depth.

The residual chromium mask is then removed in hot oxygen plasma (STS deposition tool) with the following parameters: oxygen flow, 150 sccm; pressure, 500 mTorr; hotplate temperature, 350°C; RF power, 100W; time, 20 minutes.

4. Oxidation

The oxidation is performed in the GaAs furnace tube, with the following parameters: bubbler bath, 95°C; nitrogen flow, 100 sccm. The furnace must be conditioned with water-vapor-saturated nitrogen for at least 1 hour. The sample is VERY slowly (~2 minutes) inserted in the furnace, left for ~20-24 minutes (equivalent to ~ 1.5µm and ~3µm of vertical and horizontal oxidation depths, respectively), and VERY slowly extracted (~2 minutes).

4. Thinning & cleaving

To protect the sample during the substrate thinning, spin a thick layer of resist (e.g., Shipley SPR220-7, 1 minute at 4000RPM; postbake on hotplate 10 minutes at 90°C) on top of the sample. Then, have the sample thinned to 110-120µm (Ginzton crystal shop). Finally the sample can be cleaved. the cleavage was performed at an outside company (OEPIC) on a professional machine (Dynatex International Dry Process dicing tool).

APPENDIX A FABRICATION RECIPE

Appendix B

Alignment technique

Because of the extremely small size of the waveguides, and, consequently, the low throughput, the alignment is a challenging task; in fact, standard alignment techniques used for regular waveguides are not suitable. Here we will discuss the technique employed to align our waveguides. This technique relies not only on performing the correct steps in the given order, but also on alignment “skills” which are developed with experience.

The first important task is the beam alignment: using pinholes mounted on the stages, the beam is first coarsely aligned to the stages by means of two mirrors (visible in Figure 5.1). Next, the lenses are put in place, and the stages’ positions are adjusted to obtain a beam perpendicular to the lenses; the output beam position with the lenses in place should coincide with the beam position without the lenses. To finely align the setup, the following procedure can be adopted. If the input beam fills up completely the clear aperture of the input lens, a variation of the distance between the lenses causes diffraction rings in the output beam. By imaging the rings on a camera, the stages and mirrors can then be adjusted so that the beam does not shift by changing the focus, and the rings appear symmetric with respect to the beam center. Precise beam

alignment is extremely important; in fact, it ensures that the smallest possible input spot size can be achieved.

Next, the sample is mounted on the stage, and the input coupling lens is aligned. Using an infrared sensitive card (IR card) placed behind the sample, the input lens is brought closer and closer to the sample. At the same time, the sample is continuously moved up and down, in such a way that the fringes produced by the sample crossing the input beam can be observed on the card. The sample is in focus when the diffraction fringes are the widest and when their position varies significantly with a slight variation of the sample position. By looking at the orientation of the fringes, it is also possible to check and correct for tilting of the sample along the x direction (see Figure 5.2).

The following step consists in checking the perpendicularity of the sample with respect to the beam. Although sub-micron waveguides are generally less sensitive to sample tilting, this condition minimizes the realignment while navigating through the sample. For this purpose, a beam splitter cube is placed on the beam path before the input lens. If the sample is in focus, the reflected beam will retrace the same path back to the beam splitter and it will be partially reflected at 90 degrees. By looking at this reflection while shifting the sample, it is possible to check for perpendicularity: if the sample is not perpendicular to the beam, the shifting will cause it to become unfocused and the reflection will disappear.

Next, while looking at the fringes on the IR card, we raise the sample up to the first fringe, and we translate it horizontally until we encounter a distortion in the fringes. This distortion is most likely produced by a waveguide ridge; however, since the light from the waveguide diffracts very rapidly, it is usually very difficult to distinguish on the IR card the waveguide light from the stray light. At this point, we place the output lens at a distance from the sample roughly equal to the focal length. Then, by slightly adjusting the sample position, a bright spot should appear on the IR card, which is placed 10-20 cm behind the output lens. The intensity of the spot is very sensitive to the sample position, and maintains a round shape if the focus of the output

lens is slightly changed. This spot is the signature of light coming from a waveguide. Fine adjustment of the alignment can now be performed by maximizing the signal from a power meter. At this point, the bright spot can be imaged with the IR camera, and the EDFA can be disconnected. Without amplification, the waveguide light is usually too faint to be seen on the IR card; therefore, the camera is necessary to navigate through the sample and to adjust the focus.

With enough experience and if the waveguides are of good quality, this technique can be performed in a few minutes, without requiring any expensive pieces of equipment (other than the electrostrictive actuators for fine tuning of the alignment).

APPENDIX B. ALIGNMENT TECHNIQUE

Appendix C

Fitting software

This Appendix contains an overview of the fitting software used to extract the mirror parameters from the FH cavity spectrum. The 1D nonlinear transfer function algorithm first calculates the FH transmission spectrum. Then, using the simulated FH field as source, the algorithm calculates the resulting SH spectrum in the non-depletion regime. The mathematical details of this algorithm are clearly explained in the paper by D.S. Bethune⁹⁰; in this Appendix we will describe our implementation.

