ACTIVE SLOW LIGHT IN SILICON PHOTONIC CRYSTALS: TUNABLE DELAY AND RAMAN GAIN

Isabella H. Rey

A Thesis Submitted for the Degree of PhD
at the
University of St. Andrews

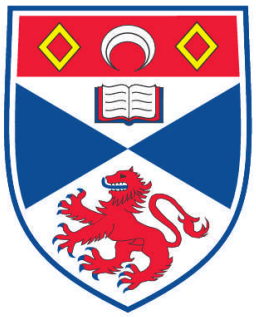

2012

Active Slow Light in Silicon Photonic Crystals: Tunable Delay and Raman Gain

Isabella H. Rey

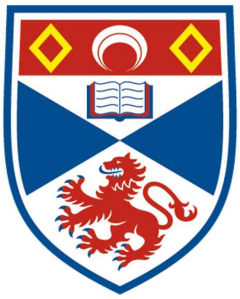

University of
St Andrews

Thesis submitted for the degree of Doctor of Philosophy
at the
University of St Andrews

September 2012

Declarations

I, Isabella H. Rey, hereby certify that this thesis, which is approximately 50,000 words in length, has been written by me, that it is the record of work carried out by me and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in March 2009 and as a candidate for the degree of Doctor of Philosophy in March 2009; the higher study for which this is a record was carried out in the University of St Andrews between 2009 and 2012.

Date: ..... Signature of Candidate: .....

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of Doctor of Philosophy in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date: ..... Signature of Supervisor: .....

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that my thesis will be electronically accessible for personal or research use unless exempt by award of an embargo as requested below,

and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. I have obtained any third-party copyright permissions that may be required in order to allow such access and migration, or have requested the appropriate embargo below.

The following is an agreed request by candidate and supervisor regarding the electronic publication of this thesis:
Access to printed copy and electronic publication of thesis through the University of St Andrews.

Date: .....

Signature of Candidate: ..... Signature of Supervisor: .....

Acknowledgements

I would like to thank my supervisor Prof. Thomas Krauss for giving me once again the opportunity of working with photonic crystals; his guidance and his enthusiasm make him an inspirational example.
A very special thank to Dr. Juntao Li, for his great help in the development of inverse tapers and the nonlinear optical setup, for very fruitful discussion and his example of hard work.
To Dr. Daryl Beggs and colleagues from the NanoOptics group of the FOM Institute AMOLF, Amsterdam, for the measurements on tunable delay, many tips and useful discussion.
To Dr. William Whelan-Curtin (Liam O'Faolain), for his invaluable presence inside and outside of the clean-room and his precious advice. Thanks also to Dr. Andrea Di Falco and Dr. Thomas White for their time and useful discussion.
A very special thank to Dr. Christopher Reardon for suggestions on the fabrication process, for his great help in many managing aspects and for lighting up the office with his unique sense of humour.
Thanks to Mr. Steve Balfour, Mr. Callum Smith and Mr. George Robb for keeping the clean-room up and running no matter what.
A great thank to all my friends and colleagues from the Microphotonics Research Group for their support and their friendship. In particular, I would like to thank Abdul Shakoor, Annett Klemm, Emiliano Rezende Martins and his wife Andrea, Kapil Debnath, Mark Scullion, Dr. Marcel Spurny, Dr. Vikash Venkataramana and Dr. Shu Chen.
The EPSRC programme UK Silicon Photonics for funding of my work.

Once again, a very big thank to (Dr.!) Paola Avetta for her subtle humour and her long and ever lasting friendship.
To my family, my father Claudio, my sisters Marcella, Alessandra and Giulia, my uncle Mario and my granny Helen, for their constant love and support. And to my mother Jane, for being always there for me.
Lastly, thanks to Dr. Karl Welna, who performed many Q-factor simulations of the photonic crystal cavities, helped with their fabrication and measured them. Most of all, a very special thank for his never ending support, patience and care.

List of publications

Journal papers

- C. Reardon, **I. H. Rey**, K. Welna, L. O'Faolain, and T. F. Krauss, "Fabrication and characterisation of photonic crystal slow light waveguides and cavities," *Journal of Visualised Experiments*, in press 2012.

  (Reardon, Rey and Welna joint first authors)

- D. M. Beggs, **I. H. Rey**, T. Kampfrath, N. Rotenberg, L. Kuipers, and T. F. Krauss, "Ultrafast tunable optical delay line based on indirect photonic transitions," *Physical Review Letters* **108**(21):213901, 2012.

  (Beggs and Rey joint first authors)

- **I. H. Rey**, Y. Lefevre, S. A. Schulz, N. Vermeulen, and T. F. Krauss, "Scaling of Raman amplification in realistic slow light photonic crystal waveguides," *Physical Review B* **84**(3):035306, 2011.

- J. Li, L. O'Faolain, **I. H. Rey**, and T. F. Krauss, "Four-wave mixing in photonic crystal waveguides: slow light enhancement and limitations," *Optics Express* **19**(5):4458–4463, 2011.

Conferences attended

- **I. H. Rey**, J. Li, L. O'Faolain, C. Xiong, C. Monat, B. J. Eggleton, and T. F. Krauss, "Si slow light photonic crystals: four wave mixing and on-chip quantum interference," *Photon 12*, Durham, Sep 2012. Invited talk.

- **I. H. Rey**, D. M. Beggs, T. Kampfrath, N. Rotenberg, L. Kuipers, and T. F. Krauss, "Ultrafast tunable optical delay based on a 2-step indirect photonic transition," *JSPS Fall School on Silicon Photonics*, Kyoto, Nov 2011. Poster presentation.

- **I. H. Rey**, D. M. Beggs, T. Kampfrath, L. Kuipers, and T. F. Krauss, "Ultrafast tunable optical delay in slow light photonic crystal wave-

guides," *GFP 2011. 8th IEEE International Conference on Group IV Photonics*, London, Sep 2011. Contributed talk.

- **I. H. Rey**, D. M. Beggs, T. Kampfrath, L. Kuipers, and T. F. Krauss, "Ultrafast tunable optical delay in engineered slow light photonic crystal waveguides," *Silicon Photonics Summer School*, St Andrews, Jul 2011. Poster presentation.

- **I. H. Rey**, D. M. Beggs, T. Kampfrath, L. Kuipers, and T. F. Krauss, "Ultrafast tunable optical delay in engineered slow light photonic crystal waveguides," *2nd Annual SU2P Symposium*, St Andrews, Mar 2011. Poster presentation.

- **I. H. Rey**, D. M. Beggs, T. Kampfrath, L. Kuipers, and T. F. Krauss, "Tunable optical delay line in engineered slow light photonic crystal waveguides," *PECS-IX 2010. 9th International Conference on Photonic and Electromagnetic Structures*, Granada, Sep 2010. Contributed talk.

- **I. H. Rey** and T. F. Krauss, "Tunable delay and Raman scattering in silicon slow light photonic crystal waveguides," *EMMP Summer School*, Edinburgh, Jun 2010. Invited talk.

- **I. H. Rey**, D. M. Beggs, T. Kampfrath, L. Kuipers, and T. F. Krauss, "Tunable optical delay line in engineered slow light photonic crystal waveguides," *ECIO 2010. 15th European Conference on Integrated Optics*, Cambridge, Apr 2010. Poster presentation.

Abstract

In the past decade, great research effort was inspired by the need to realise active optical functionalities in silicon, in order to develop the full potential of silicon as a photonic platform. In this thesis we explore the possibility of achieving tunable delay and optical gain in silicon, taking advantage of the unique dispersion capabilities of photonic crystals.

To achieve tunable optical delay, we adopt a wavelength conversion and group velocity dispersion approach in a miniaturised engineered slow light photonic crystal waveguide. Our scheme is equivalent to a two-step indirect photonic transition, involving an alteration of both the frequency and momentum of an optical pulse, where the former is modified by the adiabatic tuning possibilities enabled by slow light. We apply this concept in a demonstration of continuous tunability of the delay of pulses, and exploit the ultrafast nature of the tuning process to demonstrate manipulation of a single pulse in a train of two pulses.

In order to address the propagation loss intrinsic to slow light structures, with a prospect for improving the performance of the tunable delay device, we also investigate the nonlinear effect of stimulated Raman scattering as a means of introducing optical gain in silicon. We study the influence of slow-down factors and pump-induced losses on the evolution of a signal beam along the waveguide, as well as the role of linear propagation loss and mode profile changes typical of realistic photonic crystal structures. We then describe the work conducted for the experimental demonstration of such effect and its enhancement due to slow light.

Finally, as the Raman nonlinearity may become useful also in photonic crystal nanocavities, which confine light in very small volumes, we discuss the design and realisation of structures which satisfy the basic requirements on the resonant modes needed for improving Raman scattering.

Non dire gatto...

Contents

Chapter 1

Introduction

In today's communication networks, optical fibres allow for fast transmission of data, but since all the components for signal manipulation operate in the electronic domain, optical to electrical (O/E) and electrical to optical (E/O) conversions at the input and output of each electronic component are required. As data transmission rates are ever-increasing [1], the latency, cost and power consumption of O/E and E/O switching become prohibitive [2]. Performing part of the signal processing operations in the optical domain [3, 4], thereby complementing electronics with photonics, has the potential for eliminating these conversion inefficiencies [5, 6], making next-generation information networks even faster.

One of the most promising platforms for large scale integration of optoelectronic systems is silicon [7–9], which can rely on the mature and reliable CMOS processing technology and has very advantageous optical properties: silicon is transparent in the near- and mid-infrared and the availability of high quality silicon-on-insulator (SOI) wafers, with the associated high index contrast between silicon and silica, allows for strong optical confinement within the silicon core, and thus for miniaturisation of photonic devices [10].

While passive optical components such as waveguides, resonators, grating couplers and arrayed waveguide gratings are well established, achieving active functionalities in silicon is more challenging. This is due to the fact that silicon is an indirect band gap semiconductor, and thus cannot emit light in traditional ways, and the crystalline structure suppresses second-order nonlinear processes such as the electro-optic or Pockels effect, which is usually employed in other materials to modulate light [7].

In addition to the three main functionalities required for all-optical signal processing — light emission, modulation and detection — other basic building blocks are also needed. One of these is the tunable optical delay line, a device that temporarily stores optical data streams and adjusts their arrival times. Such a device could be employed in a great variety of applications requiring time domain operation, including optical buffers and

memories, optical packet switches at network nodes, data synchronization, time division multiplexing and optical calculation of the convolution integral [11, 12]. At present, the most practical solution of tunable delay for laboratory demonstrations of optical packet switches is based on the use of fibre delay lines of different lengths combined with optical switches [13]. However, this bulky solution becomes impractical when scaling to realistic size networks, and can generate only discrete delays, whereas fine and continuous tuning of the delay in a miniaturised device would be preferred for real applications.

Different schemes have been proposed for generating tunable delays in integrated systems, mainly involving tuning chains of coupled microresonators [14–16] or trapping and releasing optical pulses in few resonators by varying their coupling with an adjacent waveguide [17,18]. However, the former case requires the non-trivial ability to precisely align the resonant frequencies of many cavities and to tune them uniformly; in the latter, the maximum achievable delay is limited by the photon lifetime in the cavities.

An alternative approach to realising tunable delay is based on wavelength conversion and dispersion [19–21]: if an input signal is shifted in wavelength and coupled into a dispersive medium, it will propagate at a different speed with respect to an unshifted pulse, and therefore will experience a different delay. Usually wavelength conversion in silicon is achieved by means of nonlinear optical effects requiring large amounts of power and a minimum travelling distance, and often relying on phase-matching conditions [22]. A simpler process is wavelength conversion due to adiabatic tuning of an optical resonator [23], a well known phenomenon in classical physics, but relatively new in photonics, as it requires tuning of the system within the photon lifetime, which was not possible to achieve before the advent of slow light waveguides [24, 25] and high quality factor nanocavities [26, 27].

This project aims to develop a miniaturised continuously tunable optical delay line based on adiabatic wavelength conversion and dispersion in silicon. Our device will be based on photonic crystals, a very versatile platform allowing fine control of the dispersion properties of waveguides in the form of engineered slow light structures [28]. Our goal is well summarised by the schematic of Fig. 1.1, where an input light pulse entering a photonic crystal waveguide is dynamically changed in wavelength (or colour) due to a control pulse, and then propagates through a waveguide section of controlled dispersion. The study for the realisation of such dynamic delay line will be discussed in detail in Chapter 3 of this thesis, following a brief review on some basic concepts about photonic crystals and recent progress on slow light waveguides, which will be presented in Chapter 2.

Slow light photonic crystal waveguides, however, are inherently limited by increased propagation loss [29–31], as represented in Fig. 1.1 by the lowering of the pulse height as we proceed towards the output. We will therefore explore the possibility of including a gain effect for compensation

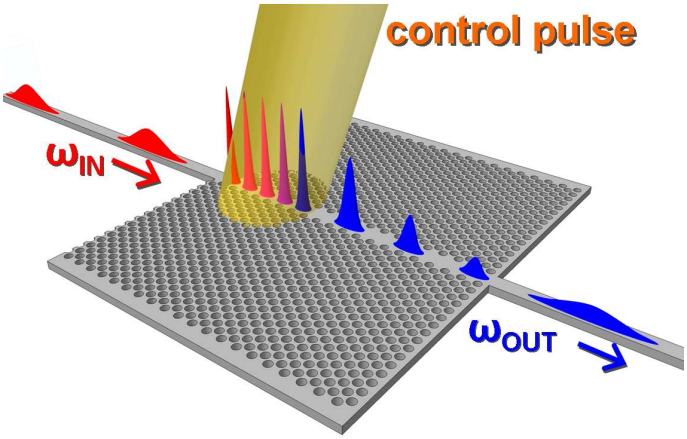

Figure 1.1: Schematic of a tunable delay line based on adiabatic wavelength conversion and group velocity dispersion in an engineered slow light photonic crystal waveguide.

of out-of-plane scattering losses. In particular, we will study stimulated Raman scattering, a third-order nonlinear optical process occurring when an input pump is scattered at different frequencies due to the interaction with the vibrational modes of the scattering medium [32].

Raman scattering has been successfully employed to achieve light amplification [33–35] and also lasing [36–38] in silicon channel waveguides. The slow light effect offered by the unique dispersion properties of photonic crystal waveguides increases light-matter interaction and thus enhances optical nonlinearities [39]. For the particular case of stimulated Raman scattering, such enhancement was theoretically demonstrated [40], and some promising experimental results on enhancement of *spontaneous* scattering have been reported [41, 42], but the demonstration of improvement due to slow light in a pump-probe experiment of *stimulated* scattering is still missing.

As the observation of nonlinear effects in compact silicon waveguides requires the fabrication of structures for efficient in- and out-coupling of light, we will begin in Chapter 4 by describing the integration of inverse taper spot-size converters with photonic crystal devices. We will then proceed in Chapter 5 by discussing the role of realistic propagation loss and mode profile effects on the amplification of an optical signal through stimulated Raman scattering; we will also present experimental results on both spontaneous and stimulated Raman scattering from the fabricated structures, with a focus on loss compensation.

The overall goal of this project is therefore summarised in Fig. 1.2: a loss-compensated photonic crystal tunable delay line, where loss compensation is achieved through the nonlinear effect of stimulated Raman scattering by the addition of a pump signal (represented in green), and efficient coupling

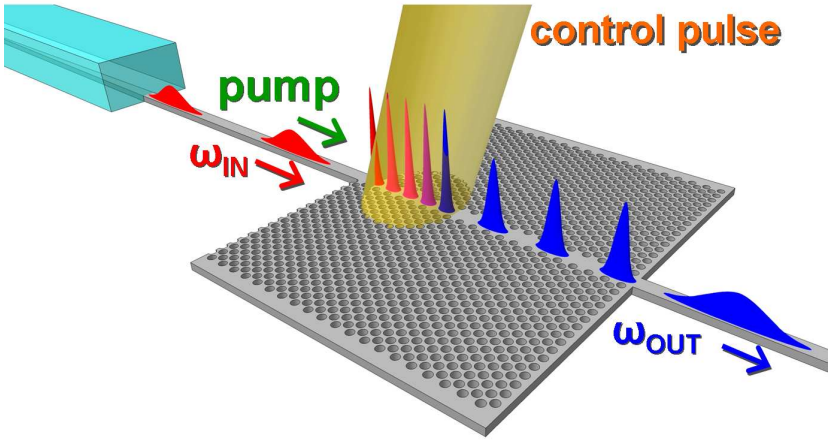

Figure 1.2: Schematic describing the aim of this project: a loss-compensated tunable delay line. With respect to Fig. 1.2, the output pulse does not decrease in intensity while propagating through the waveguide. Loss compensation is achieved by addition of a pump signal (green) which amplifies the probe signal through stimulated Raman scattering. Efficient in-coupling is achieved by integrating an inverse taper spot-size converter (on the left).

is mediated by inverse taper spot-size converters.

Finally, slow light waveguides are not the only structures able to generate high optical intensities: photonic crystal nanocavities can confine and hold light in very small volumes, and therefore they are also suitable for enhancing the Raman nonlinearity [43]. Useful application of Raman scattering in silicon optical nanocavities, however, requires first of all the realisation of structures supporting two resonant modes at a very specific frequency spacing: an example of such structures has been proposed numerically [44], but fabrication tolerances cause the resonances to shift in real devices, and to our knowledge there has been no experimental demonstration of a photonic crystal cavity exhibiting two modes at the correct frequency spacing for Raman scattering in silicon. In Chapter 6 we will show both numerically and experimentally how the position of the holes can be employed as a tuning parameter to adjust a cavity's resonances to the desired spacing in the 1550 nm wavelength range.

Chapter 2

Photonic crystals and slow light: background

Photonic crystals have emerged in the past few decades as an exciting tool to control light in an integrated format [45], opening a wide range of new possibilities in fields from telecommunications and optical signal processing, to biosensing and quantum computing, simply by playing on the geometry of their structure. One of the most promising features enabled by photonic crystals is slow light, which generates from the resonant interaction of photons with the periodic pattern defined in a dielectric, at an operating wavelength that only depends on the lattice constant of the pattern, instead of being set by a material resonance such as in electromagnetically-induced transparency in atomic vapours [46] or coherent population oscillation in semiconductors [47, 48]. By definition, slow light allows to delay and temporarily store optical signals, but is also accompanied by strong light-matter interaction, which is beneficial for reducing the size, and ultimately the operating power, of both linear and nonlinear devices.

This chapter begins by introducing some fundamental concepts concerning photonic crystals in general and slow light waveguides in particular, the platform on which the topics of tunable delay and stimulated Raman scattering discussed in the following chapters are based. We then proceed by describing the typical fabrication process of a photonic crystal waveguide and the measurement technique used to characterise slow light.

2.1 Photonic crystals and photonic crystal waveguides

The dynamics of electromagnetic fields are governed by Maxwell's equations, which in a dielectric medium with dielectric function $\varepsilon(\mathbf{r})$ can be solved to determine the electric and magnetic fields, $\mathbf{E}$ and $\mathbf{H}$ respectively. The stationary solutions of the system may be determined by expressing the

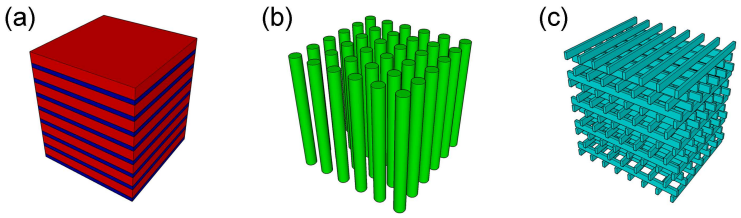

Figure 2.1: Example schematics of (a) one-, (b) two-, and (c) three-dimensional photonic crystals. (a) A multilayer film. (b) A square lattice of dielectric columns. (c) A woodpile crystal.

fields in the form of harmonic modes, i.e. as the product of a spatial pattern or mode profile with a complex exponential:

$$\mathbf{E}(\mathbf{r},t) = \mathbf{E}(\mathbf{r})\, e^{-i\omega t},$$
$$\mathbf{H}(\mathbf{r},t) = \mathbf{H}(\mathbf{r})\, e^{-i\omega t},$$

with ω the angular frequency. Knowing $\varepsilon(\mathbf{r})$ one can solve Maxwell's equations, which can be written in the form of an eigenvalue problem, to determine the spatial mode profiles $\mathbf{H}(\mathbf{r})$ (the system's eigenfunctions) and the corresponding frequencies ω (the eigenvalues).

Photonic crystals, in particular, are characterised by discrete periodic changes of $\varepsilon(\mathbf{r})$ in one, two or three dimensions (Fig. 2.1), with a period a of the order of the optical wavelength in the material. Such a periodic modulation acts on photons in close analogy to the way an atomic crystalline potential acts on electrons: the discrete periodicity of $\varepsilon(\mathbf{r})$ in one direction results in a dependence of $\mathbf{H}(\mathbf{r})$ in that direction that is a plane wave modulated by a function $\mathbf{u_k}(\mathbf{r})$ periodic on the lattice [45]:

$$\mathbf{H_k}(\mathbf{r}) = e^{i(\mathbf{k}\cdot\mathbf{r})}\mathbf{u_k}(\mathbf{r})\,, \tag{2.1}$$

where $\mathbf{k}$ is the wavevector. In analogy with solid state physics, the modes in the form (2.1) are known as Bloch states, and are classified by the value of $\mathbf{k}$. All the information about $\mathbf{H_k}(\mathbf{r})$ is given by $\mathbf{k}$ and $\mathbf{u_k}(\mathbf{r})$.

Due to the periodicity, the eigenvalue problem may be restricted to a single unit cell of the photonic crystal, leading to a discrete spectrum of frequencies (eigenvalues) $\omega_n(\mathbf{k})$ labelled by a band index n, for each $\mathbf{k}$. Being $\mathbf{k}$ a continuous parameter, each band $\omega_n(\mathbf{k})$ also varies continuously with $\mathbf{k}$, defining the band structure, or dispersion diagram, of the photonic crystal. The periodicity and symmetries of the photonic crystal in real space result in the band structure $\omega_n(\mathbf{k})$ to have the same periodicity and symmetries in the reciprocal k-space, so that the band diagram needs to be calculated only in a limited region, known as the first (irreducible) Brillouin zone (Fig. 2.2, left inset).

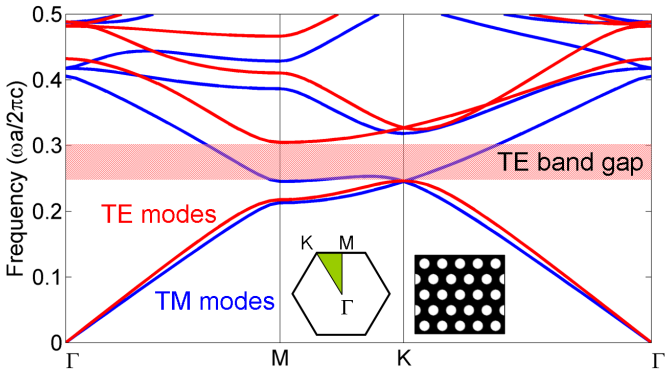

Figure 2.2: Two-dimensional photonic band structure for a triangular lattice of air holes in dielectric, with $\varepsilon = 8.24$ and $r = 0.286a$. The left inset shows the first Brillouin zone, with the irreducible Brillouin zone highlighted in green. The right inset shows a cross-sectional view of the photonic crystal structure.

In practice, the band structure and mode profiles are calculated numerically, for instance through use of MIT Photonic Band (MPB) [49], a free software package that applies a plane wave expansion method to solve the eigenvalue problem: the unknown function $\mathbf{u_k}(\mathbf{r})$, periodic on the crystal lattice, is expanded in a Fourier series. The eigenvalue problem therefore turns into a set of linear equations for the unknown Fourier coefficients, which are solved with an iterative method.

An example of photonic band diagram calculated with MPB is shown in Fig. 2.2 for a two-dimensional photonic crystal: a triangular lattice of air holes in dielectric. The red curves represent the modes of the transverse-electric (TE) polarisation, i.e. with the electric field confined to the plane of symmetry and the magnetic field normal to the plane, whereas the blue curves are the modes of the transverse-magnetic (TM) polarisation, i.e. with the magnetic field lying on the plane and the electric field perpendicular to the plane. As can be seen from Fig. 2.2, the TE band diagram of this type of crystal is characterised by a frequency region in which no mode can propagate *in plane* inside the structure, regardless of the wavevector. This region of forbidden frequencies is called a photonic band gap, in analogy with solid state physics.

In fact, the band gap is a region where there are no *extended* states, i.e. modes of the form (2.1) with purely real $\mathbf{k}$: the modes have complex wavevector instead, and they decay exponentially due to the imaginary part of $\mathbf{k}$. These evanescent states cannot be excited in a perfect infinite crystal, but they may be sustained by the presence of a defect: such modes decay exponentially away from the defect, where periodicity is restored, and are therefore localised to the defect itself. In the triangular lattice of Fig. 2.2,

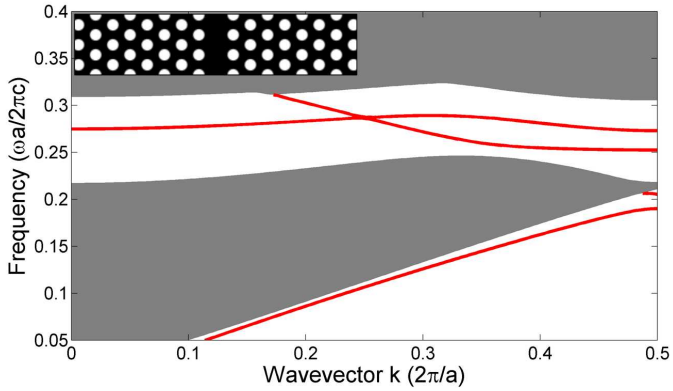

Figure 2.3: Two-dimensional band structure of a W1 waveguide, formed by removing a row of holes from the triangular lattice of air holes of Fig. 2.2 ($\varepsilon = 8.24$, $r = 0.286a$), plotted versus the wavevector component along the defect. Two defect modes (red) appear inside the band gap. The inset shows a cross-sectional view of the photonic crystal waveguide.

for instance, a *point defect* may be created by removing a single hole, or by replacing it with one of different size, shape or dielectric constant, therefore creating an optical cavity which traps light in the plane.

A *line defect*, obtained for example by the removal of an entire row of holes in a triangular lattice (Fig. 2.3, inset), creates an optical waveguide (called a W1 waveguide): light cannot penetrate on the sides of the defect due to the band gap, but it can propagate along the line defect. As periodicity is still preserved in the direction of the waveguide, we can still define a wavevector k in this direction and plot a waveguide band diagram, as shown in Fig. 2.3. The continuous regions are the result of projecting the (TE) band structure of the infinite crystal on the waveguide k axis, and any combination of wavevector and frequency in these regions corresponds to an allowed extended state of the crystal. Light in these states can leak through the sides of the waveguide and therefore cannot be guided. The presence of the defect, however, introduces new discrete bands (red curves in Fig. 2.3) inside the band gap, which correspond to guided modes.

Now, the band diagram computed for a two-dimensional photonic crystal applies in three dimensions only to a structure of infinite extent in the third direction z, and only for light propagating in the plane. Light can be localised in all three directions through use of a three-dimensional photonic crystal (Fig. 2.1c), which is the true optical equivalent of an atomic crystal. These type of structures, however, are difficult to fabricate. A widely used alternative is given by photonic crystal slabs, which approximate many desirable properties of more complicated three-dimensional photonic crystals, but are much easier to fabricate. A photonic crystal slab consists of a high

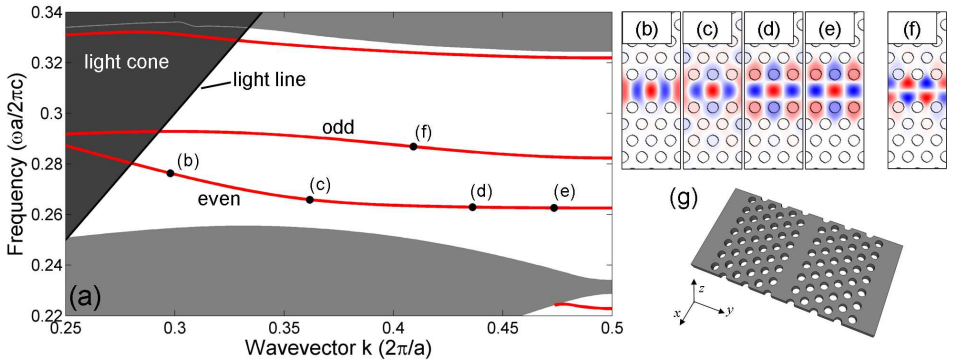

Figure 2.4: (a) Three-dimensional band structure of a W1 slab waveguide ($\varepsilon = 12.11$, $r = 0.286a$). The dark shaded region indicates the light cone. (b-f) H_z field cross-sections of the even (b-e) and odd (f) modes, calculated at the points marked by black dots in (a). (g) Schematic of the slab photonic crystal waveguide.

index layer with finite thickness in which the periodic pattern is defined, as shown in Fig. 2.4g. Confinement in the vertical direction is achieved by means of index guiding, i.e. total internal reflection. The band structure of a slab photonic crystal waveguide differs from that of a two-dimensional structure in that it presents a light cone, indicated by the dark shaded region of Fig. 2.4a. The points inside the light cone represent states where light is incident at the interface between the dielectric and the surrounding low index background at an angle lower than the critical angle for total internal reflection. The lowest boundary of the light cone is called the light line. To be guided, a waveguide mode must lie inside the band gap, or it will leak through the sides of the waveguide, *and* below the light line, or it will leak in the background.

Finally, a word on the field profiles. Since the TE modes have just one component of the magnetic field, they can be visualised by plotting H_z in the $z = 0$ plane (Figs. 2.4b-f). The usual convention is to classify them in even and odd modes as to whether H_z has even or odd symmetry in the y-direction, respectively. The lowest mode appearing in the band gap of a W1 waveguide is even (Figs. 2.4b-e), whereas the second mode is odd (Fig. 2.4f). Even modes are generally easier to excite than odd modes.

2.2 Slow light in photonic crystal waveguides

Another important information that may be deduced from the band diagram $\omega(k)$ of a photonic crystal waveguide is the group velocity v_g of light, that is, the speed at which the electromagnetic energy travels through the waveguide. The group velocity is given by the local slope of the dispersion

curve:

$$v_g = \frac{\partial \omega}{\partial k}.$$

Equivalently, the speed of light may be expressed in terms of group index $n_g = c/v_g$.

As can be seen from Fig. 2.4a, the group velocity becomes zero at the edge of the Brillouin zone $k = \pi/a$. This is due to the fact that light propagating along the waveguide is coherently backscattered at each unit cell of the crystal [39,50]; at the band edge, the forward- and backward-propagating waves agree in phase and amplitude and form a standing wave, i.e. an interference pattern with zero group velocity. The slow light regime of the waveguide occurs just before the band edge, where the slope of the dispersion curve is low but non-zero: here the forward- and backward- propagating waves are not exactly in phase any more, and the resulting interference pattern moves slowly in the forward direction.

The use of slow light in photonic crystal waveguides is of great interest because it allows to increase light-matter interactions simply by playing on the geometrical properties of the structure, with potential benefit for a number of applications [24]. Linear effects, such as phase sensitivity to refractive index changes, scale with the group index of light [50,51], allowing to reduce the size, and consequently the operating power, of devices such as switches [52], modulators [53–56] and wavelength converters [57]. Nonlinear effects are also enhanced due to slow light [40, 58–62], as an optical pulse entering a slow light waveguide is compressed in space in the propagation direction, and therefore becomes more intense [24].

The band edge slow light of a basic W1 waveguide (Fig. 2.4a) is not ideal to work with from a practical point of view [39]. First of all, the group velocity dispersion is large, due to the parabolic nature of the dispersion curve $\omega(k)$ near the band edge; for many applications it is instead desirable to reduce the group velocity dispersion to a minimum over a broad bandwidth. Secondly, the slow light regime occurs in vicinity of the mode cut-off, which in a real device may vary locally due to fabrication tolerances [63], resulting in extremely large propagation loss [30]. To avoid these issues, engineered slow light photonic crystal waveguides have been investigated, in which a controlled slow light region away from the band edge may be generated by suitably acting on the geometry of the crystal [64], for example by changing the size [65] or position [28, 66] of the first rows of holes adjacent to the waveguide, or by chirping the waveguide properties [67].

Most of these dispersion engineering techniques rely on the fact that the photonic crystal W1 waveguide supports not only gap-guided modes, but also index-guided ones [45], due to the central part of the waveguide having higher average refractive index than the surrounding lattice. The even guided mode in the band gap is actually a combination of these two types, resulting from an anti-crossing between a gap-guided mode and a

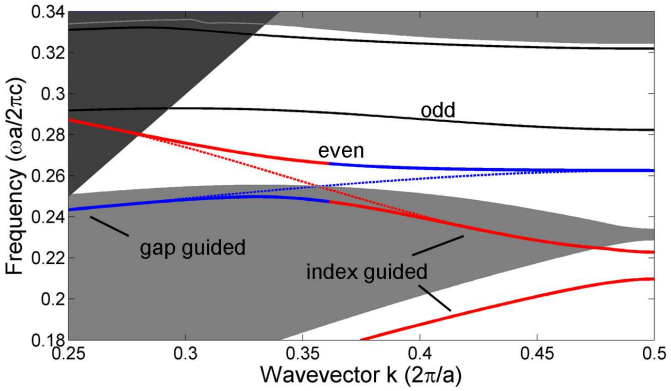

Figure 2.5: Schematic of anti-crossing occurring in a W1 photonic crystal waveguide. The fundamental index-guided mode (lowest red curve) is folded back at the edge of the first Brillouin zone and couples to a gap-guided mode (blue). The dotted lines illustrate hypothetical non-interacting modes. The two modes split at the crossing point, forming an hybrid band in which the field profile gradually evolves from an index-guided to a gap-guided profile (Figs. 2.4b-e).

folded index-guided mode [45,68,69], as shown in Fig. 2.5. The local shape — and thus group velocity — of the final mode is determined by the interaction between the two original modes, which in turn depends on the geometry of the crystal: at low k the mode behaves as an index-guided mode and its energy is concentrated inside the waveguide core (Fig. 2.4b), whereas when entering the slow light regime at higher k the mode penetrates in the crystal cladding (Figs. 2.4d-e). Therefore, the mode properties in the index-guided regime mainly depend on the first row of holes, but in the slow light regime they are also influenced by the second row.

Figure 2.6 shows how the position of the first two rows of holes may be changed to engineer the slow light dispersion curve, as reported by Li *et al.* in Ref. [28], where s_1 and s_2 indicate the displacements of the first and second row of holes, respectively, relative to the unmodified lattice. Shifting the first row of holes outwards (s_1 negative) makes the waveguide wider and increases the effective index of the mode in the entire k-range, lowering it in frequency, whereas shifting the second row inwards (s_2 positive) has the opposite effect, but is felt only at high wavevectors. By combining these two effects, a controlled region of the dispersion curve may be obtained, in which the group index is high and almost constant. A similar effect can be achieved by varying the radius of the first two rows of holes [65], but this is technologically more difficult to control in an accurate and reproducible manner.

The preservation of a useful bandwidth with low group velocity dispersion, even at high group index, is generally regarded as an advantage of slow light waveguides with respect to optical cavities [50], which are also

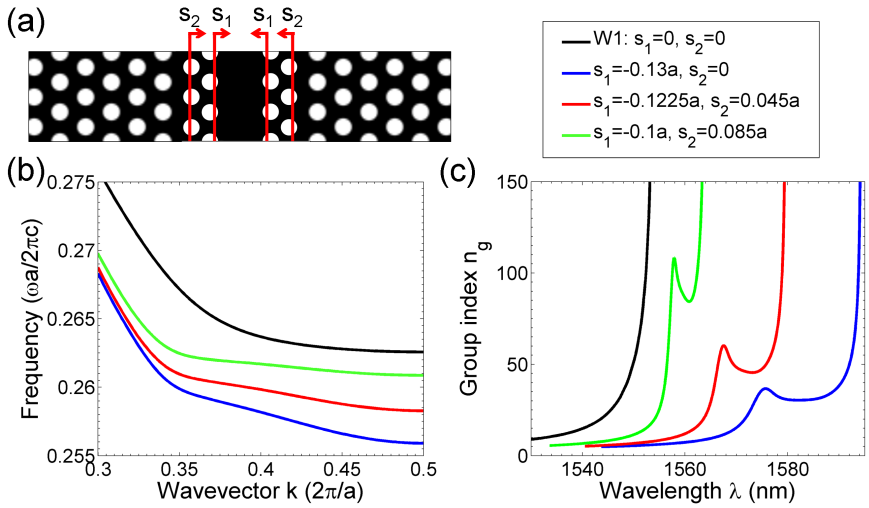

Figure 2.6: Engineered slow light waveguides as proposed in Ref. [28], with hole radius $r = 0.286a$. (a) Schematic design of the engineered W1 waveguide: the first and second rows of holes are displaced symmetrically about the waveguide axis. The displacements relative to the unmodified lattice (red lines) are expressed by s_1 and s_2, where shifts toward the waveguide centre are defined to be positive. (b) Dispersion curves and (c) group indices for the fundamental mode of engineered W1 waveguides. The mode of a W1 waveguide (black) is also presented for comparison.

extensively studied for enhancing optical fields: a good bandwidth allows for processing of shorter signals, and makes the device less sensitive to temperature and fabrication fluctuations, which is a great challenge for optical cavities. The available bandwidth, however, does decrease for higher group indices, which is clear from Figs. 2.6b-c. This delay-bandwidth product limitation is unavoidable in static devices, but in principle it may be overcome by adiabatically tuning a photonic crystal structure after insertion of a light pulse [70].

Propagation loss is also an important aspect that cannot be neglected when discussing slow light devices, as all the exciting advantages of slow light waveguides become pointless if only a tiny fraction of light is transmitted through the structure. Photonic crystal waveguides are intrinsically lossless, because the modes of interest lie inside the band gap and below the light line. All the loss in practical devices is due to fabrication imperfections, such as side-wall roughness and disorder in the position and size of the holes, which have also a direct influence on the very dispersion properties of the waveguide, especially in the slow light region [71]. The effect of disorder is amplified by the high index contrast between the silicon and the air cladding, leading to out-of-plane loss and backscattering loss [31, 72]. The former arises from scattering of light into radiation modes above the light line and scales as n_g [72, 73]; the latter is due to coupling into the

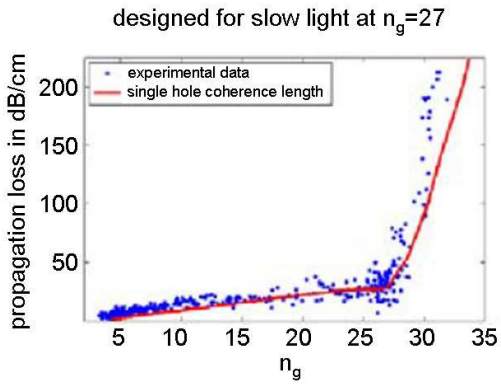

Figure 2.7: Propagation loss as a function of group index for a waveguide designed to exhibit a constant group index of $n_g \sim 27$ over a $\Delta\lambda \sim 10$ nm wavelength range ($a = 410$ nm, $r = 0.27a$, $s_1 = -52$ nm, $s_2 = 0$ (from [31]. Reprinted with permission).

backward-propagating mode and scales as n_g^2 [29, 74].

Now, as the state of the art fabrication technology already yields very good quality photonic crystal devices, with a record 2 dB/cm loss demonstrated the fast regime back in 2007 [75], great debate was generated in recent years on whether the quadratic dependence of backscattering on group index represents a major limitation to the usefulness of slow light. It is now well understood that the mode shape variation with group index (Figs. 2.4b-e) plays a significant role in determining the loss [31]: both out-of-plane and backscattering loss contributions are weighted by terms containing the mode shape, as this determines how much light is available for scattering on the side-wall of the holes. Engineered slow light waveguides of the type discussed in Ref. [28] (Fig. 2.6) turn out to successfully suppress the backscattering component and exhibit a linear dependence of loss on group index up to the region of constant n_g [31], as shown in Fig. 2.7. After this point, however, backscattering and multiple scattering become dominant [71], the loss rapidly increases and it is not advantageous to further slow down light. For this reason, we will always operate in the linear regime, which does not usually extend to very high n_g (a linear dependence up to $n_g \sim 60$ has been recently demonstrated in Ref. [76]).

Finally, the shape of the mode profile is to be taken into account also in determining the coupling efficiency of light into a slow mode: in the fast index-guided regime, the field profile is similar to that of a nanowire (Fig. 2.4b), and therefore easy to excite, but in the slow light regime the field profile is different and penetrates beyond the first row of holes (Figs. 2.4d-e), causing mode mismatch with that of an access ridge waveguide. This mismatch can be partially addressed by playing on the geometry of the photonic crystal waveguide termination [78]. More importantly, however, the coupling

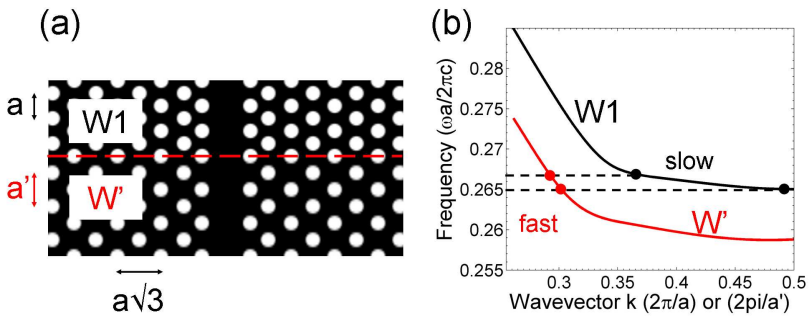

Figure 2.8: Interface between two photonic crystal waveguides supporting slow (W1) and fast (W′) Bloch modes at the same frequencies [77]. (a) Schematic design of the interface: the period of W′ is elongated along the propagation direction, with $a' > a$. (b) Dispersion diagram of the fundamental guided Bloch modes of the waveguides.

issues in the slow light regime are due to the fact that the slow mode is the result of the interference between a forward- and a backward-propagating wave, but at the boundary with a ridge waveguide only the forward component is excited [50]. An interface or transition region is therefore needed for the slow mode to be successfully excited. This may be achieved by introducing a taper to adiabatically change the group index [79, 80] or, more simply, by introducing a short section of a fast photonic crystal waveguide, where coupling is mediated by the evanescent components of the field at the interface with the slow light waveguide [81]. Such a fast section may be created by locally lowering the guided mode in the band diagram, for example by reducing the hole size [82] or by increasing the period (by ~7.2 percent) in the propagation direction [77] (Fig. 2.8). The latter approach results in very short (few periods long) fast couplers for efficient injection at moderately high group indices ($n_g < 100$) [77], and it is the solution generally adopted in our research group and in the work of this thesis.

To conclude, when compared to standard W1 waveguides, dispersion engineered structures not only can display a region of high group index with low dispersion, but they also exhibit lower propagation loss than a W1 for the same group index [31], which has allowed for useful application of slow light in the experimental demonstration of optical modulation [55], delay [83], wavelength conversion [57], self-phase modulation [62, 66], third-harmonic generation [59] and four-wave mixing [60, 61]. However, the considerations above on propagation loss, coupling efficiency, and even the dispersion engineering method itself, warn us of the importance of the mode shape in determining many key features of slow light photonic crystal waveguides, which should not be neglected when attempting to quantify the benefit of slowing down on nonlinear effects, as we shall discuss more in detail for

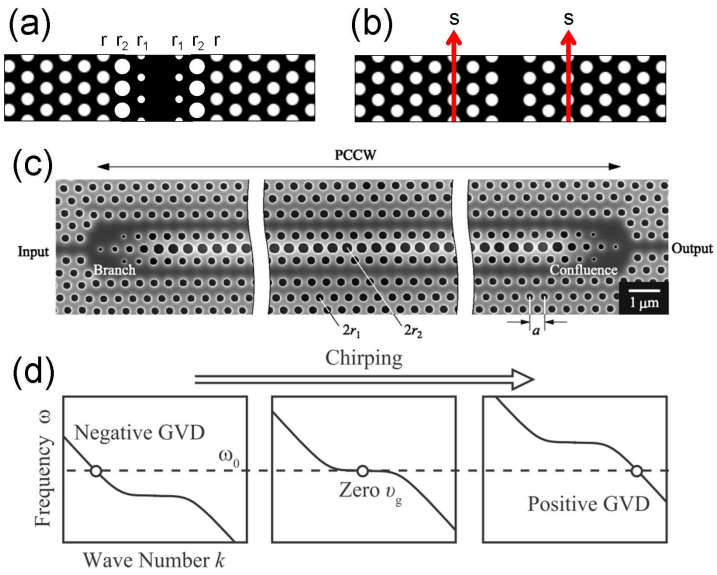

Figure 2.9: (a-b) Schematics of zero-dispersion slow light waveguides obtained (a) by varying the size of the first two rows of holes adjacent the waveguide (as in Ref. [65]), and (b) by shifting the third row of holes along the propagation direction (as in Ref. [66]). (c) SEM image of a dispersion-compensated slow light structure consisting of two coupled waveguides (from [85] © 2010 IEEE). (d) Schematic of the dispersion curve of a photonic crystal coupled waveguide exhibiting an inflection point: the curve is gradually chirped towards higher frequencies due to chirping of the hole radius; the dashed line indicated the frequency of the incident light (from [86]. Reprinted with permission).

Raman scattering in Chapter 5.

2.2.1 Other examples of slow light photonic crystals

The type of slow light waveguides described in the previous section is only one example of what are also called zero-dispersion waveguides: the dispersion of the basic W1 waveguide is flattened by playing on the coupling between two modes, which is varied simply by introducing small modifications to the first rows of holes. Early examples of this type of approach involved narrowing the waveguide width (W0.x) to improve the dispersion [69, 84] but, in real devices, narrower waveguides are usually associated with higher scattering loss [74], as the optical mode has greater overlap with the air holes.

As already mentioned, the most common approaches to flatten the mode dispersion rely on the alteration of the first rows of holes by changing their radius [65,87] (Fig. 2.9a) or position, either along the normal [28] (Fig. 2.6a) or the longitudinal [66,88] (Fig. 2.9b) direction. These three mechanisms all achieve similar delay-bandwidth products [89], but for practical applications

the holes position can be controlled more accurately [30] than the radius [90], because the hole size is influenced by many parameters during the fabrication process.

Other examples, mainly numerical, of zero-dispersion techniques that have been proposed in the literature involve introducing high-index pillars inside the holes [91], or infiltrating them with liquid [92] or other dielectric material [93]. Some authors proposed breaking the symmetry of the waveguide [94, 95] so that the dispersion can be controlled through an anti-crossing of the even mode with the odd mode. Other proposals involve altering the basic lattice itself, for instance by using elliptical [96] or even ring-shaped [97] holes, or by changing the lattice from triangular to oblique [98].

A second type of approach to suppress higher-order dispersion in slow light devices is that of dispersion-compensated structures: instead of designing the waveguide with a flattened dispersion curve, a dispersive slow light waveguide is followed by one with opposite dispersion, so that the two compensate each other. The transition from one waveguide type to the other is not abrupt but smooth, and it is obtained by chirping one of the design parameters — usually the hole radius — along the device, resulting in a gradual shift of the band diagram [99, 100]. Dispersion-compensated devices usually do not consist of a single waveguide, but of two parallel waveguides, coupled to each other. In some early studies, the two parallel waveguides were designed to exhibit opposite dispersion [101–103], but one important requirement was for the waveguides to match exactly their band edges, making the final device very sensitive to fabrication tolerances.

A more established design makes use of two identical coupled waveguides (Fig. 2.9c): the resulting even super-mode exhibits an inflection point in the dispersion curve, providing slow light [104] (Fig. 2.9d, left). This method of dispersion engineering is again based on the interaction of two modes, but in this case the two original modes belong to two separate waveguides. By chirping of the hole radius along the propagation direction, the inflection point is gradually shifted in frequency (Fig. 2.9d, middle and right), broadening the slow light condition [67,86], and different wavelengths are therefore slowed down in different regions of the device. This type of structure has been very successfully employed to demonstrate tunable slow light devices by introducing an additional thermally-controlled variable chirp [85, 105, 106], as we will discuss in more detail in Chapter 3.