The software is written in Matlab7, using the Graphical User Interface (GUI) platform, which allows a real-time visualization of any modification applied to the structure. Figure C.1 shows a screenshot of the GUI window.

In the top left corner we specify the file containing the experimental spectrum, the control (or structure) structure file (explained later) and the files containing the dispersion curves used for the waveguide section (simulated using the 2D frequency domain algorithm). An option allows the fitting of waveguides containing a single mirror instead of a cavity. When the “Load files” button is pressed, the software simulates the FH transmission spectrum of the structure, as specified in the control file. Both the experimental and simulated data are plotted in the two graphs: the spectra are plotted on the top and their Fourier transforms on the bottom. The parameters values

(right portion of the window) are initially updated according to the structure file. These parameters include the length and refractive index of the waveguide portions; the pitch, the duty-cycle, the index contrast, the phase shift and the number of pairs of the mirror. The FH loss of each section, the noise baseline and the normalization constant can also be specified.

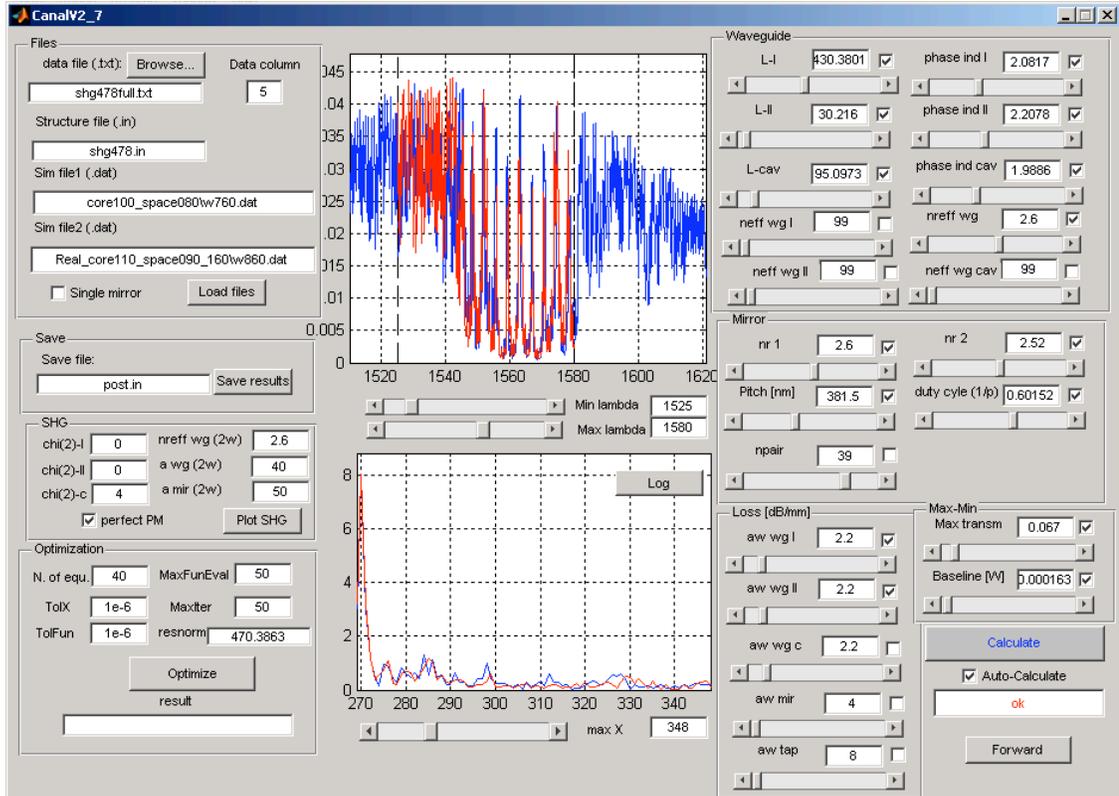


Figure C.1 Screenshot of the user interface.

Each parameter can be changed using the sliders or by editing its value directly. The simulated spectrum is updated in real time, such that a first rough fitting can be easily done manually. If the “Auto-calculate” option is off, the spectrum can be updated manually with the button “Calculate”. The sliders under the top graph are used to limit the data range to be simulated (and fitted). The button “Forward/Backward” is used to flip the input and output directions. This option was used in early experiments to compare SHG spectra with the cavity near the beginning and near the end of the waveguide. When a rough fitting of the spectrum has been

obtained, an optimization algorithm (“Optimize” button) is used for the fine tuning of the parameters. The fitting parameters can be individually specified using the checkbox on the right of the values fields.

Once the FH fitting is satisfactory, the SHG spectrum can be plotted in a separate window (“Plot SHG” button). The nonlinear coefficient and the SH loss in the waveguide and the mirror (which accounts for the SH mirror transmission) are manually specified and are used to fit the experimental SH spectrum. Since not all the required SH fitting parameters have been implemented (e.g., the mirror phases), the SH fitting is usually not as accurate as the FH one.

```

c n0 n1 a1[dB/mm] a2[dB/mm] nr length[um] x(2)[pm/V] %comment line
1 0 0 0 1 1 1 0 % air layer
1 0 0 0 1 1 1 0 % air layer
99 0 2.2 40 2.6 2.6 430.380100 0 %1st section of waveguide
2.081700 0 8 0 2.6 2.6 1 0 % taper section + mirror phase
99 0 4 40 2.6 2.6 0.229480 0 %mirror layers
99 0 4 40 2.520000 2.6 0.152020 0 % (r x y=repeat previous x layers y
r 2 39 % times)
1.988600 0 8 0 2.6 2.6 1 0 % taper section + mirror phase
99 0 2.2 40 2.6 2.6 95.097300 4 %2nd section of waveguide
99 0 8 0 2.6 2.6 1 0 % taper section
99 0 4 40 2.6 2.6 0.229480 0 %mirror layers
99 0 4 40 2.52 2.6 0.152020 0 %
r 2 39 %
2.207800 0 8 0 2.6 2.6 1 0 % taper section + mirror phase
99 0 2.2 40 2.6 2.6 30.216000 0 %3rd section of waveguide
1 0 0 0 1 1 1 0 % air layer
1 0 0 0 1 1 1 0 % air layer
maxt 0.067000 %normal. constant
base 0.000163 %noise baseline
f1 core100_space080\w760.dat %1st dispersion curve sim. file
f2 core100_space080\w1000.dat %2nd dispersion curve sim. file
-1 %end of file tag

```

Figure C.2 Sample of control file.

The optimal fitting parameters can then be saved in a control file (Figure C.2). Each line of this file corresponds to a 1D layer, and the numbers indicate, respectively: (1) the refractive index of the medium (n_0) and (2) its first derivative in respect to wavelength (n_1 , such that $n(\lambda) = n_0 + n_1 \lambda$). If $n_0=98$ or 99 , the frequency-dependent, simulated effective indices are used, as specified by f_1 and f_2 , respectively; (3) the FH and (4) SH propagation loss; (5) the FH refractive index used for the calculation of the Fresnel reflection at the boundary. This index is allowed to be different from the

propagation index to take into account the three-dimensionality of the real structure; (6) the length and (7) the nonlinear coefficient (scaled to take into account coupling efficiencies and other real-life effects). After the layer structure, a few more parameters need to be specified: *maxt* and *base* characterize the normalization constant and noise baseline; f_1 and f_2 are the files containing the 2D simulations results of real waveguide structures.

This implementation is designed to work exclusively with the waveguide devices described in this work and with the data files obtained with our optical setup. However, we point out that this software can be easily modified to simulate any kind of complex waveguide (or 1D multilayer) nonlinear devices, making it an extremely powerful analysis and design tool in a much broader spectrum of applications.

References

1. Green, P., *Progress in optical networking*. IEEE Communications Magazine, **39**(1): p. 54-61 (2001).
2. Yoo, S.J.B., *Wavelength conversion technologies for WDM network applications*. Journal of Lightwave Technology, **14**(6): p. 955-66 (1996).
3. Charlet, G., E. Corbel, J. Lazaro, A. Klekamp, R. Dischler, P. Tran, W. Idler, H. Mardoyan, A. Konczykowska, F. Jorge, S. Bigo. *WDM transmission at 6 Tbit/s capacity over transatlantic distance, using 42.7Gb/s differential phase-shift keying without pulse carver*. in Optical Fiber Communication Conference (OFC), 23-27 Feb. 2004, Los Angeles, CA, USA. Washington, DC, USA : Opt. Soc. America, 2004.(2004)
4. Menendez, R.C., P. Toliver, S. Galli, A. Agarwal, T. Banwell, J. Jackel, J. Young, S. Etemad, *Network applications of cascaded passive code translation for WDM-compatible spectrally phase-encoded optical CDMA*. Journal of Lightwave Technology, **23**(10): p. 3219-31 (2005).
5. Saleh, A.A.M. *Transparent optical networks for the next-generation information infrastructure*. in OFC '95 Optical Fiber Communication. Summaries of Papers Presented at the Conference on Optical Fiber Communication. Vol.8. 1995