The design of chirped photonic crystal coupled waveguides, however, is more complicated than for the simple zero-dispersion structures, as it requires optimisation of many parameters. In addition, while this approach is very suitable for delay devices, it cannot be employed for enhancement of nonlinear effects: the structure itself is bigger than a single W1 waveguide and, due to the initial dispersion, which is then later recovered, an optical pulse is broadened, rather than compressed, in space. In fact, suppression of optical nonlinearities has been highlighted as a positive feature of these

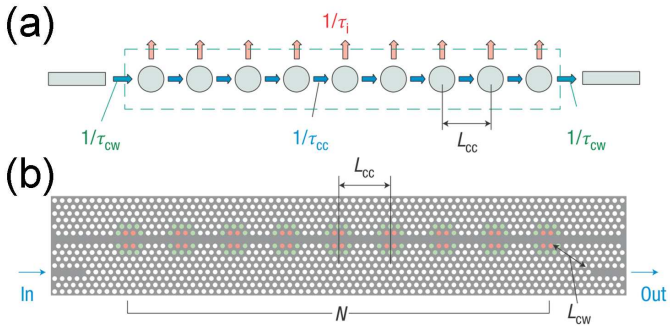

Figure 2.10: Examples of CROWs (from [107]). (a) Schematic of a generic CROW. (b) Design of a CROW based on width-modulated cavities. (Reprinted by permission from Macmillan Publishers Ltd: Nature Photonics [107], copyright 2008.)

devices [25], which is desirable when employing slow light structures for linear applications.

Finally, the third type of approach that has been proposed for achieving structural slow light in photonic crystals is the use of a sequence of coupled cavities (coupled resonators optical waveguide, CROW, Fig. 2.10): as light propagates through the structure, it is localised at each cavity, introducing a time delay [108]. Numerous works have reported studies on the dispersion of these slow light waveguides, both theoretical/numerical [109–112] and experimental [113–116]. The most impressive demonstrations of CROWs in photonic crystals are probably those reported by Notomi and co-workers (Fig. 2.10b), who realised large arrays of up to 100 [107] and 200 [117] ultrahigh-Q cavities, which were also employed to demonstrate enhancement of nonlinear effects such as four-wave mixing [117].

One issue of CROWs in photonic crystals, however, is that the regular spacing of the cavities introduces a second periodicity on the lattice, corresponding to an additional folding of the bands in the k-space, and as a consequence light can easily couple to the radiation modes above the light line [25]. More importantly, the precise control of both the inter-cavity coupling and the optical resonances of each cavity (which should be identical) at the same time is a huge challenge in photonic crystals, once again due to disorder and non-idealities introduced during the fabrication process [118]: even in some of the best reported results [107], oscillations in the transmission band on the order of 20–30 dB are unavoidable. This type of slow light structures is therefore probably more suited for use with ring resonators [119–122], where the control of individual resonances is more straightforward and reliable [123] than in photonic crystals.

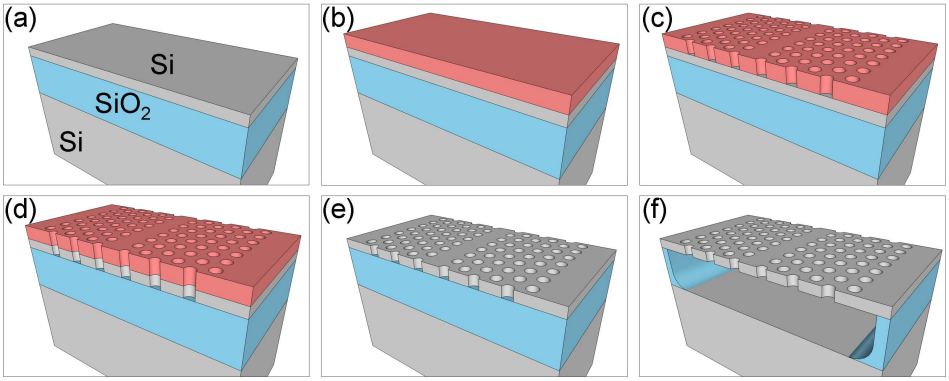

Figure 2.11: Schematic of the standard fabrication steps for photonic crystal waveguides. (a) The SOI substrate. (b) The e-beam resist ZEP is spun on the sample. (c) The photonic crystal pattern is defined on the ZEP layer through e-beam writing. (d) The pattern is transferred from the ZEP mask to the silicon layer through the dry-etch process RIE. (e) The ZEP mask is removed. (f) The silica layer is removed by wet-etch in HF solution, releasing the silicon photonic crystal membrane.

2.3 Standard fabrication of photonic crystals

The realisation of the sub-micrometric features of photonic crystals requires high-precision processing techniques. The standard fabrication process of a photonic crystal waveguide, shown schematically in Fig. 2.11, consists of three fundamental steps: first, a resist mask is defined on the chip by means of electron-beam (e-beam) lithography; the mask pattern is then transferred to the silicon through a dry-etching process; finally, the photonic crystal membrane is released by wet-etch. These steps, which are described in more detail in the following, define a fabrication recipe for silicon photonic crystals that is well established within the research group [124]. Any modification or additional step that was especially developed for the work of this thesis is described where appropriate in the following chapters.

Our samples are fabricated on SOITEC silicon-on-insulator (SOI) wafer (Fig. 2.11a), consisting of a 220 nm-thick silicon layer on top of a 2 μm-thick silica layer on silicon substrate. After sample cleaning, a ~350 nm-thick layer of ZEP-520A resist is spun and baked 10 min at 180°C (Fig. 2.11b). ZEP-520A is a high-resolution e-beam resist, used to realise the soft mask in which the photonic crystal pattern is first defined. It is a positive resist, meaning that the features to be removed (i.e. the holes) are the ones to be exposed.

The exposure takes place in a modified LEO Gemini 1530 SEM with Raith Elphy Plus pattern generator, at the maximum acceleration voltage setting of 30 kV. In the design pattern the hole radius is reduced by 20 nm with respect to the target radius, to correct for broadening of the electron beam diameter due to forward scattering of electrons throughout the re-

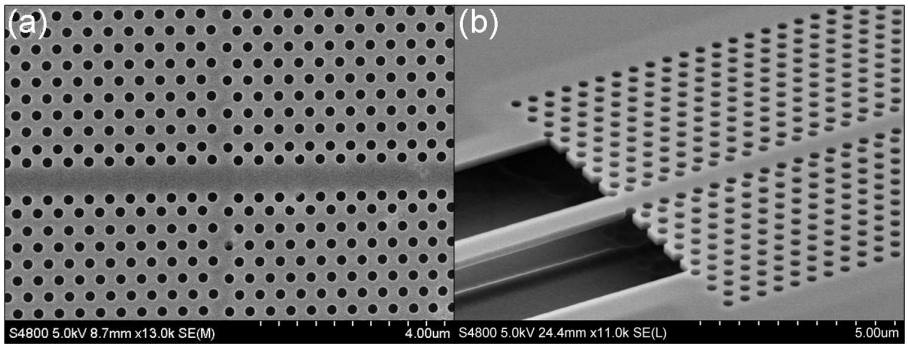

Figure 2.12: SEM images of silicon photonic crystal waveguides. (a) Example of a particularly bad stitching error; in standard devices such imperfections are not so obvious under SEM inspection. (b) A membraned photonic crystal.

sist thickness. Another non-ideality, known as the proximity effect, occurs when beam electrons penetrate the silicon underlying the resist and undergo backscattering, therefore returning into the resist and enlarging the exposed region; as a consequence, the energy received by one feature of the pattern is affected by electrons scattering from other features nearby. This effect is taken into account by running a correction software on the photonic crystal pattern. Finally, when the pattern area is bigger than the scanning area of the electron beam, known as the writefield (100 μm $\times$ 100 μm used in this work), the pattern is written through sequential exposure of adjacent areas by shifting step-by-step the sample stage. Since these mechanical shifts have finite resolution, positioning errors, also called stitching errors, may occur (Fig. 2.12a). If the photonic crystal is shorter than the writefield size, it will not be affected by stitching.

After exposure, the resist is developed in Xylene at 23°C for 45 s, then rinsed in IPA. The developer removes the exposed areas, leaving the photonic crystal pattern defined in the ZEP mask (Fig. 2.11c).

The pattern is then transferred from the resist on the underlying silicon layer by means of reactive ion etching (RIE), a dry-etch process in which ions and free radicals generated in a radio frequency (RF) plasma remove the exposed areas of silicon, while the remaining areas are protected by the ZEP mask. The etching is the combination of the physical damage of the accelerated ions (which is anisotropic) with the chemical action of the reactive neutral radicals (which is isotropic). Our RIE process is performed with a SF_6:CHF_3 1:1 gas mix at a pressure of 5×10^{-2} mBar; the plasma is sustained for 2 min at an RF power of 19 W. As a result, the silicon layer not covered by ZEP is completely etched (Fig. 2.11d). The ZEP layer is then removed with solvents, leaving the silicon photonic crystal visible (Fig. 2.11e).

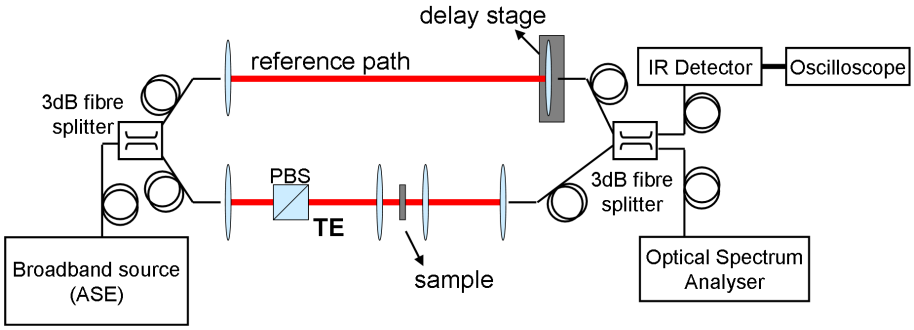

Figure 2.13: Schematic of the free-space Mach-Zehnder interferometric setup used to measure group index and transmission curves in slow light photonic crystal waveguides.

As a last step, the oxide layer beneath the photonic crystal is selectively removed by soaking the sample in dilute hydrofluoric (HF) acid (1:10 with water) for 15 to 17 min, releasing a silicon membrane (Fig. 2.11f and Fig. 2.12b). To protect the access waveguides from the HF, a 1.8 μm-thick layer of S1818 resist (Shipley) is spun on the sample prior to the HF step, and windows are opened above the photonic crystal regions by UV exposure and development in MF319 for 40 s. The S1818 mask is then removed with acetone and the sample cleaned in piranha solution (H_2SO_4:H_2O_2 3:1) for 5 min. Finally, the sample is cleaved manually to a total length of a few millimetres.

2.4 Characterisation: transmission and group index curves

The measurement of group index curves in photonic crystal waveguides is performed using a Fourier transform spectral interferometry setup, which at the same time can be used for transmission measurements. A schematic of our free-space end-fire coupling setup is shown in Fig. 2.13. Just as the fabrication steps in section 2.3, the following briefly describes a measurement technique that is well established within the research group [125].

The setup consists of a Mach-Zehnder interferometer (MZI): light in the ~1550 nm wavelength range from an amplified spontaneous emission (ASE) broadband source is separated into a sample arm and a reference arm, through a 3 dB fibre splitter. In the sample arm, light is TE-polarised by a polarisation beam splitter (PBS) and coupled in and out of the sample by means of two aspheric lenses mounted onto stages with three translational degrees of freedom. The output light is then coupled into a second fibre splitter, where it recombines with that coming from the reference arm. One of the outputs of the fibre splitter feeds light into an infrared (IR) de-

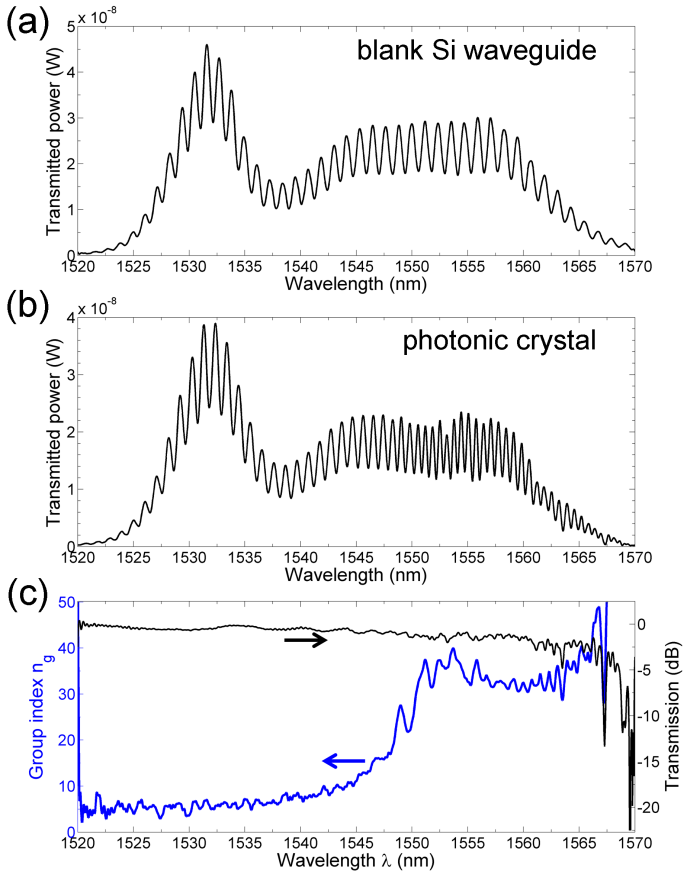

Figure 2.14: (a-b) Spectra of the MZI output measured for (a) a blank waveguide and (b) a photonic crystal waveguide on the same chip. (c) Group index (blue) and transmission (black) of the photonic crystal waveguide. The transmission is normalised to the blank waveguide. The interferograms of (a-b) and the transmission curve of (c) have been filtered numerically to remove Fabry-Perot fringes.

tector connected to an oscilloscope and is used to optimise the optical setup alignment, whereas the other output is connected to an optical spectrum analyser (OSA).

The wavelength spectrum measured by the OSA exhibits fringes (Figs. 2.14a-b) due to constructive and destructive interference between light coming from the two arms of the interferometer, with a fringe period inversely proportional to the difference in their optical path lengths. This means that the fringe spacing in the interferogram is constant if the sample waveguide has no dispersion, for example in the case of a blank strip waveguide (Fig. 2.14a), but they converge when entering the slow light regime of a photonic crystal waveguide (Fig. 2.14b), because the optical path of the sample arm becomes progressively longer.

Here we briefly describe how the group index n_g is extracted from the measured interferogram, but more details can be found in Ref. [125]. The output interferogram can be expresses as

$$I(\omega) = S(\omega) + R(\omega) + \sqrt{S(\omega)\,R(\omega)}\left[e^{i\Phi(\omega)-i\omega\tau} + \text{c.c.}\right],$$

where $S(\omega)$ and $R(\omega)$ are the spectral densities of the sample and reference outputs, respectively, which are measured separately by blocking one arm or the other. The sample output $S(\omega)$ can also be used to calculate the waveguide transmission (Fig. 2.14c, black curve), by normalising it to the same output of a blank waveguide on the same chip, or to the fibre-to-fibre spectrum of the sample arm. The delay τ is set by a delay stage in the reference arm and is initially chosen to give about 1 fringe/nm or less in the interferogram of a blank waveguide, and in such a way that the reference arm is always the shortest: thus the slow light always increases the difference in path length and always results in denser fringes. Once set, τ is kept constant for all the measurements on the same sample.

The information on the dispersion is contained in the phase difference term $\Phi(\omega) = \Phi_S(\omega) - \Phi_R(\omega)$, which is extracted from $I(\omega)$ by subtracting the non-interfering terms $S(\omega)$ and $R(\omega)$, then Fourier transforming and filtering one of the crossed terms, and transforming it back to frequency domain. Finally, the term $\Phi(\omega) - \omega\tau$ is differentiated with respect to ω to obtain the difference in group delay $\Delta\tau_g$ between the two arms, which is related to the group index n_g as $n_g = \Delta\tau_g c/L$, with L the length of the slow light region, i.e. the length of the photonic crystal without the fast coupling sections. In the calculation, the delay contributions from the access waveguides and every other element of the optical setup are accounted for by subtracting the data from a calibration measurement taken on the blank waveguide, which sets the zero delay. An example of measured group index curve is shown in Fig. 2.14c, blue curve.

Chapter 3

Tunable optical delay in slow light waveguides

Controlling the timing of optical signal pulses represents an important step towards the realisation of all-optical signal processing systems, with a variety of potential applications requiring time domain operation, including accurate synchronisation for bit-level interleaving and switching, time division multiplexing, optical calculation of the convolution integral and even optical buffering [11, 12]. Many different approaches for achieving tunable optical delay have been proposed in the literature, involving both linear and nonlinear effects in optical fibres, photonic crystals and coupled cavities systems. Each configuration has its own benefits and drawbacks in terms of maximum achievable delay, operating wavelength and speed, pulse distortion and complexity.

In this chapter we will review some of the main results reported to-date in the area of tunable optical delay lines and propose a new design based on wavelength conversion and dispersion in engineered slow light photonic crystal waveguides. We will then apply such design in an experimental demonstration of continuous fine tuning of the output timing of a signal, and demonstrate the ultrafast temporal manipulation of a single bit in a train of two pulses.

3.1 State of the art

A tunable optical delay line is a device in which the output time of an incoming optical data signal may be tuned between a minimum t_1 and a maximum t_2, depending on an external control signal. The generic schematic of a black-box tunable delay line is shown in Fig 3.1.

The performance of the device is first of all characterised by the maximum tunable delay, or storage time, that can be applied to a signal: this is indicated by Δt in Fig. 3.1, and it refers to a change in propagation time of

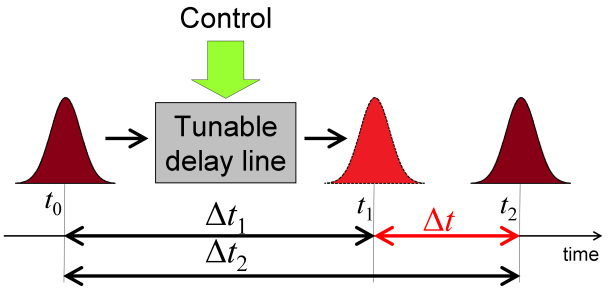

Figure 3.1: Schematic of a generic tunable delay line.

the data signal through the delay line due to a change in the control input. Depending on the application, one may be more interested in the absolute value of Δt or on the delay referred to the signal pulse duration $\Delta t/T_0$, a delay expressed as number of bits. Another important feature is the available bandwidth with low second-order dispersion, which determines the shortest pulse duration that may be transmitted with low distortion, and is therefore related to the sustainable bit rate. Finally, a tunable delay line might also be characterised by a hold-off time, a time interval after loading data during which the device cannot accept any new incoming signal. The hold-off time also limits the rate at which bits/packets may be loaded into the device.

The simplest way to generate variable optical delay is to make the signal propagate through paths of different lengths. Such a scheme has been demonstrated using optical fibres combined with optical switches to achieve buffering capability in some experiments [13, 126], but it can generate only a finite set of *discrete* time delays. Investigation about the development of optical delay lines that are *continuously* tunable has been boosted by the great research effort undertaken in the past decade in the area of slow light [12]. Many demonstrations have been reported of tunable delay resulting from direct modulation of the group velocity, relying for instance on electromagnetically-induced transparency in atomic media [127], optical parametric amplification [128] and stimulated Brillouin [129, 130] and Raman [131] scattering in optical fibres. Fibre-based tunable delay lines, however, are extremely bulky and not suitable for integration of optoelectronic devices.

Tunable delay has been demonstrated using stimulated scattering also in SOI waveguides [132] and semiconductor optical amplifiers [133, 134], but here we will mainly focus on examples of miniaturised devices in which slow light is achieved through a structural resonance, rather than a material resonance. As we have introduced in Chapter 2 for the case of photonic crystal waveguides, engineered devices offer great controllability over the group velocity of light, but static structures simply reduce the group velocity by a fixed amount, resulting in a fixed delay. In order to realise a variable

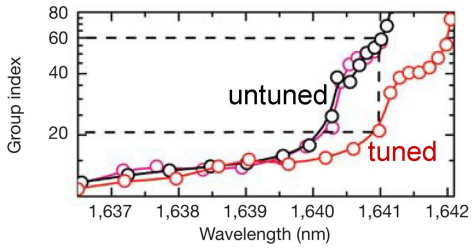

Figure 3.2: Tuning of a photonic crystal group index dispersion curve (in semi-logarithmic scale), as reported by Vlasov *et al.* (from [53]. Reprinted by permission from Macmillan Publishers Ltd: Nature [53], copyright 2005).

delay line we need to tune the material parameters with an external control.

Integrated tunable delay lines may be divided in two main categories [12]: those in which the group velocity remains constant while signals propagate through them, and those in which the group velocity may be dynamically changed in time. In the following, we will review some of the main examples reported in the literature for each of the two types. We will then discuss a third category of solutions, which falls in between the first two, based on wavelength conversion and dispersion.

3.1.1 Time-independent tunable delay lines

With the expression time-independent tunable delay lines we refer to those structures whose parameters do not vary during the time interval in which an optical signal propagates into them. The optical delay is controlled by varying these parameters when no signal is present, or the tuning mechanism is much slower than the bit propagation time.

In a photonic crystal waveguide the simplest way to vary the group velocity is to change the refractive index of the slab material. This causes the whole band diagram to shift, due to the scaling properties of photonic crystals [45], the relative frequency shift $\Delta\omega/\omega$ being equal to the relative refractive index change $\Delta n/n$. Therefore, signal frequencies corresponding to slow modes before the tuning, may correspond to fast modes after the band diagram shift. Such a tuning of the group velocity was experimentally demonstrated by Vlasov *et al.* [53] in a silicon W1 photonic crystal waveguide, in which the refractive index of silicon was varied by the thermo-optic effect. This resulted in tuning of the group index from 20 to 60 (Fig. 3.2), demonstrating the possibility of actively influencing slow light.

Experimental observation of true time shifts of optical signal pulses in an engineered photonic crystal was first reported by Baba *et al.* [105]. The proposed structure was a silicon photonic crystal coupled waveguide [67] (Fig. 2.9c), where chirping of the hole radius along the propagation direction allowed for high group index to be achieved (Fig. 2.9d) over a bandwidth

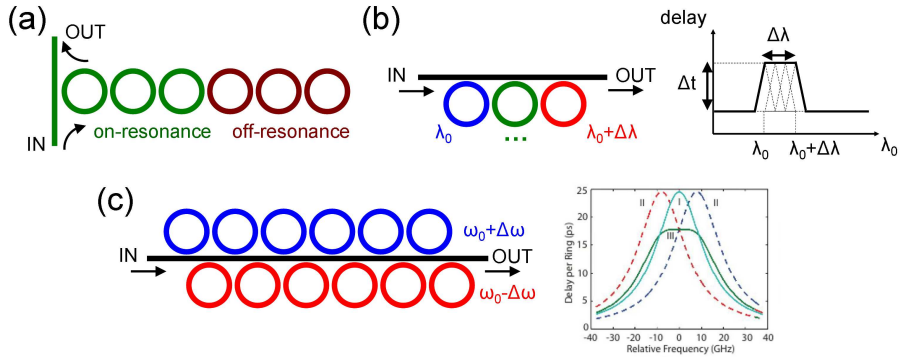

Figure 3.3: (a) Schematic of the reconfigurable reflective CROW delay line reported by Melloni *et al.* [16]. (b) Schematic of an all-pass type microring array (left) and the delay envelope spectrum (right) resulting from detuning of the rings' resonant wavelengths [135]. (c) Schematic of a balanced SCISSOR tunable delay line [136] (left) and delay spectra (right, from [137]. Reprinted with permission) for: (I) all rings aligned, (II) red- and blue-shifted rings, (III) complete device. In all the types of structures, the resonance of the rings is controlled thermally by integration of heaters.

wide enough to accommodate ~3.3 ps-long pulses. By illuminating a 250 μm-long sample with a laser spot, a refractive index profile was created along the waveguide, therefore adding also a thermally variable chirp to the fixed chirp determined by the radius. This allowed for tuning of the delay up to 23 ps (~7 pulse-widths). This figure was improved in more recent works by the same authors [85,106], demonstrating a 72 ps tunable delay for 2 ps-long pulses in a 400 μm-long waveguide [106], which was then applied as a delay scanner in an optical correlator with a 2 kHz scanning frequency.

A number of other experimental demonstrations of on-chip tunable delay are based on ring resonators, arranged in various configurations and realised in different materials. Melloni *et al.* [16] employed a 6-rings silicon oxynitride coupled resonator optical waveguide (CROW) in reflective configuration (Fig. 3.3a) and achieved continuously tunable delay of up to 8 bits (320 ps) at 25 Gb/s. Higher bandwidths and greater fabrication tolerance can be achieved adopting an all-pass filter configuration [135,138–140] (Fig. 3.3b). Tunable delays of 7.3 bits [135] and 8.2 bits [140] have been demonstrated at 40 Gb/s using this configuration. One disadvantage is the high number of rings that need to be tuned simultaneously (50 in Ref. [135] and 32 in Refs. [139,140]), whereas in the reflective CROW [16] only few of rings need to be tuned at any one time. A variation of the all-pass filter configuration is the balanced side-coupled integrated spaced sequence of resonators (SCISSOR) [136] (Fig. 3.3c), a promising concept offering good bandwidth while requiring a reasonable number of rings, but to-date less than 2 bits delay was experimentally demonstrated at 10 Gb/s [137]. Schemes based on

the use of one ring only have also been proposed [141, 142], but suffer from very low bandwidth and high losses.

All of these devices (except for [142]) are based on thermo-optic tuning, which is too slow to be useful in many applications [143]: the tuning time reported by Melloni *et al.* [16] is about 150 μs, whereas Vlasov *et al.* [53] report tuning times of the order of hundreds of nanoseconds. Even if this is remarkable for the thermo-optic effect in silicon, much faster tuning may be achieved by the free-carrier plasma effect [143], that is, the change in refractive index associated with a change in the carrier concentration [144], due to either optical or electrical pumping. However, if the whole structure or most of it is to be tuned, as in many of the reviewed works, including the photonic crystal examples by Vlasov *et al.* [53] and Baba *et al.* [105], a high carrier concentration would dramatically increase also propagation loss due to free-carrier absorption [144]. In summary, the ideal tunable device would be realised in silicon to allow for compactness and ease of integration, and it would rely on a fast tuning mechanism such as the free-carrier plasma effect; the tuning region should also be as small as possible, to reduce power consumption and minimise optical loss.

Finally, time-independent tunable delay lines in general are limited by the delay-bandwidth product [12,145,146]: the larger the group delay achievable through an optical resonance, the narrower the bandwidth within which the delay occurs. Thus, for a fixed device length, the delay can be increased only by decreasing the bandwidth — and thus the bit rate — and vice-versa. The delay-bandwidth product constraint can be overcome by use of dynamic structures.

3.1.2 Time-dependent tunable delay lines

We will now consider dynamic tunable delay lines, structures whose properties are modulated *while* an optical signal is stored inside them. Dynamic structures are not fundamentally limited by a delay-bandwidth product [12], because they rely on the adiabatic modulation of the light spectrum.

Extensive numerical investigation of adiabatic modulation in photonic crystals was performed by Yanik and Fan [15, 70, 147, 148]. The idea is to start with a system whose initial bandwidth is large enough to accommodate an incident pulse (Fig. 3.4a, left). Once the pulse has entered the photonic crystal, the refractive index is tuned uniformly and adiabatically: a uniform modulation, preserving periodicity, conserves the k components, and so *the system* can be described in terms of a time-varying band structure; if the tuning is adiabatic, no scattering of light to other bands occurs, and the frequency of each k component of *the signal pulse* itself follows the band diagram. By dynamically tilting the dispersion curve, the pulse can be spectrally compressed (Fig. 3.4a, right) and its group velocity greatly reduced, overcoming the delay-bandwidth product limit. The reverse releases

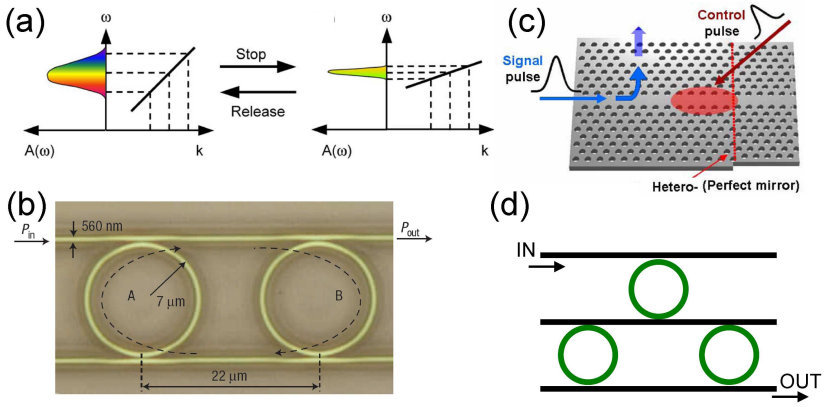

Figure 3.4: (a) General condition for stopping a light pulse: adiabatic transition from a large-bandwidth state used to accommodate an incident light pulse (left), to a narrow-bandwidth state used to hold the pulse (right) (from [15] © 2006 IEEE). (b) The device fabricated by Xu *et al.*, consisting of two ring resonators coupled to an access waveguide (from [17]. Reprinted by permission from Macmillan Publishers Ltd: Nature Physics [17], copyright 2007). (c) Schematic of the structure proposed by Upham *et al.*, consisting of a photonic crystal cavity coupled to an access waveguide, and a mirror (from [18]. Reprinted with permission). (d) Schematic of the three-rings device demonstrated by Elshaari *et al.* [150].

the pulse again.

In practice, however, the achievable relative refractive index shift in silicon is very small, of the order of 10^{-3}–10^{-4} [143], and therefore one needs to design structures in which the spectral feature is sensitive to these small variations. For example, the structure investigated by Fan *et al.* consists of a waveguide laterally coupled on both sides to two kinds of cavities, whereas Khurgin proposed the use of a CROW, in which the slope of the dispersion curve is changed through adiabatic zone folding [14]. A third example is given by engineered slow light photonic crystal waveguides: very recently Beggs *et al.* [149] applied the dispersion-tilting concept to demonstrate experimentally the adiabatic spectral compression of a light pulse.

One issue of a tunable delay line based on the adiabatic change of the dispersion curve slope is the poor storage density [12, 14], due to the fact that an optical signal entering the line at relatively high group velocity is not compressed in space when slowed down, which is instead true for static slow light devices [24]. In addition, fabrication of a great number of cavities — up to 100 — as proposed by Fan requires precise alignment and it is not straightforward to achieve in practice [107].

The adiabatic tuning principle, however, has been applied in some experimental demonstrations overcoming the traditional delay-bandwidth product through trapping and releasing of light pulses in only few resonators. In such schemes, the system is initially kept in a low-Q condition favouring in-

coupling of light, and after loading of a signal pulse it is switched to a high-Q condition through ultrafast (that is, within the cavity photon lifetime) refractive index tuning due to free-carriers injection. This functionality is usually achieved by breaking the coupling to an access waveguide, and the signal pulse can be trapped for a time interval longer than the lifetime of the low-Q static system, therefore breaking the delay-bandwidth product limit. A second ultrafast tuning event restores the system in a low-Q condition, releasing the trapped pulse.

Such an approach was first demonstrated by Xu *et al.* [17] in a compact structure consisting of two ring resonators (Fig. 3.4b). Tuning was accomplished by injecting free carriers in the same region where light was stored, inducing very high loss and limiting the maximum achievable delay to less than 100 ps for a 14 ps pulse. Two very elegant alternatives were proposed, minimising the overlap of the tuning region with the storage region. The first solution was implemented by Upham *et al.* [18, 151] in a photonic crystal cavity-waveguide-mirror system (Fig. 3.4c) whose Q-factor depends on whether constructive or destructive interference occurs between light outputting the cavity and light reflected back from the mirror. The phase difference between the two beam paths is controlled by changing the refractive index of silicon in the waveguide section between the cavity and the mirror. So the overall Q-factor can be controlled without acting directly on the cavity, where most of the light is stored. The second solution was reported by Elshaari *et al.* [150] in a ring resonator system similar to that of Xu *et al.*, but with the addition of a third ring to act as an input gate (Fig. 3.4d), which is to be tuned to trap light in the other two rings.

Note, however, that in such trap-and-release schemes once a pulse has been loaded no other pulse can be accepted until the release of the first one. Therefore, the maximum storage time defining the performance of the device represents also its intrinsic limitation on the hold-off time. Furthermore, it is still challenging to experimentally achieve a good balance between the initial low-Q condition (bandwidth) for pulse loading and the high-Q condition for pulse storage [150, 151].

3.2 Tunable delay based on wavelength conversion and dispersion

An alternative approach for realising tunable delays is based on wavelength conversion and group velocity dispersion [19]: the input signal is shifted in frequency and then injected into a dispersive medium, where different frequencies propagate at different speed. The shifted signal will therefore accumulate a delay with respect to an unshifted one. The achievable delay is proportional to the wavelength shift and the group velocity dispersion and length of the dispersive medium [152]. The signal wavelength can then be

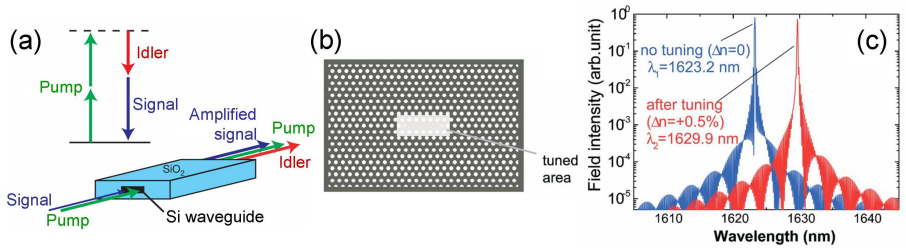

Figure 3.5: Examples of wavelength conversion processes in silicon. (a) Schematic of four-wave mixing: two pump photons are converted to a signal photon and an idler photon (from [22]. Reprinted by permission from Macmillan Publishers Ltd: Nature [22], copyright 2006). (b-c) Adiabatic wavelength conversion in a photonic crystal cavity (from [23]. Copyright 2006 by The American Physical Society) consisting of four missing air holes in a triangular lattice (b); the refractive index of the gray region is tuned. (c) FDTD calculations of wavelength spectra before (blue) and after (red) index tuning.

reconverted to its original value in a subsequent wavelength conversion step.

Several demonstrations of continuously tunable delay based on this technique have been reported, mainly using optical fibres as a dispersive medium, allowing for very high optical delays — up to few microseconds [153–155] — to be observed. Despite the large achievable delays, however, these solutions are extremely bulky, and not suitable for device integration. We may therefore investigate the possibility of realising a tunable delay line based on wavelength conversion and dispersion in more compact photonic structures such as silicon photonic crystals.

3.2.1 Wavelength conversion in silicon waveguides

Wavelength conversion in silicon may be achieved with various effects. One important example is four-wave mixing [156], a third-order nonlinear process in which two pump photons are converted into a signal and idler photons (Fig. 3.5a). The information carried by the signal can therefore be transferred also to the idler. Efficient conversion over a wavelength range of several tens of nanometres in a SOI channel waveguide was reported by Foster *et al.* [22]. Conversion efficiency and tunability range, however, come at a price: the need for very powerful pumps (11 W peak power in [22]), a minimum light travelling distance (6 to 17 mm in [22]), and phase-matching. At more realistic (continuous-wave) pump powers of around 100 mW, even in low-loss silicon nanowires the best conversion efficiencies range from −40 dB [61] to less than −10 dB [157], depending also on the waveguide length (almost 2 cm in Ref. [157]), whereas the best reported value to-date for slow light photonic crystals is −24 dB at $n_g \sim 30$ [61].

Different authors have therefore explored the possibility of realising wavelength conversion i) with high efficiency, ii) for weak light and iii) in compact

photonic structures. Winn *et al.* [158] showed theoretically that simultaneous spatial and temporal modulation of the refractive index promotes inter-band indirect transitions of photons, by imparting frequency and wavevector shifts. These transitions can be interpreted as the optical analogue to *electronic* indirect transitions in a semiconductor, in which the absorption of a photon and a phonon causes an electron to change both its energy and momentum. Refractive index tuning of this kind, however, is not trivial to accomplish [159], and to-date the only experimental demonstration has been very recently reported by Lira *et al.* in a slotted waveguide [160].

A simpler wavelength conversion effect is instead based on adiabatic tuning, the same mechanism explored by Fan *et al.* [15] for dynamically stopping a light pulse (section 3.1.2). The results of Fan *et al.* suggest the possibility of dynamically controlling the wavelength of light through adiabatic tuning of the dispersion curve, but explicit numerical investigation of such controllability was first reported by Notomi *et al.* [23] for a photonic crystal cavity (Fig. 3.5b). The authors showed that if the resonant frequency of the cavity is tuned *while* light is stored into it, the wavelength of the light follows the changes in the cavity resonant wavelength (Fig. 3.5c). In fact, adiabatic wavelength conversion is the optical equivalent of the tuning of a classical resonator [75]: by plucking the string of a guitar, we generate a sound at a certain frequency, but if we slide the finger along the string before the sound dies away, the tone of the sound changes, following the resonant frequency of the string.

When compared to four-wave mixing, which typically allows for conversion efficiencies of less than -20 dB and over long waveguides, the adiabatic wavelength conversion has the big advantage of an efficiency *close to unity*, providing that losses from other sources are negligible [23]; furthermore, it does not depend on the initial light intensity nor on phase-matching conditions, and occurs over lengths on the order of the physical length of a light pulse. When compared to inter-band indirect photonic transitions [158,160], adiabatic tuning involves only one photonic state and the wavelength shift is proportional to the magnitude of the refractive index shift, not to the frequency of the modulation [146]. The tuning should maintain the structure periodicity in order to conserve the wavevector, and it should be adiabatic to avoid coupling to other modes [15].

Modulation of the material properties in the presence of photons requires the use of high-Q cavities or slow light media to hold the light during the tuning process. Adiabatic wavelength conversion was first observed by Preble *et al.* [161] in a silicon ring resonator and by Tanabe *et al.* [162] in a photonic crystal cavity. It was then experimentally demonstrated also in the slow light regime of a silicon W1 photonic crystal waveguide by Kampfrath *et al.* [57]: a probe light pulse travelling through the waveguide is blue-shifted when the silicon refractive index is changed by illuminating the material with a pump pulse. This is shown in Fig. 3.6, where the output spectrum of the probe

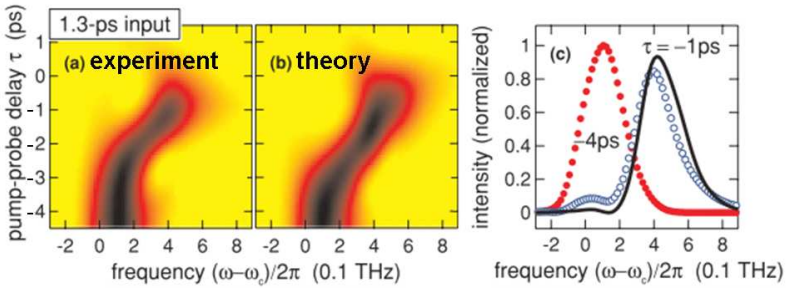

Figure 3.6: Ultrafast adiabatic frequency shifting as reported by Kampfrath *et al.* (from [57]. Copyright 2010 by The American Physical Society). (a) Experimentally determined and (b) calculated intensity spectra of a W1 waveguide response to a 1.3 ps-long input probe pulse at center frequency $\omega_c/2\pi = 202$ THz, for different pump-probe delays τ. (c) Measured and calculated output spectra at $\tau = -4$ ps and 1 ps.

pulse is reported as a function of the pump-probe delay τ. If the probe arrives before the pump ($\tau < -4$ ps), it propagates in the unshifted waveguide. If the probe arrives after the pump ($\tau > 1$ ps), its frequency falls in the band gap of the blue-shifted band diagram and is reflected. Around $\tau = -1$ ps pump and probe have maximum overlap, and the probe spectrum is adiabatically blue-shifted by 0.39 THz. With respect to configurations based on nanocavities, the use of slow light waveguide modes increases the conversion efficiency to more than 80%. This is due to the fact that in a waveguide the probe pulse *propagates* through the tuned region rather than being held inside it, and is therefore affected by free-carrier absorption only in the time necessary for wavelength conversion.

On-the-fly wavelength conversion in a silicon photonic crystal waveguide was also reported by Upham *et al.*, [163], but the probe pulse was not propagating in a slow mode; as a result, the observed frequency shift was less than the induced refractive index change, due to the probe not being fully contained within the tuning spot during the entire index change. This highlights the importance of the use of slow light for this type of effect.

3.3 Photonic crystal tunable delay line: working principle and design considerations

For our purpose of achieving tunable optical delay in integrated format, the scheme based on adiabatic wavelength conversion and dispersion is very promising. The relative frequency shift achievable through adiabatic tuning, proportional to the relative index change, is only a few nanometres in the 1550 nm range. This is sufficient if the dispersion of the waveguide is designed appropriately, through use of the tools for dispersion engineering

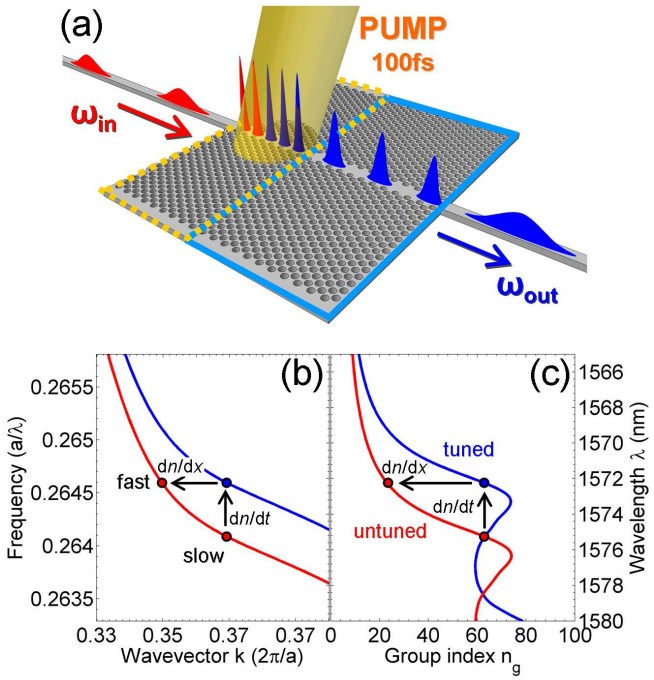

Figure 3.7: Working principle of a tunable delay line based on wavelength conversion and dispersion. (a) Schematic of the device: an input probe pulse (red) at frequency ω_{in} enters the tuning region of the waveguide (outlined in yellow), where it is adiabatically blue-shifted through the generation of free carriers by an ultrafast pump pulse; the probe pulse at the new frequency ω_{out} (blue) travels faster in the untuned delay section (outlined in blue) (picture courtesy of Dr. C. Reardon). (b-c) Calculated dispersion (b) and group index curve (c) of an untuned (red) photonic crystal waveguide of refractive index $n = 3.48$, period $a = 416$ nm, thickness $h = 220$ nm, radius $r = 120$ nm and shifts $s_1 = -46$ nm and $s_2 = 28$ nm. The dispersion of the tuned waveguide, with $\Delta n = -7 \times 10^{-3}$, is indicated in blue. The dots and arrows indicate the two-step adiabatic transition of the probe from slow to fast light.

described in section 2.2.

As shown in Fig. 3.7a, the proposed structure for achieving tunable optical delay is a single engineered slow light waveguide described by the red dispersion curve of Figs. 3.7b-c; the waveguide acts as a delay medium and, in its initial section, also as a frequency conversion component. As an input probe pulse at frequency ω_{in} enters the structure in the slow light regime, the first section of the waveguide (yellow box in Fig. 3.7a) is illuminated with a pump pulse, generating free carriers and hence blue-shifting the waveguide mode (blue curve in Figs. 3.7b-c). This causes the probe pulse to experience a change of refractive index $\mathrm{d}n/\mathrm{d}t$ *in time*, and therefore it is blue-shifted to a new frequency ω_{out} through adiabatic tuning, whereas the wavevector k is conserved [163]. The probe propagates then through the waveguide and

enters the untuned region (blue box in Fig. 3.7a), where its new frequency corresponds to a faster mode: the experienced change of refractive index $\mathrm{d}n/\mathrm{d}x$ *in space*, at the interface between the tuned and the untuned region, corresponds to a change in the pulse wavevector k, with conservation of frequency ω [163].

The overall adiabatic conversion from the slow to the fast state can be viewed also as a form of intra-band indirect photonic transition, since it involves a change in both the frequency and the wavevector of the probe pulse. Unlike the photonic indirect transitions discussed in section 3.2.1 (which are the true optical equivalent to indirect transitions of electrons), however, the frequency and wavevector transformations are performed here in two separate steps. Our approach can be thought of as the analogue of an electron experiencing a dynamic variation of the crystalline atomic potential (which is equivalent to the refractive index in photonic crystals), and allows us to adopt an index tuning profile that is much easier to achieve than the complicated index modulation implemented in Ref. [160]. Furthermore, in the case of a true photonic indirect transition, the achievable frequency shift is limited by the modulation speed (10 GHz achieved experimentally in [160]), and is therefore much smaller than that obtainable through adiabatic tuning (390 GHz demonstrated in [57]).

Our scheme for achieving tunable delay is still limited by the delay-bandwidth product in the dispersive section of the waveguide, just as time-independent configurations; however, unlike the solutions discussed in section 3.1.1, it necessitates tuning of only a small portion of the waveguide — the tuning region needs to be only as long as the probe pulse — rather than the whole structure, thus limiting the power consumption and the loss from free-carrier absorption. If longer delays are needed for the same frequency shift, we simply need a longer static delay section (blue box in Fig. 3.7a), with no extra penalty arising from the tuning process, as the tuning region remains exactly the same; in a scheme where instead the entire structure is to be tuned, increasing the waveguide length for longer delays incurs in even higher carrier-induced loss and energy requirements. These are advantages of separating the tuning section from the delay section, as also implemented in the trap-release schemes by Upham *et al.* [18] and Elshaari *et al.* [150] (Figs. 3.4c-d).

The choice of the waveguide design to be used as a delay line must take into account realistic wavelength shifts achievable with adiabatic tuning, as well as constraints given by propagation loss and group velocity dispersion. In the following we assume a wavelength shift of 3 nm, as reported by Kampfrath *et al.* in Ref. [57] (Fig. 3.6).

Linear propagation loss imposes a constraint on the maximum length of the waveguide and therefore limits the maximum achievable delay. Figure 3.8a shows an indication of the propagation loss for the design of Fig. 3.7. The loss is obtained here through use of a freely available code [31] developed

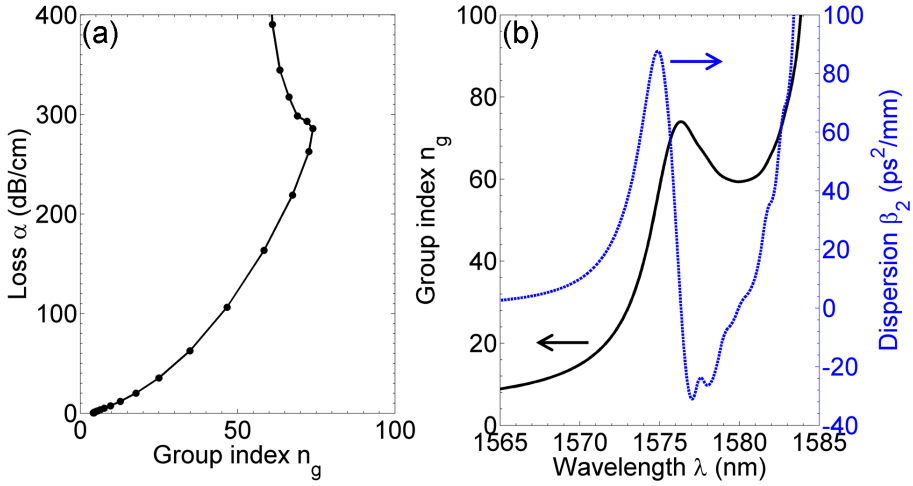

Figure 3.8: (a) Calculated linear loss as a function of group index for the waveguide design of Fig. 3.7. (b) Calculated group index curve (black) and group velocity dispersion parameter (blue) as a function of wavelength for the same waveguide.

by D. M. Beggs and S. A. Schulz, which implements a theoretical model well matching experimental results, as reported by O'Faolain *et al.* in Ref. [31]. According to the theoretical prediction, for a group index $n_{g,in} = 63$ at wavelength $\lambda_{in} = 1575.25$ nm we may expect the linear loss to be $\alpha \sim 190$ dB/cm. If we choose to keep the overall loss under a threshold of 10 dB, the maximum length of the device will be limited to $L_{loss} = 526$ μm.

The group velocity dispersion must also be considered, since it causes a pulse of finite bandwidth to distort while propagating through the device. For a given waveguide length, the minimum pulse duration $T_{0,min}$ that can propagate with acceptable distortion is given by [164]:

$$T_{0,min} = \sqrt{L\,|\beta_{2,max}|}, \tag{3.1}$$

where $|\beta_{2,max}|$ is the maximum (absolute) value of the group velocity dispersion parameter [164] in the 3 nm wavelength range of interest, that is, between 1572.25 and 1575.25 nm:

$$\beta_2 = \frac{\partial^2 k}{\partial \omega^2} = \frac{1}{c}\frac{\partial n_g}{\partial \omega}.$$

For the design of Fig. 3.7, the dispersion parameter β_2 is shown in Fig. 3.8b (blue curve), and we have $|\beta_{2,max}| \sim 87.4$ ps^2/mm at $n_g = 56$. Therefore, through use of Eq. (3.1) we obtain that the minimum pulse duration that can propagate in a waveguide of length L_{loss} is $T_{0,min} = 6.8$ ps.

The device length, together with the group index difference Δn_g between the original and the blue-shifted signals, gives us the maximum delay Δt_{max}

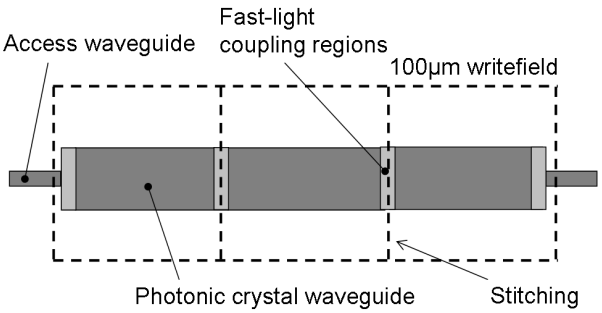

Figure 3.9: Fast-light coupling sections are included at the writefield boundaries of a long (i.e. longer than one writefield) photonic crystal waveguide [76], to minimise the optical loss from stitching errors.

achievable for the chosen design:

$$\Delta t_{max} = \frac{L_{loss}}{c} \Delta n_g.$$

A wavelength shift of 3 nm corresponds here to a change of group index from $n_{g,in} = 63$ to $n_{g,out} = 23.1$, thus $\Delta t_{max} = 69.9$ ps, or 10.3 pulse-widths.

The same band structure allows for further reduction of the minimum time duration $T_{0,min}$ of the signal, provided that the length of the device is also decreased, according to Eq. (3.1). A shorter length results also in a smaller delay which, however, we can set to a standard value of 8 pulse-widths (that is, one byte). Our band structure can then support a 5.2 ps-long pulse over a length $L = 316$ μm.

The effective rate at which signal pulses may be sent into the delay line will be ultimately determined by the carrier lifetime in the tuning section: once a pulse has been wavelength-shifted due to carriers injection, before sending a new pulse we need the carriers to recombine to bring the system back to its initial state. Carrier lifetimes in bulk silicon are very long, on the order of microseconds [165–167], but much lower values of a few hundreds of picoseconds have been reported for nanowire-based photonic devices [168], and even lower for photonic crystal structures [169, 170], due to favourable surface recombination processes and carrier diffusion [171]. However, this limitation of effective bit rate is solely related to the wavelength conversion mechanism, not to the band structure of the photonic crystal waveguide, and is determined by the available technology.

3.4 Demonstration of tunable delay

The slow light photonic crystal waveguide design described in the previous section was fabricated on SOI wafer through e-beam lithography and

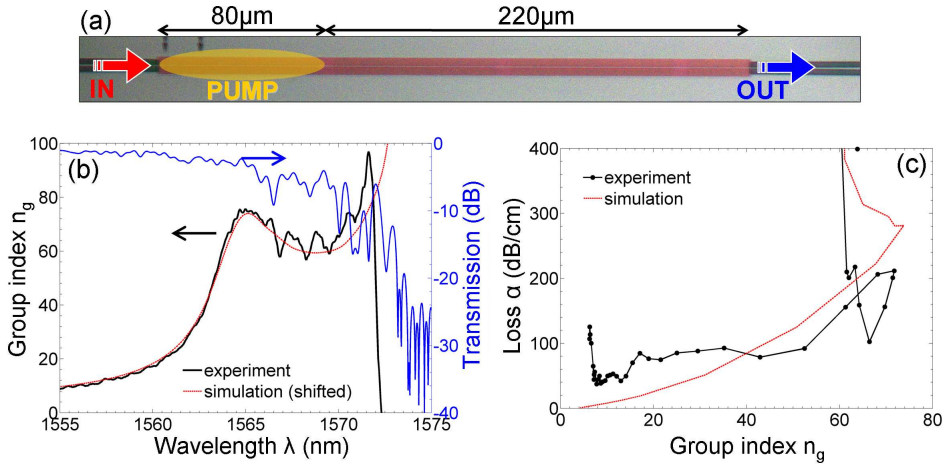

Figure 3.10: (a) Microscope image of the fabricated slow light photonic crystal waveguide. (b) Measured transmission (blue) and group index curve (black); the simulated group index curve (red), shifted in wavelength, is also reported for comparison. (c) Experimental (black) and theoretical (red) estimate of propagation loss for this particular design.

RIE dry-etch, and membraned in a dilute HF solution, as detailed in section 2.3. The e-beam writing was performed using a writefield of 100 μm $\times$ 100 μm, and therefore waveguides of greater length are inevitably affected by stitching errors (Fig. 2.12), on the order of 30–50 nm. To minimise the optical effect of stitching errors, short waveguide sections with elongated period along the propagation direction are used at the writefield boundaries, as schematically shown in Fig. 3.9. Such regions operate in the fast-light regime, which is much less sensitive to the effects of stitching errors [76], and are the same as the fast-light interfaces used to efficiently couple light into slow modes (see section 2.2).

Figure 3.10a shows a microscope picture of the fabricated waveguide, of length 300 μm (excluding the fast-light coupling and stitching sections). The first 80 μm of the waveguide are used as the tuning section where adiabatic wavelength conversion takes place, whereas the remaining 220 μm-long section acts as the dispersive medium. According to Eq. (3.1) and accounting for the highest calculated dispersion $|\beta_{2,max}| \sim 87.4$ ps^2/mm, we estimate the minimum pulse duration that can propagate through the 220 μm-long delay section to be 4.38 ps. The transmission and group index curve of the waveguide were characterised using the setup described in section 2.4, and are shown in Fig. 3.10b; the calculated curve from Fig. 3.8b is repeated here (shifted in wavelength) for comparison. Note that we are mostly interested in the shape of the group index curve, rather than its exact wavelength: the operating wavelength of the final fabricated device differs from the simulated

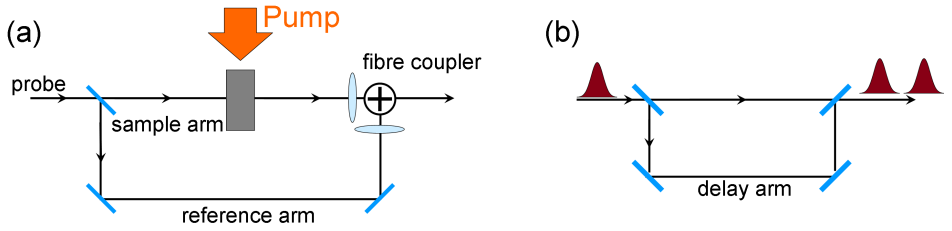

Figure 3.11: (a) Schematic of the setup used for the characterisation of the tunable delay line [172]. An input broadband probe pulse enters the device from left, and the resulting output is detected via interference with a reference pulse. The arrival of a pump laser pulse makes the optical properties of the device time dependent. The timing between the probe and the pump pulses is controlled via a mechanical stage. (b) Schematic of the delay line used to generate a train of two input probe pulses for demonstrating manipulation at the single bit level.

curve — and may vary for different waveguides on the same chip — due to many factors, including the real refractive index of silicon, and the actual slab thickness and hole radius. The propagation loss was also estimated, by comparing the transmission of the 300 μm-long waveguide to that of waveguides of shorter lengths (80, 150 and 220 μm) on the same chip, and is shown in Fig. 3.10c. The experimental curve follows a trend similar to the group index curve itself, which becomes more clear when comparing Fig. 3.10c to Fig. 3.10b: the loss increases steadily up to the local maximum of the group index $n_g \sim 75$ at 1565 nm; then, as the group index decreases, the loss is also reduced, before increasing abruptly due to backscattering becoming the dominant loss mechanism (see also section 2.2). This trend explains the loop appearing in the loss curve when plotted as a function of the group index, as in Fig. 3.10c.

The tunable delay experiments were performed by our collaborators from the NanoOptics group of the FOM Institute AMOLF, Amsterdam, through use of a spectral interferometry setup, shown schematically in Fig. 3.11a. A Ti:sapphire laser of centre wavelength 810 nm, duration 100 fs FWHMI and repetition rate 80 MHz generates the pump pulses, which are focussed onto the sample surface using cylindrical lenses; a slit in the far-field ensures the homogeneous excitation of the tuning region and prevents the delay region from being illuminated. The pump energy absorbed inside the silicon is estimated as ~2.5% of the incident pump energy, and can be varied up to a maximum of ~50 pJ. The free-carrier lifetime was measured as 200 ps. Pulses from the Ti:sapphire laser are also passed in an optical parametric oscillator to generate synchronised *broadband* probe pulses of centre wavelength 1550 nm and duration 180 fs FWHMI. Part of the input probe pulse is coupled into the waveguide, whereas the remaining part travels through a reference path, and the two recombine and interfere at the output. The intensity spectra of the reference pulse, of the waveguide output pulse, and

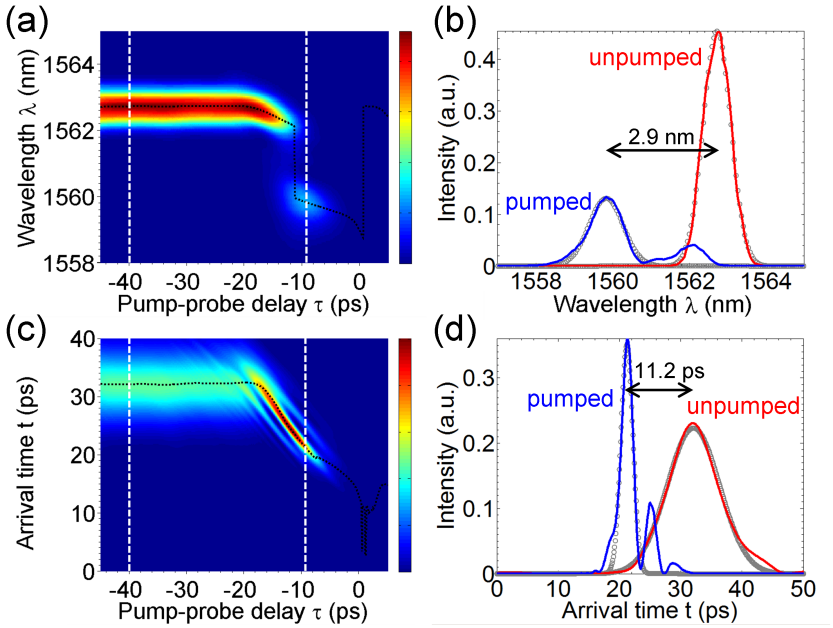

Figure 3.12: Response of the tunable delay line to an input probe pulse of centre wavelength 1563 nm and duration 4.7 ps. (a) Intensity spectrum and (c) arrival time of the output signal for different pump-probe delays τ. The black dotted curves indicate the peak position of Gaussian fits to the data, for each τ. (b) and (d) show wavelength spectra and arrival times along the τ-cuts indicated by white dashed lines in (a) and (c), respectively. Red curves ($\tau = -40$ ps) correspond to the untuned probe pulse; blue curves ($\tau = -9.3$ ps) correspond to the tuned probe pulse. The tuned pulse is blue-shifted by 2.9 nm and exits 11.2 ps earlier than its unpumped counterpart. Grey circles are Gaussian fits.

of their superimposition are measured with an OSA (resolution 0.1 nm) and are used to retrieve the complex-valued output electric field in the frequency domain [173]. The probe energy is maintained at low levels (~10 pJ) to avoid nonlinear effects. The measurement is repeated varying the timing between the pump and the probe pulses, which is set by a mechanical stage.

From these measurements we can characterise the time-dependent linear transfer function of the waveguide over the 60–70 nm bandwidth of the probe, which in turn allows us to determine the linear response of the system to any input probe pulse whose bandwidth is contained within the measured spectrum. A detailed description of this measurement technique is reported in Ref. [172]. The information on the delay of the waveguide output pulse with respect to the reference pulse is contained in their phase difference, which can be extracted in a similar way to what described in section 2.4 for the case of group index measurement.

The response of our tunable delay line to an input probe pulse of duration 4.7 ps FWHMI and centre wavelength 1563 nm, with an absorbed

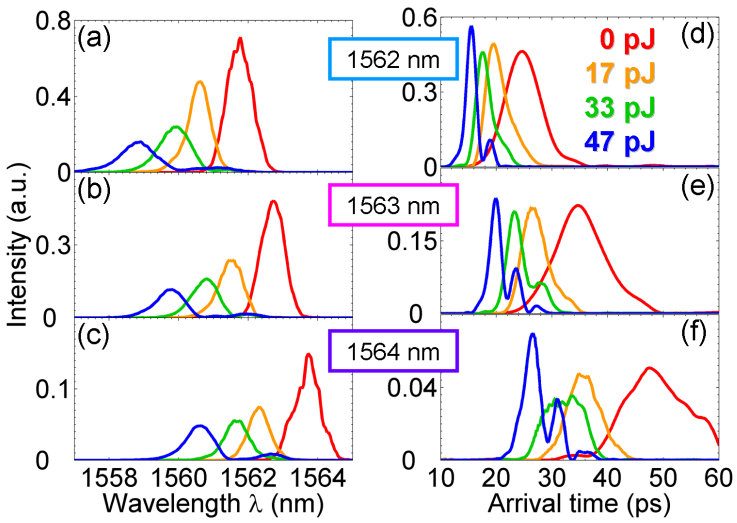

Figure 3.13: Continuous tunability of the delay. (a-c) Wavelength spectra and (d-f) arrival times of 4.7 ps-long probe pulses for absorbed pump energies 0 (unpumped), 17, 33 and 47 pJ. The input centre wavelength varies from 1562 nm (a,d), to 1563 nm (b,e) and 1564 nm (c,f). Frequency shifts and delay values are summarised in Fig. 3.14.

pump energy of ~47 pJ, is shown in Fig. 3.12. Figure 3.12a shows the output wavelength spectra for different pump-probe delays τ. For $\tau < -20$ ps the signal pulse travels through the tuning region before the arrival of the pump, and exits the waveguide untuned, as also indicated by the red spectrum of Fig. 3.12b. The ~0.2 nm blue-shift of the central wavelength results from spectral filtering due to higher propagation loss at longer, slower wavelengths. The transmission through the waveguide at this wavelength (Fig. 3.10b) was measured as -3 dB, due to the ~100 dB/cm propagation loss in the slow light regime (Fig. 3.10c). For $\tau > 0$ ps the probe arrives after the pump and is rejected, as the free carriers are still present inside the tuning region. At $\tau = -9.3$ ps the probe and pump pulses have maximum overlap in the tuning region, and the probe is adiabatically blue-shifted by 2.9 nm to a new centre wavelength of 1559.9 nm (Fig. 3.12b, blue line). The transmission of the converted pulse was -6.5 dB; accounting for -1.5 dB due to the ~50 dB/cm linear loss in the fast light regime, we estimated the efficiency of the adiabatic conversion to be around -5 dB.

Figure 3.12c shows the probe pulse arrival time for different pump-probe delays τ, and demonstrates the operating principle of the tunable delay line. The untuned pulse (red curve of Fig. 3.12d) propagates through the waveguide with a group index $n_{g,in} = 39$ and exits with a delay of 32.3 ps with respect to the reference pulse. The blue-shifted pulse (blue curve of Fig. 3.12d), however, travels faster, at a group index $n_{g,out} = 17$, and exits 11.2 ps

earlier than the unconverted pulse. Therefore, the observed wavelength conversion translates in the time domain into a ~2.4 pulse-widths change of delay.

The wavelength conversion process does not alter the spectral content of the signal, save for the frequency conversion: the widths of the red and blue pulses in Fig. 3.12b are the same in the wavelength domain. In the time domain, however, the higher group velocity dispersion at the original wavelength (Fig. 3.8b) causes the slow red pulse to be wider than the tuned pulse. Such broadening could be addressed by adding suitably engineered waveguide sections to operate dispersion compensation.

Note that our approach based on tuning of the signal rather than of the waveguide helps to balance the losses of tuned and untuned pulses. The converted pulse is subject to free-carrier absorption in the tuning section, but travels in a low-loss fast mode through the delay section. On the other hand, the unconverted pulse propagates in the lossier slow light regime, but does not suffer from conversion losses.

Finally, a fraction of the pumped pulse remains unconverted or partially converted, as can be seen from the residual peak around 1562 nm in the blue curve of Fig. 3.12b. This is light that was not contained in the tuning region during the entire pumping time: for example, the head and tail of the signal which, respectively, exit and enter the tuning region while the conversion is taking place, and therefore do not experience the full variation of refractive index in time. The interference between the converted and the partially converted light in the frequency domain translates into the presence of modulation fringes in the blue pulse of Fig. 3.12d in the time domain. Full conversion could be achieved by increasing the length of the tuning region or by compressing the input signal pulse by making it slower, but the present configuration is a good compromise between the amount of converted light with energy requirements, loss and dispersion.

The delays in our structure can be tuned continuously by adjusting the pump power, as demonstrated in Fig. 3.13. Figures 3.13a-c show the spectral response of the delay line to input probe pulses of duration 4.7 ps FWHMI and centre wavelengths 1562, 1563 and 1564 nm, for different absorbed pump energies. The corresponding outputs in the time domain are shown in Figs. 3.13d-f. The measured frequency shifts and tunable delays as a function of absorbed pump energies are summarised in Fig. 3.14. As can be seen from Fig. 3.14a, the frequency shift depends linearly on the absorbed pump energy and is independent of the input signal wavelength, as expected: the refractive index change causing the frequency shift is proportional to the density of generated free carriers, which in turn is proportional to the pump energy [144]. By fitting the data of Fig. 3.14a we obtain a frequency shift of ~7.7 GHz/pJ. The amount of tunable delay for varying pump energy, instead, depends on the group index dispersion, and is therefore different for different input wavelengths (Fig. 3.14b). The highest tunable delay is

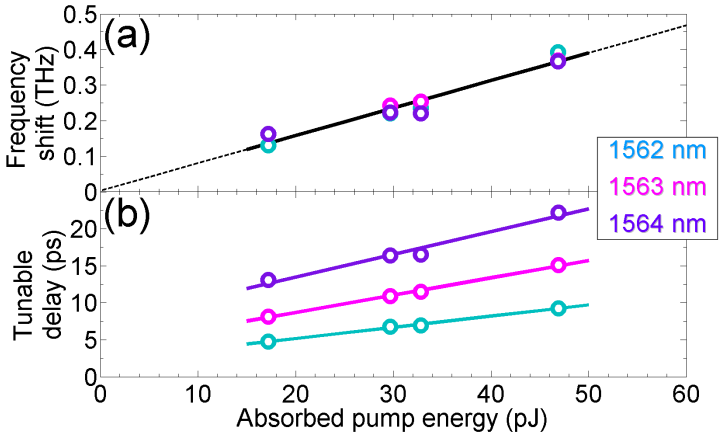

Figure 3.14: (a) Measured frequency shifts and (b) tunable delays as a function of absorbed pump energy for input probe wavelengths 1562, 1563 and 1564 nm. Circles correspond to the peak position of the pulses in Fig. 3.13. Lines are linear fits to the data. The frequency shift is directly proportional to the absorbed pump energy.

achieved here for input wavelength 1564 nm ($n_{g,in} \sim 75$) and absorbed pump energy 47 pJ, and is equal to 22.2 ps, or 4.7 pulse-widths.

3.5 Manipulation at single bit level

In section 3.4 we have applied the wavelength conversion and dispersion approach to implement a two-step photonic indirect transition and demonstrate the dynamic delay of optical pulses and its fine tunability. We will now focus on the ultrafast nature of the controlling mechanism (the generation of free carriers by the 100 fs pump pulse), in contrast to the quasi-static thermo-optic structures discussed in section 3.1.1, and apply its potential by operating on individual pulses in a pulse stream.

The experiment was performed in the same way as described in section 3.4, but using two broadband probe pulses of equal amplitude separated by 30 ps. Such a pulse train was generated through use of a beam splitter and a delay arm, as schematically shown in Fig. 3.11b. The response of the system to two probe pulses of centre wavelength 1563 nm, duration 4.7 ps, and 30 ps apart was then determined from the measured transfer function, and is shown in Fig. 3.15 for different values of the pump-probe delay τ. The absorbed pump energy was estimated as 33 pJ. For $\tau = -60$ ps both pulses travel through the tuning region unpumped (Fig. 3.15b), and exit from the delay section still separated by 30 ps (Fig. 3.15f). The modulation fringes in the wavelength spectrum result from the interference between the two pulses in the time domain. As discussed previously, the pulses are broadened in the

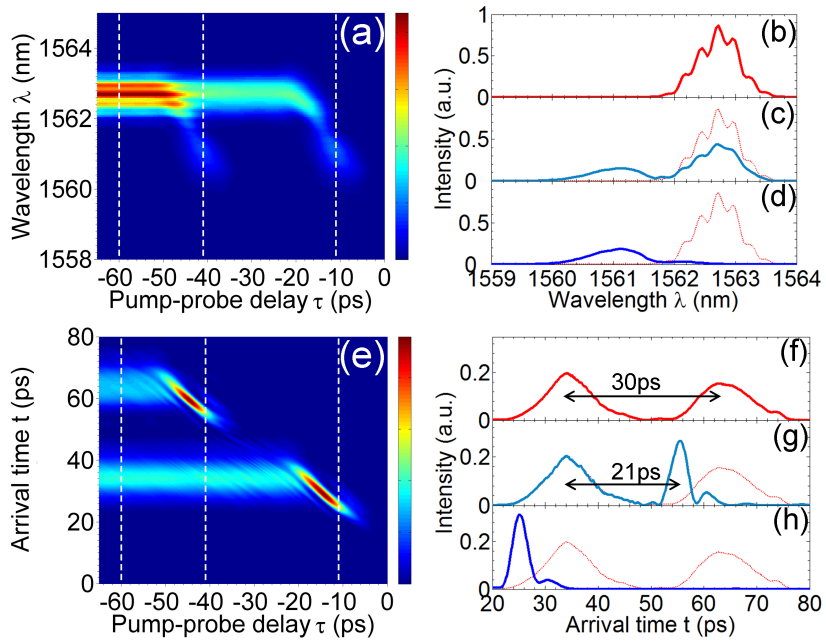

Figure 3.15: Response of the tunable delay line to a train of two input probe pulses of centre wavelength 1563 nm and duration 4.7 ps, with 30 ps separation. (a) Intensity spectrum and (e) arrival time of the output for different pump-probe delays τ. (b-d) and (f-h) show wavelength spectra and arrival times along τ-cuts indicated by white dashed lines in (a) and (c), respectively. Red curves ($\tau = -60$ ps) correspond to the untuned probe pulses: both pulses exit at the original wavelength and still separated by 30 ps. Cyan curves ($\tau = -41$ ps) correspond to the second probe pulse being blue-shifted and exiting the waveguide only 21 ps after the first pulse. Blue curves ($\tau = -11$ ps) correspond to the first probe pulse being tuned; the second pulse is rejected due to the hold-off time of the device.

time domain due to the group velocity dispersion in the slow light regime; the peak intensity of the second pulse is lower than that of the first one due to smaller cross-correlation with the reference pulse.

For $\tau = -41$ ps the pump pulse arrives when the second probe pulse is inside the tuning region. As can be seen from Figs. 3.15c,g, the first pulse is not altered and exits the waveguide in the ground state, whereas the second pulse is blue-shifted and travels faster in the delay section, catching up with the first pulse: the separation between the two pulses is now only 21 ps. The second pulse is therefore manipulated independently of the first.

At $\tau = -10$ ps, the pump operates on the first pulse of the train (Figs. 3.15d,h). The second pulse in this case is rejected, as the free carriers have not yet dissipated. The ~200 ps free-carrier lifetime is in fact the major limitation on the device hold-off time, but note that no special care was taken on the sample to reduce it. The carrier lifetime could be reduced by sweeping out the carriers through integration of a reverse-biased pn-

junction [56], or by appropriate treatment of the silicon slab [170, 174, 175]. Also note that while this limitation on the hold-off time is technological, in principle the device can operate bit-by-bit on individual pulses, or even on multiple bits at the time. In the case of the trap-and-release schemes discussed in section 3.1.2, the hold-off time will still be limited by their very own performance (the maximum delay) even without carrier lifetime constraints.

Finally and importantly, the adiabatic wavelength conversion, which is the core of our tuning process, can be achieved also by electrical carrier injection, as demonstrated by Tanabe *et al.* for the case of a photonic crystal nanocavity [176]. Thus the external pump laser is not a necessity, and the scheme is compatible with integration of optoelectronic devices.

3.6 Summary

A fine control of optical pulses in the time domain is a desirable feature for various types of applications in the all-optical signal processing framework, and a great research effort has been undertaken in recent years to develop such functionality, particularly in the area of slow light. Many slow light tunable delay lines have been demonstrated in the literature, based on different effects and implemented in a variety of platforms, but in this chapter we have focussed our attention on silicon integrated devices based on structural slow light. We have distinguished between time-independent and time-dependent delay lines: the former are quasi-static, relying on the slow thermo-optic effect, and require tuning of most of the structure; the latter are dynamic, relying on ultrafast tuning, and the main examples consist of trap-and-release schemes in optical microresonators.

We have introduced the concept of a silicon integrated tunable delay line based on a two-step photonic indirect transition, i.e. involving a change in both frequency and momentum of an optical signal pulse, implemented in the form of a wavelength conversion and group velocity dispersion approach. Wavelength conversion is achieved through adiabatic tuning, whereas a controlled dispersion is introduced through use of photonic crystal waveguide dispersion engineering. We have then applied this concept in a proof of principle experiment and demonstrated fine tuning of the delay up to 22.2 ps for a 4.7 ps-long pulse in a 300 μm-long device. We have also demonstrated the potential of the ultrafast approach by manipulating a single pulse in a train of two pulses only 30 ps apart. Table 3.1 compares our results with some of the literature works discussed in section 3.1.

Our approach has the advantage of being scalable and flexible, as only a small part of the waveguide needs to be tuned, and greater delays may be achieved by increasing the length of the delay section, which remains static. The main constraint on the bit rate is given by the hold-off time, which is

limited by the 200 ps free-carrier lifetime, but there are various techniques that may be implemented to improve this figure, as discussed in section 3.5. The maximum delay is related to the length of the static delay section, and is mainly limited by the tolerable propagation loss. If the loss could be reduced further, either by playing on the waveguide design [31] or by introducing optical gain, greater delays could be achieved. For example, assuming a loss $\alpha = n_g$ dB/cm (that is, reducing the simulated loss to about a third for Fig. 3.8a), the design of Fig. 3.8b would allow for a 209 ps delay, corresponding to 18.7 pulse-widths for a dispersion-limited 11.16 ps-long pulse over a length of 1.4 mm and with a 3 nm blue-shift from $\lambda_{in} = 1575.7$ nm ($n_g = 70$) to $\lambda_{out} = 1572.7$ nm ($n_g = 26$). One possible way to introduce optical gain in silicon and therefore compensate for the propagation loss relies on stimulated Raman scattering, as we will discuss in Chapter 5.

Table 3.1: Summary of figures of merit for the reviewed methods for achieving tunable optical delay.

Delay line	Configuration	Max tunable delay	Pulse duration or bit rate	Tunable delay (bits)	Device length	Tuning mechanism	Tuning region	Complexity
Baba *et al.* [105]	chirped PhC coupled wg	23 ps	~3.3 ps	~7	250 μm	thermo-optic	whole wg	fabrication of different r
Adachi *et al.* [85]	chirped PhC coupled wg	74 ps	~3.4 ps	~22	280 μm	thermo-optic	whole wg	fabrication of different r
Ishikura *et al.* [106]	chirped PhC coupled wg	72 ps	~2 ps	~36	400 μm	thermo-optic	whole wg	fabrication of different r
Melloni *et al.* [16]	reflective CROW	320 ps	25 Gb/s	8	~7 mm	thermo-optic	1 to 6 rings	no
Shinobu *et al.* [135]	all-pass filter CROW	184 ps	40 Gb/s	~7.3	1.06 mm	thermo-optic	up to 50 rings	tuning of many rings
Fontaine *et al.* [140]	all-pass filter CROW	205 ps	40 Gb/s	~8.2	~2 cm	thermo-optic	up to 32 rings	tuning of many rings
Cardenas *et al.* [137]	SCISSOR	133 ps	10 Gb/s	~1.33	250 μm	thermo-optic	8 rings	no
Xu *et al.* [17]	2 rings dynamic	97 ps	14 ps	~7	~36 μm	carrier injection	2 rings	no
Elshaari *et al.* [150]	3 rings dynamic	~300 ps	~5 ps	~60	~174 μm	carrier injection	2 rings	no
This work	λ conversion & dispersion	22.2 ps	4.7 ps	4.7	300 μm	carrier injection	80 μm	no

Chapter 4

Efficient light injection: inverse taper couplers

The experimental study of nonlinear effects such as stimulated Raman scattering in silicon waveguides requires the ability to couple sufficient pump power into the fabricated structures.

In this chapter we will explore the possibility of using well-known inverse taper spot-size converters to efficiently couple light into silicon. After reviewing the working principle of these structures, we will describe the challenges posed by their integration with photonic crystals and by the operating wavelength of the Raman signal, and how the solution to these issues led to the development of a reliable fabrication recipe.

4.1 Working principle and simulations

The erbium doped fibre amplifier (EDFA) that we use as a pump source for Raman scattering delivers up to 600 mW of continuous-wave light in the wavelength range 1540–1563 nm. Such a high pump power, however, is wasted if no special care is taken to optimise the in-coupling of light

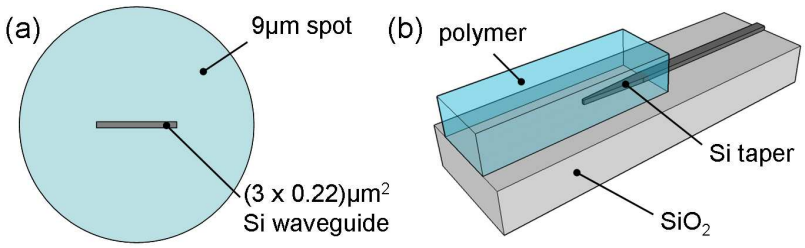

Figure 4.1: (a) In-scale cross section comparison between a typical input light spot and a silicon access waveguide. (b) Schematic of an inverse taper spot-size converter (not to scale).

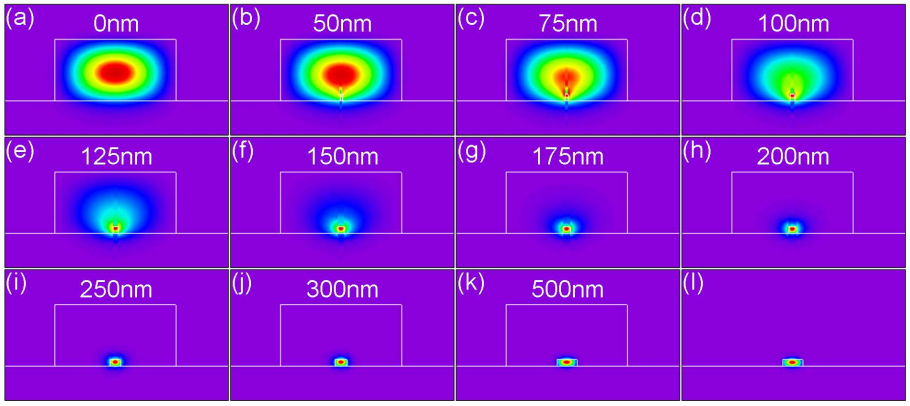

Figure 4.2: Working principle of an inverse taper spot-size converter. Light from a polymer waveguide (a) couples evanescently to the silicon taper, and is progressively confined within it as the taper widens (b-k). The taper width varies here from a 50 nm tip (b) to a 500 nm wire (k), as indicated in each figure. When the polymer waveguide is terminated (l), all the light is confined within the silicon. In this simulation, the polymer waveguide has width 3 μm, height 2 μm and refractive index 1.575. The refractive indices of the silicon taper, the silica bottom layer and the air cladding are 3.48, 1.44 and 1, respectively. The simulation is performed here for TE polarisation.

inside the fabricated devices. Figure 4.1a shows an in-scale schematic of the cross-section difference between the 9 μm mode field size typical of an optical fibre or of the focal point of an objective, and the standard 3 μm-wide silicon strip waveguide that we normally use to feed light into photonic crystals: the mode-size mismatch affecting coupling efficiency is clear. In addition, the high refractive index of silicon (~3.48) results in significant reflection at the interface with air, and in the presence of strong Fabry-Perot fringes in the transmission spectrum of any silicon waveguide. As a consequence, the coupling loss to our strip access waveguide is typically around 8–10 dB per facet, and can be as high as 20 dB per facet in single mode nanowires [177, 178].

A possible solution for improving coupling is the integration of inverse taper spot-size converters [177–181]. In such structures, a silicon nanowire is tapered at its input facet to a width much smaller than that required to support a propagating mode (Fig. 4.1b). The taper is then overlaid with a large polymer waveguide, which can by itself support guided modes as its refractive index, lower than silicon, is higher than the underlying silica ($n_{SiO_2} = 1.44$). The input light is initially fed into the polymer waveguide with high coupling efficiency, due to both its large cross section and low refractive index. Then, as the high-index taper adiabatically widens, light is progressively and evanescently confined within the silicon (Fig. 4.2). The same structure is then repeated at the output facet to improve out-coupling

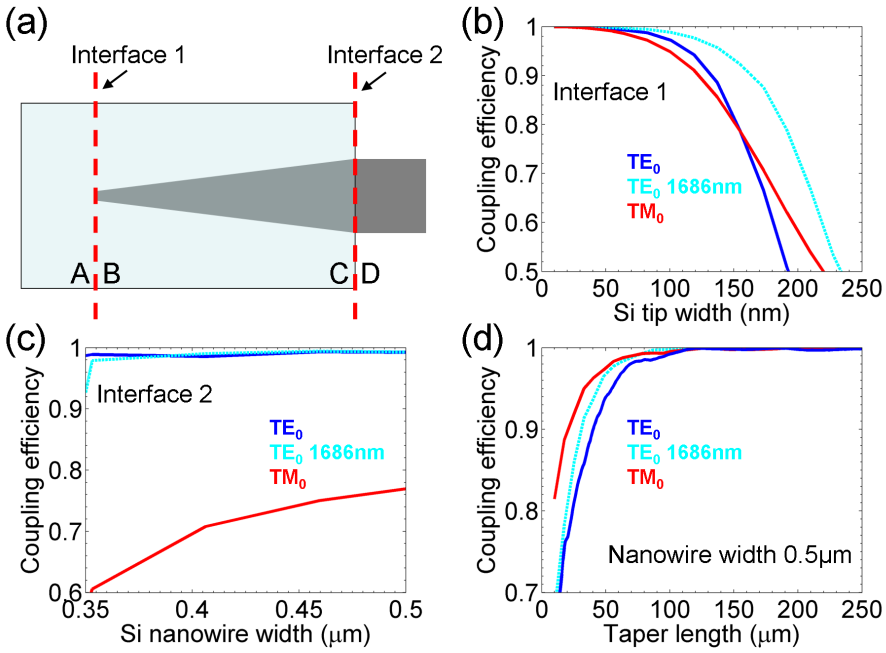

Figure 4.3: (a) The coupling efficiency of light is simulated by dividing the spot-size converter into three sections: i) interface 1 between sections A and B, ii) interface 2 between sections C and D, iii) the taper between sections B and C. (b-d) Coupling efficiency for the fundamental TE (blue) and TM (red) modes at wavelength 1550 nm, and for the fundamental TE mode at 1686 nm (cyan), at (b) interface 1, (c) interface 2, and (d) along the taper. In (d) the taper tip width is 50 nm and the final nanowire width is 500 nm. In all simulations the polymer waveguide is 3 μm wide and 2 μm high.

of the weak Raman signal.

As we shall describe in more detail in Chapter 5, in our study of stimulated Raman scattering the pump and the Raman signal will have two different polarisations and wavelength separation of more than 130 nm. Given such requirements, inverse tapers converters are probably the best choice for efficient coupling, for example when compared to grating couplers [182,183], as they can perform over a very broad bandwidth and they work reasonably well for both polarisations at the same time [181].

The taper tip must be small and the taper long enough to ensure the adiabatic transition of light from the fundamental mode of the polymer waveguide to that of the silicon nanowire, avoiding coupling to non-guided or higher-order modes which would disperse the optical power. To gain a better understanding of the role of the various geometric parameters in the coupling mechanism, we performed three-dimensional simulations of the mode converters using FIMMPROP, a commercial simulation tool based on an eigenmode expansion algorithm for analysis of two- and three-dimensional

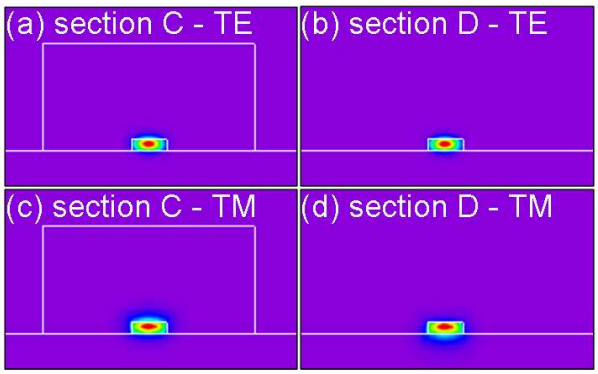

Figure 4.4: Intensity profiles of the fundamental TE (a-b, same as in Figs. 4.2k-l) and TM (c-d) modes in sections C (a,c) and D (b,d), for a 500 nm-wide silicon nanowire. The TM mode exhibits a mode mismatch in the vertical direction which may be responsible for the lower coupling efficiency at interface 2 (Fig. 4.3c).

propagation phenomena in waveguides. In the simulations, the mode converter system was divided into three sub-systems, as illustrated in Fig. 4.3a: i) a first interface between the polymer waveguide (section A) and the taper tip (section B), ii) a second interface where the polymer waveguide is terminated (between sections C and D), iii) the taper itself (between sections B and C). The polymer and silicon refractive indices were set to 1.575 (the refractive index of SU-8 [180, 184]) and 3.48, respectively.

Figure 4.3b shows the coupling efficiency between the fundamental modes on the two sides of interface 1 as a function of the taper tip width, for both TE and TM polarisations at the Raman pump wavelength 1550 nm, and for TE polarisation at the Raman signal wavelength 1686 nm. The two polarisations have similar levels of coupling efficiency, consistently with the results reported in Ref. [181] for our silicon thickness. Provided that the tip width is around 100 nm or less, the coupling is better than 97% and 95% for TE and TM polarisations respectively. The polymer waveguide is here 3 μm wide and 2 μm high, and only a very weak dependence on its dimensions was found at this interface.

At interface 2, TE and TM polarisations behave differently (Fig. 4.3c): the TE mode is transferred with a coupling efficiency of ~99% as soon as the nanowire is wider than 350 nm, but only 77% of the TM mode is coupled for a 500 nm-wide nanowire. The lower TM coupling might be due to a mode profile mismatch between sections C and D (Figs. 4.4c-d), which does not affect the better-confined TE mode (Figs. 4.4a-b). Although not shown here, we have verified by simulation that the TM coupling efficiency at this interface increases for wider nanowires, and can surpass 90% for widths larger than 800 nm. At these widths, however, there is a risk of TM-to-TE conversion inside the taper, due to coupling of the fundamental TM

mode to the first higher-order TE mode appearing in the waveguides, whose dispersion crosses that of the TM mode for silicon widths around 650 nm.

Regarding the optimal taper length, between sections B and C, 50 μm is already sufficient to respect adiabaticity for both polarisations and for nanowire widths up to ~600 nm. As shown in Fig. 4.3d, coupling efficiencies of 97% and 94% are predicted for TE and TM polarisations, respectively, for taper length 50 μm and nanowire width 500 nm. Therefore, the complete system allows in theory for about 93% coupling efficiency (−0.3 dB) of light from the polymer waveguide into the silicon nanowire for TE polarisation, and at least 80% (−1 dB) for TM polarisation. Note that this estimate does not include the coupling from air to the polymer waveguide: we will see in section 4.2.2 that the measured transmission from air into the polymer tends to be around 70% (−1.5±0.5 dB), thus bringing the overall in-coupling efficiency to ~66% (−1.8 dB) for TE polarisation and ~56% (−2.5 dB) for TM polarisation.

4.2 Fabrication and characterisation of improved coupling

The fabrication of spot-size converters requires the definition of good tapers in the silicon layer, and the addition of polymer waveguides precisely aligned to the tapers. The silicon tapers and the trenches to host the polymer waveguides are defined in the same e-beam writing and etch step as the photonic crystals.

4.2.1 Silicon inverse tapers

Particular care must be taken in the design of the tapers: the proximity error correction software (see section 2.3) divides the area to be exposed on the sides of the taper into small polygon sections (as highlighted with the red dashed box in Fig. 4.5a), each with its own exposure dose, resulting in an unwanted modulation of the taper width. The initial taper design therefore included a line 15 nm from the exposed areas (Fig. 4.5a) in order to smooth out such modulations, which however were still visible in the final structure, as shown in Fig. 4.5c. Results were improved by writing two lines 50 nm apart from each other (Fig. 4.5b): SEM images of smoother tapers are shown in Figs. 4.5d-e. In the design file, the taper tip has width between 150 nm and 200 nm (between the two internal lines) to correct for electron scattering and corresponding spot broadening inside the resist.

4.2.2 ZEP polymer waveguides

The polymer waveguides were initially realised in ZEP-520A, the same polymer used as an e-beam resist to define photonic crystals, as proposed by

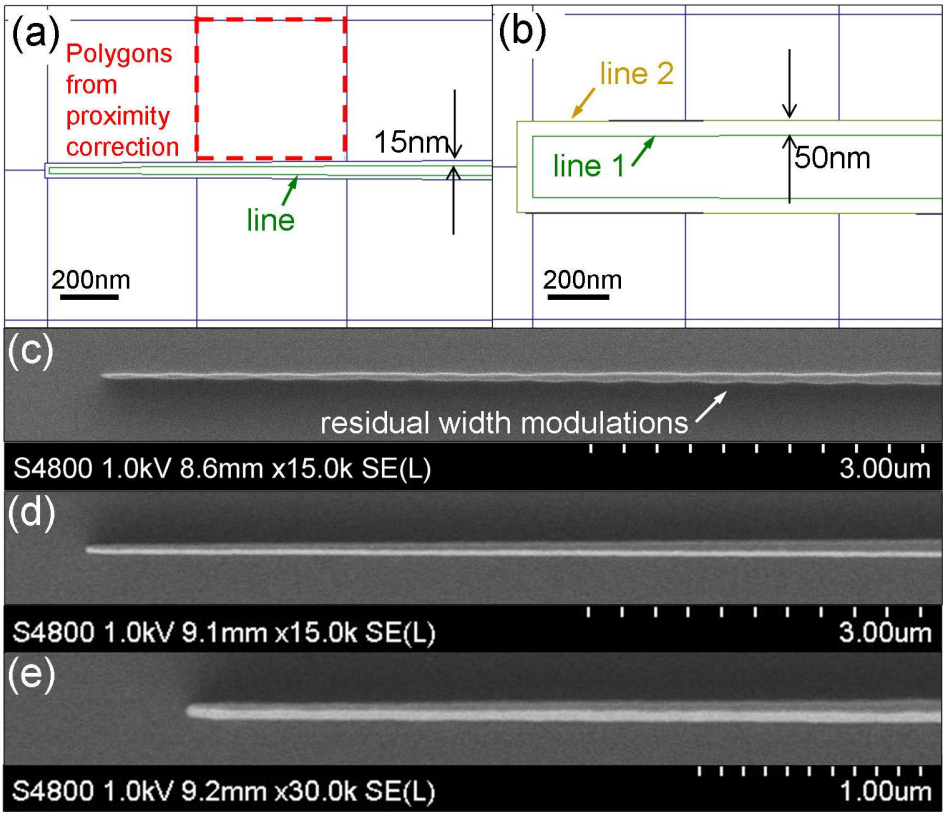

Figure 4.5: (a) Initial design of an inverse taper: the areas to be written are divided into small rectangles by the proximity correction software; a line is included in the design in an attempt to smooth out imperfections arising from the proximity correction. (b) Improved taper design, in which two lines are written, with spacing 50 nm; the taper tip is also wider. (c) SEM image of an inverse taper fabricated using the initial design of (a): a modulation of the taper width is visible. (d-e) SEM images of tapers fabricated with the improved design of (b).