- Technical Digest Series. Conference Edition, 26 Feb.-3 March 1995, San Diego, CA, USA. Washington, DC, USA : Opt. Soc. America, 1995.(1995)
6. Chou, M.H., K.R. Parameswaran, M.M. Fejer, I. Brener, *Multiple-channel wavelength conversion by use of engineered quasi-phase-matching structures in LiNbO₃ waveguides*. Optics Letters, **24**(16): p. 1157-9 (1999).
 7. Franken, P.A., G. Weinreich, C.W. Peters, A.E. Hill, *Generation of optical harmonics*. Physical Review Letters, **7**(4): p. 118-119 (1961).
 8. Armstrong, J.A., N. Bloembergen, J. Ducuing, P.S. Pershan, *Interactions between Light Waves in a Nonlinear Dielectric*. Physical Review, **127**(6): p. 1918 - 1939 (1962).
 9. Giordmaine, J.A., *Mixing of Light Beams in Crystals*. Physical Review Letters, **8**(1): p. 19 LP - 20 (1962).
 10. Feisst, A., P. Koidl, *Current induced periodic ferroelectric domain structures in LiNbO₃ applied for efficient nonlinear optical frequency mixing*. Applied Physics Letters, **47**(11): p. 1125-7 (1985).
 11. Thompson, D.E., J.D. McMullen, D.B. Anderson, *Second-harmonic generation of GaAs 'stack of plates' using high-power CO₂ laser radiation*. Applied Physics Letters, **29**(2): p. 113-15 (1976).
 12. Tien, P.K., R. Ulrich, R.J. Martin, *Optical second harmonic generation in form of coherent Cerenkov radiation from a thin-film waveguide*. Applied Physics Letters, **17**(10): p. 447-& (1970).
 13. Fejer, M.M., G.A. Magel, D.H. Jundt, R.L. Byer, *Quasi-phase-matched second harmonic generation: Tuning and tolerances*. IEEE Journal of Quantum Electronics, **28**(11): p. 3631-2654 (1992).
 14. Yamawaku, J., H. Takara, T. Ohara, K. Sato, A. Takada, I. Morioka, O. Tadanaga, H. Miyazawa, M. Asobe, *Simultaneous 25 GHz-spaced DWDM wavelength conversion of 1.03 Tbit/s (103*10 Gbit/s) signals in PPLN waveguide*. Electronics Letters, **39**(15): p. 1144-5 (2003).
 15. Chou, M.H., I. Brener, G. Lenz, R. Scotti, E.E. Chaban, J. Shmlovich, D. Philen, S. Kosinski, K.R. Parameswaran, M.M. Fejer, *Efficient wide-band and tunable*

- midspan spectral inverter using cascaded nonlinearities in LiNbO₃ waveguides*. IEEE Photonics Technology Letters, **12**(1): p. 82-4 (2000).
16. Vodopyanov, K.L., O. Levi, P.S. Kuo, T.J. Pinguet, J.S. Harris, M.M. Fejer, B. Gerard, L. Becouarn, E. Lallier, *Optical parametric oscillation in quasi-phase-matched GaAs*. Optics Letters, **29**(16): p. 1912-14 (2004).
 17. Dallesasse, J.M., N. Holonyak, A.R. Sugg, T.A. Richard, N. Elzein, *Hydrolyzation oxidation of Al_xGa_{1-x}As-AlAs-GaAs quantum well heterostructures and superlattices*. Applied Physics Letters, **57**(26): p. 2844-6 (1990).
 18. Grabherr, M., R. Jager, R. Michalzik, B. Weigl, G. Reiner, K.J. Ebeling, *Efficient single-mode oxide-confined GaAs VCSEL's emitting in the 850-nm wavelength regime*. IEEE Photonics Technology Letters, **9**(10): p. 1304-6 (1997).
 19. Macdougall, M.H., P.D. Dapkus, V. Pudikov, H.M. Zhao, G.M. Yang, *Ultralow threshold current vertical-cavity surface-emitting lasers with AlAs oxide-GaAs distributed Bragg reflectors*. IEEE Photonics Technology Letters, **7**(3): p. 229-31 (1995).
 20. Luo, Y., D.C. Hall, L. Kou, L. Steingart, J.H. Jackson, O. Blum, H. Hou, *Oxidized Al_xGa_{1-x}As heterostructure planar waveguides*. Applied Physics Letters, **75**(20): p. 3078-80 (1999).
 21. Ripin, D.J., K.Y. Lim, G.S. Petrich, P.R. Villeneuve, S.H. Fan, E.R. Thoen, J.D. Joannopoulos, E.P. Ippen, L.A. Kolodziejski, *One-dimensional photonic bandgap microcavities for strong optical confinement in GaAs and GaAs/Al_xO_y semiconductor waveguides*. Journal of Lightwave Technology, **17**(11): p. 2152-60 (1999).
 22. Eyres, L.A., P.J. Turreau, T.J. Pinguet, C.B. Ebert, J.S. Harris, M.M. Fejer, L. Becouarn, B. Gerard, E. Lallier, *All-epitaxial fabrication of thick, orientation-patterned GaAs films for nonlinear optical frequency conversion*. Applied Physics Letters, **79**(7): p. 904-6 (2001).
 23. Yu, X., L. Scaccabarozzi, J.S. Harris, P.S. Kuo, M.M. Fejer, *Efficient continuous wave second harmonic generation pumped at 1.55 μm in quasi-phase-matched AlGaAs waveguides*. Optics Express, **13**(26): p. 10742-10753 (2005).