Grande *et al.* for the demonstration of vertical directional couplers [185]. ZEP is spun on the sample at 800 rpm for 1 min and hard-baked 10 min at 180°C. The same spin and baking steps are repeated a second time, resulting in a ~1.8 μm-thick double layer of ZEP. The 3 μm-wide polymer waveguides are defined in an e-beam lithography step involving a three-point alignment procedure: precise alignment to the silicon inverse tapers (Fig. 4.6c), while less critical for successful transmission of the fundamental TE mode, is very important for the TM mode to avoid TM-to-TE conversion. After development 45 s in Xylene at 23°C, the ZEP waveguides have a boat-like cross-section (Fig. 4.6a), which can be partially corrected by baking 2 to 4 min at 130°C before cleaving (Fig. 4.6b). Both the temperature and the baking time should be controlled precisely; in particular, if the temperature

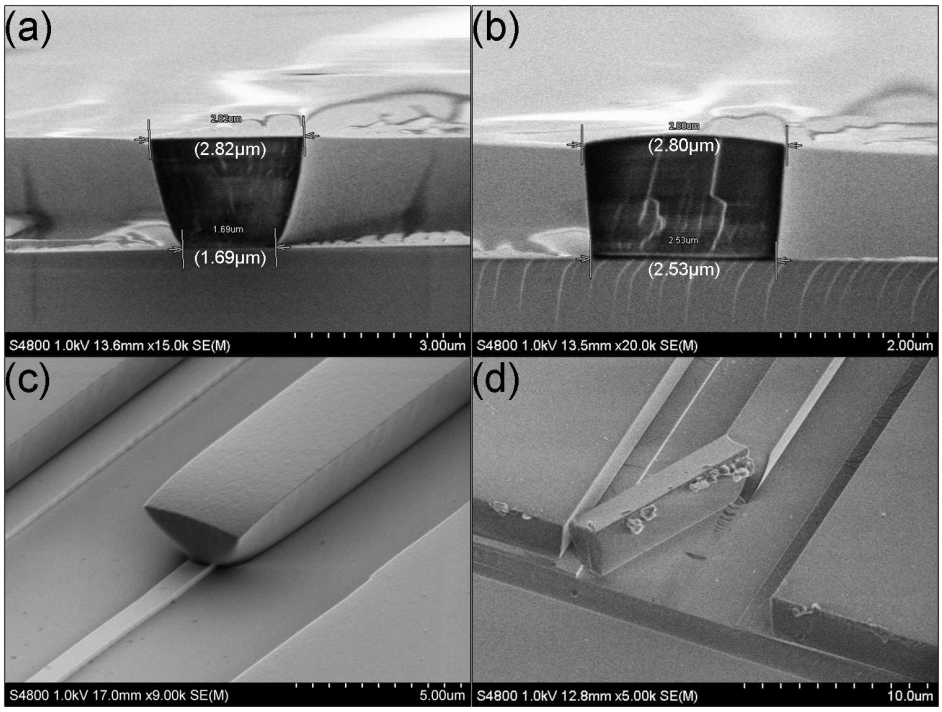

Figure 4.6: SEM images of fabricated ZEP waveguides. (a) Cleaved facet of a ZEP waveguide. (b) Cleaved facet of a ZEP waveguide re-flowed by baking 4 min at 130°C to correct the shape of the cross-section. (c) A ZEP waveguide aligned to an inverse taper; this waveguide was baked 2 min at 130°C. (d) Facet of a ZEP waveguide that broke while cleaving.

is below 120°C, it will not be sufficient to re-flow the waveguides, but at 150°C or higher ZEP will melt.

Such spot-size converters fabricated with ZEP waveguides result already in very satisfactory coupling efficiencies, as shown in Fig. 4.7, where the measured total insertion loss for waveguides of a 2.6 mm-long sample is plotted in the wavelength range of our EDFA pump source. Such loss includes both in- and out-coupling from air, as well as some propagation loss and, in the case of a structure comprising spot-size converters, also in- and out-coupling between the ZEP waveguide and a short (200 μm length) 680 nm-wide silicon nanowire through inverse tapers. For this particular sample we measured in TE polarisation (Fig. 4.7a) a total loss of 3 dB for a blank ZEP waveguide, and 5.6 dB for a waveguide with spot-size converters. By dividing these figures by 2 (and neglecting propagation losses), we roughly estimate that the one-way coupling loss from air into the silicon nanowire is at most 2.8 dB, with ~1.5 dB loss from air into the polymer and ~1.3 dB loss from the polymer into the nanowire, a significant improvement with respect to the 7.5–8 dB loss of the blank silicon waveguide. The spot-size converters

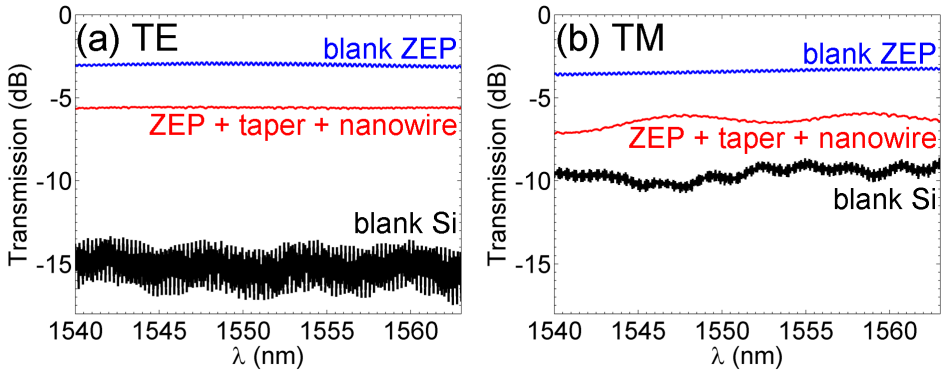

Figure 4.7: Measured total insertion loss for (a) TE and (b) TM polarisation, from a 3 μm-wide blank silicon channel waveguide (black), a ~2.8 μm-wide blank ZEP waveguide (blue) and a structure with inverse tapers couplers on both input and output sides (red). The total insertion loss includes both in- and out-coupling into the waveguides from free-space, as well as propagation loss. The total sample length is 2.6 mm. The structure with inverse tapers clearly outperforms the blank silicon waveguide.

coupling efficiencies for TM polarisation (Fig. 4.7b) are not very different from TE polarisation; the silicon blank waveguide, however, performs better in TM rather than TE polarisation, reducing (but not filling) the gap with the inverse taper structure.

The estimated coupling efficiencies are lower than the best experimental values reported in the literature [177–181], but they fall within a good average [22] (see also Table 4.1). During the experiment we observe scattering of light at the termination of the polymer waveguides, and therefore we believe that good part of the loss is due to excitation of higher-order modes in the polymer waveguides, which cannot couple to the small tapers. The literature values refer to experiments performed with use of taper-lensed fibres for in- and out-coupling, and these may also be more effective in exciting the fundamental mode of the polymer waveguides than the aspheric lenses in our free-space setup, which have higher numerical aperture.

Despite the good performance of the measured spot-size converters, however, ZEP waveguides are very soft: this represents a major disadvantage for their mechanical stability and practicality, as they can easily break in proximity of the facet during cleaving or even just while handling the sample (Fig. 4.6d). We therefore decided to explore a tougher solution, and opted for SU-8.

4.2.3 SU-8 polymer waveguides

With respect to ZEP, which is a positive resist, meaning that the features to be removed (i.e. the lateral trenches) are those to be exposed, SU-8 is a

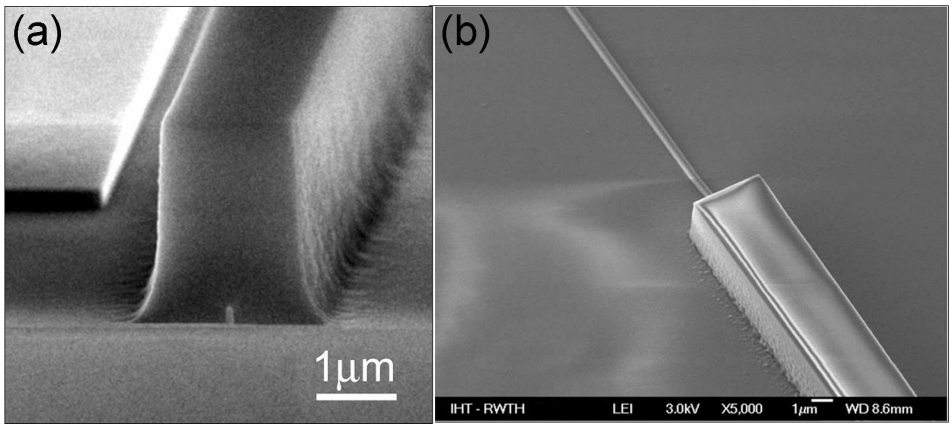

Figure 4.8: SEM images of fabricated SU-8 waveguides taken from the literature: (a) from [180] (Reprinted with permission) and (b) from [186] (Reprinted from [186]. Copyright 2009, with permission from Elsevier).

negative resist, so one needs to directly expose the part that will remain, i.e. the waveguide itself. SU-8 is often reported by other authors as the polymer of choice for this type of structures [178, 180, 186], and results in very sharp features when exposed with high e-beam acceleration voltages, on the order of 100 kV (Fig. 4.8). In our e-beam system, however, the maximum voltage that can be set is 30 kV, meaning that the electrons do not have enough energy to penetrate deep into the silica layer underlying the SU-8, and easily return into the polymer when backscattered. As SU-8 is very sensitive, the backscattered electrons are sufficient to expose it outside the region that is directly written, leaving after development a significant amount of unwanted material on the sides of the waveguide (Fig. 4.9a).[1] The shape of the SU-8 waveguide can be improved by etching in an oxygen plasma which, however, damages the photonic crystals by excessively smoothing the edges (Fig. 4.9b).

As we did not have the option to increase the e-beam acceleration voltage, we found that a good solution was to decrease it instead, and use the very exposure from the scattered electrons to define the waveguide shape. As shown in Fig. 4.9c, we write directly only a width of 1 μm with an e-beam voltage of 15 kV on a ~2.3 μm-thick SU-8 layer, and obtain a ~6 μm-wide SU-8 waveguide with a rounded cross-section (Fig. 4.9d). The part swept directly by the electron beam is visible as the flat stripe on top of the waveguide, whereas the rest is the result of the combined exposure from the electrons scattered in the resist and those backscattered from the substrate. In the design file the waveguide is defined as a set of lines along the propa-

[1] Note that the same effect — but the other way round — is responsible for the boat-shaped ZEP waveguide of Fig. 4.6a: due to backscattered electrons the exposed trench is wider on the bottom.

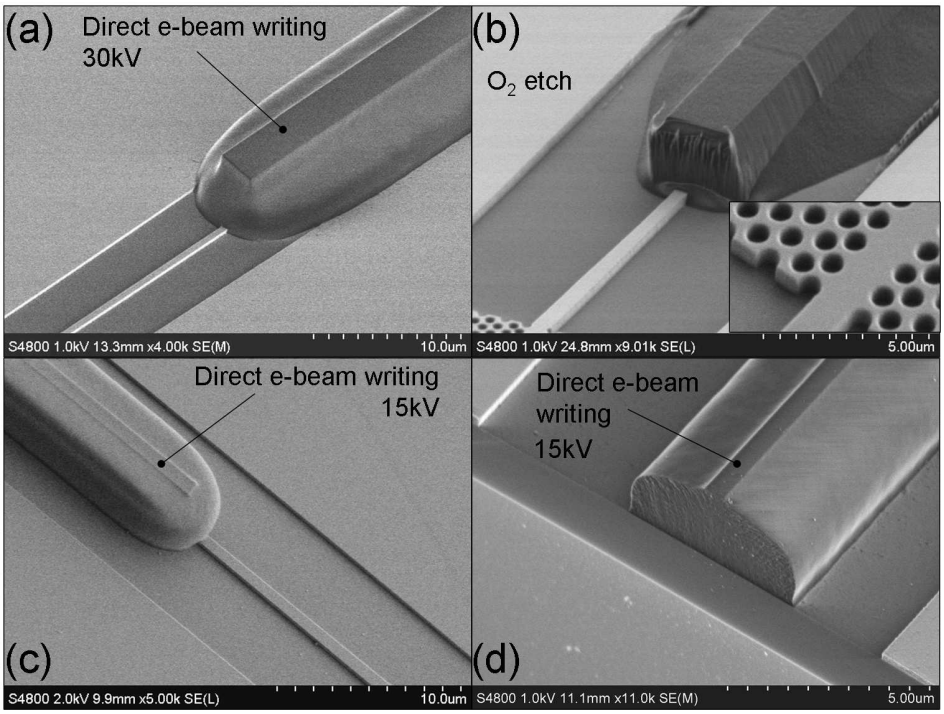

Figure 4.9: SEM images of fabricated SU-8 waveguides. (a) A waveguide exposed at 30 kV acceleration voltage; the residual polymer on the sides of the directly exposed region is the result of the exposure from backscattered electrons. (b) An attempt to correct the shape of the SU-8 waveguide by etching in an oxygen plasma for 30 min; this etch compromises the quality of the photonic crystal (inset, 20 min oxygen etch), by smoothing out the edges. (c-d) A SU-8 waveguide aligned to a taper (c) and a cleaved SU-8 facet (d) for e-beam exposure at 15 kV; the directly exposed portion can have width between 0.5 and 1 μm, depending on the desired final width of the waveguide, which is determined by the backscattered electrons.

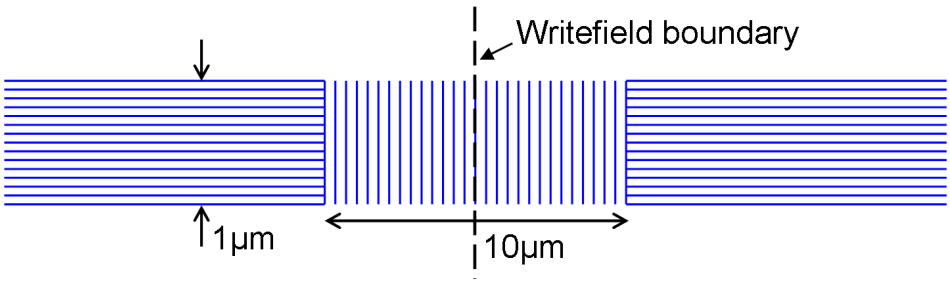

Figure 4.10: Schematic of the design file of the SU-8 waveguide (not to scale). The waveguides are written by e-beam exposure as a set of lines along the waveguide direction with 50 nm spacing. At the writefield boundary the lines are tilted by 90° to slow down the beam and improve stitching errors.

gation direction with 50 nm spacing (Fig. 4.10), rather than a rectangular area: due to the high sensitivity of SU-8, the exposure dose (and thus the exposure time) needed is low, and when exposing areas the electron beam sweeps very quickly across the pattern, generating errors at the writefield boundaries or at points where the beam reverts its direction. The lines are set to be perpendicular to the waveguide direction at the writefield boundaries to further slow down the beam and improve stitching errors.

In summary, the fabrication of SU-8 waveguides starts by spinning a 2.3 μm-thick layer of SU-8 on the sample, followed by a soft-bake 3 min at 65°C and 3 min at 95°C. The waveguides are defined by e-beam writing at 15 kV acceleration voltage as described above, with a three-point alignment procedure. After exposure, the sample is baked again 3 min at 65°C and 3 min at 95°C, developed in EC solvent for 20 s and rinsed in IPA. The sample is then hard-baked for at least 5 min at 180°C, resulting in very stable waveguides, both mechanically and chemically. The large cross-section favours in-coupling of light from free-space, but also lowers the intensity level for a given input pump power, thus reducing the risk of burning the polymer when pumping.

Finally, the cleaving step is always accompanied by the formation of unwanted silicon dust particles, which are impossible to remove if they stick to the SU-8 waveguides; these particles may act as scattering points for light, again with the risk of burning the waveguides when operating at high powers. Therefore, to protect the SU-8 waveguides, we spin an additional layer of SU-8 prior to the cleaving step, and the silicon particles are simply washed away when removing this layer with EC solvent after cleaving. The SU-8 waveguides — already hard-baked — are not affected.

4.2.4 Polymer light absorption at the Raman wavelength

Our inverse taper spot-size converters work well in the 1550 nm wavelength range for both TE and TM polarisations, and we achieve coupling efficiencies similar to those measured with ZEP waveguides (see also the summary in Table 4.1 at the end of this chapter). When testing the coupling efficiency at the Raman wavelength around 1686 nm, however, we found a significant drop in the transmission spectrum of blank SU-8 waveguides, as shown in Fig. 4.11a (black curve). The same behaviour was also observed for ZEP waveguides (blue curve). We therefore tested the transmission of unpatterned polymer layers on glass substrate with a spectrophotometer, and could replicate the dip in transmission not only for SU-8, but also for S1818 and PMMA, as shown in Fig. 4.11b. In fact, it appears from the relevant literature that in this wavelength range we might be exciting a C–H vibrational mode[2] (stretching first overtone) [187, 188]. If this is truly the

[2]Note in Fig. 4.11b also the small dip in transmission around 1435 nm: this is where we would need to place the pump wavelength to obtain Raman emission at 1550 nm!

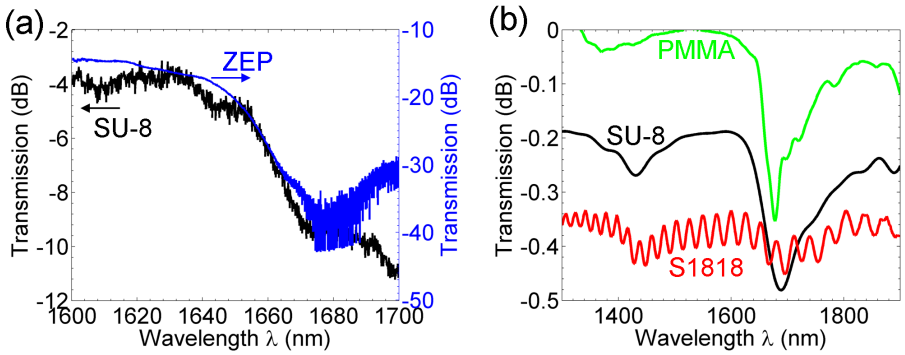

Figure 4.11: (a) Transmission curves of a ~3 mm-long SU-8 waveguide (black) and a ~2.6 mm-long ZEP waveguide (blue): both waveguides exhibit a drop in transmission around 1680 nm due to material absorption (note that the curves should not be compared in terms of values on the y-axis: the SU-8 data is normalised to that of the transmission line without sample, whereas the ZEP curve is normalised to the source). (b) Comparison of the transmission of three unpatterned polymer layers: 135 μm-thick SU-8, 30 μm-thick S1818, and ~75 μm-thick PMMA; material absorption around 1680 nm is evident.

case, we are dealing with a fundamental limitation, and any polymer would probably exhibit the same issue at this wavelength.

One possible solution would be to replace the polymer with a low-index inorganic material such as SiON [178], but we lack in-house fabrication facilities for this type of process. More simply, we can make the SU-8 waveguides very short, between 100 to 200 μm in length past the taper tip. We still need, however, to manually cleave the ~750 μm-thick substrate, which becomes very difficult if the target sample length is 1 mm or less. Therefore, we initially explored the possibility of adding a standard 3 μm-wide access waveguide between the photonic crystals and the tapers (Fig. 4.12b), which allowed us to reduce dramatically the length of SU-8 without reducing the length of the sample. These silicon access waveguides, however, produce Raman signal themselves and also generate free carriers (with the associated loss), masking the effect of the photonic crystals alone and making the results more difficult to interpret.

The final solution was to thin the substrate, thus making very short samples easier to cleave reliably (Fig. 4.12c). The thinning process involves attaching the top side of the sample to a metal rod using dental wax, then manually grinding ~400 μm off the bottom side using a 5 μm alumina particles paste mixed to water. The grinding step takes place just *after* e-beam writing and RIE etching, not before, as the sample surface must be as clean as possible when defining the photonic crystal structures. Prior to the grinding, the top side of the sample is covered with a ~1.5 μm-thick layer of PMMA (950K PMMA, spun at 900 rpm for 1 min and hard-baked 5 min

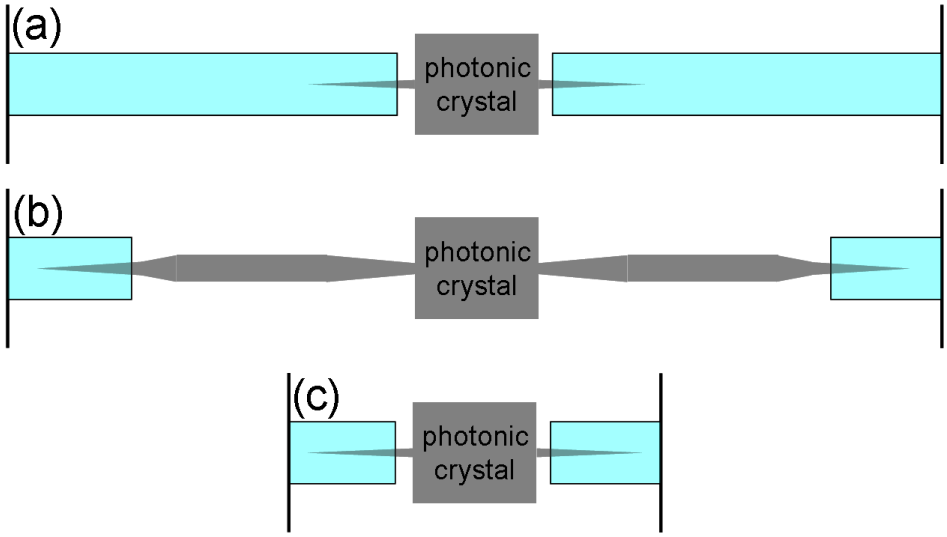

Figure 4.12: Schematic of different access waveguides with integrated spot-size converters. (a) The access waveguides are made of SU-8 and the total sample length is a few millimetres; this configuration is affected by SU-8 absorption around 1680 nm wavelength, as shown in Fig. 4.11. (b) The photonic crystals are connected to standard 3 μm-wide silicon access waveguides that terminate with inverse tapers and very short SU-8 waveguides; in this configuration, Raman signal and free-carriers are generated in the access waveguides, masking the effect of the photonic crystal. (c) Same as (a), but the substrate is thinned by ~400 μm to be able to cleave very short samples, hence reducing the SU-8 waveguide length without the need for silicon access waveguides.

at 180°C), which here has the double function to protect the top surface from the dirt of the wax and to absorb mechanical stress during the grinding: we found that around 11% of tapers break during grinding without the protective PMMA layer, whereas none of the tapers break when the PMMA layer is used. After grinding, the sample is detached from the metal rod by melting the wax, and both the wax and PMMA are removed by cleaning in trichloroethylene, followed by acetone and IPA.

4.3 Integration with photonic crystals

The main challenge we encountered when integrating spot-size converters with photonic crystal waveguides was related to the 15 min HF wet etch needed for membraning. In the standard fabrication process, as described in section 2.3, the silicon photonic crystal itself acts as a mask for membraning, and the HF (diluted 1:10 with water) penetrates through the holes to attack the underlying silica. The trenches defining the access waveguides, where silica is also exposed, are protected by a 1.8 μm-thick layer of S1818 resist

(spun at 5000 rpm for 1 min and baked 1 min at 100°C), in which windows are opened on top of the photonic crystals by UV lithography (30 to 40 s UV exposure and development in MF319 for 30 s).

As can be seen from the SEM image of Fig. 4.13a, in the transverse direction the under-etch of the silicon layer is only about 1 μm. Longitudinally, however, the only barrier to HF above the exposed access waveguide trenches is the S1818: the acid below the photonic crystal can easily penetrate under the resist and etch part of the silica up tens or even few hundreds of micrometres away from the photonic crystal (Fig. 4.13b). Even very far away from the photonic crystal a fraction of HF can infiltrate from top *through* the resist layer and attack the silica, leaving a tiny gap visible under SEM inspection below the edges of silicon (Fig. 4.13c). While this is not an issue for standard access waveguides, which from the photonic crystal widen *up* to 3 μm (Fig. 4.13b), it becomes a major problem for spot size converters, where the silicon instead tapers *down* to 100 nm or less: even the slightest under-etch causes the taper to detach (Fig. 4.13d), and the waveguide cannot be used.

We therefore needed to find a way of membraning the photonic crystal while at the same time preventing any under-etch of the tapers. We have explored the use of other resists as HF mask, such as ZEP and AZ 4562, and also the use of a vapour HF technique developed by Spurny *et al.* [189], without much success, as the acid would always somehow penetrate through the resist (Figs. 4.14a-c). Nevertheless, we found that the tapers would not detach when breaking the standard membraning process into three separate 5 min HF steps (rather than a single 15 min etch), each with the definition of a fresh S1818 mask undergoing a post-development hard-bake for 15 min at 100°C to fill any gap between the resist and the silica; following this procedure, the tapers are only partially under-etched, and still in place, as shown in Fig. 4.14d.

A better and less time-consuming option is to switch to a buffered HF solution by mixing the commercial dilute HF (1:1 in water) 1:6 with ammonium fluoride NH_4F. The mask is a thicker (~5 μm) layer of S1818 spun on the sample at 900 rpm and baked 1 min at 100°C; windows are opened on top of the photonic crystals by UV exposure for 2 min 30 s and development 30 to 40 s in MF319. The buffered HF solution, while more aggressive towards silica than the standard dilute HF, penetrates less under the resist mask when this has been hard-baked for at least 15 min at 100°C. As a result, we can obtain much sharper silica slopes, as shown in Fig. 4.14e for the case of a sample in which the initial part of the access waveguide trench was not covered by resist. If the mask window is a few micrometres smaller than the photonic crystal, the taper is not under-etched at all (Fig. 4.14f). Nevertheless, the wet etch should not last longer than 5 to 7 min, or the acid will penetrate from top *through* the resist and detach the taper tips.

With this procedure we could reach our goal of membraning the photonic

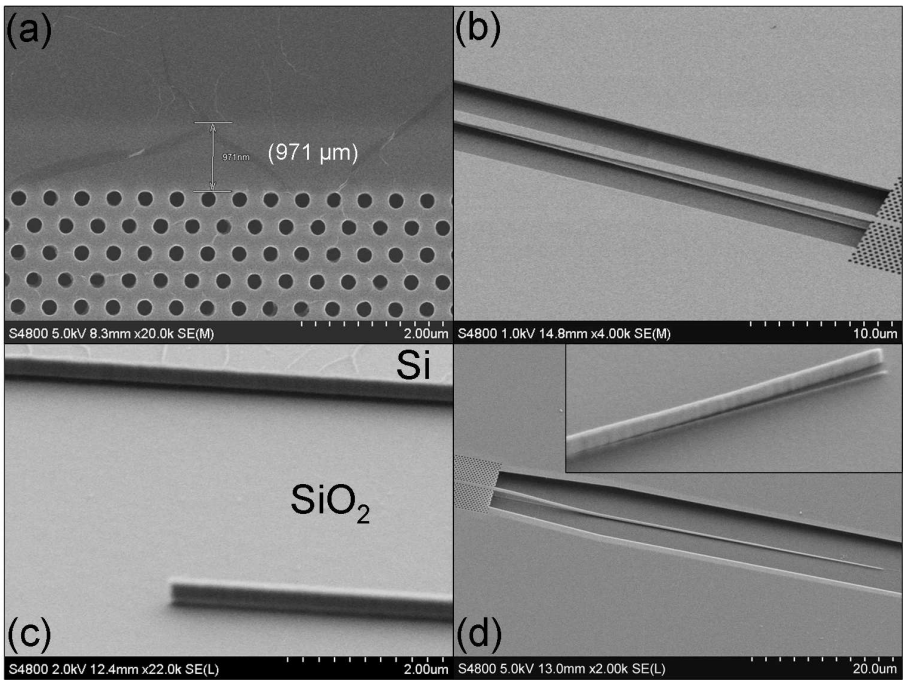

Figure 4.13: SEM images of samples after membraning in HF. (a) The under-etch of silicon is visible in the SEM as a lighter stripe around the photonic crystal. (b) Despite the use of a S1818 protective layer, the silica in the trenches of the access waveguides is partially removed by the HF penetrating under the S1818 from the photonic crystal; as a result, the access waveguides may be partially membraned. (c) SEM image of a taper tip in a sample of the type shown in Fig. 4.12b: the taper is very far from the photonic crystal, but a small gap between the silica and the silicon is clearly visible, due to some HF penetrating through the S1818 protective layer. (d) Example of a membraned and bent inverse taper. Inset: a taper tip that detached although no HF window was opened on the waveguide, indicating that the HF penetrated through the S1818 mask: the original position of the taper is still visible in the silica.

crystals without affecting the small silicon tapers, but we soon encountered repeatability issues, mainly related to the definition of the S1818 mask. We found that in many cases the membraning process is not uniform, which shows under the optical microscope and under the SEM as a rough, rather than straight, under-etch profile of the silicon in the transverse direction to the waveguide. The results can be very variable even from waveguide to waveguide within the same sample, and we think it might be due to the presence of some residual S1818 on top of the photonic crystal after development and hard bake.

In addition, the high thickness of the S1818 mask makes it difficult to align precisely and reliably the HF windows to the edges of the photonic crystal waveguide, and often a small portion of the access silicon nanowire

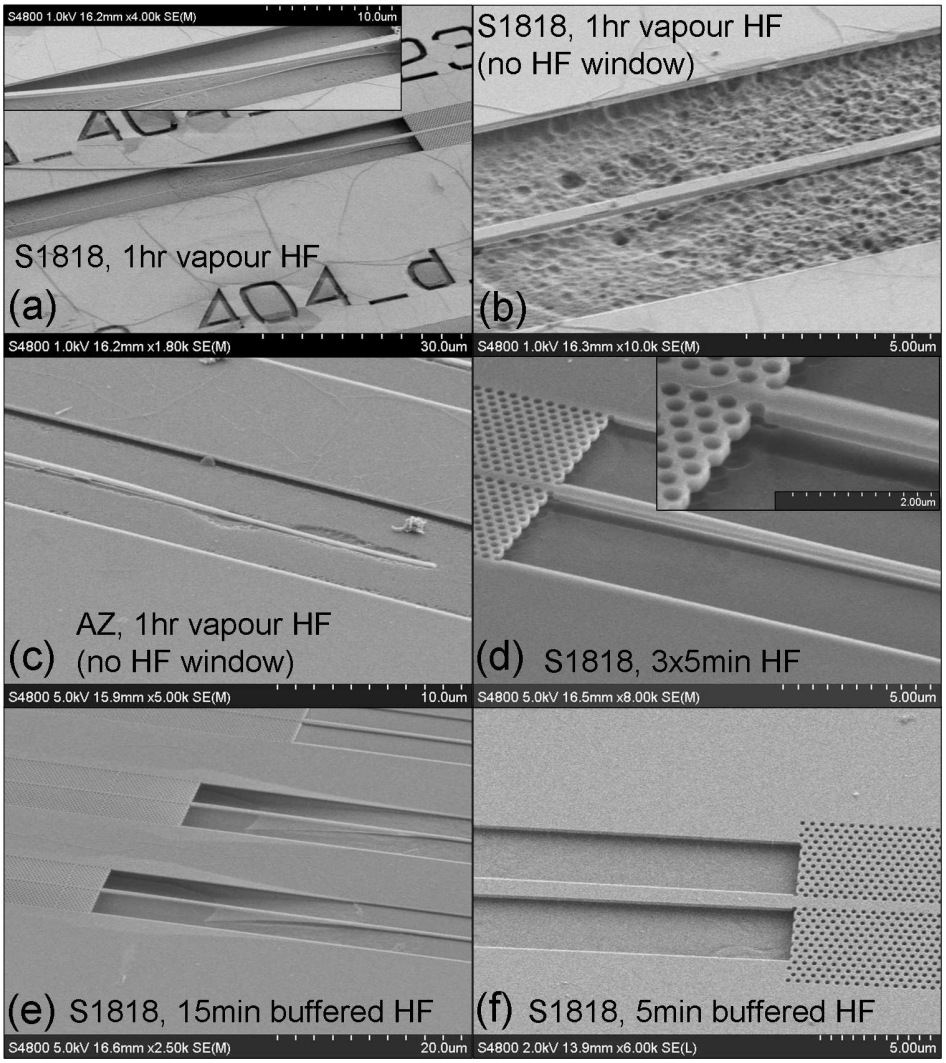

Figure 4.14: (a) Example of a detached inverse taper after 1 hr vapour HF etch, using S1818 as a mask; inset: close-up of the same sample. (b) An inverse taper of a different waveguide on the same sample as (a), in which no window was opened on top of the photonic crystal: any attack of the silica is due to HF that has penetrated through the mask. (c) Similar to (b), but for a sample covered by a thick (~10 μm) layer of AZ resist. (d) Example of membrane obtained with three 5 min HF (wet) etches, each time with a fresh and hard-baked S1818 mask: the inverse tapers do not detach; inset: close-up of the same sample. (e) Example of membrane obtained with a 15 min buffered HF etch, using a hard-baked S1818 mask (with window longer than the photonic crystal), exhibiting much steeper silica slope than with the normal HF solution. (f) Waveguide membraned with a 5 min buffered HF etch, using a hard-baked S1818 mask few micrometres shorter than the photonic crystal.

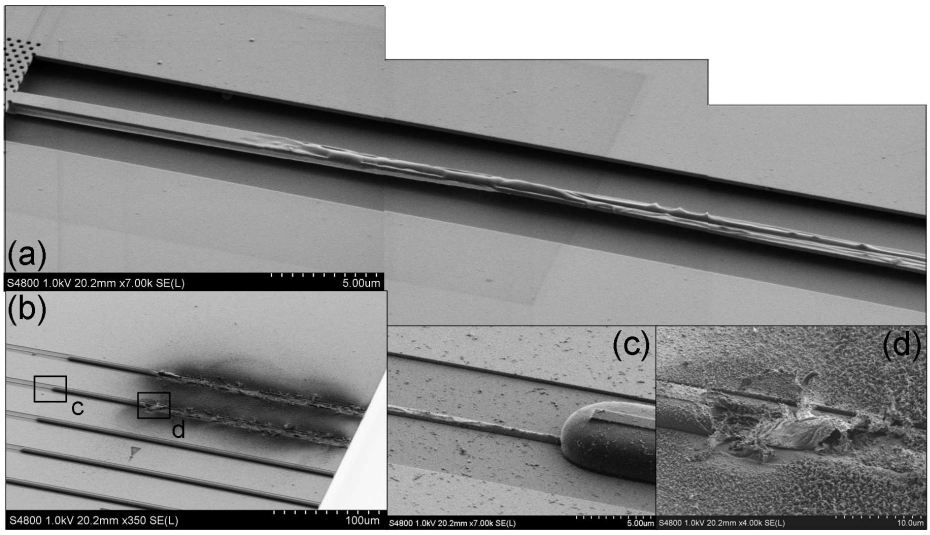

Figure 4.15: SEM images of burnt waveguides of the type shown in Fig. 4.12b. If the access waveguide or taper is at any point suspended (a), it will burn even at relatively low powers when pumping in TM polarisation, initiating a catastrophic optical damage. The failure point travels backwards, towards the source, and in some cases can reach the SU-8 waveguide (c, note that the nanowire before the SU-8 is no longer smooth due to the damage) and blow it up (d). Note that in (b) the triangular marker etched in the silicon indicates the position of the taper tip: the SU-8 right on top of the taper is not affected at all, as the light there is confined within the silicon.

is suspended just before the photonic crystal. These sections of suspended nanowire act as scattering points for TM-polarised light. When pumping in TM polarisation the silicon in these points is very likely to melt even at relatively low pump powers (Fig. 4.15a) and initiate a catastrophic optical damage: the failure point propagates backwards and can reach the SU-8 waveguide, blowing it up (Figs. 4.15b-d), and leaving us with very picturesque SEM images but a completely useless structure. On the contrary, nanowires that were not suspended gave no problems even at much stronger TM pumps. This issue highlights the importance of a very good and repeatable alignment of the HF window to the edge of the photonic crystal to prevent any membraning of the input nanowire.

Both the problems of under-etch uniformity and alignment of the HF windows were improved by replacing the S1818 with a ~1.5 μm-thick layer of PMMA, in which windows are defined through e-beam lithography in a three-point alignment procedure, rather than through the less-precise UV technique used for S1818. After exposure, the PMMA is developed 1 min in a solution of water and IPA 3:7, and hard-baked 5 min at 180°C and 1 hr at 200°C. Through use of PMMA as an HF mask we obtained more uniform and more reliable membraning. The PMMA is then removed by rinsing in

acetone and IPA.

4.4 Summary: complete fabrication procedure

The overall fabrication process required for the definition of photonic crystal waveguides with integrated spot-size converters involves three e-beam steps, grinding, dry and wet etch, and is summarised as follows:

1. The photonic crystal pattern and inverse tapers are defined in ZEP resist through e-beam lithography and transferred into silicon by RIE, as described in section 2.3. The inverse tapers are designed as described in section 4.2.1. After etching, ZEP is removed by soaking in dimethylformamide (DMF), acetone and IPA.

2. The sample is covered with a protective 1.5 μm-thick PMMA layer and thinned by 400 μm through manual grinding. The PMMA and the wax used for grinding are removed by soaking in trichloroethylene, acetone and IPA. Note that the grinding step is not a general requirement for integration of spot-size converters: it becomes necessary only due to absorption of SU-8 around the 1686 nm wavelength range of interest for Raman scattering with a 1550 nm pump. There is no need to grind the sample when operating at other wavelengths.

3. The sample is again covered with a 1.5 μm-thick PMMA layer, and windows are opened by e-beam writing, following a three-point alignment procedure. After development and hard-bake of the PMMA mask, the photonic crystals are membraned by soaking 5 to 7 min in a buffered HF solution. The PMMA layer is then removed with acetone and IPA.

4. The sample is cleaned in piranha solution for 5 min, before spinning a ~2.3 μm-thick SU-8 layer, which is exposed by e-beam writing, following a second three-point alignment procedure. After development, the SU-8 is hard-baked for at least 5 min at 180°C. The addition of the SU-8 waveguides should always be left as the last step before cleaving: the less fabrication steps are performed *after* the definition of the SU-8 waveguides, the less likely they are to accumulate unwanted dirt particles, which may scatter light and burn the structure when pumping.

5. The sample is finally spun with another 2.3 μm-thick SU-8 layer, to protect the polymer waveguides from silicon dust particles generated while cleaving. After cleaving, this SU-8 layer is removed by soaking in EC solvent, acetone and IPA.

Table 4.1: Summary of spot-size converters coupling efficiencies from various works.

Spot-size converter	Tip width (nm) Wire width (nm)	Taper length (μm) SOI thickness (nm)	Low n material Dim. $w \times h$ (μm^2)	Coupl. TE (dB) Coupl. TM (dB)	Coupling mode
Shoji *et al.* [177]	60 300	200 300	polymer 3×3	-0.8 –	4.3 μm diam. high-index fibre
Almeida *et al.* [179]	100 470	40 270	air –	-3 -1.65	5 μm diam. high-index fibre
McNab *et al.* [180]	75 450	150 220	epoxy 2×3	-0.5 -1	4.2 μm diam. tapered-lensed fibre
Tsuchizawa *et al.* [178]	80 300 or 400	200–300 μm 300 or 200	SiON 3×3	-0.5 –	4.3 μm diam. high-index fibre
as above	as above	as above	as above	-2.5 –	9 μm diam. conventional fibre
Foster *et al.* [22]	– –	– –	air –	-12.8 –	tapered-lensed fibre
Pu *et al.* [181]	12 452	300 238	SU-8 3.4×3.4	-0.36 -0.66	2.9 μm diam. tapered fibre
This work	100 750	100 220	SU-8 6×2.3	-2.8 -3.17	free-space

Chapter 5

Raman scattering in slow light waveguides

In Chapter 3 we have explored the possibility of using an engineered slow light photonic crystal waveguide to realise a dynamic optical delay line. One of the main limitations of our proposed structure, and of slow light devices in general, is the increase of propagation loss with group index. The suppression of such loss, or at least its compensation by introducing optical gain, would allow for longer slow light devices, and thus longer (tunable) delays, virtually lossless.

Achieving optical gain in silicon, however, is one of the greatest challenges in silicon-based photonics [10], since bulk silicon is an indirect band gap semiconductor and therefore has a very low light emission efficiency via carrier recombination. Due to crystal symmetry, second-order nonlinear optical effects are highly suppressed [32], but third-order nonlinear phenomena do take place, with a number of potential practical applications [190].

One important example is stimulated Raman scattering: due to the single-crystal structure, the Raman gain coefficient in silicon is several orders of magnitude higher than in amorphous glass fibres [191], which have already established as a platform for Raman amplifiers and lasers; in addition, the large index contrast and the sub-micrometre scaling possibilities offered by SOI substrates allow for tight optical confinement and consequently high optical intensities. Stimulated Raman scattering may therefore become strong enough to be used as a means of achieving optical gain in silicon waveguides, and in fact it has been employed to demonstrate optical amplifiers [33] and even lasers [37, 38] in silicon.

In this chapter we will start by introducing the fundamental concepts about spontaneous and stimulated Raman scattering, and reviewing some of the main results reported to-date regarding Raman optical amplification and lasing in silicon waveguides. Then, starting from the existing literature studies on slow-light enhancement of the Raman interaction, we will discuss

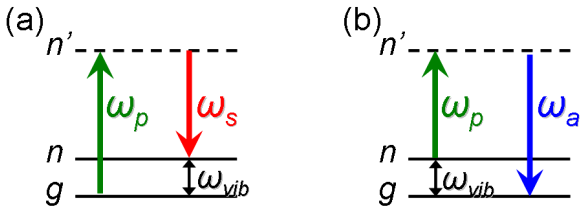

Figure 5.1: Energy level diagrams describing (a) Raman Stokes scattering from the ground state g to an excited vibrational level n and (b) Raman anti-Stokes scattering from n to g. Here n' indicates an intermediate virtual state.

a simple model describing the total signal gain dependence on the slowdown factor, which we will also expand to include the effect of realistic propagation losses and mode profiles in photonic crystal waveguides. Finally, we will present some experimental results on both spontaneous and stimulated Raman scattering from engineered slow light waveguides.

5.1 Spontaneous vs stimulated Raman scattering

Spontaneous Raman scattering in a generic material system is the result of the interaction of light with vibrational modes, and can be understood through use of energy level diagrams, as shown in Fig. 5.1 [32]. The absorption of an incident pump photon at frequency ω_p promotes the transition of the system from the ground state g to an excited vibrational state n by means of an intermediate transition to a virtual level n' and emission of a red-shifted photon at frequency ω_s (Fig. 5.1a). Alternatively, a transition from n to g may occur, with n' serving as the intermediate level and with emission of a blue-shifted photon at frequency ω_a (Fig. 5.1b). The two processes are called Raman Stokes and Raman anti-Stokes scattering, respectively. The former is usually much stronger than the latter because, in thermal equilibrium, the population of level n is smaller than the population of level g.

Equivalently, the Raman effect *in crystals* can be described in terms of an interaction with optical phonons[1] [191,192]: Stokes scattering corresponds to annihilation of an incident pump photon with emission of an optical phonon and a Stokes photon (Fig. 5.2a), whereas anti-Stokes scattering corresponds to emission of an anti-Stokes photon by means of absorption of an optical phonon (Fig. 5.2b). Energy conservation requires the frequency difference between the pump and the Raman photon to be equal to the frequency of the phonon. In addition, since the wavevector of photons at our wavelengths of interest is 3–4 orders of magnitude smaller than that of phonons, momentum conservation imposes that only phonons from the centre of the Brillouin zone

[1]Scattering from acoustical phonons is known as Brillouin scattering.

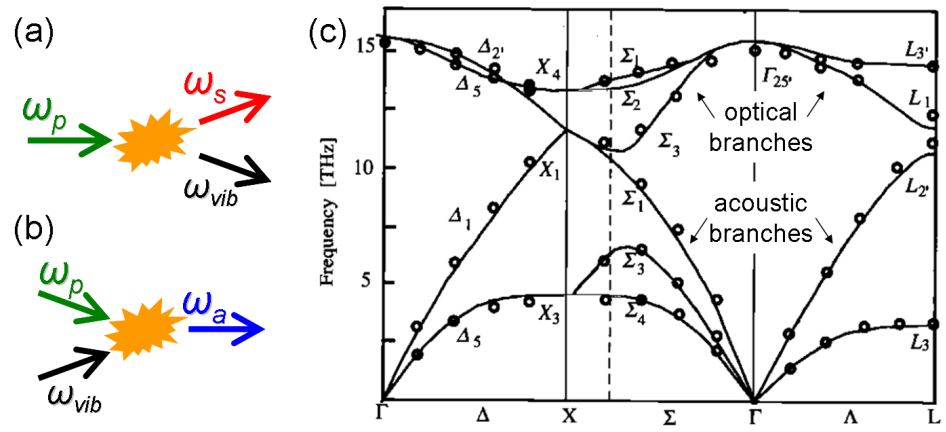

Figure 5.2: (a-b) Schematic of Raman scattering in crystals: an optical phonon is (a) emitted in Stokes scattering and (b) absorbed in anti-Stokes scattering. (c) Phonon spectrum of silicon; the three optical branches are degenerate in the Brillouin zone centre Γ at frequency 15.6 THz, which determines the Raman shift of silicon (from [192]. Reprinted with permission).

can participate in the scattering. In the particular case of crystalline silicon, the phonon spectrum consists of three optical and three acoustical branches, as shown in Fig. 5.2c [192]. The degeneracy of the three optical branches at the zone centre means that there is only one strong channel for Raman scattering at 15.6 THz.

The spontaneous Raman effect is a quite weak process, but under excitation by an intense pump beam, efficient scattering to the Stokes frequency can occur as a result of its stimulated version, a third-order nonlinear effect. While the detailed classical model of stimulated Raman scattering may be found in Ref. [32], here we will limit ourselves to a brief qualitative description of the process. With reference to the notation of Fig. 5.1a, the material vibrations at frequency ω_{vib} give rise to Raman scattering, generating Stokes components starting from an incident pump at frequency ω_p. The Stokes field beats with the pump field, resulting in a modulation of the total intensity at frequency $\omega_p - \omega_s = \omega_{vib}$, which coherently excites vibrations at the resonance frequency ω_{vib}. Therefore, the two processes reinforce each other: the beating of the pump and Stokes signals leads to stronger vibrations, which lead to a stronger Stokes field, which in turn leads to stronger vibrations, and so on. This positive feedback results into exponential growth of the amplitude of the Stokes wave, which may have been generated spontaneously in the material (resulting in amplified spontaneous emission) or injected externally as a seed signal to be amplified.

We are particularly interested in the case of a probe signal subject to Raman gain inside a silicon waveguide. The evolution of the signal along the waveguide will be determined by the superimposition of a positive gain

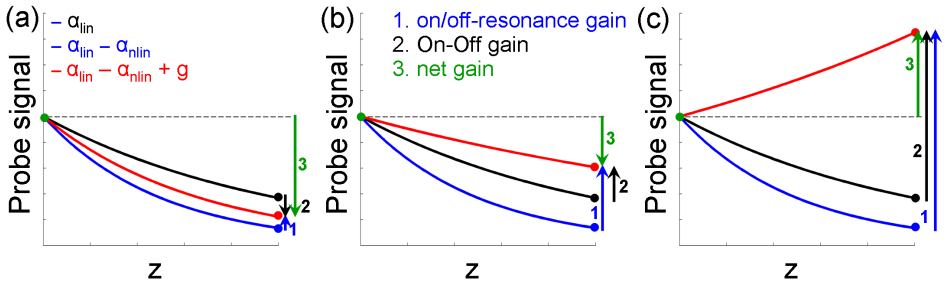

Figure 5.3: Comparison of three different definitions of gain in three possible situations. The black curve indicates the signal evolution along the waveguide propagation direction (z) subject only to the intrinsic linear propagation loss of the system $-\alpha_{lin}$ (i.e. without pump). The blue curve indicates the evolution of an off-resonance signal in the presence of the pump: the signal is subject to both linear and nonlinear losses $-\alpha_{lin} - \alpha_{nlin}$. The red curve indicates the evolution of an on-resonance signal in the presence of the pump: the signal is subject to Raman gain in addition to linear and nonlinear losses $-\alpha_{lin} + g - \alpha_{nlin}$. The on/off-resonance gain compares the red curve output with the blue curve output. The On-Off gain compares the red and the black outputs. The net gain compares the red curve output to the input. (a) The pump introduces more loss than gain: the on/off-resonance gain is positive, but both On-Off and net gains are negative. (b) The pump introduces more gain than loss: the output with pump is greater than without pump, but still lower than the input; the On-Off gain is positive, but the net gain is negative. (c) The Raman gain also compensates for propagation loss and the system performs as an amplifier: the net gain is positive.

term g arising from the Raman effect with the negative contributions due to the linear propagation loss α_{lin} intrinsic to the system and the nonlinear loss α_{nlin} induced by the presence of the pump: $-\alpha_{lin} + g - \alpha_{nlin}$. It is important to identify these terms here in order to better understand the different definitions of gain reported by different authors in the literature (Fig. 5.3).