24. Vanderziel, J.P., M. Ilegems, *Multilayer GaAs-Al_{0.3}Ga_{0.7}As dielectric quarter wave stacks grown by molecular beam epitaxy*. Applied Optics, **14**(11): p. 2627-30 (1975).
25. Fiore, A., V. Berger, E. Rosencher, P. Bravetti, J. Nagle, *Phase matching using an isotropic nonlinear optical material*. Nature, **391**(6666): p. 463-6 (1998).
26. Bravetti, P., A. Fiore, V. Berger, E. Rosencher, J. Nagle, O. Gauthier-Lafaye, *5.2-5.6- μm source tunable by frequency conversion in a GaAs-based waveguide*. Optics Letters, **23**(5): p. 331-3 (1998).
27. Fiore, A., S. Janz, L. Delobel, P. Van Der Meer, P. Bravetti, V. Berger, E. Rosencher, J. Nagle, *Second-harmonic generation at $\lambda = 1.6 \mu\text{m}$ in AlGaAs/Al₂O₃ waveguides using birefringence phase matching*. Applied Physics Letters, **72**(23): p. 2942-4 (1998).
28. Moutzouris, K., S.V. Rao, M. Ebrahimzadeh, A. De Rossi, V. Berger, M. Calligaro, V. Ortiz, *Efficient second-harmonic generation in birefringently phase-matched GaAs/Al₂O₃ waveguides*. Optics Letters, **26**(22): p. 1785-1787 (2001).
29. Ashkin, A., G.D. Boyd, J.M. Dziedzic, *Resonant optical second harmonic generation and mixing*. IEEE Journal of Quantum Electronics, **QE-2**(6): p. 109-124 (1966).
30. Myers, L.E., W.R. Bosenberg, *Periodically poled lithium niobate and quasi-phase-matched optical parametric oscillators*. IEEE Journal of Quantum Electronics, **33**(10): p. 1663-72 (1997).
31. Vodopyanov, K., *Pulsed Mid-IR optical parametric oscillators*. Solid-State Mid-Infrared Laser Sources, **89**: p. 141-178 (2003).
32. Fejer, M.M., *Nonlinear optical frequency conversion*. Physics Today, **47**(5): p. 25-32 (1994).
33. Xu, C.Q., K. Shinozaki, H. Okayama, T. Kamijoh, *Three wave mixing using a fiber ring resonator*. Journal of Applied Physics, **81**(3): p. 1055-62 (1997).
34. Mookherjea, S., A. Yariv, *Coupled resonator optical waveguides*. IEEE Journal of Selected Topics in Quantum Electronics, **8**(3): p. 448-56 (2002).

35. Van, V., T.A. Ibrahim, P.P. Absil, F.G. Johnson, R. Grover, P.T. Ho, *Optical signal processing using nonlinear semiconductor microring resonators*. IEEE Journal of Selected Topics in Quantum Electronics, **8**(3): p. 705-13 (2002).
36. Nakagawa, S., N. Yamada, N. Mikoshiba, D.E. Mars, *Second-harmonic generation from GaAs/AlAs vertical cavity*. Applied Physics Letters, **66**(17): p. 2159-61 (1995).
37. Pellegrini, V., R. Colombelli, I. Carusotto, F. Beltram, S. Rubini, R. Lantier, A. Franciosi, C. Vinegoni, L. Pavesi, *Resonant second harmonic generation in ZnSe bulk microcavity*. Applied Physics Letters, **74**(14): p. 1945-7 (1999).
38. Berger, V., *Second-harmonic generation in monolithic cavities*. Journal of the Optical Society of America B (Optical Physics), **14**(6): p. 1351-60 (1997).
39. Liscidini, M., L.C. Andreani, *Second-harmonic generation in doubly resonant microcavities with periodic dielectric mirrors*. Physical Review E, **73**(1): p. 016613 (2006).
40. Simonneau, C., J.P. Debray, J.C. Harmand, P. Vidakovic, D.J. Lovering, J.A. Levenson, *Second-harmonic generation in a doubly resonant semiconductor microcavity*. Optics Letters, **22**(23): p. 1775-7 (1997).
41. Lodenkamper, R., M.L. Bortz, M.M. Fejer, K. Bacher, J.S. Harris, Jr., *Surface-emitting second-harmonic generation in a semiconductor vertical resonator*. Optics Letters, **18**(21): p. 1798-800 (1993).
42. Janz, S., Y. Beaulieu, A. Fiore, P. Bravetti, V. Berger, E. Rosencher, J. Nagle, *Surface emitted second-harmonic generation from a quasi-phase matched waveguide in an Al_xGa_{1-x}As/Al₂O₃ microcavity*. Optics Express, **2**(12) (1998).
43. Regener, R., W. Sohler, *Efficient second-harmonic generation in Ti:LiNbO₃ channel waveguide resonators*. Journal of the Optical Society of America B (Optical Physics), **5**(2): p. 267-77 (1988).
44. Arbore, M.A., M.M. Fejer, *Singly resonant optical parametric oscillation in periodically poled lithium niobate waveguides*. Optics Letters, **22**(3): p. 151-3 (1997).