Some works report the on/off-resonance gain, that is, a gain calculated as the ratio between the on-resonance and off-resonance output probe (i.e. the probe wavelength matches or not the correct spacing from the pump wavelength, as dictated by the 15.6 THz Raman shift), with the pump present in both cases: this definition compares the Raman output ($-\alpha_{lin} + g - \alpha_{nlin}$) to a signal undergoing both linear and nonlinear losses but no gain ($-\alpha_{lin} - \alpha_{nlin}$), and therefore will always result in positive values[2] (Fig. 5.3a). Such gain differs from the On-Off gain, which is instead the ratio between the output probe power while the pump is on and the output probe power while the pump is off, with the probe being always on-resonance: this second definition compares the Raman output to a signal

[2]Red- or blue-shifts of a waveguide's dispersion due to thermal heating and carrier injection, respectively, may complicate further this simple picture.

undergoing only the *intrinsic* propagation loss ($-\alpha_{lin}$), and gives a better indication on whether the system performs well when pumped. For example, if the pump introduces -5 dB of nonlinear loss and $+1$ dB of Raman gain (Fig. 5.3a), the on/off-resonance gain is $+1$ dB: this definition is useful to isolate the positive Raman contribution, and coincides with the Raman gain itself. In our example, however, the On-Off gain will be -4 dB, meaning that the pump brings more harm that benefit. If instead the On-Off gain is positive, it is indeed convenient to pump the system (Fig. 5.3b). As the On-Off gain measures the balance between the positive and negative nonlinear contributions, it is sometimes referred to as net nonlinear gain [193] or net internal gain [194].

An even better measure of the system's performance is ultimately given by the net gain, i.e. the ratio between the Raman output signal to the input probe signal, which takes into account also the linear loss: if the On-Off gain is $+4$ dB but the propagation loss amounts to -6 dB, the net gain will be -2 dB (Fig. 5.3b). A positive net gain (Fig. 5.3c) is therefore the indication that the system performs as an optical amplifier. Even better, if the net gain is reported as fibre-to-fibre, then the Raman gain overcomes not only the linear and nonlinear losses, but also the in- and out-coupling losses.

If we want to achieve loss compensation in a silicon waveguide, we need to aim at least to the situation illustrated in Fig. 5.3b.

5.2 Ten years of Raman in silicon waveguides

The first experimental observation of *spontaneous* Raman scattering in SOI waveguides was reported in 2002 by Claps *et al.* [195]: the authors measured a few picowatts Stokes signal at 1542 nm generated from a 5×5 μm^2 rib waveguide, with continuous-wave coupled pump powers between 63 and 588 mW at 1430 nm (Fig. 5.4a). The following year the same authors reported the first observation of *stimulated* Raman scattering from a 1.8 cm-long SOI waveguide in a pump-probe experiment [196], recording a 0.25 dB On-Off signal gain at 975 mW continuous-wave coupled pump power. Their experiment was then followed by measurement of a 2.3 dB on/off-resonance gain by Rong *et al.* of Intel Corporation [197] from a smaller (1.52×1.45 μm^2) rib waveguide of length 4.8 cm, at lower input pump power (270 mW in-coupled).

Raman conversion at even lower pump powers was observed by further shrinking the size of the silicon waveguide to a 220×445 nm^2 nanowire, as the smaller cross-section provides stronger optical confinement and enhances the pump intensity. Dadap *et al.* reported measurements of spontaneous Raman scattering from silicon nanowires for coupled pump powers between 5 and 35 mW [198], as well as a 0.7 dB On-Off gain at 29 mW coupled pump power over only 4.2 mm length [199]. An On-Off gain of 3.1 dB from a 7 mm-long

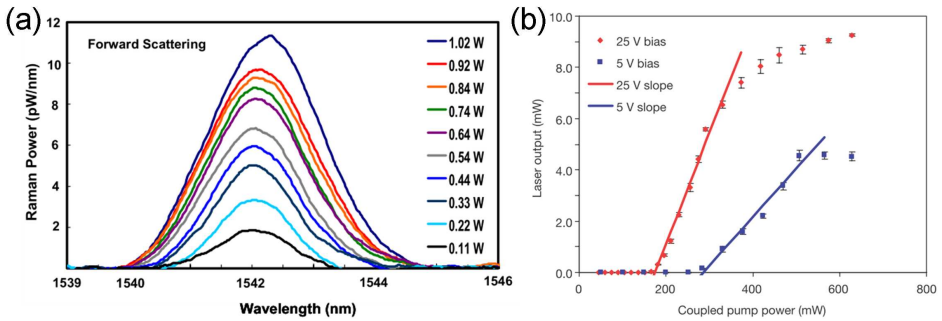

Figure 5.4: (a) First observation of spontaneous Raman scattering in silicon waveguides: measured Raman spectra for different pump powers; the pump power values are measured before the coupling optics (from [195]. Reprinted with permission). (b) First demonstration of a continuous-wave silicon Raman laser: laser output power (1686 nm) as a function of input pump power (1550 nm) at reverse bias of 5 and 25 V (from [38]. Reprinted by permission from Macmillan Publishers Ltd: Nature [38], copyright 2005).

nanowire was also observed by Xu *et al.* in pulsed operation, with 2.8 W coupled peak power [194], which was then increased to 6 dB at the expense of 20.7 W peak pump power [193].

It soon became clear, however, that the amount of achievable net gain in silicon channel waveguides via Raman scattering is limited by the nonlinear loss associated with two-photon absorption (TPA). Even though linear optical absorption at telecom wavelengths is small in silicon, due to the energy band gap being higher than the energy of a single photon, two-photon absorption may take place, resulting in pump depletion. This effect, if considered alone, is relatively weak [200], but due to the long carrier recombination lifetime (of the order of tens of nanoseconds [197, 201]) a noticeable free-carrier population builds up under high power pumping conditions, and the resulting free-carrier absorption (FCA) can largely surpass the Raman gain [197,201–207]. Therefore, reduction of the free-carrier lifetime is crucial to achieve net optical gain in silicon [207].

Net optical gain due to stimulated Raman scattering in a SOI waveguide was first reported in 2004 by Liu *et al.* of Intel Corporation [33]. The authors used a pulsed pump with a pulse-width narrower than the carrier lifetime, which significantly reduced the free-carrier generation due to TPA, resulting in a net gain of 2 dB from a 4.8 cm-long waveguide with a 470 mW coupled peak pump power. Soon after, Liang and Tsang reported a 7.8 dB net Raman gain, corresponding to 6.8 dB net fibre-to-fibre gain for 55 W coupled peak power [34], followed just a week later by 11.4 dB net fibre-to-fibre gain for 40 W coupled peak power reported by Boyraz and Jalali [35].

In less than a year Intel Corporation also achieved the first demonstration of net optical gain for continuous-wave pumping [208]. In this case, nonlinear

optical loss was successfully reduced by introducing a 25 V reverse-biased p-i-n diode across the waveguide to sweep carriers out, thus reducing their effective lifetime to 1 ns and generating a net gain of more than 3 dB in a 4.8 cm-long waveguide at 700 mW coupled pump power.

Demonstrations of silicon lasers soon followed the observation of net optical gain: the very first Raman laser emission from silicon was reported in late 2004 by Boyraz and Jalali [36], through use of a pulsed pump, with a silicon waveguide being the gain medium in a 8 m-long optical fibre ring cavity. Early in 2005 Rong *et al.* demonstrated pulsed Raman lasing in a single silicon waveguide acting as a Fabry-Perot cavity [37], establishing the viability of integrated silicon active elements; their structure included a reverse-biased p-i-n junction to reduce carrier lifetime. Only a month later, the same authors also demonstrated continuous-wave lasing operation in a similar device [38] (Fig. 5.4b).

All these pioneering experiments — with only two years from the first observation of spontaneous Raman scattering to the demonstration of lasing — show that Raman scattering gives the possibility to use silicon as an ultracompact on-chip gain medium at the desired telecommunications frequencies: since the Raman shift depends only on the frequency of the optical phonons, light emission may be tuned over all wavelength ranges of interest, as long as the corresponding pump wavelength is available and silicon is still transparent.

In the following years, the most important experimental results on Raman scattering as a means of achieving amplification and lasing in silicon were again dominated by Rong *et al.*, with the demonstration of ring-resonator Raman lasers with thresholds as low as 20 mW [209,210], as well as a cascaded Raman laser [211] and a 2.3 dB On-Off gain at 40 Gb/s [212]. A silicon Raman amplifier at mid-infrared wavelengths was also demonstrated by Raghunathan *et al.* [213]. At the same time, a number of theoretical and numerical studies have been presented which provide a better understanding of the role of the many different parameters involved in determining the performance of silicon Raman lasers [214, 215], pulsed [207, 216–218] and continuous-wave [206, 219–222] Raman amplifiers, and third-order optical phenomena in SOI waveguides in general [223].

Despite the initial rapid advance, however, the development of silicon Raman amplifiers and lasers has seen no major improvement since 2007 [210] and various authors have moved their efforts away from amplification and lasing, with Raman scattering being investigated as a means of achieving other functionalities, most notably all-optical isolation [224–228] and Raman-induced slow light [132, 229], but also all-optical logic gates [230], switches [231] and polarisers [232].

One critical issue is that the potential advantage of Raman amplification being an all-optical effect turns into a disadvantage when it requires high pump powers to be achieved: the initial estimates of the Raman gain

coefficient g_R in silicon were excitingly promising, with reported values up to 76 cm/GW [195], but various experimental results suggest lower values, in the range 4.2–57 cm/GW [198,225,233].[3] The need for high pump power lasers accompanying the lower-than-expected values of g_R, together with the issue of free-carrier absorption, mean that silicon Raman sources and amplifiers still cannot find a practical use, at least in the framework of integrated photonic chips, unless the required pump powers and operating lengths can be significantly reduced. Such a challenge has motivated researchers to investigate ways to enhance the Raman interaction, in particular through use of photonic crystal slow light waveguides and photonic crystal nanocavities (we will briefly discuss cavities in Chapter 6).

5.2.1 Slow light enhancement of the Raman interaction

As already mentioned in section 2.2, slow modes in photonic crystal waveguides may be used to access the nonlinear optical properties of the material at lower input pump powers: when light enters a slow mode, the small group velocity increases the light-matter interaction time [39], resulting in a higher probability of a Raman scattering event.

Enhancement of the Raman nonlinearity at the band-edge of generic photonic crystal structures was reported in some early theoretical and numerical studies for both spontaneous [234,235] and stimulated scattering [236,237], but the effect of slow light on stimulated Raman scattering for the specific case of SOI photonic crystal slab waveguides was first theoretically investigated by McMillan *et al.* [40]. In their important contribution, the authors explicitly demonstrate an inverse dependence of the Stokes mode gain on the group velocities of *both* the pump and the Stokes modes and predict an enhancement of the Raman interaction by a factor on the order of 10^4. This prediction, however, corresponds to using group indices of around 6×10^3: group indices of almost 10^3 have indeed been demonstrated [53], but values in excess of 100 are usually already too high to be useful from a practical point of view, mainly due to the associated large propagation loss (see section 2.2). In addition, the study by McMillan *et al.* considers the Raman interaction alone, but not the competing processes of TPA and FCA, which are also enhanced due to slow light.

A simple model to evaluate the scaling of the total signal gain with slowdown factor by taking into account both losses and Raman gain was proposed by Krause *et al.* [238] and further developed by Rukhlenko *et al.* [239]. Both works indicate that, ultimately, the highest gains are achievable when slowing down the signal while the keeping the pump in a relatively fast mode, due to the fact that slowing down the pump incurs into higher free-carrier density and thus higher associated loss. Krause *et al.* predict signal gains

[3] The reason for such a large variation in reported values is not clear.

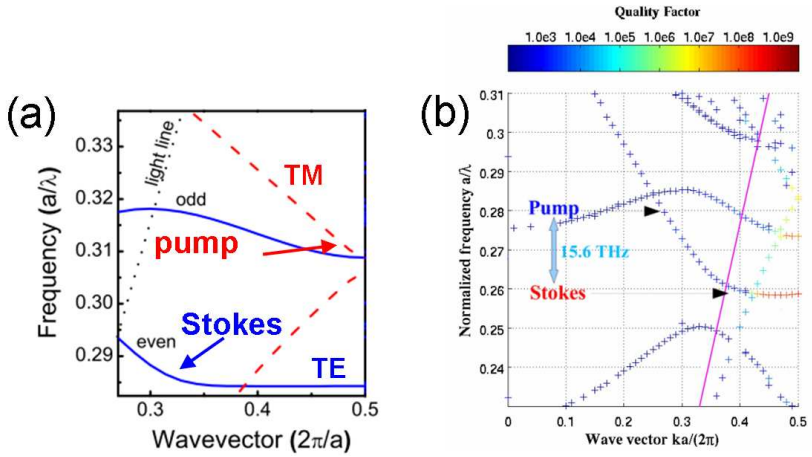

Figure 5.5: Band structure configurations adopted for the demonstration of enhanced spontaneous Raman scattering in the experiments by (a) McMillan *et al.* (slow pump, fast signal) (from [41]. Reprinted with permission from [41]. Copyright 2008, American Institute of Physics.) and (b) Checoury *et al.* (fast pump, slow signal) (from [42]. Reprinted with permission).

up to 13 dB for 200 mW continuous-wave pump power, whereas Rukhlenko *et al.* estimate gains higher than 100 dB for much stronger pump intensities, on the order of 1 GW/cm^2 (corresponding to 5 W for a 0.5 μm^2 effective area). As we will see in sections 5.3.1 and 5.3.2, however, the model developed in Refs. [238, 239] does not account for realistic linear propagation losses and mode areas in photonic crystal waveguides: while useful to understand the interplay of the different loss and gain terms with varying slowdown factor, it cannot be directly applied for quantitative estimates of the achievable Raman gains in this particular type of slow light structures.

From an experimental point of view, band-edge slow light enhancement of the spontaneous Raman scattering in photonic crystal waveguides was demonstrated in 2008 by Inoue *et al.* [240] in an AlGaAs W1 waveguide with a *fast pump* and *slow signal* configuration and by McMillan *et al.* [41] in a silicon W1 with *slow pump* and *fast signal* (Fig. 5.5a). In both cases the waveguides were 1 mm long. The following year Checoury *et al.* [42] reported similar results in a non-membraned 50 μm-long SOI W1 waveguide, despite extremely high pump propagation loss (~700 dB/cm) due to the pump mode being above the silica light line (Fig. 5.5b). The same authors reported in 2010 observation of stimulated Raman scattering in the form of amplified spontaneous emission from a membraned W0.66 silicon photonic crystal waveguide [241]: the waveguide, only 50 μm long, acted as a Fabry-Perot cavity doubly resonant at the pump and signal wavelengths. Their work, however, did not report any study of slow light enhancement and to-date there has been no record of stimulated Raman scattering in the form

of a pump-probe experiment in slow light waveguides.

Our aim is to complement the theoretical understanding of stimulated Raman scattering in slow light photonic crystal waveguides in order to establish the level of performance that can *realistically* be achieved by these structures, in particular in terms of loss compensation of a slow mode, and then to demonstrate loss compensation in a pump-probe experiment.

5.3 Raman gain in slow light photonic crystals

Understanding the interplay of the various quantities involved in stimulated Raman scattering is the first step to establish the viability of Raman gain as a loss compensation effect in slow light waveguides. In this section we will discuss the scaling of the total signal gain with the slowdown factor of both the pump and the Stokes waves. We will begin in section 5.3.1 by considering the simplified model proposed in Ref. [238], which we will analyse from a new angle, as it will help us to better appreciate the properties of realistic photonic crystals discussed in section 5.3.2.

5.3.1 Simplified model

Consider the case of a continuous-wave signal at frequency ω_s co-propagating in a silicon waveguide with a continuous-wave pump at frequency ω_p (Fig. 5.6), where the frequency difference between the pump and the signal exactly corresponds to the Raman shift of silicon $\omega_p - \omega_s = \Omega_R = 2\pi \times 15.6$ THz. The evolution of the pump and signal powers $P_p(z)$ and $P_s(z)$, respectively, along the waveguide propagation direction z obeys the following equations [207, 238]:

$$\frac{\partial P_p}{\partial z} = -\alpha_p P_p - \frac{\beta_{\mathrm{TPA}}}{A_{eff,p}} P_p^2 - \frac{\sigma_p \beta_{\mathrm{TPA}} \tau_c}{2\hbar\omega_p A_{eff,p}^2} P_p^3, \tag{5.1}$$

$$\frac{\partial P_s}{\partial z} = -\alpha_s P_s + \frac{g_R - 2\beta_{\mathrm{TPA}}}{A_{eff,p}} P_p P_s - \frac{\sigma_s \beta_{\mathrm{TPA}} \tau_c}{2\hbar\omega_p A_{eff,p}^2} P_p^2 P_s, \tag{5.2}$$

where the first term on the right-hand side describes linear propagation loss $\alpha_{p,s}$; the second term accounts for TPA of two pump photons in Eq. (5.1) through the TPA coefficient β_{TPA} and, in Eq. (5.2), crossed-TPA of a signal and a pump photon and stimulated Raman scattering through the Raman gain coefficient g_R, with $A_{eff,p}$ the effective area of the pump mode; the third term describes FCA induced by free carriers generated through TPA of the pump, with $\sigma_{p,s}$ the FCA coefficients for pump and signal and τ_c the free-carrier lifetime. Equations (5.1) and (5.2) are written under the strong pump assumption $P_p \gg P_s$: in this approximation, pump depletion through the Raman process is negligible, as well as the effect of signal-generated free carriers [206, 207, 238].

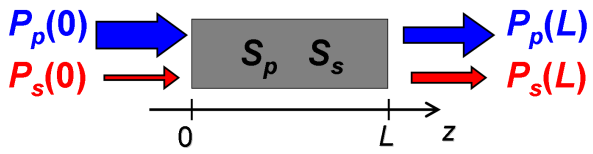

Figure 5.6: Schematic of the model for stimulated Raman scattering: a signal beam is amplified by co-propagating with a pump in a silicon slow light waveguide.

When considering propagation in a slow light waveguide, we need to modify Eqs. (5.1) and (5.2) to include the slowdown factors S_p and S_s of the pump and signal, respectively. While the slowdown factor is often defined as the ratio of the group index n_g over the phase index [39], we will use here the definition adopted by Monat *et al.* [62], $S = n_g/n_{\mathrm{Si}}$, with $n_{\mathrm{Si}} = 3.48$ the refractive index of silicon; this definition is more appropriate in our case, as we aim to scale quantities that are defined for bulk silicon [239]. This will become clearer in section 5.3.2. The effect of slow light in our case can be accounted for with simple scaling rules [62, 238]. If the linear loss is dominated by out-of-plane scattering, we can assume the first term on the right-hand side of Eqs. (5.1) and (5.2) to scale linearly with the slowdown factor [31], $\alpha_{p,s} = \kappa_{p,s} S_{p,s}$, with κ the loss per unit slowdown factor. Such a linear scaling has been recently demonstrated for group indices up 60 [76]. TPA and the Raman interaction scale quadratically with the slowdown factor [40, 238], whereas FCA scales with its third power. Therefore, in a slow light waveguide, Eqs. (5.1) and (5.2) are modified as follows:

$$\frac{\partial P_p}{\partial z} = -\kappa_p S_p P_p - \frac{\beta_{\mathrm{TPA}}}{A_{eff,p}} S_p^2 P_p^2 - \frac{\sigma_p \beta_{\mathrm{TPA}} \tau_c}{2\hbar\omega_p A_{eff,p}^2} S_p^3 P_p^3, \tag{5.3}$$

$$\frac{\partial P_s}{\partial z} = -\kappa_s S_s P_s + \frac{g_R - 2\beta_{\mathrm{TPA}}}{A_{eff,p}} S_p S_s P_p P_s - \frac{\sigma_s \beta_{\mathrm{TPA}} \tau_c}{2\hbar\omega_p A_{eff,p}^2} S_p^2 S_s P_p^2 P_s. \tag{5.4}$$

The strong pump assumption in this case is valid as long as $S_p P_p \gg S_s P_s$ [239].

For convenience, we shall express Eq. (5.4) in terms of a local signal gain function G_s:

$$\frac{\partial P_s}{\partial z} = S_s G_s(S_p P_p)\, P_s, \tag{5.5}$$

where

$$G_s(S_p P_p) = -\kappa_s + \frac{g_R - 2\beta_{\mathrm{TPA}}}{A_{eff,p}} S_p P_p - \frac{\sigma_s \beta_{\mathrm{TPA}} \tau_c}{2\hbar\omega_p A_{eff,p}^2} (S_p P_p)^2 \tag{5.6}$$

is a quadratic function of the product $S_p P_p$. Figure 5.7 shows how the overall local gain $S_s G_s$ varies as a function of P_p for different values of S_s and S_p. Other parameters are summarised in Table 5.1.

Table 5.1: Values of the parameters used in the calculations of signal gain.

Parameter name	Symbol	Value	Reference
Stokes wavelength	λ_s	1550 nm	–
Input pump power	P_{p0}	250 mW	–
Loss per unit S	$\kappa_{s,p}$	4 dB/cm/S	[31]
		= 1.15 dB/cm/n_g	
Effective area	$A_{eff,p}$	0.4 μm^2	[59]
Raman gain	g_R	20 cm/GW	[195, 196, 199]
TPA coefficient	β_{TPA}	0.5 cm/GW	[207, 242, 243]
Carrier lifetime	τ_c	200 ps	[169, 244]
FCA coefficient	$\sigma_{s,p}$	1.45×10^{17}	[206]
		$\times(\lambda_{s,p}/1550\mathrm{nm})^2\mathrm{cm}^2$	

It is easy to see from Eqs. (5.4–5.6) that while S_p and S_s have the same impact on the Raman gain term alone [40], they play very different roles on the total signal evolution when also loss scaling is taken into account: whether the signal experiences gain or loss is determined by the sign of G_s, which is only a function of S_pP_p, and does not depend on S_s. The signal slowdown factor S_s is an overall multiplying factor to the local gain, and therefore only enhances the already experienced gain or loss: this is shown in Fig. 5.7a as a scaling along the y-axis, and it also implies that S_s has no influence on the two zeros of the parabola, which represent the transparency pump threshold P_{th} (lowest zero) and the optimal pump power P_{opt} (highest zero) [238]:

$$\frac{P_{th,opt}}{A_{eff,p}} = \frac{g_R - 2\beta_{\mathrm{TPA}}}{S_p}\frac{\hbar\omega_p}{\sigma_s\beta_{\mathrm{TPA}}\tau_c}\left[1 \pm \sqrt{1 - \frac{2\kappa_s}{(g_R - 2\beta_{\mathrm{TPA}})^2}\frac{\sigma_s\beta_{\mathrm{TPA}}\tau_c}{\hbar\omega_p}}\right]. \tag{5.7}$$

Slowing down the pump by a factor S_p, on the other hand, corresponds to scaling the local gain function G_s along the x-axis, as shown in Fig. 5.7b. The transparency threshold P_{th} and optimal pump power P_{opt} scale as $1/S_p$ (Eq. (5.7)), and therefore the pump slowdown factor determines for which pump powers the local gain is positive.

Let us now consider how the shape of G_s translates into the evolution of the signal power $P_s(z)$ along the waveguide. Figure 5.8b shows the case for fast pump and signal, $S_p = S_s = 1$. We assume that the input pump power $P_{p0} \equiv P_p(0)$ is such that the initial local gain $G_s(P_{p0})$ is positive (Fig. 5.8a), and therefore the signal starts by being amplified. As we proceed along z, $P_p(z)$ decreases due to losses (Eq. (5.3)), progressively lowering also the local gain G_s. When the pump falls at the transparency threshold P_{th}, the local gain vanishes and the signal reaches its peak value $P_{s,max}$. After this point, G_s is negative and the signal decreases.

As already highlighted, slowing down the signal simply increases the

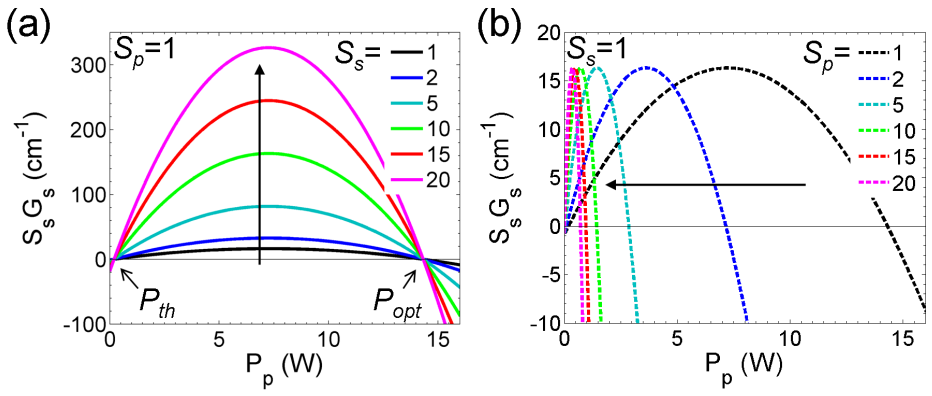

Figure 5.7: Local signal gain function $S_s G_s(P_p)$ computed using the parameters given in Table 5.1 for different values of (a) S_s and (b) S_p. The transparency threshold P_{th} and optimal pump power P_{opt} are indicated in (a).

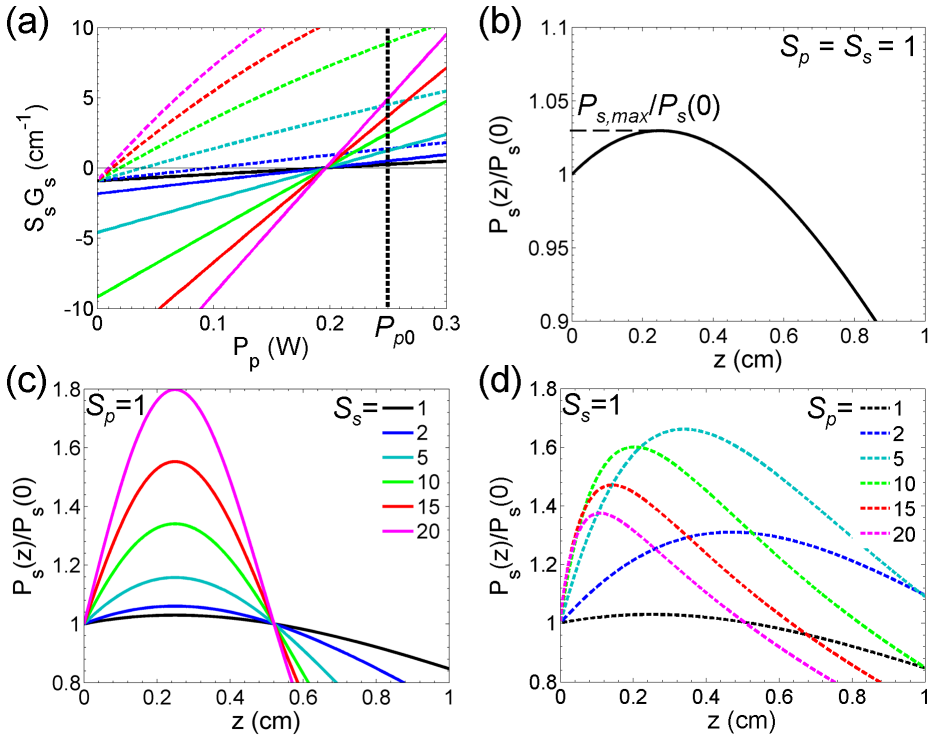

Figure 5.8: (a) Close-up of the local gain $S_s G_s$ from Figs. 5.7(a-b) for P_p up to 300 mW; the vertical dotted line indicates the input pump power $P_{p0} = 250$ mW used for the plots in (b-d). (b-d) Evolution of the signal power along the waveguide (z direction) for (b) $S_p = S_s = 1$, and for the same values of (c) S_s and (d) S_p as in Fig. 5.7.

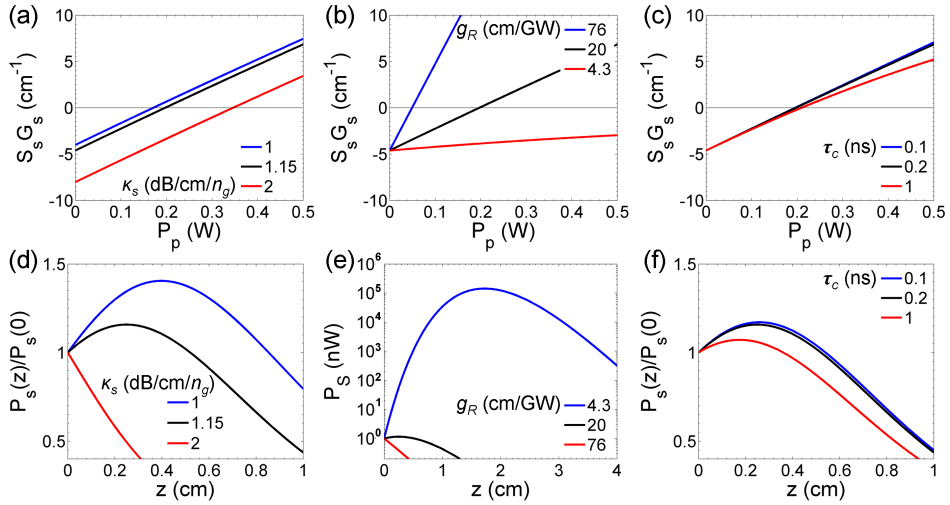

Figure 5.9: Variations of (a-c) the local gain S_sG_s and of (d-f) the signal power for different values of (a,d) the signal linear loss per unit slowdown factor κ_s, (b,e) the Raman gain coefficient g_R (note that in (e) the y-axis is displayed in logarithmic scale) and (c,f) the free-carrier lifetime τ_c, within the range of typically reported values. Slowdown factors are set to $S_p = 1$ and $S_s = 5$. The black curves in each plot are the same as those for $S_s = 5$ in Figs. 5.7a and 5.8a,c. If not specified, other parameters are as in Table 5.1.

local gain or loss, and this is reflected directly on the signal evolution (Fig. 5.8c). The effect of slowing down the pump, however, is more complicated: on one hand the transparency threshold is reduced and the local gain is higher for the same pump power (Fig. 5.8a), but on the other hand the pump power itself decays faster along the waveguide due to increased pump linear and nonlinear losses (Eq. (5.3)). This results in the existence of an optimum pump slowdown factor S_p for a given input pump power P_{p0} (Fig. 5.8d), after which further increasing S_p worsen the performance, as the signal peak $P_{s,max}$ is lowered. Nevertheless, an increase of S_p above the optimum may be advantageous over short lengths.

If the aim is to achieve the best signal amplification, the waveguide length L should be chosen so that the output signal corresponds to the peak signal, $P_s(L_{opt}) = P_{s,max}$, and therefore $P_p(L_{opt}) = P_{th}$. The best possible amplification would be obtained for waveguide length and input pump power chosen so that $P_{p0} = P_{opt}$ and $P_p(L_{opt}) = P_{th}$, in order to cover the full range of pump powers for which G_s is positive. Details of this type of optimisation may be found in Ref. [239].

Note from Fig. 5.9 that the shape of the local gain function G_s is very sensitive to the values of the many parameters involved. Linear propagation loss (Fig. 5.9a) depends on the technology and on the waveguide design, but typically ranges around 1–2 dB/cm/n_g [31,76]. The range of reported values

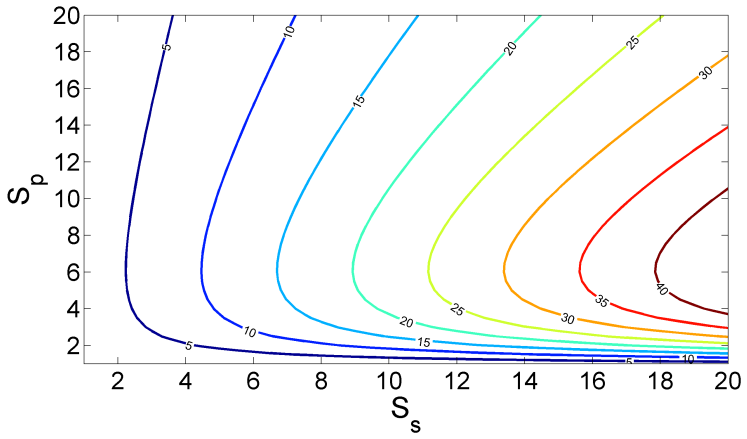

Figure 5.10: Peak signal gain $P_{s,max}/P_s(0)$ (in dB) as a function of the pump and signal slowdown factors S_p and S_s, respectively. All other parameters are as indicated in Table 5.1.

for the Raman gain coefficient g_R (Fig. 5.9b) is very large, 4.2–76 cm/GW [198, 225, 233], and we have chosen here a conservative value (Table 5.1), as adopted also by other authors [239]. Also the free-carrier lifetime τ_c (Fig. 5.9c) can largely vary depending on technology; the value of 200 ps adopted here is based on the measurements on our samples performed by our collaborators from the NanoOptics group of the FOM Institute AMOLF, Amsterdam, for the purpose of tunable delay (see section 3.4 and Ref. [244]), but even lower values have been reported [169]. The resulting final signal power profile $P_s(z)$ heavily depends on the combination of these parameters (Figs. 5.9d-f).

In general, it is even possible for the local gain function G_s never to reach positive values. From inspection of Eq. (5.7) it is straightforward to see that in order to allow for positive values of G_s, the linear loss and free-carrier lifetime must satisfy the condition $\tau_c\kappa_s < (g_R - 2\beta_{\mathrm{TPA}})^2\hbar\omega_p/(2\sigma_s\beta_{\mathrm{TPA}})$ [238], or $\tau_c\kappa_s < 15$ dB/cm $\cdot$ ns for our choice of parameters (Table 5.1). For κ_s as in Table 5.1, we obtain a condition on the free-carrier lifetime, $\tau_c < 3.75$ ns.

Note that for the pump powers and free-carrier lifetime considered here, the main limitation that the Raman gain g_R must overcome to result into signal amplification is not FCA, but the linear propagation loss of the signal: $\alpha_s = S_s\kappa_s$ is the y-intercept in Fig. 5.8a, and the gain must compensate for it before reaching the transparency threshold P_{th}.

Finally, if the structure could be designed to slow down both the pump and the signal at the same time, it would be possible to combine the benefits of i) gain enhancement with S_s and ii) lowering of the transparency threshold with S_p, with much higher signal gains achievable. This is shown in Fig. 5.10, where the peak signal gain $P_{s,max}/P_s(0)$ is plotted as a function of S_p

and S_s. Once more we stress that the final figures have a strong dependence on the combination of parameters used, and lower [238] or higher [239] gains may be simulated with a different choice. The authors of both Refs. [238] and [239] however conclude that ultimately, if there are no restrictions on the available input pump power P_{p0}, the highest signal amplification should be achieved for fast to moderately slow pump modes, due to the increase of pump nonlinear losses with S_p as described above.

5.3.2 Influence of realistic linear loss and mode profiles

The model described in section 5.3.1 is useful to understand the fundamental difference in the roles played by the slowdown factors of pump and signal, but cannot be *directly* applied to the study of photonic crystal devices to gain insight into the gain or loss compensation achievable with real-life structures. First of all, the propagation loss α does not necessarily scale linearly with S, and the exact dependence varies with the waveguide design [31]; we have already highlighted in Figs. 5.9a,d how the exact value of the linear loss can have a strong effect on the final signal evolution.

Secondly, and more importantly, we have so far assumed the effective area $A_{eff,p}$ appearing in Eqs. (5.3) and (5.4) to be constant with S (see Table 5.1), as also assumed in Ref. [238]. However, we have already shown in Fig. 2.4 that this is not the case, and we have discussed in section 2.2 the importance of the role played by the mode profile in our method for slow light dispersion engineering [28], and in the accurate determination of propagation loss [31], coupling efficiency [81] and the enhancement of nonlinear effects such as four-wave mixing [61]. Note that in Eqs. (5.3) and (5.4) only the effective area $A_{eff,p}$ of the pump mode appears, and there is no dependence on the mode profile of the signal: in reality we must account for both mode profiles, because the overlap between pump and Stokes fields changes the strength of the Raman interaction,[4] as already indicated by McMillan *et al.* [40].

In summary, our aim in this section is to improve the simplified model of section 5.3.1 by i) including a realistic dependence of the propagation loss α on the slowdown factor S, ii) applying mode profile effects to the correct scaling of the Raman interaction in photonic crystals (as defined by McMillan *et al.* [40]), and iii) extending such scaling of the Raman gain term also to the TPA and FCA terms. We will then apply the more complete model to investigate some specific photonic crystal designs in section 5.3.3.

The expression of the Raman gain term in photonic crystal waveguides can be obtained through coupled-mode analysis, by following the derivation detailed by Panoiu *et al.* in Ref. [245] up their Eq. (14), and then proceeding

[4] This also means that the mode profiles cannot be ignored even when expressing Eqs. (5.3) and (5.4) in terms of intensities $I_{p,s}$ rather than powers $P_{p,s}$ [40, 239].

along the guidelines indicated by McMillan *et al.* in Ref. [40] for continuous-wave operation. Without re-stating the derivation, the main result of Ref. [40] can be expressed as

$$\frac{\partial P_s}{\partial z} = -\frac{3\omega_s a}{v_{g,p} v_{g,s} \varepsilon_0} \chi^R \frac{\Re\left[\int_{V_{\mathrm{Si}}} \mathbf{E}_s^*(\mathbf{r})\, \hat{\xi}^R \vdots \mathbf{E}_p(\mathbf{r})\, \mathbf{E}_p^*(\mathbf{r})\, \mathbf{E}_s(\mathbf{r})\, \mathrm{d}V\right]}{\left[\int_{V_{\mathrm{tot}}} n_p^2(\mathbf{r})\, |\mathbf{E}_p(\mathbf{r})|^2 \mathrm{d}V\right] \left[\int_{V_{\mathrm{tot}}} n_s^2(\mathbf{r})\, |\mathbf{E}_s(\mathbf{r})|^2 \mathrm{d}V\right]} P_p P_s, \tag{5.8}$$

where a is the length of the photonic crystal unit cell in the direction of propagation (and in our case is equal to the photonic crystal period), $v_{g,p}$ and $v_{g,s}$ are the group velocities of the pump and signal modes, respectively, ε_0 is the vacuum permittivity and $\mathbf{E}_{p,s}(\mathbf{r})$ are the pump and signal mode profiles. In addition, V_{tot} is the volume of one photonic crystal unit cell and V_{Si} is the silicon volume in the unit cell. Regarding the Raman susceptibility $\hat{\chi}^R$, we have here used the fact that in silicon all its non-zero components have the same magnitude $\chi^R = \Im\left(\chi^R_{1221}\right)$ [246]; following the same notation as in Ref. [241], we have extracted χ^R and defined a normalised susceptibility $\hat{\xi}^R$ to keep track of the anisotropy of the Raman effect. The components of $\hat{\xi}^R$ equal 0 when the corresponding components of $\hat{\chi}^R$ are zero, and 1 otherwise. In particular, in the coordinate system of the silicon crystal principal axes we have [246]:

$$\xi^R_{ijji} = \xi^R_{ijij} = 1, \quad i,j = 1,2,3, \quad i \neq j. \tag{5.9}$$

If we now compare Eq. (5.8) to the expression for the evolution of the Stokes wave in bulk silicon [40, 43, 247]

$$\frac{\partial I_s}{\partial z} = g_R I_p I_s = -\frac{3\omega_s \chi^R}{\varepsilon_0 c^2 n_{p,\mathrm{Si}} n_{s,\mathrm{Si}}} I_p I_s, \tag{5.10}$$

we can re-write Eq. (5.8) as

$$\frac{\partial P_s}{\partial z} = \frac{g_R}{A_R} S_p S_s P_p P_s, \tag{5.11}$$

where we have used $v_{g,i} = c/n_{g,i}$ and $S_i = n_{g,i}/n_{i,\mathrm{Si}}$, $i = p, s$ (justifying our choice for the definition of the slowdown factor in terms of the silicon refractive index), and we have defined the Raman interaction effective area A_R as

$$A_R = \frac{\left[\int_{V_{\mathrm{tot}}} n_p^2(\mathbf{r})\, |\mathbf{E}_p(\mathbf{r})|^2 \mathrm{d}V\right] \left[\int_{V_{\mathrm{tot}}} n_s^2(\mathbf{r})\, |\mathbf{E}_s(\mathbf{r})|^2 \mathrm{d}V\right]}{a n_{p,\mathrm{Si}}^2 n_{s,\mathrm{Si}}^2 \Re\left[\int_{V_{\mathrm{Si}}} \mathbf{E}_s^*(\mathbf{r})\, \hat{\xi}^R \vdots \mathbf{E}_p(\mathbf{r})\, \mathbf{E}_p^*(\mathbf{r})\, \mathbf{E}_s(\mathbf{r})\, \mathrm{d}V\right]}. \tag{5.12}$$

This definition of A_R is equivalent to V_R/a, with V_R the Raman volume defined by Checoury *et al.* [241].[5]

[5] Yang and Wong use a similar definition of V_R for photonic crystal cavities, but without $\hat{\xi}^R$ [43].

Equation (5.11) therefore appears in the same form as the Raman gain term of Eq. (5.4), with $A_{eff,p}$ simply replaced by A_R. This alteration, however, is of great importance, as A_R is conceptually very different from $A_{eff,p}$: the Raman effective area A_R takes into account the overlap of the pump and signal modes with each other and with the silicon fraction of the total volume (as the Raman interaction occurs only in silicon, not in air), *and* accounts for the Raman selection rules (see Ref. [233]) through $\hat{\xi}^R$. As a consequence, for the same photonic crystal design and the same mode profiles, A_R may vary with the waveguide orientation with respect to the silicon crystal principal axes, as we will discuss in section 5.3.4.

Now that we have defined the correct scaling factor for the Raman gain term in a photonic crystal waveguide, we can extend the result also to the TPA and FCA terms. The TPA nonlinearity is of the same order as the Raman scattering and scales in the same way. Therefore, we can define a TPA effective area A_{TPA} in the same form as A_R:

$$A_{\mathrm{TPA},pi} = \frac{\left[\int_{V_{\mathrm{tot}}} n_p^2(\mathbf{r})\,|\mathbf{E}_p(\mathbf{r})|^2 \mathrm{d}V\right]\left[\int_{V_{\mathrm{tot}}} n_i^2(\mathbf{r})\,|\mathbf{E}_i(\mathbf{r})|^2 \mathrm{d}V\right]}{a n_{p,\mathrm{Si}}^2 n_{i,\mathrm{Si}}^2 \Re\left[\int_{V_{\mathrm{Si}}} \mathbf{E}_i^*(\mathbf{r})\,\hat{\xi}^{\mathrm{TPA}} \vdots \mathbf{E}_p(\mathbf{r})\,\mathbf{E}_p^*(\mathbf{r})\,\mathbf{E}_i(\mathbf{r})\,\mathrm{d}V\right]}, \quad i = s, p \tag{5.13}$$

where $A_{\mathrm{TPA},pp}$ is to be used for the TPA term appearing in the pump rate equation, Eq. (5.3), and $A_{\mathrm{TPA},ps}$ is to be used in the crossed-TPA term of the signal equation, Eq. (5.4). Equivalent integrals were also defined by Osgood *et al.* in the study of nonlinear effects in silicon nanowires [248]. In the coordinate system of the silicon crystal principal axes, the normalised TPA susceptibility $\hat{\xi}^{\mathrm{TPA}}$ is defined as [223]

$$\xi_{ijkl}^{\mathrm{TPA}} = \frac{2}{1+\rho}\left[\frac{\rho}{3}\left(\delta_{ij}\delta_{kl} + \delta_{ik}\delta_{jl} + \delta_{il}\delta_{jk}\right) + (1-\rho)\,\delta_{ijkl}\right], \quad i,j,k,l = 1,2,3$$

with $\rho = 1.27$.

Finally, for the terms describing FCA due to carriers generated by TPA of the pump beam, we can define an effective FCA area A_{FCA}:

$$A_{\mathrm{FCA},pi}^2 = \frac{1}{a^2 n_{p,\mathrm{Si}}^4 n_{i,\mathrm{Si}}^2} \times \frac{\left[\int_{V_{\mathrm{tot}}} n_p^2(\mathbf{r})\,|\mathbf{E}_p(\mathbf{r})|^2 \mathrm{d}V\right]^2\left[\int_{V_{\mathrm{tot}}} n_i^2(\mathbf{r})\,|\mathbf{E}_i(\mathbf{r})|^2 \mathrm{d}V\right]}{\Re\left[\int_{V_{\mathrm{Si}}} |\mathbf{E}_i(\mathbf{r})|^2 \mathbf{E}_p^*(\mathbf{r})\,\hat{\xi}^{\mathrm{TPA}} \vdots \mathbf{E}_p(\mathbf{r})\,\mathbf{E}_p^*(\mathbf{r})\,\mathbf{E}_p(\mathbf{r})\,\mathrm{d}V\right]}, \quad i = s, p. \tag{5.14}$$

A similar definition is also adopted in Ref. [241], but without the normalised TPA susceptibility $\hat{\xi}^{\mathrm{TPA}}$.

In summary, the evolution of the pump and signal beams in a slow light photonic crystal waveguide in the presence of the Raman interaction and

nonlinear losses can be expressed by re-writing Eqs. (5.3) and (5.4) as

$$\frac{\partial P_p}{\partial z} = -\alpha_p(S_p)\,P_p - \frac{\beta_{\mathrm{TPA}}}{A_{\mathrm{TPA},pp}} S_p^2 P_p^2 - \frac{\sigma_p \beta_{\mathrm{TPA}} \tau_c}{2\hbar\omega_p A_{\mathrm{FCA},pp}^2} S_p^3 P_p^3, \tag{5.15}$$

$$\frac{\partial P_s}{\partial z} = -\alpha_s(S_s)\,P_s + \left(\frac{g_R}{A_R} - \frac{2\beta_{\mathrm{TPA}}}{A_{\mathrm{TPA},ps}}\right) S_p S_s P_p P_s - \frac{\sigma_s \beta_{\mathrm{TPA}} \tau_c}{2\hbar\omega_p A_{\mathrm{FCA},ps}^2} S_p^2 S_s P_p^2 P_s, \tag{5.16}$$

where the nonlinear interaction effective areas A_R, $A_{\mathrm{TPA},pi}$ and $A_{\mathrm{FCA},pi}$ are defined in Eqs. (5.12), (5.13) and (5.14), respectively.

5.3.3 Design comparison

We will now apply the model described in section 5.3.2 to evaluate the Raman gains realistically achievable in a few specific photonic crystal designs. In particular, we will compare a W1 waveguide to four structures engineered for slow light, as described in section 2.2 [28].

All the dispersion curves and complex field amplitudes $\mathbf{E}(\mathbf{r})$ are calculated using MPB [49], whereas the propagation loss $\alpha(S)$ is estimated with a model based on the loss engineering code developed by O'Faolain *et al.* in Ref. [31] (see also section 2.2), as summarised in Fig. 5.11. Figure 5.11a shows the experimental loss data (circles and crosses) and corresponding code prediction (solid blue and black lines) for two designs of engineered slow light waveguides, as reported in Ref. [31]. The model clearly fits well the experimental data, but it underestimates the loss for group indices below 10–15, as can also be seen from the close-up of Fig. 5.11c (see also Ref. [76]). Therefore, when we estimate the propagation loss for a third, different design (red curve in Figs. 5.11a-c), we better represent reality if we correct the curve at low n_g with a polynomial fit to the experimental data (green curve in Figs. 5.11b-c). Such procedure is justified by the fact that the code predicts that all the designs under consideration have almost equal losses for $n_g < 15$ (and we have verified this also for the other engineered designs considered in this chapter), and this is also confirmed by the available experimental data (circles and crosses in Fig. 5.11c). Therefore, in the following, the propagation loss for the engineered designs will always be estimated by merging the result of the loss code for $n_g > 20$ with a polynomial fit to the experimental data of Fig. 5.11a for $n_g < 10$. Finally, Fig. 5.11d shows that for W1 waveguides the prediction from the loss code (black curve) is not as accurate as for engineered waveguides, and for our purposes we found that a polynomial fit (red curve) would better represent the experimental data from Ref. [31] (black circles). As such, we will assume that the loss curve for W1 waveguides is given by the red fit of Fig. 5.11d.