45. Hofmann, D., G. Gschreiber, W. Grundkotter, R. Ricken, W. Sohler. *Mid-infrared continuous-wave singly resonant optical parametric oscillator with periodically poled Ti:LiNbO₃ waveguide*. in Digest. 2000 Conference on Lasers and Electro-Optics Europe, 10-15 Sept. 2000, Nice, France. Piscataway, NJ, USA : IEEE, 2000.(2000)
46. Schreiber, G., R. Ricken, K. Rochhausen, W. Sohler. *Doubly resonant near-infrared optical parametric oscillator with periodically poled Ti:LiNbO₃ waveguide*. in on Lasers and Electro-Optics (CLEO 2000). Technical Digest. Postconference Edition. TOPS Vol.39, 7-12 May 2000, San Francisco, CA, USA. Salem, MA, USA : Opt. Soc. America, 2000.(2000)
47. Schreiber, G., D. Hofmann, W. Grundkoetter, Y.L. Lee, H. Suche, V. Quiring, R. Ricken, W. Sohler, *Nonlinear integrated optical frequency converters with periodically poled Ti:LiNbO₃ waveguides*. Proceedings of the SPIE - Integrated Optics Devices V; Jan 23-25 2001; San Jose, CA, United States, **4277**: p. 144-60 (2001).
48. Saleh, B.E.A., M.C. Teich, *Fundamentals of photonics*. Series in Ppure and Applied Optics, ed. Wiley. New York: John Wiley & Sons, Inc. (1991).
49. Boyd, G.D., D.A. Kleinman, *Parametric interaction of focused gaussian light beams*. Journal of Applied Physics, **39**(8): p. 3597-3639 (1968).
50. Kogelnik, H., *Theory of optical waveguides*. Guided-Wave Optoelectronics, ed. T. Tamir. New York: Springer-Verlag. 7-88 (1990).
51. Marcuse, D., *Theory of Dielectric Waveguides*. San Diego: Academic Press (1991).
52. Bortz, M.L., S.J. Field, M.M. Fejer, D.W. Nam, R.G. Waarts, D.F. Welch, *Noncritical quasi-phase-matched second harmonic generation in an annealed proton-exchanged LiNbO₃ waveguide*. IEEE Journal of Quantum Electronics, **30**(12): p. 2953-60 (1994).
53. Franken, P.A., J.F. Ward, *Optical harmonics and nonlinear phenomena*. Reviews of Modern Physics, **35**(1): p. 23-39 (1963).

54. Rao, S.V., K. Moutzouris, M. Ebrahimzadeh, *Nonlinear frequency conversion in semiconductor optical waveguides using birefringent, modal and quasi-phase-matching techniques*. Journal of Optics A: Pure and Applied Optics, **6**(6): p. 569-584 (2004).
55. Moutzouris, K., S.V. Rao, M. Ebrahimzadeh, A. De Rossi, M. Calligaro, V. Ortiz, V. Berger, *Second-harmonic generation through optimized modal phase matching in semiconductor waveguides*. Applied Physics Letters, **83**(4): p. 620-622 (2003).
56. De Rossi, A., N. Semaltianos, E. Chirlias, B. Vinter, V. Ortiz, V. Berger, *Third order mode optically pumped semiconductor laser*. Applied Physics Letters, **80**(25): p. 4690-2 (2002).
57. Stern, M.S., *Semivectorial polarized finite-difference method for optical waveguides with arbitrary index profiles*. IEE Proceedings-J Optoelectronics, **135**(1): p. 56-63 (1988).
58. Johnson, S.G., S. Fan, A. Mekis, J.D. Joannopoulos, *Multipole-cancellation mechanism for high-Q cavities in the absence of a complete photonic band gap*. Applied Physics Letters, **78**(22): p. 3388-90 (2001).
59. Cho, A.Y., *Molecular beam epitaxy from research to manufacturing*. MRS Bulletin, **20**(4): p. 21-28 (1995).
60. Scherer, A., J.L. Jewell, Y.H. Lee, J.P. Harbison, L.T. Florez, *Fabrication of microlasers and microresonator optical switches*. Applied Physics Letters, **55**(26): p. 2724-6 (1989).
61. Kim, J.H., D.H. Lim, G.M. Yang, *Selective etching of AlGaAs/GaAs structures using the solutions of citric acid H_2O_2 and de-ionized H_2O buffered oxide etch*. Journal of Vacuum Science & Technology B, **16**(2): p. 558-560 (1998).
62. Bacher, K., J.S. Harris, *A wet etching technique for accurate etching of GaAs/AlAs distributed Bragg reflectors*. Journal of the Electrochemical Society, **142**(7): p. 2386-8 (1995).
63. Clawson, A.R., *Guide to references on III-V semiconductor chemical etching*. Materials Science and Engineering: R: Reports, **31**(1-6): p. 1-438 (2001).
64. Perozziello, E., *Personal communication*. 2004.