Figure 5.12a shows the band structure of a silicon W1 waveguide, as an example to illustrate the type of modes configuration that we will adopt

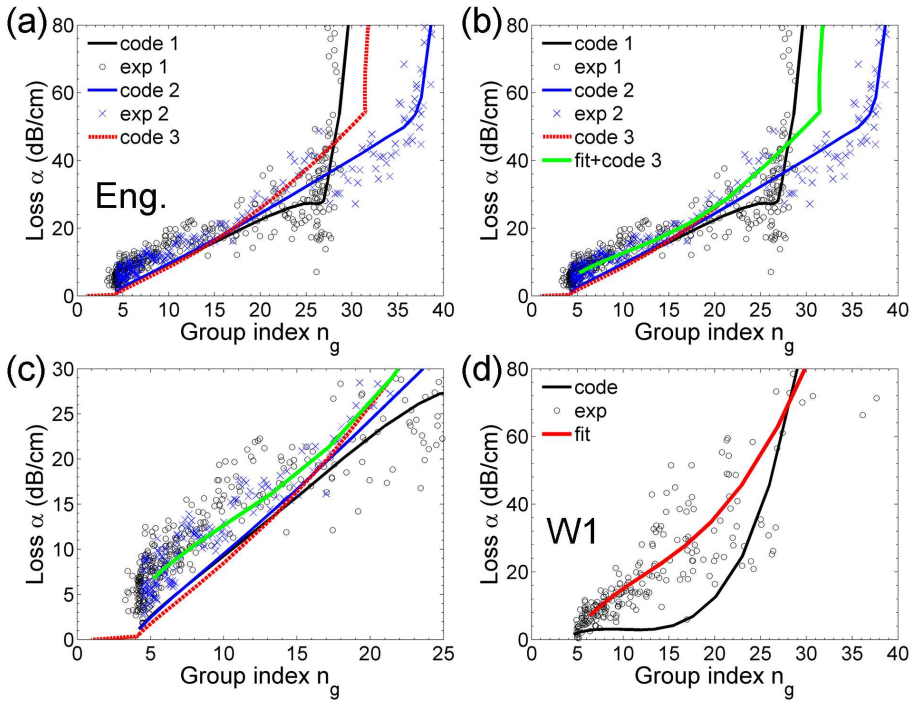

Figure 5.11: Propagation loss α as a function of the group index n_g for different waveguide designs. (a-b) Comparison of the experimental data (circles and crosses) with the loss code prediction (black and blue curves) for two designs reported by O'Faolain *et al.* (data from Ref. [31]). To better represent reality, the loss code prediction for a third design (red curve) is corrected with a polynomial fit to the available experimental data at low group indices (green curve), where the code tends to underestimate propagation loss. (c) Close-up of (b) for $n_g \leq 25$. (d) Comparison of the experimental data (circles) with the code prediction (black curve) for the W1 waveguide reported by O'Faolain *et al.* (data from Ref. [31]); the red curve is a polynomial fit to the experimental data.

in our study. The dispersion diagram is calculated by assuming a photonic crystal structure defined in a membraned slab of thickness $h = 220$ nm, with period $a = 412$ nm and radius $r = 0.286a$, chosen so that the fundamental TE mode lies in the ~1550 nm range (middle shaded band). We have considered the possibility of placing the pump onto this mode and slowing it down, but the only guided mode available at the corresponding Stokes wavelength ~1686 nm (bottom shaded band) is the fundamental TM mode (see also the TM band diagram in Fig. 5.14a): at these frequencies the TM mode has too large effective area to result in any useful gain (Fig. 5.14b). The second option is to place the signal onto the fundamental TE mode and use the fast TM mode to host the pump at ~1435 nm (top shaded band) — the odd TE mode is also available at these frequencies, but it would be more difficult to excite experimentally.

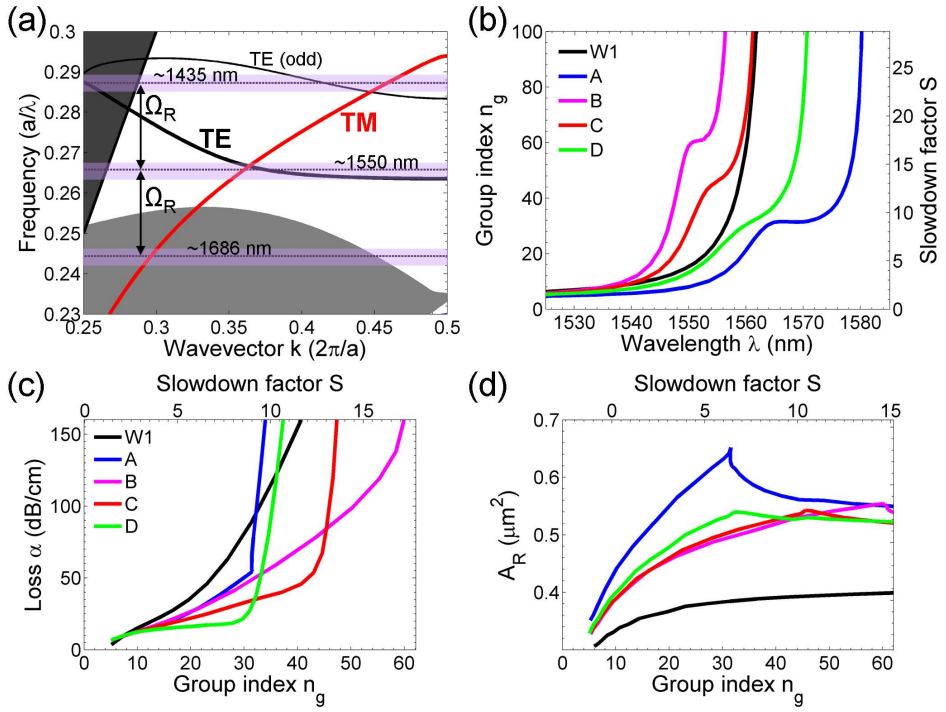

Figure 5.12: (a) TE band diagram of a membraned W1 waveguide with period $a = 412$ nm and hole radius $r = 0.286a$ (3D calculation). The fundamental index-guided TM mode is also shown in red (see also the full TM band structure in Fig. 5.14a). The middle shaded band indicates the ~1550 nm wavelength range; the top and bottom shaded bands are, respectively, blue- and red-shifted from the middle band by the silicon Raman shift Ω_R. (b) Group index n_g and (c) propagation loss α of the W1 waveguide and engineered designs A, B, C, and D. (d) Raman effective areas A_R calculated for waveguides aligned along the standard [110] direction of crystalline silicon (Fig. 5.14d).

Figures 5.12b-c show the group index $n_{g,s}$ and propagation loss α_s calculated for the W1 waveguide of Fig. 5.12a, together with the curves of four engineered slow light waveguides, which we will indicate as A, B, C, and D. The period of the engineered waveguides is set to $a = 404$ nm, so that the slow light falls around the same wavelengths as for the W1. The cross-sectional view of the different types of designs is also represented in Fig. 5.13. With respect to the W1 waveguide (Fig. 5.13, top), waveguides A and B (Fig. 5.13, middle) have been modified simply by shifting the first row of holes outwards (s_1 negative) and the second row of holes inwards (s_2 positive), following the rules for dispersion engineering as explained in section 2.2. In particular, the dispersion of waveguide A has been engineered to exhibit a ~7 nm bandwidth of constant group index around $n_g = 30$, by setting $r = 0.282a$, $s_1 = -0.12a$, and $s_2 = 0$: including this design in our

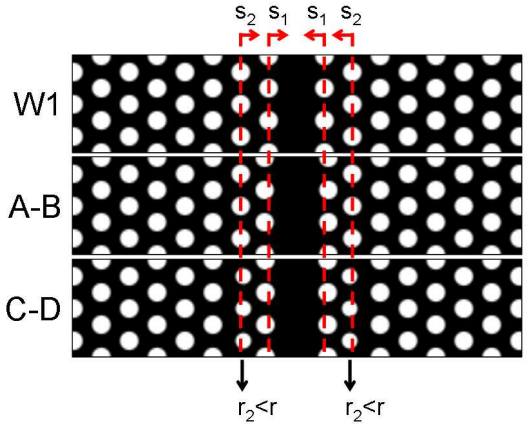

Figure 5.13: Cross-sectional view schematic of the different types of design analysed for stimulated Raman scattering in this study. Top: a standard W1 photonic crystal waveguide. Middle: engineered photonic crystal waveguides of the type discussed in section 2.2, with s_1 and s_2 the lateral shifts of the first and second rows of holes adjacent the waveguide, respectively; shifts are defined positive towards the waveguide centre. Designs A and B are of this type. Bottom: engineered photonic crystal waveguides in which not only the first two rows of holes are shifted, but the second row of holes is also reduced in radius. Designs C and D are of this type.

analysis is a natural choice, as this type of structure has already been used in various works to demonstrate enhancement of nonlinear effects [61, 62, 249]. Waveguide B ($r = 0.286a$, $s_1 = -0.1a$, $s_2 = 0.05a$), on the other hand, has been selected because is exhibits an almost linear dependence of the loss up to relatively high group indices $n_g \sim 60$ (see also section 2.2), as has also been very recently confirmed by experiment [76]; this very design was originally suggested in Ref. [31] as a possible good candidate for nonlinear applications. Finally, in structures C and D, in addition to the usual shifts s_1 and s_2 of the first two rows of holes, we also reduce the radius of the second row of holes, $r_2 < r$ (Fig. 5.13, bottom): as we have described in sections 2.1 and 2.2, the mode profile in the slow light regime partially extends inside the crystal cladding, and is strongly influenced by the geometry of the second row of holes; by reducing the radius r_2 we would expect the field to be less affected by the corresponding side-wall roughness, and therefore the loss to be improved. We include structure C ($r = 0.286a$, $s_1 = -0.1a$, $s_2 = 0.08a$, $r_2 = 0.26a$) in our analysis because, having an almost linear loss curve $\alpha_s(S_s)$, it comes very close to the simplified model of section 5.3.1. Structure D (same as C, but $r_2 = 0.24a$) appears in this study because its loss is almost constant with S_s; in fact it has been highlighted elsewhere [89] that, as the purpose of slow light is indeed to increase the delay or light-matter interaction time over a fixed length, the loss per unit length is not a suitable measure for this type of devices: in other words, as we increase n_g,

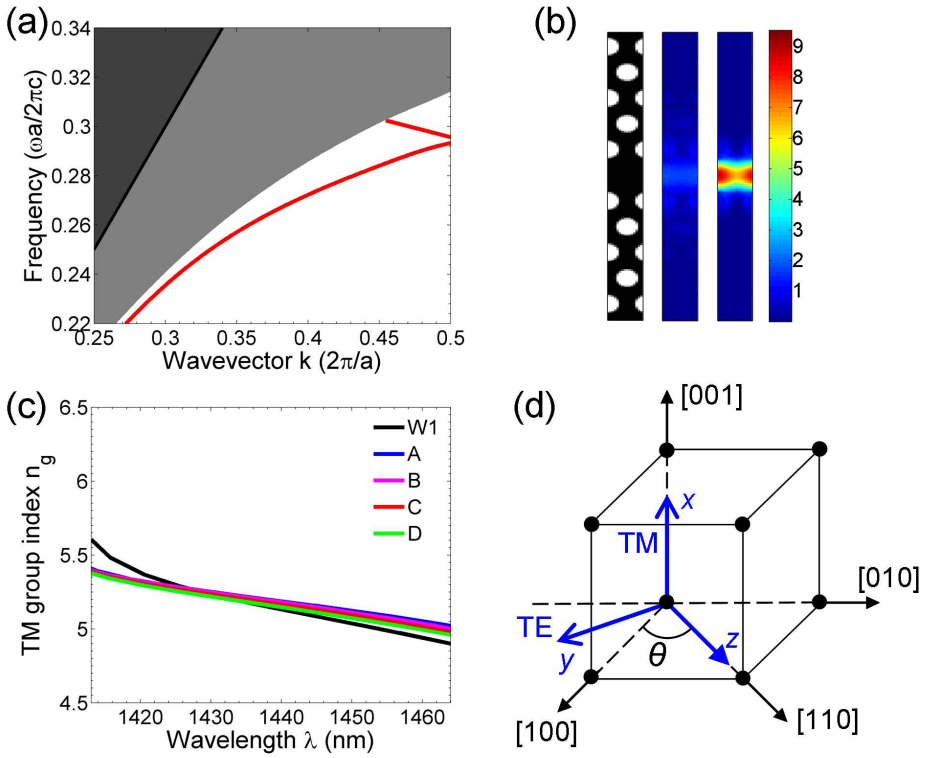

Figure 5.14: (a) TM band diagram of the W1 waveguide; the fundamental index-guided mode is indicated in red (3D calculation) (b) Cross-sectional view of the W1 waveguide (left), together with the corresponding intensity profiles (arbitrary units) of the TM mode at wavelengths ~1686 nm (middle) and ~1435 nm (right). (c) Group index of the fast TM mode in the pump wavelength range for the same designs considered in Fig. 5.12. (d) Crystal axes in a cubic unit cell as used to describe waveguide orientation in the diamond cubic crystal structure of silicon: in practical devices, waveguides are typically aligned along the [110] direction (blue), due to favourable cleaving properties; the corresponding TE and TM mode polarisation directions are also indicated.

not only the unwanted loss increases, but also the desired time delay. As a consequence, the loss per unit time has been proposed as a better measure [89], and design D was highlighted in Ref. [55] as one where such loss is particularly low. Our aim with this analysis is to compare the Raman gains that we may expect to achieve with established realistic devices (W1, A, B), as well as with more advanced, novel designs (C, D).

The effective Raman area A_R calculated with Eq. (5.12) is plotted in Fig. 5.12d for each of the considered designs. Here we have assumed that the waveguide is oriented along the [110] direction of crystalline silicon (Fig. 5.14d), as is standard practice in real fabricated devices due to favourable cleaving properties in this direction: in the calculation of the denominator

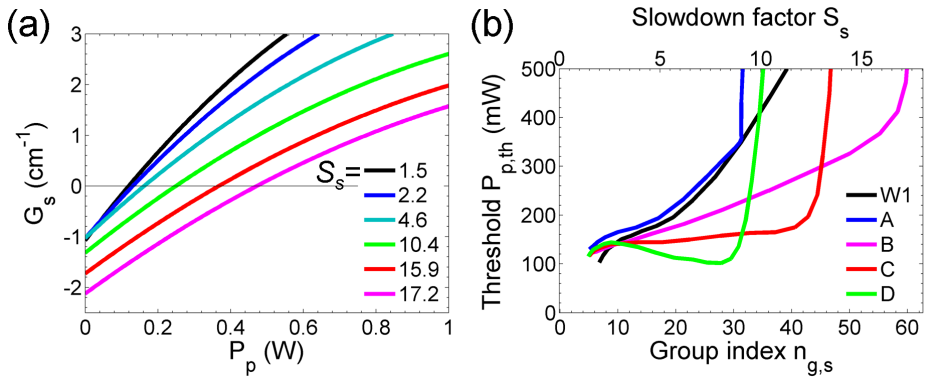

Figure 5.15: (a) Local signal gain G_s as a function of pump power P_p for design B, calculated for different values of slowdown factor S_s; the signal slowdown factor S_s affects the shape of G_s through the linear loss κ_s by translating the curve to lower values, and through the Raman area A_R by lowering the slope. (b) Transparency pump threshold P_{th} for all the designs considered in this section: the transparency threshold has in general a dependence on S_s, which cannot be accounted for by the simplified model.

of A_R we take $\hat{\xi}_R$ as defined in Eq. (5.9) and we express the mode profiles $\mathbf{E}_p$ and $\mathbf{E}_s$ in the coordinate system of the silicon crystal principal axes, rotated in the plane by 45° with respect to the coordinate system of the waveguide.

For each design we change the signal slowdown factor S_s by sweeping the wavelength along the n_g curves of Fig. 5.12b, and we also take into account the dependence of the Raman gain coefficient on wavelength as $g_R(\lambda_s) = g_R(\lambda_{s0}) \times (\lambda_{s0}/\lambda_s)$ (see Eq. (5.10) and Ref. [223]), with $\lambda_{s0} = 1550$ nm and $g_R(\lambda_{s0}) = 20$ cm/GW (Table 5.1). The group index of the pump TM mode takes values between 5 and 5.5 (Fig. 5.14c), and we assume the corresponding propagation loss α_p to be 5 dB/cm, comparable to that of a good fast TE mode [124]. The pump loss α_p can indeed be considered constant, since the group index $n_{g,p}$ does not vary significantly in the wavelength range of interest. Parameters β_{TPA}, $\sigma_{p,s}$ and τ_c are as defined in Table 5.1.

To illustrate the effect of realistic loss α_s and Raman area A_R, we plot in Fig. 5.15a the local signal gain G_s for design B (which allows for the highest slowdown factors S_s). The local gain G_s is defined here as a variation of Eq. (5.6):

$$G_s(S_s, S_p P_p) = -\kappa_s(S_s) + \left(\frac{g_R}{A_R(S_s)} - \frac{2\beta_{\mathrm{TPA}}}{A_{\mathrm{TPA},ps}(S_s)} \right) S_p P_p - \frac{\sigma_s \beta_{\mathrm{TPA}} \tau_c}{2\hbar\omega_p A^2_{\mathrm{FCA},ps}(S_s)} (S_p P_p)^2, \quad (5.17)$$

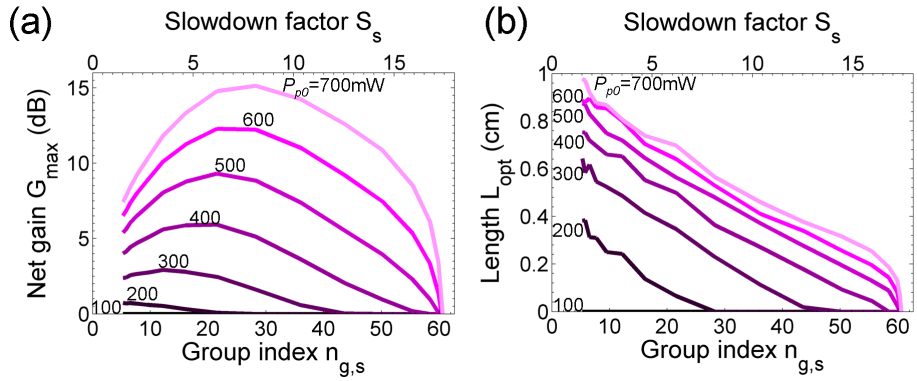

Figure 5.16: (a) Peak net signal gain G_{max} and (b) optimal waveguide length L_{opt} for waveguide B as a function of signal slowdown factor S_s for various input pump powers P_{p0}; the optimal signal slowdown factor depends on the input pump power and the optimal length is not constant with S_s.

with $\kappa_s(S_s) = \alpha_s(S_s)/S_s$, so that Eq. (5.16) can be re-written as

$$\frac{\partial P_s}{\partial z} = S_s G_s(S_s, S_p P_p)\, P_s,$$

in the same way as described in section 5.3.1. In the simplified model of section 5.3.1, since κ_s and $A_{eff,p}$ were constant with S_s, the local gain G_s had no dependence on S_s neither: if we would have plotted G_s alone, we would have always obtained the black curves of Figs. 5.7a and 5.8a (i.e. the case $S_s = 1$), for any value of S_s. In real photonic crystals, however, κ_s and the nonlinear areas, specifically A_R, depend on S_s, affecting the shape of G_s: for design B, Fig. 5.15a shows that the increase of κ_s with S_s (i.e. the loss α_s increases more than linearly with S_s) results in a rigid translation of G_s towards lower values, whereas the increase of A_R lowers the slope. The immediate result is a raise of the transparency pump threshold P_{th} with S_s (Fig. 5.15b, purple curve), which instead was constant in the simplified model (Fig. 5.8a). Mathematically, if to first approximation we neglect the TPA and FCA terms in Eq. (5.17), we obtain

$$P_{th}(S_s) = \frac{\kappa_s(S_s)\, A_R(S_s)}{g_R S_p}.$$

Therefore, in real-life structures, the signal slowdown factor S_s *does* influence for which pump powers the waveguide can exhibit positive gain.

The consequence on the signal evolution is the existence of an optimal signal slowdown factor, which is a function of the input pump power P_{p0}. One of the main results of the simplified model was that, provided the product $S_p P_p$ enabled an initial positive gain condition, slowing down the signal could only make things better (Fig. 5.8c). Now this is not the case:

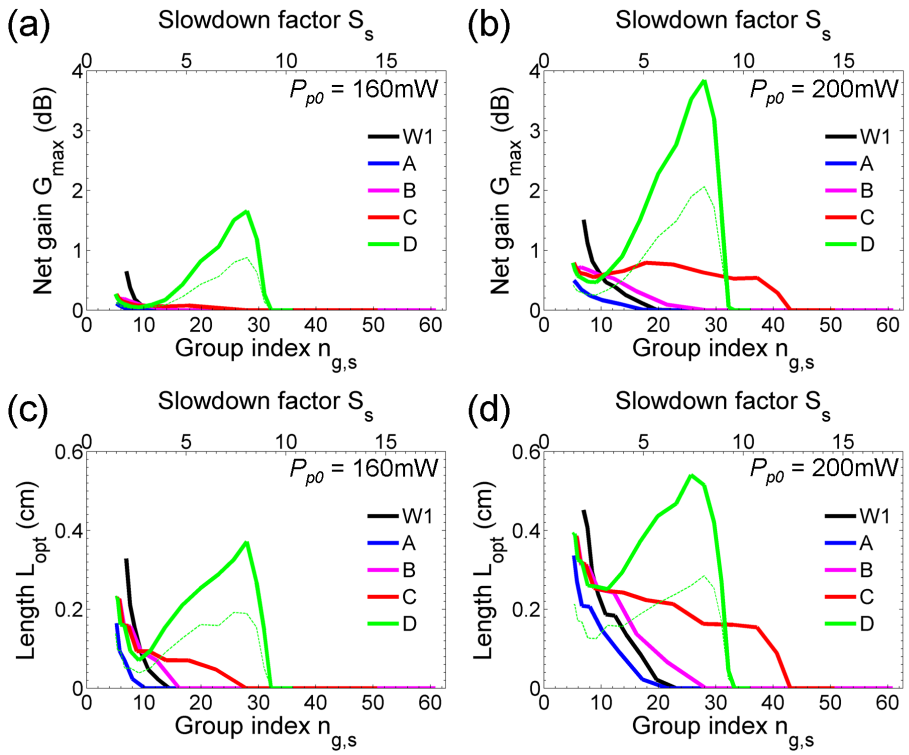

Figure 5.17: (a-b) Peak net signal gain G_{max} and (c-d) optimal waveguide length L_{opt} as a function of the signal slowdown factor S_s for all the waveguide designs under consideration, at input pump power $P_{p0} = 160$ mW (a,c) and 200 mW (b,d). The best performance is that of design D at $S_s \sim 8$, with a gain of 1.66 dB and 3.8 dB at pump powers 160 and 200 mW, respectively. The thin green curves are obtained for design D by increasing the TM pump propagation loss from 5 to 10 dB/cm.

from the plot in Fig. 5.16a of the peak net signal gain $G_{max} = P_{s,max}/P_s(0)$ as a function of S_s, we see that for design B whether it is better to slow the signal depends on the input pump power P_{p0}. Therefore, just as there exists a threshold for achieving positive gain, there is also a threshold for slow light to become advantageous with respect to the fast regime *of the same design*. Note also from Fig. 5.16b the dependence of the optimal length L_{opt} (such that $P_s(L_{opt}) = P_{s,max}$) on S_s for each input pump power, whereas L_{opt} did not depend on S_s in the simplified model (Fig. 5.8c).

Figure 5.15b displays the transparency threshold P_{th} also for the other designs under consideration, and Fig. 5.17 compares the peak net signal gain[6] G_{max} (Figs. 5.17a-b) and corresponding optimal length L_{opt} (Figs.

[6]Note that the net signal gain G_{max} curve should be read in conjunction with the corresponding optimal length L_{opt} curve. If $G_{max} = 0$ dB *and* $L_{opt} = 0$, the maximum net signal gain occurs at the input point $z = 0$: the waveguide exhibiths net loss.

5.17c-d) for all waveguides at low input pump powers, $P_{p0} = 160$ mW and 200 mW. The W1 waveguide exhibits the highest loss, but its A_R is the lowest (Figs. 5.12c-d), allowing for a small gain to be achieved at moderate n_g. Waveguide A has lower loss than the W1 over the entire slow light region, but the much higher Raman area A_R prevents it from showing appreciable net gain: with reference to the G_s curve of Fig. 5.15a, a lower α_s for design A means there is less loss needing compensation than for the W1, but a greater A_R means that the rate (slope) of the compensation is lower. With similar losses but smaller Raman area A_R, design B overall performs better than design A. As for waveguide C, its loss curve $\alpha_s(S_s)$ is very close to linear (Fig. 5.12c), but the Raman area increases with S_s: as a result, the moderate net gain G_{max} remains fairly constant within the slow light region, rather than increasing with S_s as we would expect from the simplified model. Finally, in design D, the increase in A_R is counteracted by a loss curve $\alpha_s(S_s)$ that is almost constant with S_s, and slowing down the signal brings a clear advantage already at low pump powers: the best G_{max} appears here for $S_s \sim 8$ ($n_g \sim 28$), and is estimated as 1.66 dB and 3.8 dB for pump power 160 and 200 mW, respectively, over lengths $L_{opt} = 3.7$ mm and 5.1 mm.

Doubling the pump propagation loss α_p from 5 to 10 dB/cm in design D almost halves the net gain G_{max} and the waveguide length L_{opt} for both the pump powers considered, as shown by the thin green curves in Fig. 5.17.

If higher pump powers can be used, slowing down the signal eventually becomes advantageous for all of the considered designs.

5.3.4 Gain dependence on waveguide orientation

We will now conduct the same analysis of section 5.3.3 but for waveguides oriented along the [100] direction of crystalline silicon, i.e. 45° to the cleaving plane — with reference to the notation of Fig. 5.14d, $\theta = 0°$. To better understand the motivation behind this additional investigation, we need to briefly introduce the selection rules dictated by the anisotropy of the Raman effect in silicon.

The definition of the silicon Raman susceptibility tensor given by Eq. (5.9) implies that the efficiency of the Raman scattering depends on the polarisation direction of the pump and Stokes signals relative to the crystallographic axis. Details on how to determine the selection rules may be found in Refs. [191, 233], and here we will limit ourselves to the final result for light propagating in the (001) plane of bulk crystalline silicon. Following a notation similar to that of Ref. [233], if we assume purely TE and TM modes, the bulk Raman gain g_R is to be multiplied by a factor C equal to

$$C^{TE-TE} = 4\sin^2\theta\cos^2\theta, \tag{5.18}$$

$$C^{TE-TM} = 1, \tag{5.19}$$

$$C^{TM-TE} = 1, \tag{5.20}$$

$$C^{TM-TM} = 0. \tag{5.21}$$

Here, the first and second superscripts describe the pump and Stokes polarisations, respectively. For TE-TE configuration (Eq. (5.18)), the efficiency of the Raman scattering varies from zero for $\theta = 0°$ or $90°$ to its maximum value for $\theta = 45°$ (standard waveguide orientation in SOI wafers). TE-TM and TM-TE configurations (Eqs. (5.19–5.20)) always correspond to the maximum scattering efficiency, regardless of the propagation direction in the plane. TM-TM is not an allowed configuration (Eq. (5.21)). In the analysis of section 5.3.2, these selection rules are included in the integral at the denominator of the Raman area A_R in Eq. (5.12), where the components of the pump and signal electric fields are differently mixed by the normalised Raman susceptibility $\hat{\xi}^R$ depending on the waveguide orientation.

In narrow and highly confined waveguides, however, the selection rules can change, as the modes are no longer purely TE or TM: in the case of the W1 waveguide analysed in section 5.3.3, Figs. 5.18 and 5.19 show that both the TE signal and TM pump modes exhibit strong longitudinal components along the propagation direction z. Checoury *et al.* [241] drew the attention to this effect by showing that in a W0.66 photonic crystal waveguide, and for the TE-TE configuration, orienting the waveguide along the [100] direction ($\theta = 0°$ in Fig. 5.14d) results in a more efficient scattering than in the standard [110] direction ($\theta = 45°$); this is indeed due to the presence of a strong longitudinal component in the TE(-like) modes, as in weakly confined waveguides the Raman interaction between a TE pump and a TE signal is forbidden in the [100] direction (Eq. (5.18)). In the formalism of section 5.3.2, this translates into a smaller Raman area.

We now extend the study from Checoury *et al.* by showing that changing the waveguide orientation to the [100] direction of crystalline silicon increases the Raman scattering efficiency also in our configuration, where pump and signal are TM- and TE-polarised, respectively. Figure 5.20a shows that the effective Raman area A_R calculated in the [100] direction (solid curves) is considerably reduced with respect to the standard [110] direction (dashed curves), for all the designs considered in our analysis. It is worth stressing that the photonic crystal properties (such as dispersion, loss and mode profiles) are exactly the same as before, and the only difference lies in the way the electric field components are selected and combined by the Raman susceptibility tensor $\hat{\xi}^R$. In particular, the reduction in A_R is to be attributed to the longitudinal component of the TM-like pump mode: the longitudinal component of the TE-like mode simply changes the orientation of the signal electric field in the plane, but we see from Eq. (5.20) that, for a purely TM pump, the scattering efficiency should be always the same, regardless of the orientation of the TE mode in the plane.

Figure 5.20b shows that the reduction of the Raman effective area results in higher peak signal gains than in the standard orientation for the same

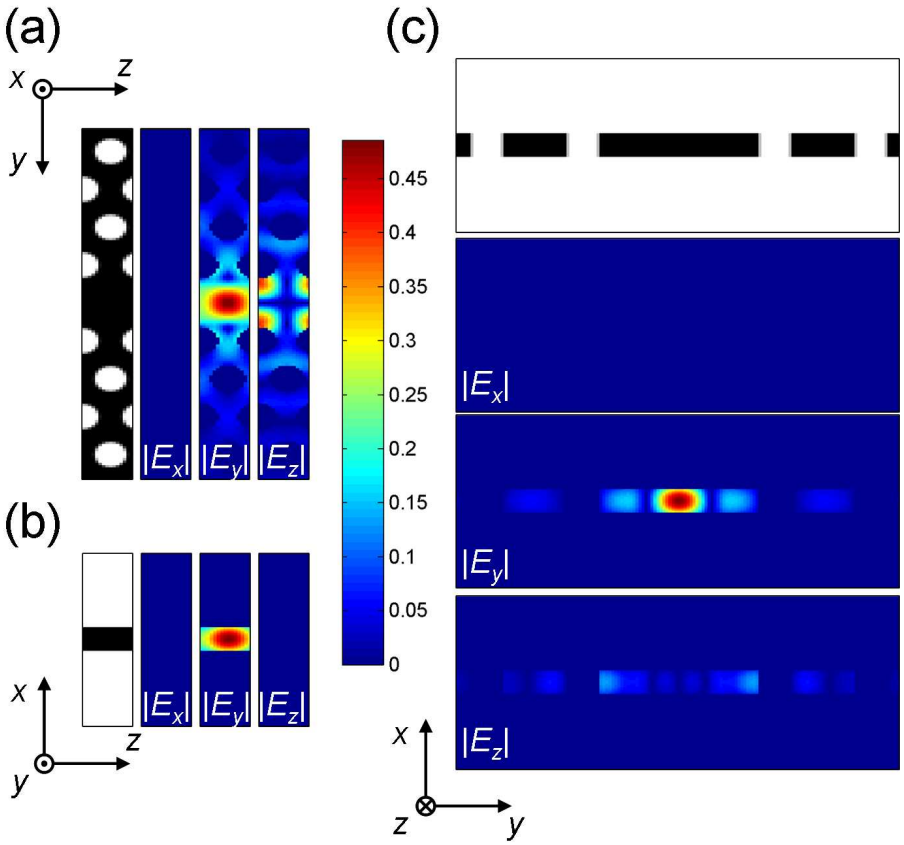

Figure 5.18: Cross-sectional cuts in the (a) x-, (b) y- and (c) z-directions of the W1 photonic crystal waveguide and of the electric field components E_x, E_y and E_z (arbitrary units) for the signal TE mode at wavelength $\lambda_s \sim 1545$ nm (with reference to the band structure of Fig. 5.12a). The cross-sections are always taken in the middle of the photonic crystal unit cell. The longitudinal component E_z is comparable to the transversal component E_y.

input pump power, and that a small gain already appears at even lower pump powers (Fig. 5.20c). Note that as the reduction in A_R is more pronounced for low group indices, the performance of the W1 waveguide in the fast light regime becomes comparable to that of design D in the slow light regime. At input pump power $P_{p0} = 160$ mW we estimate $G_{max} = 3.56$ dB at length $L_{opt} = 7.16$ mm for the fast W1 waveguide ($S_s \sim 2$) and $G_{max} = 3.33$ dB at length $L_{opt} = 5.26$ mm for design D in the slow light regime ($S_s \sim 7.4$). These figures well summarise the performance comparison between a low-loss and low-A_R waveguide on one side, and a slow waveguide on the other side. Still, our estimate identifies a window where slow light can be useful for signal amplification through stimulated Raman scattering at low pump powers, which looks even more encouraging when compared to the

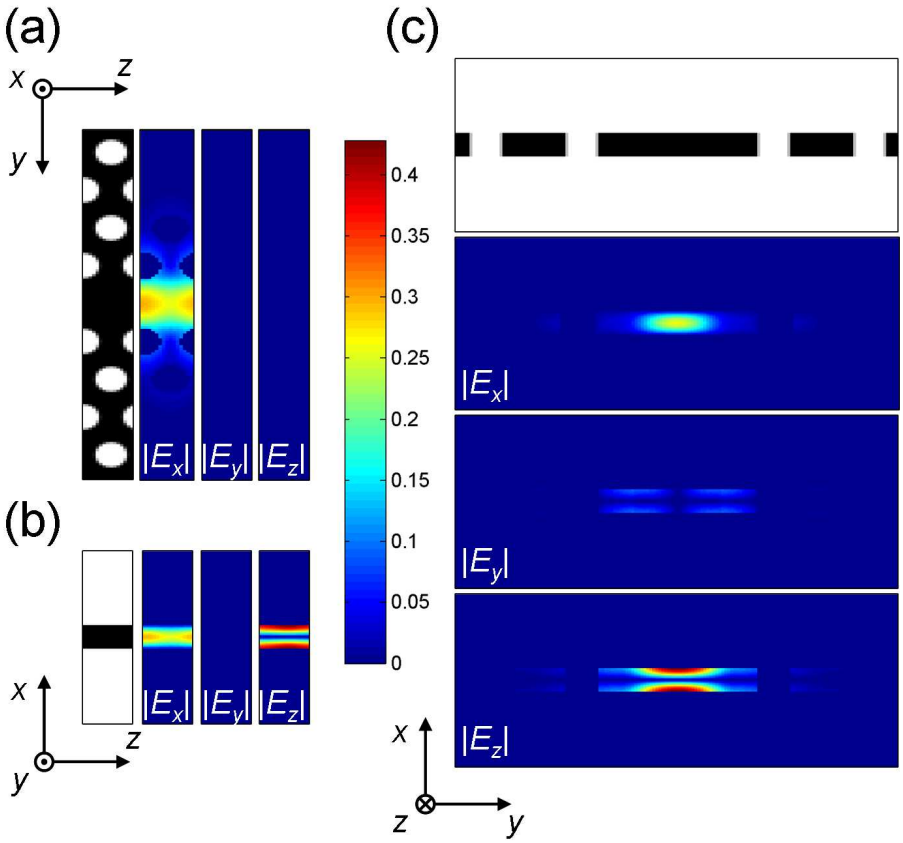

Figure 5.19: Cross-sectional cuts in the (a) x-, (b) y- and (c) z-directions of the W1 photonic crystal waveguide and of the electric field components E_x, E_y and E_z (arbitrary units) for the pump TM mode at wavelength $\lambda_p \sim 1435$ nm (with reference to the band structure of Fig. 5.12a). The cross-sections are always taken in the middle of the photonic crystal unit cell. The longitudinal component E_z is comparable to the transversal component E_x.

current record for continuous-wave stimulated Raman gain, which is 3.7 dB for 724 mW input pump power, over a 4.8 cm-long rib waveguide and with a 25 V reverse-biased p-i-n junction to reduce free-carrier lifetime [233].

In conclusion, we started from the promise in the literature of slow light being a very powerful tool enabling extremely high Raman gains in silicon [40, 238, 239]. Throughout our study we have reconsidered such promise and we have shown that, even when accounting for realistic loss and mode profile effects, slow light has indeed the potential for achieving similar signal gains to those already reported, but at lower pump power *and* shorter device length, and without the need of special measures for carrier lifetime reduction. We would expect even higher gains (or similar gains at lower pump powers) to be possible if both pump and signal beams could be slowed down

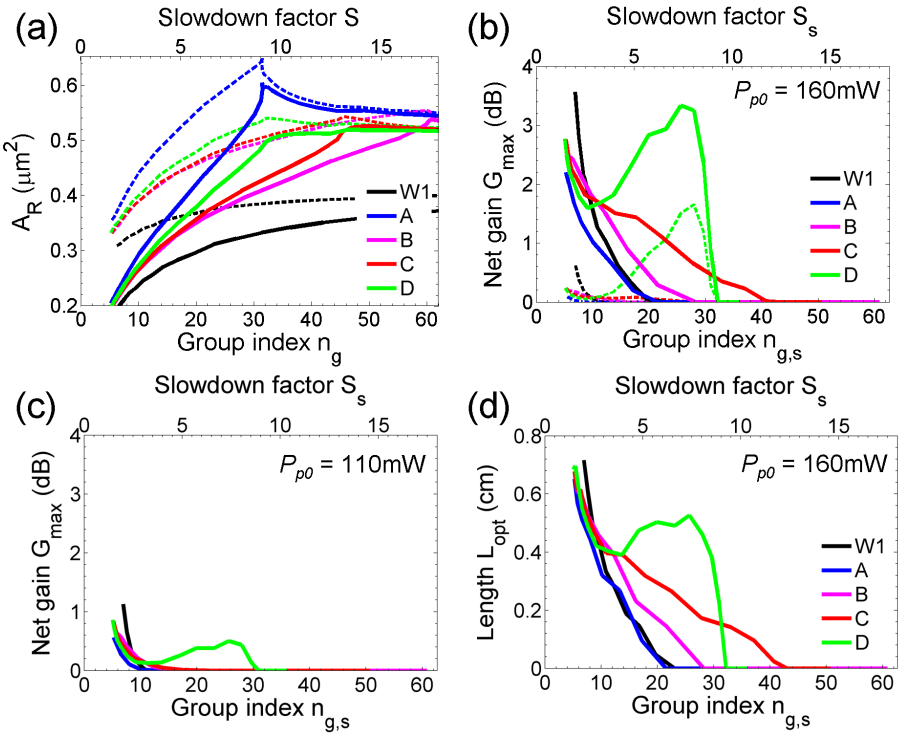

Figure 5.20: (a) Raman effective areas A_R for all the designs under consideration, calculated for waveguides aligned along the [100] direction of crystalline silicon (solid curves) and along the standard [110] direction (dashed curves); see also Fig. 5.14d. (b-c) Peak net signal gain G_{max} and (d) optimal waveguide length L_{opt} as a function of the signal slowdown factor S_s, at input pump power $P_{p0} = 160$ mW (b,d) and 110 mW (c); dashed curves in (b) are for waveguides aligned along the standard [110] direction, for comparison. At 160 mW, design D reaches 3.33 dB gain in the slow-light regime at length 5.26 mm; the W1 waveguide reaches 3.56 dB gain at length 7.16 mm.

at the same time, thus truly enabling useful operation of silicon Raman amplifiers at low pump powers, but the design and realisation of photonic crystal structures supporting two slow modes at the correct frequency spacing is extremely challenging [41, 241].

5.3.5 Prospects for loss compensation at 1686 nm

In the previous sections we have found that, if suitably designed, engineered slow light photonic crystal waveguides can achieve signal amplification through stimulated Raman scattering at reasonably low pump powers even when realistic detrimental effects are taken into account. Before being interested in signal amplification, however, we simply aim at compensation of the linear propagation loss in the slow light regime, and therefore our

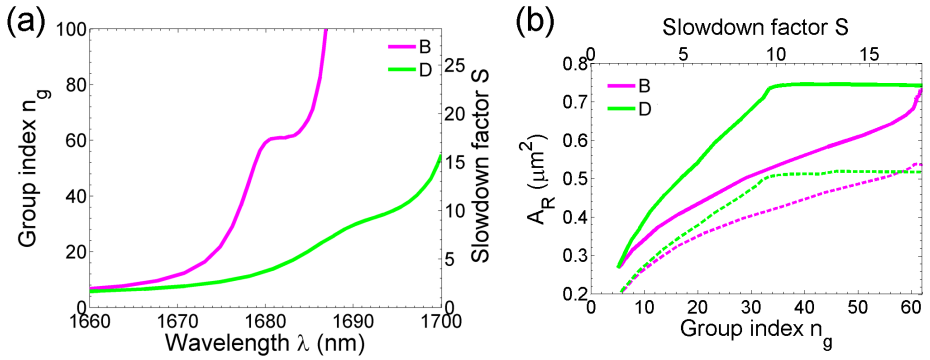

Figure 5.21: (a) Group index n_g for engineered designs B and D scaled at a larger period, $a = 448$ nm. (b) Effective Raman area A_R for the same designs (solid curves) calculated for waveguides oriented along the [100] direction of crystalline silicon ($\theta = 0°$ in the notation of Fig. 5.14d); the effective Raman areas for same waveguide orientation but period $a = 404$ nm are repeated here for comparison (dashed curves).

estimate is even more encouraging.

All calculations in the previous sections were performed for pump beam in the 1435 nm wavelength range and Stokes signal in the 1550 nm range. As already mentioned in Chapter 4, however, in the actual experiment our best pump source is a continuous-wave tunable laser amplified by an EDFA, and the available pump wavelengths span in the range 1540–1563 nm, thus placing the Stokes signal at 1674–1701 nm. We must therefore scale the designs analysed so far to larger periods in order to shift the modes to longer wavelengths. In particular, from now onwards (and also in the experiment) we will only consider the engineered designs B and D: we choose design D because in principle it allows for the best Raman performance in the slow light regime, but since we have no experimental data available to confirm the peculiar behaviour of the linear loss predicted by the loss code (Fig. 5.12c), we will still consider also design B, for which the loss prediction has already been confirmed [76].

Figure 5.21a shows the group index curves of designs B and D scaled to period $a = 448$ nm. Figure 5.21b shows the effective Raman area calculated along the [100] direction ($\theta = 0°$) for the new, enlarged structures (solid curves), as compared to the effective Raman area in the same waveguide orientation but for period $a = 404$ nm (dashed curves), which is repeated here for comparison. The effective Raman area is now much larger due to the bigger period, and as such we will expect a reduction in the scattering efficiency.

This is indeed confirmed by Fig. 5.22a, where we plot the peak signal gain G_{max} (solid curves) calculated in the same conditions and with the same input pump power P_{p0} as in Fig. 5.20b: the gain is now very small. We note

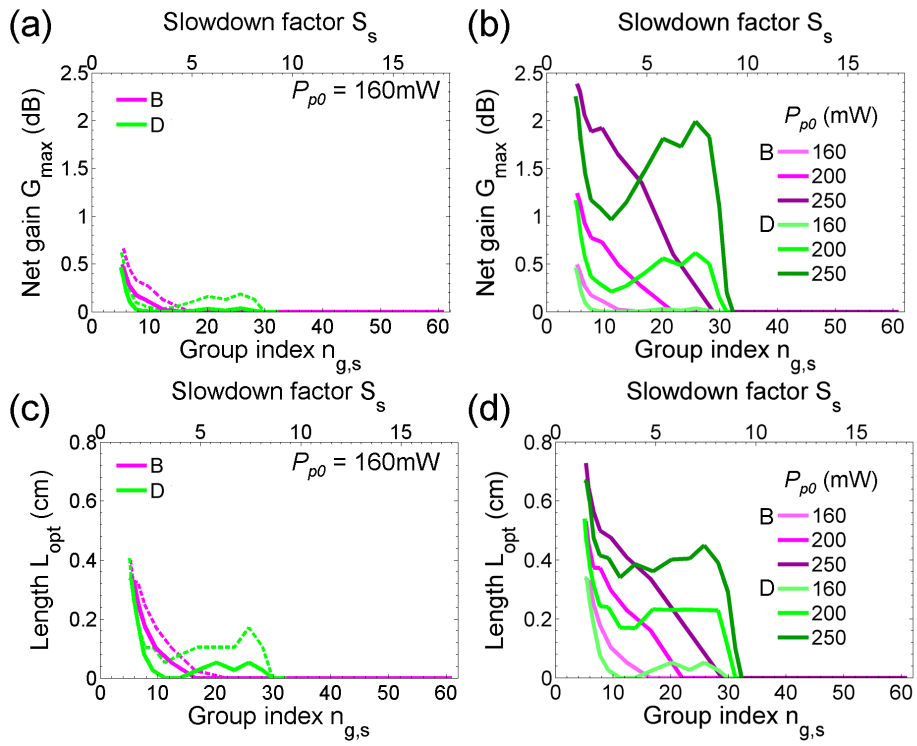

Figure 5.22: (a-b) Peak net signal gain G_{max} and (c-d) optimal waveguide length L_{opt} as a function of the signal slowdown factor S_s; the input pump power P_{p0} is fixed at 160 mW in (a,c) and varied up to 250 mW in (b,d). Dashed curves in (a,c) are calculated by neglecting the inverse dependence of the Raman gain coefficient g_R on the signal wavelength (see Eq. (5.10)). In all cases the waveguides are aligned along the [100] direction of crystalline silicon.

that there are in fact two contributions to the reduction of the overall signal gain: i) the larger Raman area and ii) the inverse dependence of the bulk Raman gain on wavelength $g_R(\lambda_s) = g_R(\lambda_{s0}) \times (\lambda_{s0}/\lambda_s)$ (see Eq. (5.10) and Ref. [223]), with $\lambda_{s0} = 1550$ nm and $g_R(\lambda_{s0}) = 20$ cm/GW (Table 5.1). This second effect, however, plays only a minor role in the reduction of the total signal gain, as we can see from the dashed curves in Fig. 5.22a, where G_{max} is plotted with the same bulk gain coefficient as at 1550 nm. Thus the main cause to the total gain reduction is indeed the increase in A_R due to the bigger period, and to achieve similar gains as before we would have to increase the input pump power (Fig. 5.22b).

However, by considering the gain G_{max} curve (Fig. 5.22a) in conjunction with the optimal length L_{opt} curve[7] (Fig. 5.22c), we see that even at these

[7]Note again that the net signal gain G_{max} curve should be read in conjunction with the corresponding optimal length L_{opt} curve. If $G_{max} = 0$ dB *and* $L_{opt} = 0$, the maximum net signal gain occurs at the input point $z = 0$: the waveguide exhibits net loss.

wavelengths it should be possible to achieve a tiny gain in the slow light region for waveguide lengths around 500 μm (design D) or, in other words, a lossless 500 μm slow light photonic crystal waveguide.

5.4 Raman scattering measurement

In this section we will describe the work conducted for the experimental observation of Raman scattering in silicon slow light photonic crystal waveguides. After a few comments regarding chip layout in section 5.4.1, we will proceed by describing the experimental setup in section 5.4.2. We will then discuss experimental results on spontaneous and stimulated Raman scattering in sections 5.4.3 and 5.4.4, respectively.

5.4.1 Sample layout considerations

The samples for measuring Raman scattering in silicon photonic crystals were fabricated following the recipe detailed in Chapter 4. We realised chips for both the waveguide orientations analysed in section 5.3, and we fabricated waveguides of length 96, 296, 496 and 696 μm (Fig. 5.23) for both designs B and D, using a writefield size of 100 μm during the e-beam writing step. Waveguides longer than the writefield size comprise fast-light coupling sections at the writefield boundaries to minimise the effect of stitching errors on the optical loss, as we have already described in Fig. 3.9 [76]. The actual slow light length is thus shorter than the total waveguide length.

Waveguides of different lengths on the same chip are aligned on the input side, as shown in the SEM images of Fig. 5.23: this way, we would expect the input pump power and input signal power to be similar from waveguide to waveguide, within fabrication tolerances. On the output side, the SU-8 waveguides have different lengths, which may have an effect on the detected signal at ~1680 nm due to material absorption at these wavelengths (section 4.2.4).