65. Scaccabarozzi, L., *Gallium Arsenide waveguides for non-linear optical applications*, Master thesis, Dept. of Nuclear Engineering, Politecnico di Milano (2001)
66. Donnelly, V.M., D.L. Flamm, *Anisotropic etching in chlorine-containing plasmas*. Solid State Technology, **24**(4): p. 161-6 (1981).
67. Franz, G., *Robust reactive ion etching processes for GaAs/AlGaAs/AlAs by application of statistical concepts*. Journal of the Electrochemical Society, **140**(4): p. 1147-51 (1993).
68. Scherer, A., H.G. Craighead, E.D. Beebe, *Gallium arsenide and aluminum gallium arsenide reactive ion etching in boron trichloride/argon mixtures*. Journal of Vacuum Science & Technology B (Microelectronics Processing and Phenomena), **5**(6): p. 1599-605 (1987).
69. Abdollahi-Alibeik, S., J. Zheng, J.P. Mcvittie, K.C. Saraswat, C.T. Gabriel, S.C. Abraham, *Modeling and simulation of feature-size-dependent etching of metal stacks*. Journal of Vacuum Science & Technology B (Microelectronics and Nanometer Structures), **19**(1): p. 179-85 (2001).
70. Ashby, C.I.H., J.P. Sullivan, K.D. Choquette, K.M. Geib, H.Q. Hou, *Wet oxidation of AlGaAs: the role of hydrogen*. Journal of Applied Physics, **82**(6): p. 3134-6 (1997).
71. Nickel, H., *A detailed experimental study of the wet oxidation kinetics of $Al_xGa_{1-x}As$ layers*. Journal of Applied Physics, **78**(8): p. 5201-3 (1995).
72. Choquette, K.D., K.M. Geib, C.I.H. Ashby, R.D. Twisten, O. Blum, H.Q. Hou, D.M. Follstaedt, B.E. Hammons, D. Mathes, R. Hull, *Advances in selective wet oxidation of AlGaAs alloys*. IEEE Journal of Selected Topics in Quantum Electronics, **3**(3): p. 916-26 (1997).
73. Durand, O., F. Wyckzisk, J. Olivier, M. Magis, P. Galtier, A. De Rossi, M. Calligaro, V. Ortiz, V. Berger, G. Leo, G. Assanto, *Contraction of aluminum oxide thin layers in optical heterostructures*. Applied Physics Letters, **83**(13): p. 2554-6 (2003).

74. Keller, R.R., A. Roshko, R.H. Geiss, K.A. Bertness, T.P. Quinn, *EBS*
measurement of strains in GaAs due to oxidation of buried AlGaAs layers.
Microelectronic Engineering, **75**(1): p. 96-102 (2004).
75. Regener, R., W. Sohler, *Loss in low-finesse Ti:LiNbO₃ optical waveguide
resonators*. Applied Physics B (Photophysics and Laser Chemistry), **B36**(3): p.
143-7 (1985).
76. Johnson, S.G., M. Ibanescu, M.A. Skorobogatiy, O. Weisberg, J.D. Joannopoulos,
Y. Fink, *Perturbation theory for Maxwell's equations with shifting material
boundaries*. Physical Review E (Statistical, Nonlinear, and Soft Matter Physics),
65(6): p. 066611-7 (2002).
77. Lee, K.K. *Revolutionary innovations in microphotonics*. in MRS Spring Meeting.
San Francisco.(2004)
78. Deri, R.J., M.A. Emanuel, *Consistent formula for the refractive index of Al_xGa_{1-x}
As below the band edge*. Journal of Applied Physics, **77**(9): p. 4667-72 (1995).
79. Sfigakis, F., P. Paddon, V. Pacradouni, M. Adamcyk, C. Nicoll, A.R. Cowan, T.
Tiedje, J.F. Young, *Near-infrared refractive index of thick, laterally oxidized
AlGaAs cladding layers*. Journal of Lightwave Technology, **18**(2): p. 199-202
(2000).
80. Knopp, K.J., R.P. Mirin, D.H. Christensen, K.A. Bertness, A. Roshko, R.A.
Synowicki, *Optical constants of (Al_{0.98}Ga_{0.02})_xO_y native oxides*. Applied Physics
Letters, **73**(24): p. 3512-3514 (1998).
81. Parameswaran, K.R., J.R. Kurz, R.V. Roussev, M.M. Fejer, *Observation of 99%
pump depletion in single-pass second-harmonic generation in a periodically
poled lithium niobate waveguide*. Optics Letters, **27**(1): p. 43-5 (2002).
82. Hall, D.C., H. Wu, L. Kou, Y. Luo, R.J. Epstein, O. Blum, H. Hou, *Refractive
index and hygroscopic stability of Al_xGa_{1-x}As native oxides*. Applied Physics
Letters, **75**(8): p. 1110-12 (1999).
83. Johnson, S.G., M.I. Povinelli, M. Soljagic, A. Karalis, S. Jacobs, J.D.
Joannopoulos, *Roughness losses and volume-current methods in photonic-crystal
waveguides*. Applied Physics B: Lasers and Optics, **81**(2-3): p. 283-293 (2005).

84. Guha, S., F. Agahi, B. Pezeshki, J.A. Kash, D.W. Kisker, N.A. Bojarczuk, *Microstructure of AlGaAs-oxide heterolayers formed by wet oxidation*. Applied Physics Letters, **68**(7): p. 906-8 (1996).
85. Twesten, R.D., D.M. Follstaedt, K.D. Choquette, R.P. Schneider, *Microstructure of laterally oxidized $Al_xGa_{1-x}As$ layers in vertical-cavity lasers*. Applied Physics Letters, **69**(1): p. 19-21 (1996).
86. Cheong, S.K., B.A. Bunker, T. Shibata, D.C. Hall, C.B. Demelo, Y. Luo, G.L. Snider, G. Kramer, N. El-Zein, *Residual arsenic site in oxidized $Al_xGa_{1-x}As$ ($x=0.96$)*. Applied Physics Letters, **78**(17): p. 2458-60 (2001).
87. Ashby, C.I., J.P. Sullivan, P.P. Newcomer, N.A. Missert, H.Q. Hou, B.E. Hammons, M.J. Hafich, A.G. Baca, *Interfacial arsenic from wet oxidation of $Al_xGa_{1-x}As/GaAs$: Its effects on electronic properties and new approaches to MIS device fabrication*. MRS Fall Meeting - Control of Semiconductor Surfaces and Interfaces, **448**: p. 285-290 (1997).
88. Shi, S.S., E.L. Hu, J.P. Zhang, Y.I. Chang, P. Parikh, U. Mishra, *Photoluminescence study of hydrogenated aluminum oxide-semiconductor interface*. Applied Physics Letters, **70**(10): p. 1293-5 (1997).
89. Joannopoulos, J.D., R. D. Meade, J.N. Winn, *Photonic Crystals: molding the flow of light*. Princeton, NJ: Princeton University Press (1995).
90. Bethune, D.S., *Optical harmonic generation and mixing in multilayer media: analysis using optical transfer matrix techniques*. Journal of the Optical Society of America B (Optical Physics), **6**(5): p. 910-16 (1989).
91. Chow, E., A. Grot, L.W. Mirkarimi, M. Sigalas, G. Girolami, *Ultracompact biochemical sensor built with two-dimensional photonic crystal microcavity*. Optics Letters, **29**(10): p. 1093-5 (2004).
92. Jugessur, A.S., P. Pottier, R.M. De La Rue, *One-dimensional periodic photonic crystal microcavity filters with transition mode-matching features, embedded in ridge waveguides*. Electronics Letters, **39**(4): p. 367-9 (2003).

93. Barwicz, T., H.A. Haus, *Three-dimensional analysis of scattering losses due to sidewall roughness in microphotonic waveguides*. Journal of Lightwave Technology, **23**(9): p. 2719-32 (2005).
94. Almeida, V.R., R.R. Panepucci, M. Lipson, *Nanotaper for compact mode conversion*. Optics Letters, **28**(15): p. 1302-4 (2003).
95. Taillaert, D., W. Bogaerts, P. Bienstman, T.F. Krauss, P. Van Daele, I. Moerman, S. Verstuyft, K. De Mesel, R. Baets, *An out-of-plane grating coupler for efficient butt-coupling between compact planar waveguides and single-mode fibers*. IEEE Journal of Quantum Electronics, **38**(7): p. 949-55 (2002).
96. Leo, G., V. Berger, A. De Rossi, M. Calligaro, X. Marcadet, A. Fiore, J. Nagle, *Parametric processes in GaAs/Al_xO_y structures*. Proceedings of the SPIE Nonlinear Materials, Devices and Applications, **3928**: p. 94-107 (2000).
97. De Rossi, A., V. Berger, G. Leo, G. Assanto, *Form birefringence phase matching in multilayer semiconductor waveguides: tuning and tolerances*. IEEE Journal of Quantum Electronics, **41**(10): p. 1293-302 (2005).
98. Berger, V., C. Sirtori, *Nonlinear phase matching in THz semiconductor waveguides*. Semiconductor Science and Technology, **19**(8): p. 964-70 (2004).
99. Capasso, F., R. Paiella, R. Martini, R. Colombelli, C. Gmachl, T.L. Myers, M.S. Taubman, R.M. Williams, C.G. Bethea, K. Unterrainer, H.Y. Hwang, D.L. Sivco, A.Y. Cho, A.M. Sergent, H.C. Liu, E.A. Whittaker, *Quantum cascade lasers: ultrahigh-speed operation, optical wireless communication, narrow linewidth, and far-infrared emission*. IEEE Journal of Quantum Electronics, **38**(6): p. 511-32 (2002).