The definition of photonic crystal waveguides oriented along the standard [110] direction (Fig. 5.24a) is quite straightforward: the pattern is defined in a virtual (u, v) coordinate system (red), which determines the orientation of the writefield boundaries (red dashed lines), but shifts from one writefield to the other occur by translating the sample stage in its own real (x, y) coordinate system (black). The movements in (x, y) are precisely aligned to (u, v) during the setup procedure of the e-beam machine, but the initial orientation of (u, v) is normally imposed by aligning to the bottom edge of the chip.

When writing a long photonic crystal waveguide 45° to the cleaving plane ($\theta = 0°$ in Fig. 5.14d), we cannot simply tilt the structure in the layout (Fig. 5.24b), or the writefield boundaries will intersect the photonic crystal along many lines and there will be no obvious fast-light coupler to minimise the

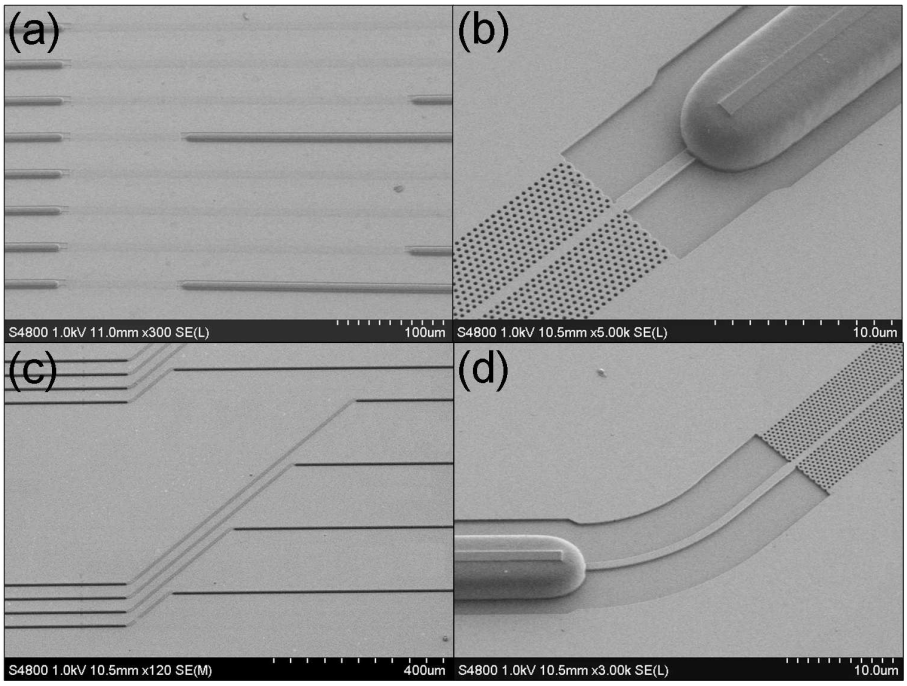

Figure 5.23: SEM images of the fabricated samples for Raman scattering measurement. The photonic crystal waveguides have lengths 96, 296, 496 and 696 μm and are oriented (a-b) along the standard [110] direction of crystalline silicon and (c-d) along the [100] direction. All waveguides are aligned on the input side (left).

effect of stitching errors. Since the photonic crystal waveguide is the feature most sensitive to stitching errors, it is best to write it along the normal orientation to the writefield boundaries.

This leaves us with two options. We can manually tilt the sample at 45° to the stage, as shown in Fig. 5.24c, maintaining the coordinate systems (u, v) and (x, y) aligned to each other; however, as we should make sure that the waveguide orientation is as close as possible to the [100] direction of crystalline silicon, this approach may not be reliable enough, especially if the sample has rectangular (rather than square) shape, which would make it difficult to define the initial orientation of (u, v) relative to the sample. The alternative is to tilt the (u, v) coordinate system by 45° instead, as illustrated in Fig. 5.24d: this method is more reliable than the previous one, as we can initially align the (u, v) system to the straight bottom edge of the sample and then impose a fixed 45° rotation. A drawback is represented by the alignment procedure between the (x, y) and (u, v) coordinate systems during setup, which takes much longer to converge.

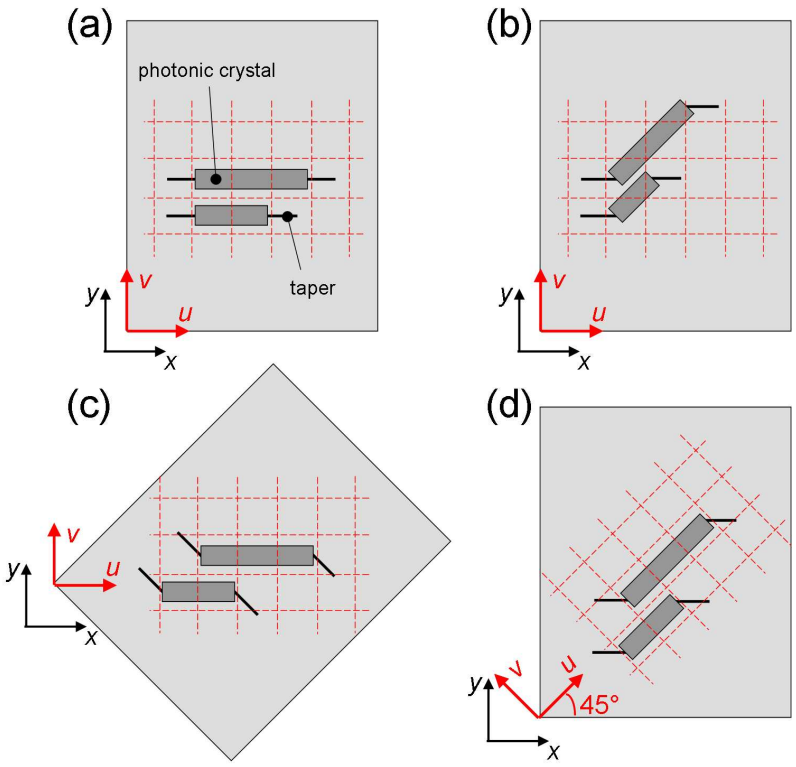

Figure 5.24: Schematic of the layout adopted during the definition of long photonic crystal waveguides during e-beam writing, for structures oriented along (a) the standard [110] direction dictated by the cleaving plane of silicon ($\theta = 45°$ in Fig. 5.14d) and (b-d) the [100] direction, 45° to the cleaving plane ($\theta = 0°$ in Fig. 5.14d). (a) The sample, the virtual (u, v) coordinates and the real stage (x, y) coordinates are all oriented in the same direction; the writefield boundaries (red dashed lines) are normal to the photonic crystal waveguide, and simple fast-light couplers can be defined to reduce the impact of stitching errors (see Fig. 3.9 and Ref. [76]). (b) The sample and coordinate systems are as in (a), but the photonic crystal pattern is tilted by 45°: the position of the writefield boundaries with respect to the photonic crystal is not trivial. (c) The sample is tilted by 45°, while the two coordinate systems are aligned to each other and the photonic crystal waveguide is normal to the writefield boundaries. (d) The sample is aligned to the stage (x, y) coordinates, but the virtual (u, v) coordinates are tilted by 45°; the photonic crystal waveguide is normal to the writefield boundaries.

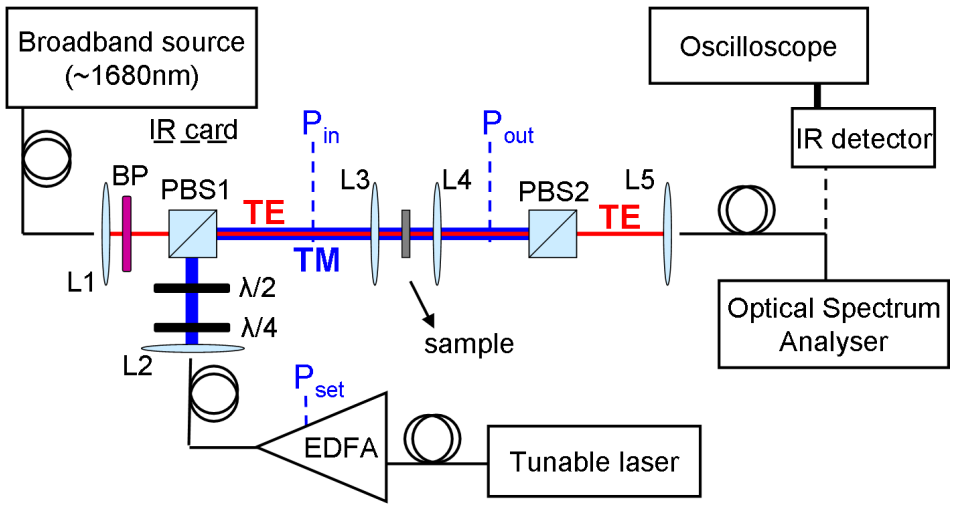

Figure 5.25: Schematic of the experimental setup used to perform Raman measurements in silicon photonic crystal waveguides.

5.4.2 Optical setup for Raman measurements

The high-power Raman scattering measurements were performed using a variation of a free-space transmission setup, as schematically shown in Fig. 5.25. Light is coupled in and out of the sample and the optical fibres through use of aspheric lenses (L1–5), each mounted on a xyz-precision stage. The pump source is a continuous-wave tunable laser with output wavelengths spanning in the range 1520–1630 nm. The laser output, set at 4 mW power, is amplified up to 600 mW by an EDFA operating in the wavelength range 1540–1563 nm. It is therefore the EDFA that limits the pump wavelength range. The pump polarisation is set to TM with a quarter- and a half-waveplates, and the pump beam (blue) is combined with the signal beam (red) through a polarisation beam splitter (PBS1).

The signal in the Raman wavelength range 1674–1701 nm is provided by a broadband source, which can be either a continuous-wave low-power LED or a pulsed super-continuum white light source. A bandpass filter (BP) in the signal arm cuts any light below 1600 nm and the signal polarisation is set to TE through PBS1. Light is in- and out-coupled to and from the sample through lenses L3–4, and then on the output side a second polarisation beam splitter (PBS2) filters out the transmitted pump power in TM polarisation. The Stokes signal is collected into a fibre through lens L5 and fed into an optical spectrum analyser (OSA).

During the alignment phase, the output fibre is connected to an infrared photo-detector rather than to the OSA; the detector is then connected to an oscilloscope and used to maximise the collected light. We manually switch the output fibre between the detector (during the alignment) and the OSA

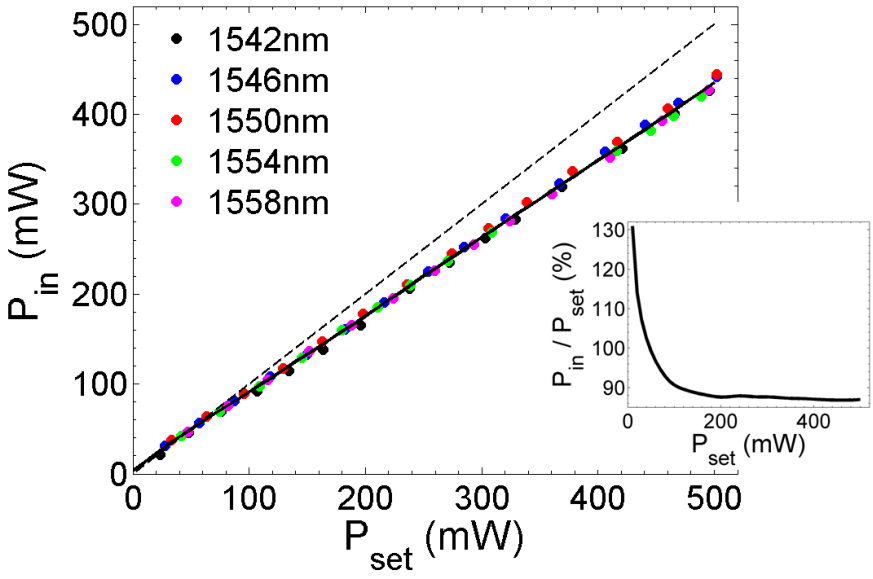

Figure 5.26: Input pump power P_{in} measured between PBS1 and L3 (see Fig. 5.25) as a function of the power P_{set} set on the EDFA for five different wavelengths. The black solid line indicates the average power. The black dashed line indicates the function $P_{in} = P_{set}$. Inset: P_{in} expressed in percentage of P_{set}.

(during the measurement), instead of using a 3 dB fibre splitter as in Fig. 2.13, because the Raman signal is very weak and we need to collect as much light as possible.

The alignment procedure is defined as follows:

1. signal ON, pump OFF: the signal arm (L1, L3–5) is optimised for maximum light coupling at the Stokes wavelength;

2. signal OFF, pump ON: PBS2 is removed, and in-coupling of the pump is optimised by adjusting lenses L2 and L3, but without moving L4 and L5;

3. signal ON, pump OFF: with PBS2 back in place, L1 only is adjusted to increase the in-coupled signal; L4 and L5 can also be re-adjusted at this stage, but L2 and L3 are kept fixed.

In this way, we ensure that the input is optimised for in-coupling of the pump beam, whereas the output is optimised for collection of the Raman signal. During these steps, when either the pump or the signal needs to be OFF, we block the corresponding beam path before PBS1 with beam blockers (not shown in Fig. 5.25).

Note that the two waveplates in the pump arm must be manually rotated to maximise the amount of light in TM polarisation *every time* the laser

wavelength is changed, and often also when varying the EDFA power at a fixed wavelength. We achieve this by placing an IR sensor card after PBS1, as indicated in Fig. 5.25, and rotating the waveplates until the spot on the card is at its minimum. This may seem a very rough way of proceeding, but in fact it is very reliable, as for a certain set power P_{set} on the EDFA, we always measure the same power in the line after PBS1, even when varying the wavelength. This is shown in Fig. 5.26, where the power P_{in} measured with a power meter between PBS1 and L3 is plotted as a function of P_{set} for five different wavelengths. The dashed line indicates $P_{in} = P_{set}$ and we see from the inset of Fig. 5.26 that, for $P_{set} > 100$ mW, the overall loss induced by L2, the two waveplates and PBS1 is around 10–13%.

5.4.3 Measurement of spontaneous Raman scattering

Figure 5.27a shows an example of the spontaneous Raman spectrum (pump on, signal off) collected from a 496 μm-long silicon nanowire at wavelength 1680 nm. The pump wavelength is 1545 nm and the pump power is set to 200 mW (black) and 300 mW (red) on the EDFA, corresponding to around 175 and 263 mW input power before lens L3 (Fig. 5.26). The measurement was taken with a resolution of 1 nm in the OSA.

To confirm that what we see is indeed Raman scattering, we have recorded various spectra while changing the pump wavelength (Figs. 5.27c-d, each spectrum normalised to its own peak value), and we have found that the detected signal lies exactly 15.6 THz below the pump. This is shown in Fig. 5.27b, where the black line represents the expected Raman wavelength dependence and the red dots indicate the recorded Raman peak wavelength from many different measurements on various waveguides and on different samples. In Fig. 5.27d the Raman spectra become more distorted and noisy towards 1700 nm, as we approach the detection limit of the OSA, even at the high sensitivity settings used for these measurements.

Figure 5.28a shows spontaneous Raman spectra measured from a 496 μm-long photonic crystal waveguide of type B, oriented along the standard [110] direction of crystalline silicon, at a fixed input power $P_{in} \sim 175$ mW. The length of the slow light region, obtained by subtracting the length of all the fast light couplers, is 430 μm. Each data point in the Raman spectra is averaged four times by the OSA during acquisition. Overlapped to the Raman spectra is the group index curve of the waveguide (black solid curve), measured with the interferometric technique described in section 2.4. The group index curve is interrupted around $n_g \sim 50$ because the waveguide was too long to resolve the interference fringes: the black dotted line represents the group index measured on a shorter neighbouring waveguide with the same design, and it is reported here to show the full dispersion. Different colours in the Raman spectra correspond to separate scans of the pump wavelength: once the setup has been aligned as described in section 5.4.2, we measure

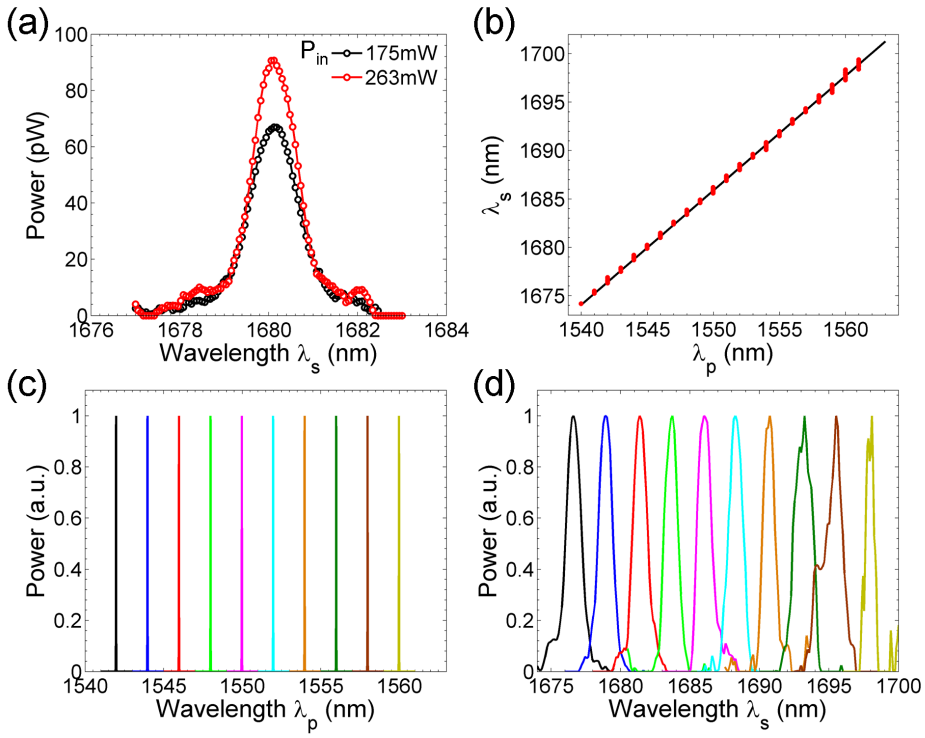

Figure 5.27: (a) Spontaneous Raman spectra measured from a 496 μm-long silicon nanowire at two different input pump powers P_{in}; the pump wavelength is 1545 nm and the measurements are taken with 1 nm resolution. (b) Peak Raman wavelength as a function of the corresponding pump wavelength: expected dependence (black line) and measured data (red dots). (c) Normalised pump spectra for different wavelengths (resolution 0.01 nm); what measured here is residual pump light passing through PBS2 and coupling into the output lens L5. (d) Normalised spontaneous Raman spectra (resolution 1 nm) corresponding to the pump wavelengths in (c).

the spontaneous spectra by changing the laser wavelength in steps (and by manually adjusting the two waveplates at each step), but the precision stages on which lenses L1–5 are mounted may start misaligning in the time it takes to acquire a few spectra at this high sensitivity. We must therefore re-align the system every 5–6 measurements before continuing with a new scan of the pump wavelength. Figure 5.28a shows that the measurement is fairly repeatable across different scans and, interestingly, the waveguide overall generates more Raman signal at longer wavelengths, corresponding to higher n_g.

To ensure that what we see is a slow light effect, and not due to non-idealities of the system, we have taken into account other possible causes that may lead us to see more Raman power at some wavelengths. We have already verified that we obtain the same input pump power P_{in} in front of lens L3 for different wavelengths, even by manually adjusting of the pump

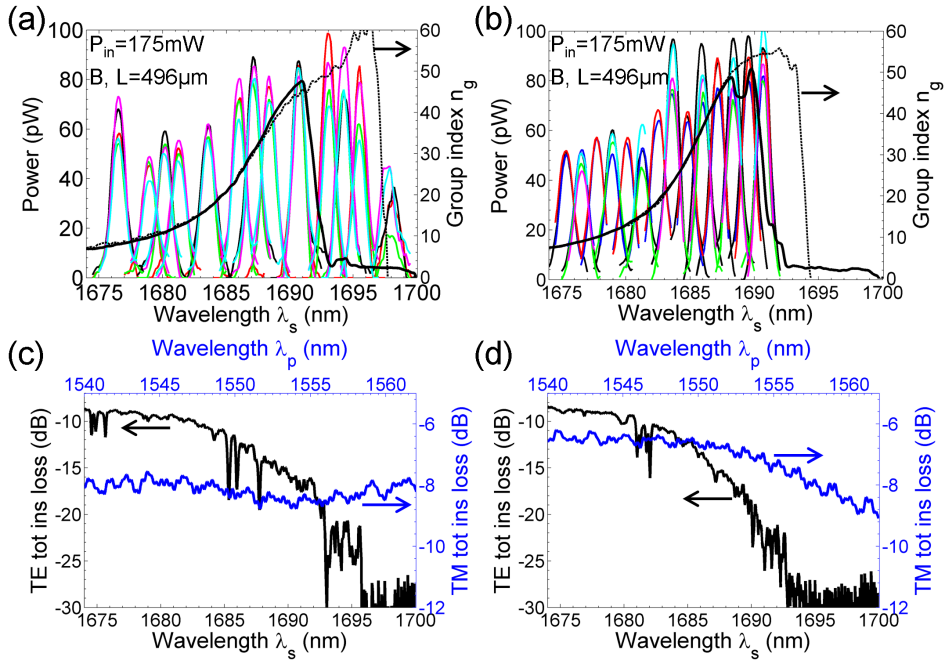

Figure 5.28: (a-b) Measured spontaneous Raman spectra (coloured curves) and group index (black solid) for two waveguides of type B and length 496 μm on the same chip (the waveguides were written with two different e-beam doses and thus have slightly different hole size and n_g curve). The dotted curves are the n_g spectra of shorter neighbouring waveguides of same design (slightly shifted in wavelength to match the solid curves). Both waveguides generate more Raman signal in the slow light regime than in the fast light regime. (c-d) Total insertion loss of the two waveguides in (a-b) for both the TE signal mode (black) and the TM pump mode (blue). The average TM insertion loss is 8.17 dB for (a,c) in the range 1540–1563 nm and 6.6 dB for (b,d) in the range 1540–1554 nm. By dividing these figures by 2 and neglecting propagation loss, the in-coupled pump power in the photonic crystal is at least 68 mW for (a) and 82 mW for (b).

polarisation with the waveplates. One possible issue might be the alignment: during the alignment phase (section 5.4.2) the pump is set to low power (~30–50 mW) and its wavelength is fixed, so that effectively the alignment is optimised for that particular wavelength, which is then changed during the scan. In Fig. 5.28a, however, black, red and green curves are obtained by initially aligning at $\lambda_p = 1550$ nm ($\lambda_s \sim 1686$ nm), whereas purple and cyan curves are obtained by aligning at $\lambda_p = 1542$ nm ($\lambda_s \sim 1676.5$ nm) and $\lambda_p = 1558$ nm ($\lambda_s \sim 1695.5$ nm), respectively: the result is the same. A third possibility is that the sample itself favours some wavelengths by coupling more pump light, but this is also ruled out by the total TM insertion loss measured for this waveguide which, as shown in Fig. 5.28c (blue curve), is quite constant in the wavelength range tested and varies only within 1 dB.

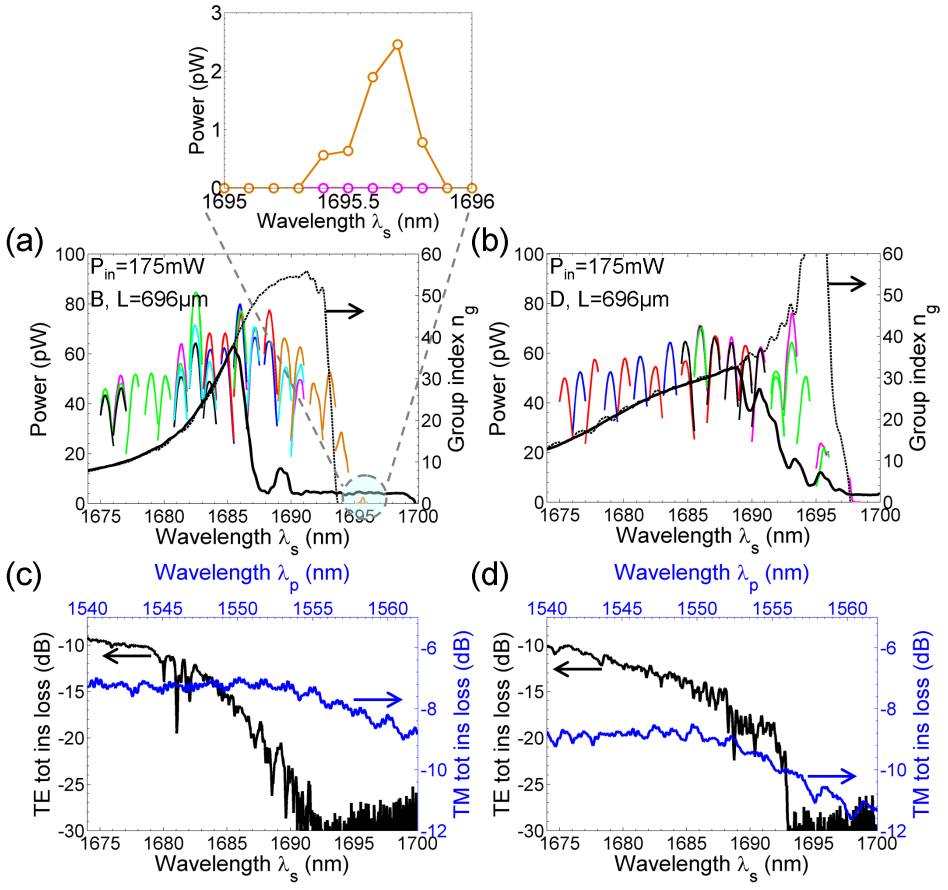

Figure 5.29: (a-b) Measured spontaneous Raman spectra (coloured curves) and group index (black solid) for waveguides of type B (a) and D (b) on the same chip as those of Fig. 5.28. Both waveguides have total length 696 μm, with effective slow light length 602 μm (a) and 607 μm (b). The inset of (a) shows the collected Raman signal when the pump wavelength λ_p corresponds to a Stokes wavelength λ_s inside the TE band gap. (c-d) Total insertion loss of the two waveguides in (a-b) for both the TE signal mode (black) and the TM pump mode (blue). The average TM insertion loss is 7.29 dB for (a,c) and 9 dB for (b,d) in the range 1540–1556 nm. By dividing these figures by 2 and neglecting propagation loss, the in-coupled pump power in the photonic crystal is at least 75 mW for (a) and 62 mW for (b).

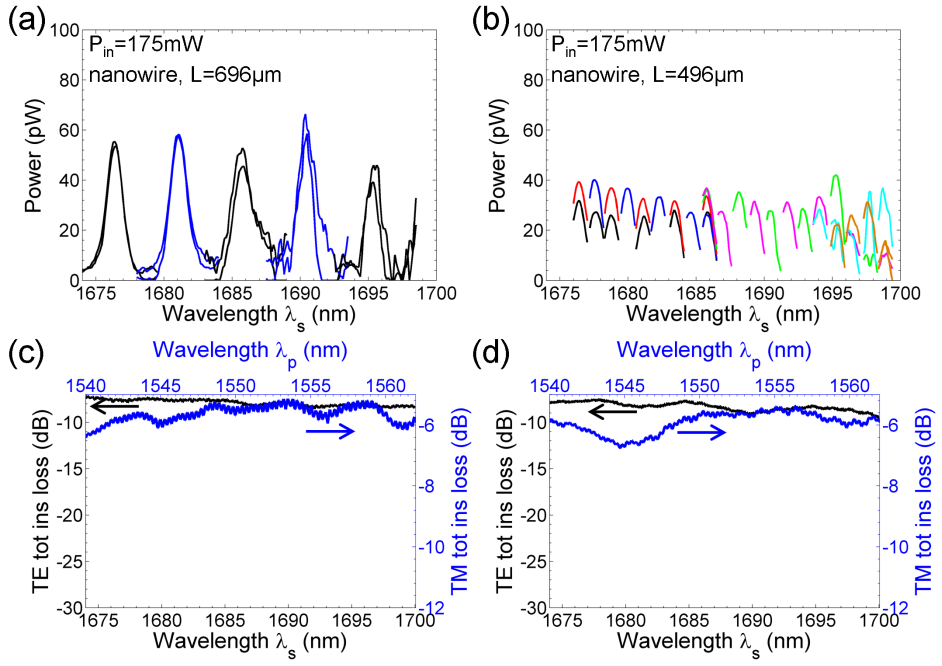

Figure 5.30: (a-b) Measured spontaneous Raman spectra for nanowires of length 696 μm (a) and 496 μm (b) on the same chip as the waveguides of Figs. 5.28 and 5.29. (c-d) Total insertion loss of the two nanowires in (a-b) for both the TE signal mode (black) and the TM pump mode (blue). The average TM insertion loss is 5.64 dB for (a,c) and 5.93 dB for (b,d) in the range 1540–1563 nm. By dividing these figures by 2 and neglecting propagation loss, the in-coupled pump power in the nanowire is at least 91 mW for (a) and 88 mW for (b).

The enhancement of the generated Raman power at slow wavelengths is also confirmed by measurements on different photonic crystal waveguides on the same sample (Fig. 5.28b and Figs. 5.29a-b), whereas for the case of nanowires there is no obvious increase of Raman signal at longer wavelengths (Figs. 5.30a-b).[8] In addition, we have also verified that the Raman contribution from the silicon inverse tapers is negligible. Figures 5.29a-b show examples for which the photonic crystal waveguides are pumped in such way that the Stokes wavelength falls inside the TE band gap: any Raman signal generated by the input taper is filtered out by the photonic crystal, and no scattering can occur inside the photonic crystal itself because there is no guided mode allowed. Any light that we collect at these wavelengths must therefore originate from the output taper, which is still pumped as the TM pump transmission is unrelated to the band gap: the Raman signal

[8]Note that in Figs. 5.29a-b and 5.30b the Raman spectra span over only 1 nm around the peak, in order to reduce the acquisition time and thus reduce the effect of misalignment of the precision stages over time.

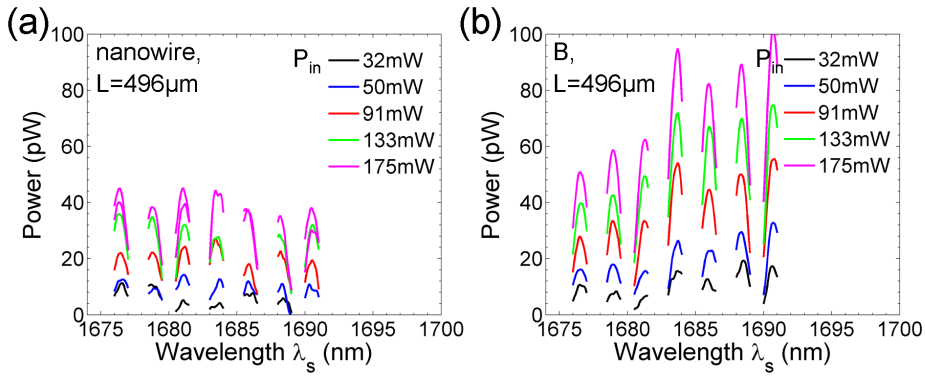

Figure 5.31: Measured spontaneous Raman spectra for (a) a nanowire of length 496 μm and (b) an engineered photonic crystal waveguide of type B of same length (this is the same waveguide as in Figs. 5.28b,d), for different input pump wavelengths and powers. The photonic crystal exhibits stronger Raman scattering at slow wavelengths at all the tested input pump powers.

generated at these wavelengths is barely detectable (Fig. 5.29a, inset), and thus the tapers contribution is negligible.

We can therefore conclude that the stronger spontaneous Raman signal detected at longer wavelengths is indeed due to a slow light effect, with no influence from issues such as alignment, coupling or Raman scattering from the access tapers. Our results are consistent with those reported by McMillan *et al.*, who reported spontaneous Raman enhancement due to slow pump [41], and with those reported by Checoury *et al.*, who reported enhancement due to slow signal, with a fast TE pump [42].

The fact that we collect a greater Raman signal from the slow light region is even more significant when we take into account the total insertion loss of the TE mode: we see from Figs. 5.28c-d and 5.29c-d that between the fast and slow light regimes there is a difference in transmission of up to 12–18 dB. If the Raman emission would depend only on loss, we would expect the spontaneous signal to decrease as we proceed towards the cut-off, as the propagation loss also increases. Instead, the Raman scattering at slow wavelengths is not only comparable to that at fast wavelengths, but actually increases, and this is a clear sign of group index dependence.

For an engineered waveguide of type D, shown in Fig. 5.29b, the increase of Raman signal at longer wavelengths is less pronounced than for type B (Figs. 5.28a-b and 5.29a), following the softer increase in group index.

In addition, in Fig. 5.31 we compare the Raman emission from a 496 μm-long nanowire to that of a photonic crystal waveguide of type B of same length, for different wavelengths and different input pump powers P_{in}. In the case of the photonic crystal (Fig. 5.31b), the generated Raman signal at slow wavelengths is better than that at fast wavelengths already at low

pump powers, whereas in the case of the nanowire (Fig. 5.31a) there is no obvious difference between wavelengths.

It is also interesting to notice that the nanowires generate Raman levels comparable to that of the fast light of the engineered waveguides, despite the total insertion loss for both TE and TM polarisations being better than for photonic crystals (the insertion loss graphs of Figs. 5.28c-d, 5.29c-d and 5.30c-d are plotted on the same scale).[9] We should also point out that from these measurements it is difficult to extract quantitative estimates on, for instance, the Raman scattering efficiency, mainly because it is not straightforward to estimate the actual loss on the collection side and we cannot fully trust the absolute powers measured by the OSA — with a different OSA we measure different powers (but of the same order) with the same resolution settings.

We can, nevertheless, spend a few more words on the qualitative dependence of the generated Raman output on the group index of the signal. When considering the model for stimulated Raman scattering we found that, at these long wavelengths and at low pump powers, slowing down the signal brings no advantage with respect to the fast light of the same design (section 5.3.5), and especially for waveguides of type B. Now, in contrast, we clearly see an increase of the generated Raman due to slow light. This suggests that the dependence of the *spontaneous* Raman scattering on group index is different than that of the *stimulated* scattering, which in fact can be predicted theoretically.

The evolution of a Stokes signal spontaneously generated through Raman scattering in a slow light photonic crystal waveguide can be described by the following equation:

$$\frac{\partial P_s}{\partial z} = -\alpha_s(S_s)\, P_s + \gamma S_p S_s P_{p0} e^{-\alpha_p(S_p) z}, \tag{5.22}$$

where we have simply included the slowdown factor dependence to the standard expression for spontaneous Raman scattering [195]. Here, γ indicates the spontaneous Raman coefficient, which is directly proportional to the gain coefficient g_R used in the equations for stimulated scattering [195], and therefore, just as g_R, will scale as the product $S_p S_s$. In Eq. (5.22) we have omitted for simplicity the TPA and FCA terms, so that the evolution of P_s along the waveguide can be expressed analytically by

$$P_s(z) = -\gamma S_p S_s P_{p0} \frac{e^{-\alpha_p(S_p) z} - e^{-\alpha_s(S_s) z}}{\alpha_p(S_p) - \alpha_s(S_s)}, \tag{5.23}$$

where we have imposed $P_{s0} = P_s(0) = 0$.

[9] However, the detected Raman powers in the various plots should not all be compared one-to-one as, for instance, the Raman spectra of Figs. 5.29a-b and 5.30b were collected at a later stage than those of Figs. 5.28a-b, 5.30a and 5.31a-b: a replacement of the output fibres in the setup resulted in generally lower collected powers in the later measurements.

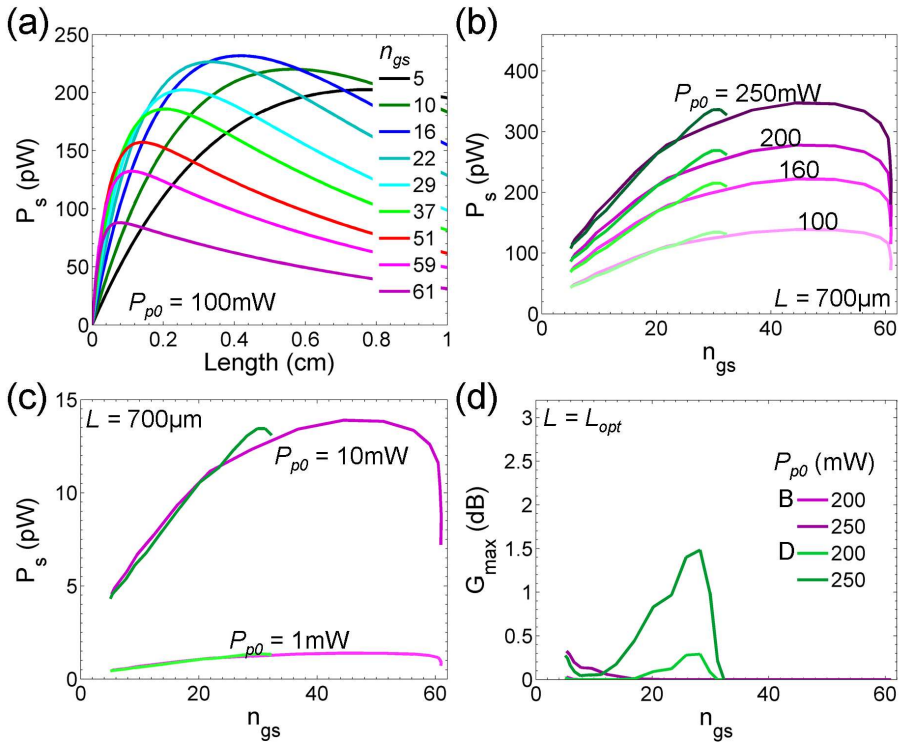

Figure 5.32: (a) Evolution of the spontaneous Raman signal in the waveguide design B, calculated using Eq. (5.23) for input pump power $P_{p0} = 100$ mW. (b-c) Spontaneous Raman signal at waveguide length $L = 700$ μm for designs B (purple) and D (green) for different input pump powers. (d) Net signal gain G_{max} due to stimulated Raman scattering calculated analytically from Eq. (5.24). (b-c) and (d) show that spontaneous and stimulated Raman scattering have different dependence on signal group index. All calculations are performed here for the standard [110] waveguide orientation.

Correspondingly, if TPA and FCA are neglected, the expression for stimulated Raman scattering Eq. (5.16) becomes

$$\frac{\partial P_s}{\partial z} = -\alpha_s(S_s)\, P_s + \frac{g_R}{A_R} S_p S_s P_{p0} e^{-\alpha_p(S_p)z} P_s,$$

whose analytical solution is

$$P_s(z) = P_{s0} e^{-\alpha_s(S_s)z + \frac{g_R}{A_R} S_p S_s P_{p0} L_{eff}}, \tag{5.24}$$

with

$$L_{eff} = \frac{1 - e^{-\alpha_p(S_p)z}}{\alpha_p(S_p)}.$$

The evolution of a Stokes signal spontaneously generated inside a slow light photonic crystal waveguide of type B is shown in Fig. 5.32a for in-coupled pump power $P_{p0} = 100$ mW for different values of the signal group

index, and it is calculated using Eq. (5.23). In the calculation we have used $\gamma = \gamma_0 \times (A_{R0}/A_R)$, with $\gamma_0 = 3.5 \times 10^{-9}\text{cm}^{-1}$ as used in Ref. [41] in the fast light region, and A_{R0} is the effective Raman area for fast light in the standard [110] direction. Other parameters are as in section 5.3.5.

Figure 5.32b shows the spontaneous Raman output at a fixed waveguide length $L = 700$ μm, for both designs B (purple) and D (green), for the same in-coupled pump powers P_{p0} as in Fig. 5.22b. Slowing down the signal increases the generated Raman output for both designs, and already for very low pump powers (Fig. 5.32c), in agreement with the experimental observation of Fig. 5.31b. In addition, the two designs generate the same amount of spontaneous Raman signal for the same group index. Figure 5.32d shows the maximum signal gain G_{max} due to stimulated Raman scattering in the standard [110] direction, which is reported here for comparison of the group index dependence with the spontaneous effect. The net signal gain is calculated using Eq. (5.24). The same quantities calculated for the [100] direction are plotted in Fig. 5.33. The net signal gain due to stimulated scattering (Fig. 5.33b) has slightly higher values than in Fig. 5.22b, as we have now neglected the nonlinear loss terms.

Our experimental results for the generation on spontaneous Raman scattering in slow light photonic crystals are therefore in agreement with the theoretical dependence on signal group index, and the dependence of stimulated Raman scattering would need to be tested separately.

In addition, Fig. 5.34a shows the measured spontaneous Raman peak power as a function of the input pump power P_{in} for the same photonic crystal waveguide as in Figs. 5.28b,d, whereas Fig. 5.34b shows the pump power outputting another waveguide (the same as in Fig. 5.28a,c), measured with a power meter between L4 and PBS2 (see Fig. 5.25). At the operating pump powers of the previous figures (green dashed line) we clearly see a deviation from the linear dependence (black dotted line), which might be indicating that the free-carrier lifetime for these samples is higher than the 200 ps estimated from measurements on the samples for tunable delay (section 3.4).

Finally, we show in Fig. 5.35 the same type of spontaneous Raman measurements for waveguides oriented along the [100] direction of crystalline silicon. The total insertion loss for both TE and TM polarisations (Fig. 5.33c-d) is higher in these waveguides than in those oriented along the standard [110] direction (Figs. 5.28c-d and 5.29c-d). However, as we have verified on separate test samples that the 45° bends of 20 μm radius (Fig. 5.23d) should not introduce significant loss, the reason for the lower transmission is not clear and would require further investigation. From the measurements of spontaneous Raman scattering we cannot distinguish any obvious increase of the detected signal in the slow light region, which is also unexpected, especially given the good results obtained for the standard orientation; unfortunately, not many waveguides were available on this particular chip for

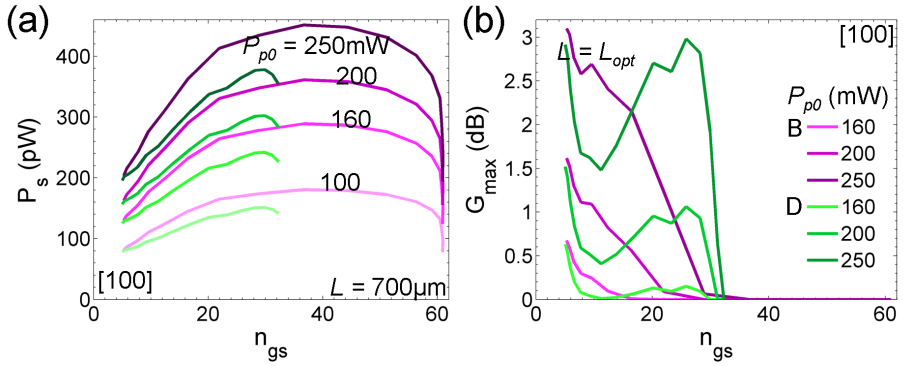

Figure 5.33: (a) Spontaneous Raman signal at waveguide length $L = 700$ μm for designs B (purple) and D (green) for different input pump powers, for waveguides oriented along the [100] direction of crystalline silicon. (b) Net signal gain G_{max} due to stimulated Raman scattering calculated analytically from Eq. (5.24).

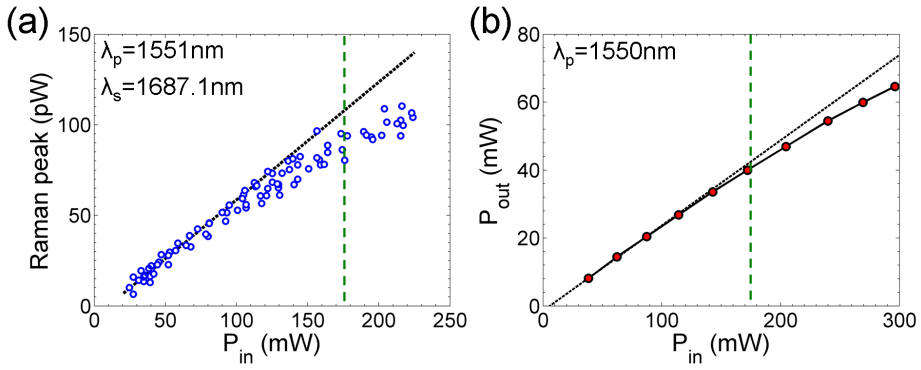

Figure 5.34: (a) Collected spontaneous Raman peak power as a function of the input pump power before lens L3, measured for the same 496 μm-long photonic crystal waveguide as in Figs. 5.28b,d and 5.31b. The black dotted line is a linear fit to the data below 80 mW pump power. (b) Output pump power measured with a power meter between L4 and PBS2 for the same 496 μm-long photonic crystal waveguide as in Figs. 5.28a,c, as a function of the input pump power before lens L3. The black dotted line is a linear fit to the first three data points. In both graphs, the green dashed line indicates the operating point of previous figures.

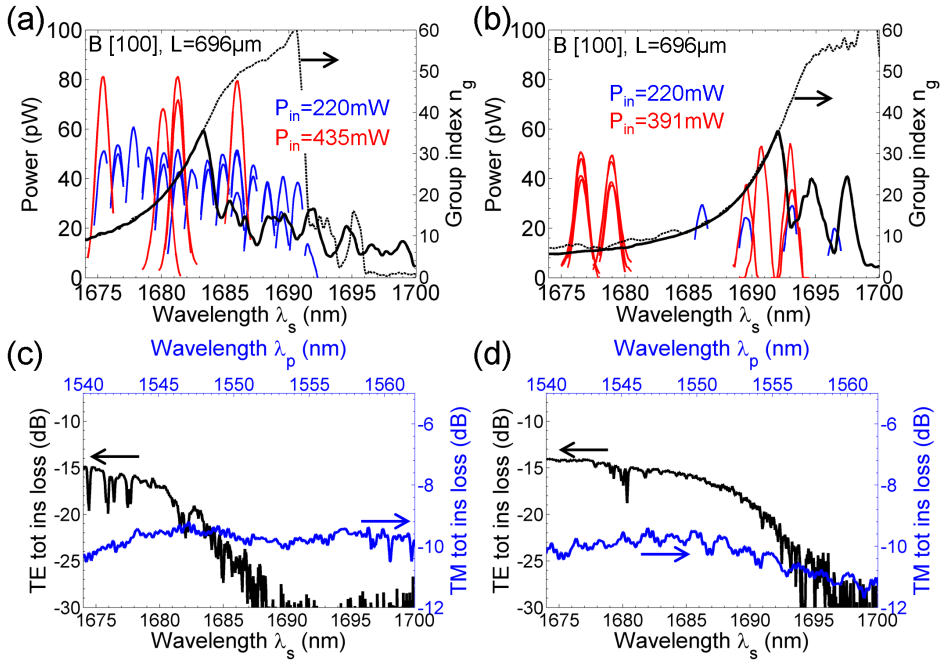

Figure 5.35: (a-b) Measured spontaneous Raman spectra (red and blue curves) and group index (black solid) for two waveguides of type B and length 696 μm on the same chip, oriented along the [100] direction of crystalline silicon, i.e. 45° to the cleaving plane. (c-d) Total insertion loss of the two waveguides in (a-b) for both the TE signal mode (black) and the TM pump mode (blue). The average TM insertion loss is 9.74 dB for (a,c) in the range 1540–1554 nm and 10 dB for (b,d) in the range 1540–1556 nm. By dividing these figures by 2 and neglecting propagation loss, the in-coupled pump power in the photonic crystal is at least 72 mW (blue) and 142 mW (red) for (a) and 70 mW (blue) and 124 mW (red) for (b).

testing. Nevertheless we notice that, at least in the fast light region, the Raman signal levels are comparable to those in the fast light region of the samples aligned along the standard [110] direction (Figs. 5.29a-b), despite the total TE insertion loss being now 5 dB worse. This might be indeed the indication of a higher Raman scattering efficiency in this orientation, and further investigation would be needed to fully understand if this is the case, but due to time restrictions it was was not possible to conclude this study.

5.4.4 Pump-probe experiment

With this section we conclude the experimental study of this chapter by showing some preliminary results for the experiment on stimulated Raman scattering. We tested some slow light photonic crystal waveguides by recording a broadband probe signal spectrum with and without the pump, and Fig. 5.36 shows the outcome of the measurements in three different wave-

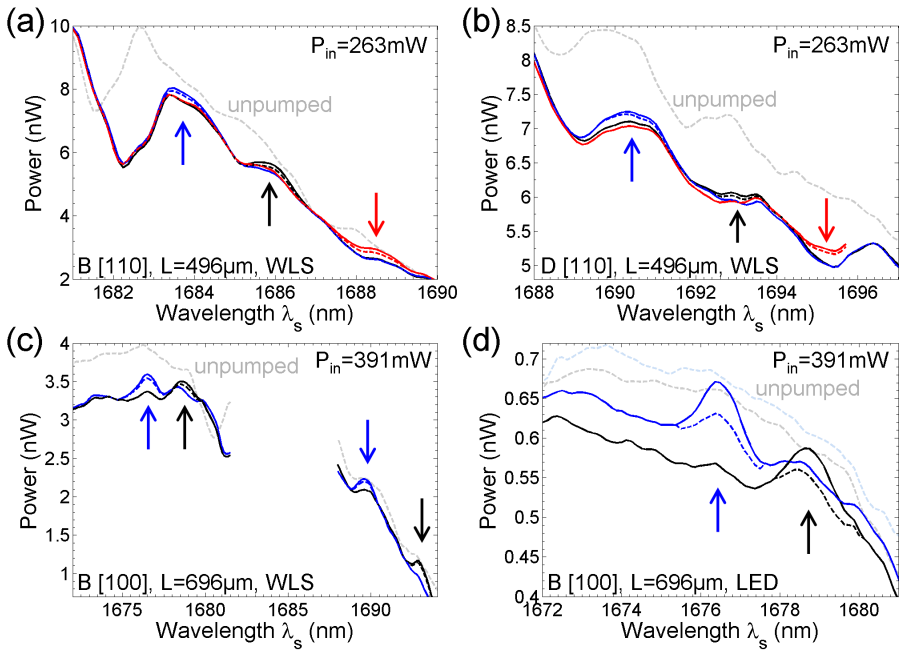

Figure 5.36: Spectra measured in pump-probe experiments for three different waveguides (data in (c-d) are for the same waveguide) aligned along the standard [110] direction of crystalline silicon (a-b) and along the [100] direction (c-d). The pump is a broadband white light source in (a-c) and a low-power LED in (d). Grey dashed curves are measured in the unpumped condition. Solid curves are recorded in the pumped condition for varying pump wavelength, and the corresponding Stokes wavelength is indicated by the arrows. The coloured dashed spectra are obtained by subtracting the spontaneous signal at the same wavelength, and indicate the occurring of stimulated scattering. In the case of (d), the blue and black spectra do not overlap when they are both off-resonance, due to misalignment of the precision stages during the measurement.

guides (the spectra in Figs. 5.36c-d are taken from the same waveguide).

In these graphs, the solid curves indicate the measured spectra in the presence of the pump for a few different pump wavelengths, and the Stokes position is indicated by the corresponding arrows. Let us consider Fig. 5.36a as an example (the waveguide is the same as for Figs. 5.28b,d). In the wavelength range indicated by the *red* arrow, the red solid curve has higher values than the blue and black spectra: at these wavelengths the blue and black curves are off-resonance with the pump, whereas the red curve is on-resonance, and the difference between spectra is due to a small Raman contribution appearing in the red curve. Furthermore, the red dashed curve is obtained by subtracting the corresponding spontaneous signal (measured by blocking the signal arm before PBS1): the red dashed curve is still higher than the solid blue and black spectra, which means that by pumping in the

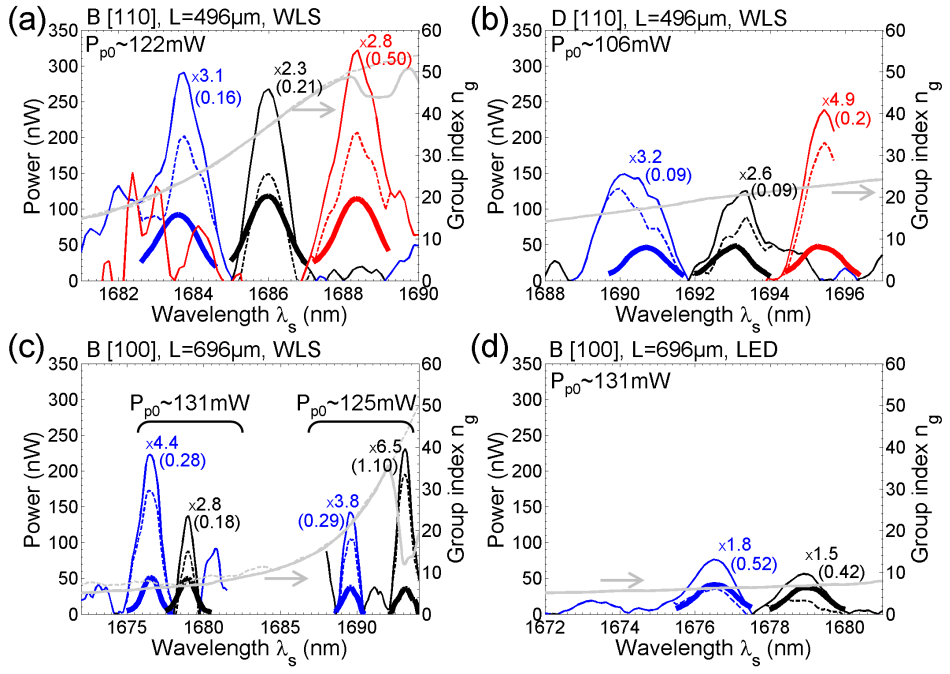

Figure 5.37: Raman output generated with (solid thin curves) and without (i.e. spontaneous, solid thick curves) probe signal, for the same waveguides as in Fig. 5.36. The solid thin curves are derived from the data of Fig. 5.36 by subtracting the off-resonance signal from the on-resonance signal, and demonstrate that in the presence of a probe the waveguides generate more Raman output than without the probe, a manifestation of stimulated scattering. The dashed curves are the difference between the on-resonance dashed curves from Fig. 5.36 and the off-resonance spectra, and therefore indicate the amount of Raman signal due only to the stimulated effect. Numbers next to each curve indicate the ratio between the peaks of the thin and thick solid curves (top) and the on/off-resonance gain in dB (bottom, in brackets). The group index is overlapped in grey. For each graph, the estimated in-coupled pump power is given by P_{p0}.

presence of the seed signal we generate more Raman output than without the signal. This is a clear indication of stimulated Raman scattering.

The fact that the waveguides generate more Raman signal with the broadband probe than without is more evident from Fig. 5.37, where the data from Fig. 5.36 are further elaborated. The thick curves indicate the spontaneous signal (i.e. the total Raman output generated *without* input probe), whereas the thin solid curves are the difference between the on-resonance and off-resonance spectra from Fig. 5.36, and therefore indicate the total Raman output generated *with* the input probe; the dashed curves are the difference between the two previous curves and, as such, they represent the amount of Raman signal due only to the stimulated scattering. The overlapped grey curves indicate the group index. Finally, the numbers

next to each curve identify the ratio between the peak Raman signal generated with and without probe (top number) and the on/off-resonance gain expressed in dB (bottom number, in brackets). So for the red curves in Fig. 5.37a, for example, in the presence of the probe the waveguide generates ~2.8 times more Raman signal than in the spontaneous case, and the on/off-resonance gain is ~0.5 dB.

In general, the on/off-resonance gain seems to increase with slowdown factor, but it is difficult from these data to identify any obvious trend. Note that the data of Fig. 5.37d are for the same waveguide and same pumping conditions as the fast light data of Fig. 5.37c, and the only difference is in the source used as a broadband probe: a low-power LED (Fig. 5.37d) rather than a more powerful white light source (which is pulsed, so that the recorded spectrum is the time-average detected by the OSA). The total amount of Raman signal is higher when the test is performed with a more powerful probe, as we would expect for stimulated scattering.

The Raman signal generated through stimulated scattering, however, is clearly not sufficient to compensate for the linear propagation loss intrinsic to the waveguide. The grey dashed curves in Fig. 5.36 are the probe spectra recorded when the pump is off, and we see that when pumping the curves are generally attenuated by up to 0.7–1 dB, most likely due to the effect of free carriers (as already mentioned in section 5.4.3, see Fig. 5.34), and also red-shifted by up to ~0.7 nm, which indicates that the sample is heating up, despite the pump propagating in a fast mode.

It seems that we are still far from loss compensation even if the estimated in-coupled pump powers (see Fig. 5.37) are not much smaller than those used for the calculations of Fig. 5.22. However, note that the best results at low pump powers in Fig. 5.22 were for a waveguide of type D oriented 45° to the cleaving plane, which we are not reporting experimentally here: we did test few of such designs in the chip with waveguides oriented at 45°, but we could not obtain any conclusive result, partially also due to instabilities of the translation stages in the optical setup.

In all the measured waveguides in general, the nonlinear loss is more significant than originally thought, and it tends to be equal to, or higher than, the introduced Raman gain (Fig. 5.36). We have already highlighted in section 5.4.3 that the free-carrier lifetime in these samples may be higher than that measured for the tunable delay samples, but it is also possible that the Raman gain coefficient is lower than the already conservative 20 cm/GW value used for the calculations in section 5.3. In addition, the pump itself is likely to be subject to higher propagation loss (both linear and nonlinear) than those assumed in the calculations, which also has a noticeable effect on the final signal gain (see Fig. 5.17b). Note, however, from Fig. 5.36, that the Raman gain appears to be compensating better for the pump-induced losses in the slow light region of the waveguide oriented along the [100] direction, which may be another indication of stronger scattering efficiency in this

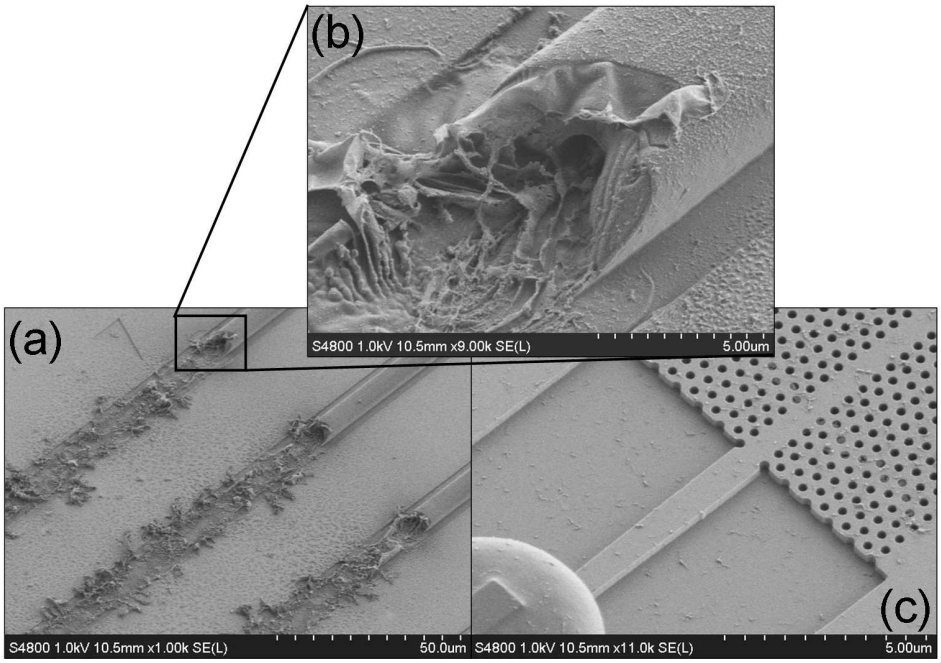

Figure 5.38: SEM images of waveguides burnt due to failure of the inverse tapers. (a) Burnt SU-8 waveguides. (b) Close-up of (a). (c) Access nanowire of the same waveguide as in (b): the nanowire, not suspended, is not the prime cause of failure of the waveguide, which was instead the case for Fig. 4.15.

direction.

We have also tried to pump at higher powers than those reported here, but then the waveguides burn due to catastrophic optical damage, as shown in Fig. 5.38. The failure point, however, does not correspond to a suspended nanowire, which was the case discussed in Chapter 4 (Fig. 4.15): now the access nanowires are not under-etched (Fig. 5.38c) and remain intact — we can in fact couple higher powers than in samples where the nanowire is suspended — but it seems that the inverse tapers burn at some point under the SU-8 (which we have verified by inspection under the optical microscope).

Overall, from these measurements performed on only few samples it does not appear that stimulated Raman scattering is a viable solution for loss compensation *in the standard waveguide orientation*, as the pump introduces more gain than loss already at these relatively low powers; the unsuitability of this waveguide orientation for the long wavelengths tested (~1686 nm) was already suggested by the numerical analysis of section 5.3.5, where even along the better 45° orientation the simulated gains were not high. The case for waveguides oriented at 45° to the cleaving plane is not as clear: on one hand, the nonlinear loss terms play a significant role, and at the pump

powers coupled inside the waveguide we would have expected to achieve a more significant gain contribution; on the other hand, it appears that in this orientation the Raman gain can compensate at least for FCA at these long wavelengths. We believe it would be worth to investigate whether improvements in the transmission of these structures are possible, and then test these designs scaled for operation with a 1435 nm pump and a 1550 nm Stokes signal, as analysed in section 5.3.4, where the model predicted achievement of significant positive net gain: given the importance of nonlinear losses, we do not expect positive net gain to be possible unless steps are taken to reduce the free-carrier lifetime, but nevertheless the smaller waveguide cross-section will correspond to a stronger Raman interaction, and it may indeed be possible to observe some loss compensation for this orientation.

5.5 Summary

The Raman scattering is a well-known optical effect resulting from the interaction of light with the vibrational modes of a material system and, in its stimulated form, has been widely demonstrated in the past decade as a possible solution to achieving optical gain in silicon photonic chips. An initially very intense research activity has allowed to overcome detrimental effects such as absorption from free carriers, and to progress in very few years from the first observation of spontaneous Raman scattering from SOI waveguides to many demonstrations of net optical gain and silicon Raman lasers. The need to further reduce the required pump powers has subsequently pushed research into looking for alternative ways of enhancing the nonlinearity; as slow light photonic crystal waveguides had already been investigated, with success, for increasing the efficiency of other nonlinear optical effects, some authors have started studying the possibility of enhancing also the Raman effect, with very promising and exciting estimates [40]. These estimates, however, were obtained by neglecting the influence of losses and by assuming group index values that are not realistically achievable in practical devices. When slowing down light, other detrimental effects compete with the enhancement of the positive Raman gain, and the need to better understand what performance can be obtained in real-life structures has motivated the study presented in the first part of this chapter.

We have approached the Raman effect from a new angle: our focus was not only on how to use slow light as a means of enhancing the Raman nonlinearity, but also on how to employ the latter as a way of introducing optical gain in lossy slow light photonic crystal waveguides, with the prospect of improving the performance of slow light devices such as the tunable delay line investigated in Chapter 3.

First of all, we needed to understand how stimulated Raman scattering depends on the slowdown factor of both the pump and the Stokes beams: in

this chapter we have discussed and extended a simple and intuitive model previously introduced by other authors. This simplified model, however, fails to describe the full picture of the behaviour of real-life photonic crystal waveguides. In particular, we have discussed how the realistic dependence of linear propagation losses and mode profiles on the slowdown factor strongly influence the final result on the net signal gain. We have also shown how, by appropriately orienting the waveguide in the (001) plane of crystalline silicon, we can take advantage of the anisotropy of the Raman effect in silicon to achieve higher scattering efficiencies, and consequently higher signal gains, without the need of changing the photonic crystal design.

We have also found that not all slow light devices are the same, and that one must pay careful attention to these realistic effects in order to choose the most appropriate design, and especially if the aim is to achieve signal amplification. After reconsidering the exceptionally high gains proposed by previous authors, we have estimated that slow light photonic crystal waveguides may still have a potential to achieve similar continuous-wave gains to those already demonstrated, but over shorter lengths and at lower pump powers. Our estimates were therefore also reasonably promising for our purpose of achieving loss compensation.

We have then discussed the experimental work conducted in fabricated silicon slow light waveguides to study both spontaneous and stimulated Raman scattering. We were able to establish a clear group index dependence of the strength of the observed spontaneous Raman scattering in waveguides oriented along the standard [110] direction of crystalline silicon, in agreement with the theoretical expectations — where again we have taken into account realistic losses and mode profile effects — and we collected some preliminary results on a possible increase of the Raman interaction between a TM pump and a TE signal for waveguides oriented 45° to the cleaving plane.

In addition, we have also shown preliminary experimental results on stimulated Raman scattering through a pump-probe experiment: we have verified that the Raman output generated in the presence of the probe is stronger than without the probe and, to our knowledge, this is the first time that stimulated Raman scattering from a pump-probe experiment is reported from a silicon photonic crystal waveguide. We have found, however, that at the power levels coupled into the photonic crystals, the detrimental effects from FCA are higher than originally estimated, whereas the Raman gain is not always sufficiently strong to compensate for such nonlinear loss, especially in the standard [110] direction. In terms of compensation of the pump-induced loss, however, the data from Fig. 5.36 appear to be more promising for the slow light region of waveguides aligned along the [100] direction than for the [110] direction. We did attempt to pump at even higher pump powers, but while very significant improvement in terms of both coupling efficiency and sustainable pump powers has been achieved with the

inverse taper couplers with respect to the first attempts (Chapter 4), this is still not enough, as we have found that the inverse tapers burn above a certain threshold.

Overall, these measurements do tell us that detrimental effects may be playing a more significant role than what estimated in the model, while the Raman gain may be even lower than that used in the calculations, possibly due to the combination of a low Raman gain coefficient g_R and pump losses higher than the those estimated. Our preliminary tests on only a very few samples, however, are not sufficient to draw hard conclusions on the possibility of achieving compensation of the linear loss, especially for the 45° orientation and if measures are taken to reduce the free-carrier lifetime. Amongst the features needing improvement, we can identify the 45° design, which currently transmits significantly less power than in the standard orientation: it is possible that tilting the virtual (u, v) coordinate system during e-beam writing (Fig. 5.24d) introduces additional disorder and non-idealities in both the photonic crystals and inverse tapers, thus increasing loss; it would be worth testing whether adopting a configuration of the type of Fig. 5.24c corrects this issue, as having the (u, v) coordinates aligned to the real (x, y) coordinates of the sample stage is the most natural operating condition of the e-beam machine. The issue of aligning reliably the (u, v) coordinates to the [100] direction of crystalline silicon, discussed in section 5.4.1, could possibly be solved by first aligning the (u, v) system to the tilted bottom edge of the sample and then imposing a 45° rotation. Improvement should also be achievable for the coupling structure: it may be possible that reducing the side-wall roughness of the tapers will also increase the pump threshold for them to fail, but alternative structures such as grating couplers should also be considered.

Finally, it would be also very interesting to explore ways of slowing down the pump at the same time as the Stokes signal, as this would reduce the pump thresholds for achieving positive gains. Although not shown here, we have started investigating a few possibilities, such as W0.x waveguides, coupled photonic crystal waveguides and the use of the odd TE mode as a pump. We did identify a few promising designs, but due to time restrictions we could not optimise these structures and investigate good couplers to achieve useful transmission, as we needed to concentrate the efforts on the more straightforward mode configuration presented here, which was more likely to produce an experimental result in shorter time scales.

Chapter 6

Photonic crystal cavities for Raman scattering

In Chapter 5 we have studied the enhancement of stimulated Raman scattering due to slowing down the light in engineered photonic crystal waveguides. Optical cavities, and in particular photonic crystal cavities, are other structures where the optical intensity is enhanced, due to light making multiple round-trips [50]. This intensity enhancement may also be used to increase nonlinear effects, and for the particular case of stimulated Raman scattering such possibility was proposed in a detailed numerical study by Yang and Wong [43]. In order to achieve a useful device, operating at low pump powers, the first important challenge to address is the definition of a photonic crystal cavity design supporting two modes at the correct spacing for Raman scattering in silicon.

In this chapter we will initially introduce some basic notions about photonic crystal cavities design, in particular the concepts of Q-factor, gentle confinement and far-field optimisation. We will also discuss the progress to-date in the study of Raman scattering in photonic crystal cavities and then we will explore both numerically and experimentally the possibility of achieving a design supporting resonances for both the Raman pump and the Stokes signal.

6.1 Light confinement in photonic crystal cavities

In section 2.1 we have discussed how the creation of defects in an otherwise perfectly periodic photonic crystal introduces confined modes inside the photonic band gap: light in these modes decays evanescently inside the perfect crystal and is therefore pinned to the defect. Throughout our study we have mainly focussed on line defects obtained by removing an entire row of holes, thus forming an optical waveguide, but if only one or few holes are removed, the defect mode is surrounded by reflective walls, and therefore an

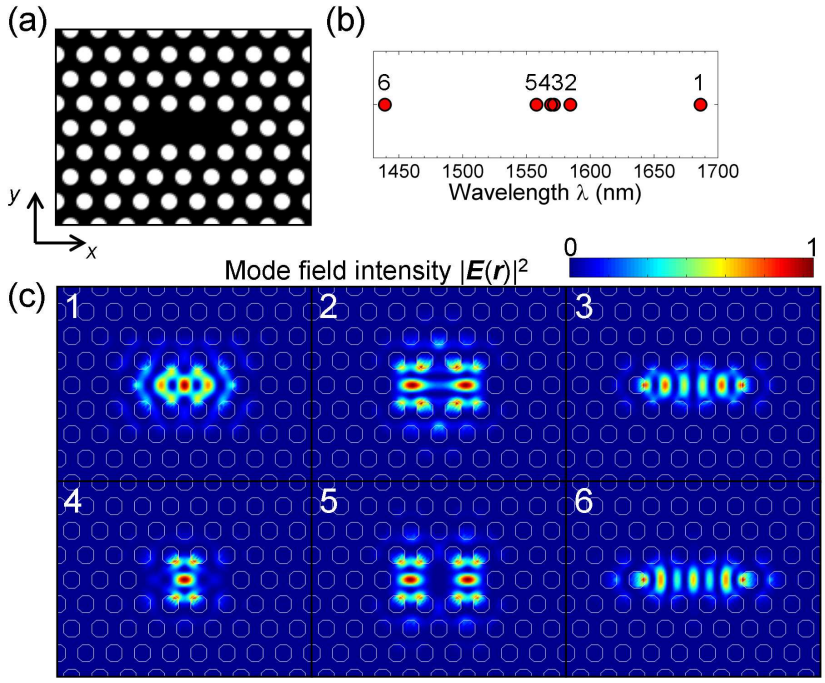

Figure 6.1: (a) Cross-sectional view of an L3 photonic crystal cavity, obtained by removing three rows of holes along the ΓK direction of a triangular lattice of air holes in silicon; the period is $a = 465$ nm and the hole radius is $r = 0.286a$. (b) Resonant frequencies of the six cavity modes, calculated through a three-dimensional MPB simulation. (c) Intensity profiles of the six cavity modes.

optical cavity is formed.

The removal of n holes along the ΓK direction of a triangular lattice results in what is called an Ln cavity. The cross-sectional view of an L3 cavity is shown in Fig. 6.1a. This type of cavity supports six modes, whose resonant frequencies and spatial intensity profiles are shown in Figs. 6.1b-c.

The ability of an optical cavity to store light is summarised by a dimensionless figure of merit called quality factor Q. Assuming no absorption from the cavity material, the stored field decays in time due to reflection losses at the interfaces of the cavity. The field decay in time can be described by a sinusoidal function modulated by an exponential [45], as $\exp\left[-i\left(\omega_0 - i/\tau\right)t\right]$, where ω_0 is the central frequency of the mode and τ a time decay constant. This time-varying field corresponds to a Lorentzian *intensity* distribution in the frequency domain, proportional to $1/\left[(\omega-\omega_0)^2 + (1/\tau)^2\right]$, so that the distribution's full width at half maximum (FWHM) is given by $\Delta\omega = 2/\tau$. The longer the time decay or lifetime τ, the narrower the peak in the frequency domain. Normalisation to the central frequency ω_0 leads to the

definition of the Q-factor:

$$Q \triangleq \frac{\omega_0}{\Delta\omega} = \frac{\tau\omega_0}{2}.$$

In addition, $1/Q$ describes a dimensionless energy decay rate and is equivalent to [45]

$$\frac{1}{Q} = \frac{P}{\omega_0 U}, \tag{6.1}$$

with P the radiated power and U the electromagnetic energy inside the cavity.

If the cavity is characterised by more than one loss mechanism, the total radiated power P is given by the sum of the power lost through each mechanism. In the case of an ideal photonic crystal cavity defined in a suspended silicon slab, the only two ways light can escape the cavity are in-plane ($\parallel$) and out-of-plane ($\perp$) leakages, so that $P = P_\parallel + P_\perp$ and thus from Eq. (6.1) follows

$$\frac{1}{Q_{theor}} = \frac{1}{Q_\parallel} + \frac{1}{Q_\perp}.$$

In a real device, loss is also originated from the imperfections of the structure due to fabrication, and therefore we will have:

$$\frac{1}{Q_{real}} = \frac{1}{Q_{theor}} + \frac{1}{Q_{fab}} = \frac{1}{Q_\parallel} + \frac{1}{Q_\perp} + \frac{1}{Q_{fab}}.$$

In-plane confinement is due to the photonic band gap, and so $Q_\parallel$ can be made arbitrarily high simply by adding more rows of holes to the photonic crystal [250]. The total theoretical Q is therefore usually limited by the vertical, out-of-plane confinement $Q_\perp$, which is determined by the total internal reflection at the interface between the silicon slab and air.

The vertical confinement of a cavity mode is usually studied by Fourier-transforming its spatial electric field profile into the k-space, as shown in Fig. 6.2 for the fundamental mode of the L3 cavity of Fig. 6.1. Those components whose in-plane $\mathbf{k}$-vector $\mathbf{k}_\parallel = \mathbf{k}_x + \mathbf{k}_y$ is smaller (in modulus) than $2\pi/\lambda_0$ (λ_0 being the wavelength of light in air) fail to satisfy the total internal reflection condition, and are therefore leaky components [251], directly related to the Q-factor [252]. The leaky region is indicated in Fig. 6.2b by a white circle of radius $2\pi/\lambda_0$, which is simply a cut of the light cone at frequency ω_0. In an L3 cavity, the fundamental mode is usually the one with the highest Q-factor [253].

As the L3 cavity is a quite basic structure, different types of designs have been explored to improve the Q-factor. Most notably, two waveguide-based designs, the hetero-structure cavity [26] and the width-modulated cavity [27], achieve the best performances in terms of light confinement, with experimentally demonstrated Q-factors up to few millions [254, 255].

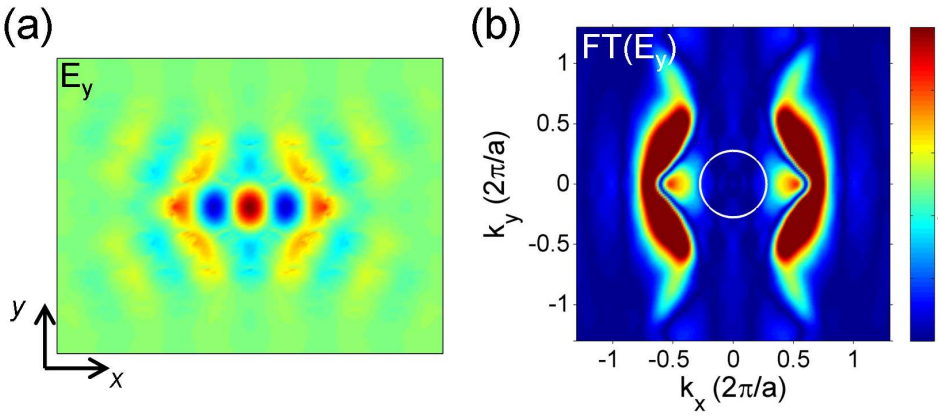

Figure 6.2: (a) Electric field profile E_y of the fundamental cavity mode shown in Fig. 6.1. (b) Two-dimensional Fourier transform of (a). The light cone is marked by a white circle: the area inside the light cone is the leaky region.

Nevertheless, the simpler L3 cavity is widely used, especially as the Q-factors of the basic structure can be improved up to values on the order of 10^5 [256, 257] by applying the simple yet powerful principle of gentle confinement [251]. Briefly, the spatial profile of the electric field inside a cavity can be expressed as a sinusoidal function of wavelength λ modulated by an envelope function whose shape depends on the type of cavity. In the reciprocal k-space, the sinusoidal corresponds to peaks at $k = \pm 2\pi/\lambda$, but the spectrum is then modified by the envelope function. If the envelope is abrupt or, in general, deviates from a smooth Gaussian curve, as in Fig. 6.3a, the k-spectrum will exhibit strong components inside the leaky region (Fig. 6.3c), even if the main peaks due to the sinusoidal are outside of this region. If the cavity structure can be modified to correspond to a Gaussian envelope, i.e. to a more gentle confinement (Fig. 6.3b), the leaky components can be highly suppressed (Fig. 6.3d). In L3 cavities this can be achieved simply by fine-tuning the position and/or size of the first holes adjacent the cavity in the x-direction (see Fig. 6.4a, purple holes, for the case of an L5), in order to increase the Q-factor while maintaining a small modal volume V, as ultimately it is the Q/V ratio that determines the strength of the light-matter interaction.

One important issue of high-Q cavities, however, is that their radiation pattern lacks directionality, making it difficult from a practical point of view to in- and out-couple light from the vertical direction. Such issue is, again, directly related to the cavity's electric field intensity profile, as its Fourier transform determines the far-field emission pattern. The verticality of the emission pattern can be improved by superimposing on the photonic crystal a second periodicity with double the period [258], as shown in Fig. 6.4d, so that the k components at π/a are folded back at $k = 0$ (Fig. 6.4b). As this

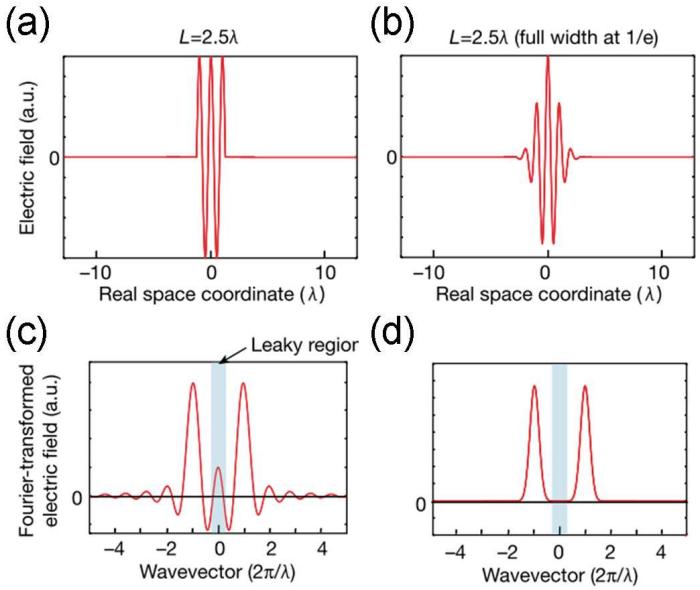

Figure 6.3: Electric field profile (a-b) and corresponding spatial Fourier spectra (c-d), for the case of an abrupt envelope (a,c) and a gentle Gaussian envelope (b,d). In the case of gentle confinement the leaky components inside the light cone (blue shaded region) are suppressed (from [251]. Reprinted by permission from Macmillan Publishers Ltd: Nature [251], copyright 2003).

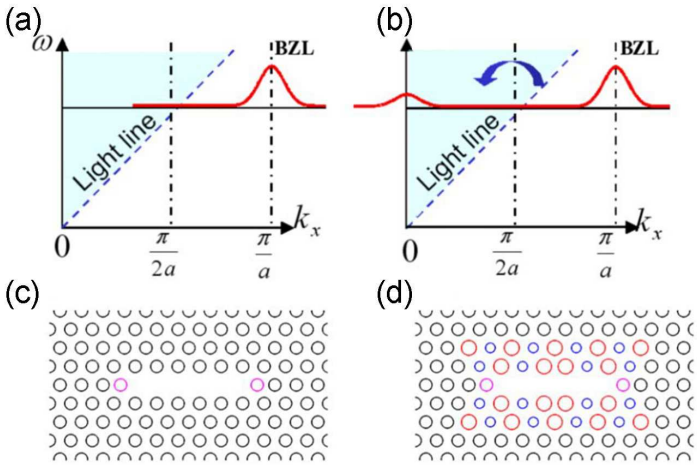

Figure 6.4: Principle of band folding for far-field optimisation: (a) schematic of the near field distribution in the reciprocal space for a cavity optimised for gentle confinement; (b) band-folded distribution. (c-d) Designs implementing the principle of band folding: (c) L5 cavity and (d) modified L5 (emphasized) (from [258]. Copyright 2009 by The American Physical Society).

corresponds to re-introducing components inside the light cone, the Q-factor will decrease, but the amount of leakage can be controlled by appropriately tuning the strength of the perturbation, and an optimum can be reached with useful coupling efficiencies while maintaining Q-factors on the order of 10^4 [259]. This type of far-field optimised structures has been in fact employed to experimentally demonstrate nonlinear effects such as continuous-wave second- and third-harmonic generation from silicon photonic crystal cavities [260].

6.2 Towards Raman scattering in photonic crystal cavities

The enhancement of stimulated Raman scattering from high-Q/V photonic crystal cavities, with prospects for Raman lasing, was first suggested by Yang and Wong in 2005 [44], who predicted scaling of the lasing threshold as $V/(Q_pQ_s)$. This dependence closely resembles that of Raman scattering in slow light waveguides, Eq. (5.16), with the Q-factors replacing the slowdown factors S. The authors proposed using a simple L5 cavity modified for gentle confinement, designed in such way that the fundamental mode, hosting the Stokes signal, has spacing 15.6 THz with one of the higher-order modes, which would host the pump. They reported fabrication of such a cavity, but no measurements, suggesting that matching of the Raman spacing between the two modes could not be achieved in practice. From calculations, they estimated that with a Q_S on the order of 10^4 and Q_p of around 1000, the coupled pump power lasing threshold should be on the order of 300 μW. We note, however, that this calculation is based on the use of a Raman gain coefficient $g_R = 70$ cm/GW; with more realistic values of 10–20 cm/GW, the coupled pump power needed for lasing increases to around 1–3 mW.

The same authors then developed a detailed coupled mode theory framework to study Raman lasing in silicon photonic crystal cavities [43], accounting also for the detrimental effects of TPA, FCA and refractive index shifts from the Kerr effect, FCA and thermal dispersion. Importantly, they note that the cavity design must reach a minimum Q_s for the signal to allow for lasing, and that the required value increases for increasing free-carrier lifetime. This suggests that when designing structures for lasing it is more important to achieve high Q for the signal rather than for the pump: a higher Q_p allows to reach the (in-coupled) lasing threshold for lower input pump powers, but the signal mode determines whether lasing is possible in the first place.

So far, however, the only progress on Raman scattering in photonic crystal cavities is an experimental result reported in 2011 by Takahashi *et al.* [261]: the authors pumped in the 1.5×10^6-Q fundamental mode of a hetero-structure cavity (Fig. 6.5a) and collected light emitted in the vertical

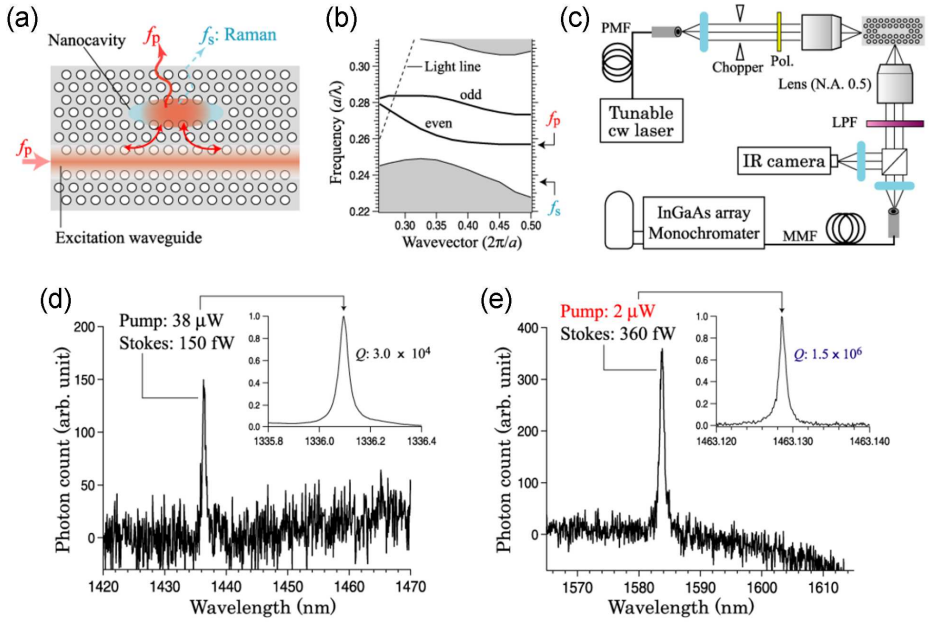

Figure 6.5: (a) Schematic of the pumping configuration employed by Takahashi *et al.* [261] to observe spontaneous Raman scattering: a photonic crystal cavity evanescently coupled to an excitation waveguide. (b) Band diagram of the line-defect waveguide. The high-Q cavity mode is formed just above the band edge of the even guided mode for the L3 cavity, and just below it for the hetero-structure cavity. The Stokes frequency corresponds to a slab mode of the two-dimensional photonic crystal. (c) Experimental setup for Raman spectroscopy; the Raman signal is detected by a liquid-Nitrogen-cooled InGaAs detector. (d) Raman scattering spectrum in an L3 cavity shifted for gentle confinement. Inset: resonant peak for the fundamental mode utilised for pumping. (e) Raman spectrum in the hetero-structure cavity with Q of 1.5 million. Inset: the fundamental resonant mode (from [261]. Reprinted with permission).

direction with a very sensitive equipment (Fig. 6.5c), detecting 360 fW of spontaneously generated Raman signal with an estimated coupled pump power of only 2 μW (Fig. 6.5e). In this experiment, which is the first demonstration of Raman scattering from silicon photonic crystal cavities, the authors estimated that the scattering efficiency from the ultra-high-Q hetero-structure cavity is 46 times higher than from an L3 cavity with Q on the order of 10^4 (Fig. 6.5d).

Note from the band structure of Fig. 6.5b, however, that by pumping in the fundamental mode of the cavity, the Stokes signal is not supported by another cavity mode, and in fact is not even inside the band gap. The numerical predictions by Yang and Wong [43] of operation of a photonic crystal cavity laser for $Q \sim 10^4$ assume that such Q is achieved for the *signal*, which is the most important for achieving lasing. Therefore, the

first step to proceed along the direction indicated by Yang and Wong is to design the cavity in such way that one of the higher-order modes matches the Raman shift of silicon with the fundamental mode. While this was proposed numerically by Yang and Wong with an L5 cavity [44], to our knowledge there has been no experimental demonstration of a photonic crystal cavity supporting two matching modes with the correct frequency spacing,[1] which is a challenge that we will investigate in the remaining of this chapter.

6.3 Matching the cavity resonances to the Raman shift

If we rely only on the basic Ln cavity, it may be difficult already in the design phase to obtain a suitable mode matching in the correct operating wavelength range. In addition, even if a design could be numerically simulated, in practice there are always fabrication tolerances that must be taken into account, such as variations of the refractive index, slab height and hole radius and position, which in the end result in the cavity peaks being shifted from the design. We see these effects also in waveguides, where fluctuations of the fabrication conditions may cause waveguides with the same design to have dispersion curves shifted in wavelength not only from sample to sample, but also on different parts of the same chip.

Therefore, we must introduce a tuning parameter, an additional degree of freedom to ease the design in the first place, but which can also be used in practice to tune the fabricated devices and account for fabrication tolerances. One powerful tool to alter the modes spacing is tuning of the position of the holes, as we have already applied in the case of slow light waveguides for dispersion engineering (see Chapters 2 and 5). When tuning the position of specific holes we can expect, in general, that all modes will be affected, but as different modes have different profiles (Fig. 6.1c), they will sense differently the alteration and thus shift at different rates.

As an example, we can start by tuning the position of the first holes adjacent the cavity in the x-direction; these holes are marked as holes 0 in Fig. 6.6a (red). Figure 6.6b shows the result of introducing a shift $s_{0,x}$ of these holes along the x-direction, where the shift is defined positive towards the cavity centre, consistently with the notation used for slow light dispersion engineering (section 2.2). The red circles indicate the wavelength of the fundamental cavity mode (mode 1), whereas the yellow circles correspond to modes 2–5 (mode 6 is too blue-shifted). The vertical dotted line, in purple, indicates our zero condition, i.e. the basic L3 cavity of Fig. 6.1. Referring to

[1] Perhaps the work by Checoury *et al.* [241] on the first demonstration of stimulated Raman scattering from silicon photonic crystal waveguides can be considered as an exception: the measured waveguide was so short that it acted as a Fabry-Perot cavity, with two modes matching the 15.6 THz Raman shift.

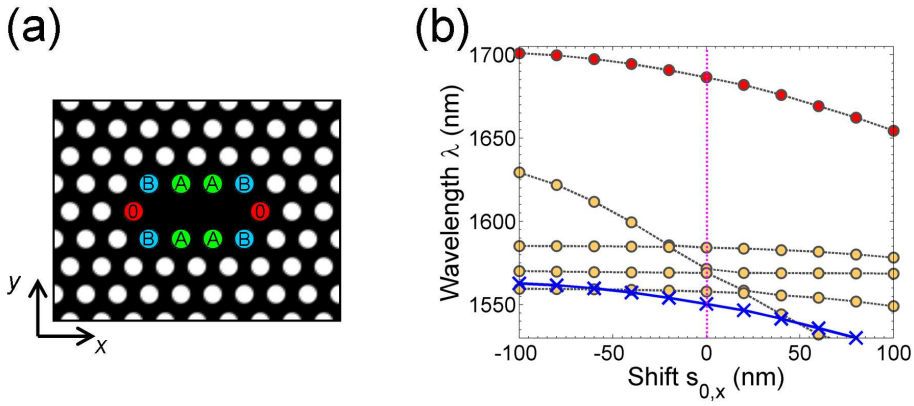

Figure 6.6: (a) Cross-sectional view of an L3 photonic crystal cavity. Holes 0, in red, are used for improving the Q-factor through gentle confinement. Holes A and B, in green and blue respectively, are used to tune the spacing between the fundamental mode and the higher order modes. (b) Example of alteration of the resonant wavelength of the L3 cavity modes by shifting holes 0 along the x-axis. The red circles indicate the fundamental mode, whereas the higher-order modes are indicated by yellow circles. The blue crosses are the projection of the fundamental mode on the corresponding pump wavelength matching the 15.6 THz Raman shift of silicon. The vertical dotted line indicates the unmodified L3 cavity from Fig. 6.1. Shifts are defined positive towards cavity centre.

the modes order of Fig. 6.1c, we can see from Fig. 6.6b that mode 3 is much more sensitive to $s_{0,x}$ than all the other modes, which is directly related to its intensity profile having hot spots at holes 0. The fundamental mode also appears to be more sensitive than the remaining three modes.

Now, the blue crosses in Fig. 6.6b indicate the projection of the fundamental mode at the corresponding pump wavelength dictated by the 15.6 THz Raman shift of silicon. In the zero condition there is no higher-order mode that matches the required pump wavelength. Since different modes are tuned differently with $s_{0,x}$, however, matches are possible with mode 3 for positive $s_{0,x}$ and with mode 5 for negative $s_{0,x}$.

Once ensured that it is possible to obtain wavelength matches, we need to establish whether some high-order modes are better than others in terms of Raman overlap with the fundamental mode. Such overlap is expressed by the effective Raman volume:

$$V_R = \frac{\left[\int_{V_{\text{tot}}} n_p^2(\mathbf{r})\,|\mathbf{E}_p(\mathbf{r})|^2 \mathrm{d}V\right]\left[\int_{V_{\text{tot}}} n_s^2(\mathbf{r})\,|\mathbf{E}_s(\mathbf{r})|^2 \mathrm{d}V\right]}{n_{p,\text{Si}}^2 n_{s,\text{Si}}^2 \Re\left[\int_{V_{\text{Si}}} \mathbf{E}_s^*(\mathbf{r})\,\hat{\xi}^R\!:\!\mathbf{E}_p(\mathbf{r})\,\mathbf{E}_p^*(\mathbf{r})\,\mathbf{E}_s(\mathbf{r})\,\mathrm{d}V\right]}, \tag{6.2}$$

which is equivalent to aA_R, with A_R the effective Raman area defined by Eq. (5.12). This definition is also consistent with those used by McMillan *et al.* [41] and Checoury *et al.* [241]. Figure 6.7 shows the Raman volume

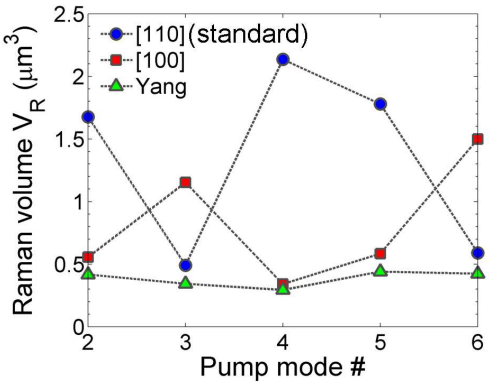

Figure 6.7: Raman volume V_R calculated between the fundamental mode and each of the higher-order modes for the basic L3 cavity of Fig. 6.1. Blue circles and red squares are calculated with Eq. (6.2) in the standard [110] direction of crystalline silicon and in the [100] direction, respectively. Green triangles are calculated according to the definition by Yang and Wong Eq. (6.3) [43], and are reported here for comparison.

V_R calculated for overlap of the fundamental mode (mode 1) with each of the higher-order modes (modes 2–6), where we have used the electric field profiles from the basic L3 design (Fig. 6.1c). The blue circles indicate V_R calculated for cavities oriented along the standard [110] direction of crystalline silicon, whereas the red squares are calculated for the [100] direction. Just as for A_R, the smaller V_R the better, and the data from Fig. 6.7 indicate that pump modes that correspond to high V_R in one orientation, allow for small V_R in the other orientation, and vice-versa. Therefore, the requirement on small V_R is not a limiting factor on the choice of the pump mode, as we can always achieve a small volume by orienting the cavity appropriately.

Note that in Fig. 6.7 the green triangles indicate the Raman volume calculated according to the definition given by Yang and Wong [43]:

$$V_{R(Yang)} = \frac{\left[\int_{V_{\text{tot}}} n_p^2(\mathbf{r})\,|\mathbf{E}_p(\mathbf{r})|^2 \mathrm{d}V\right]\left[\int_{V_{\text{tot}}} n_s^2(\mathbf{r})\,|\mathbf{E}_s(\mathbf{r})|^2 \mathrm{d}V\right]}{n_{p,\text{Si}}^2 n_{s,\text{Si}}^2 \left[\int_{V_{\text{Si}}} |\mathbf{E}_p(\mathbf{r})|^2 |\mathbf{E}_s(\mathbf{r})|^2 \mathrm{d}V\right]}, \tag{6.3}$$

which is similar to Eq. (6.2) except for the omission of the Raman susceptibility tensor $\hat{\xi}^R$. With this definition, not only the anisotropy of silicon is neglected (and so the [110] and [100] directions are equivalent), but also there is not much difference in V_R between the various higher-order modes. It is therefore very important to consider also the effects of $\hat{\xi}^R$ when choosing a design, as it may account for up to a factor 3–4.

Finally, when evaluating which higher-order modes are most suitable to act as a pump, we should also take into account how easy it will be to couple to the mode in the experiment. For an objective lens with given numerical

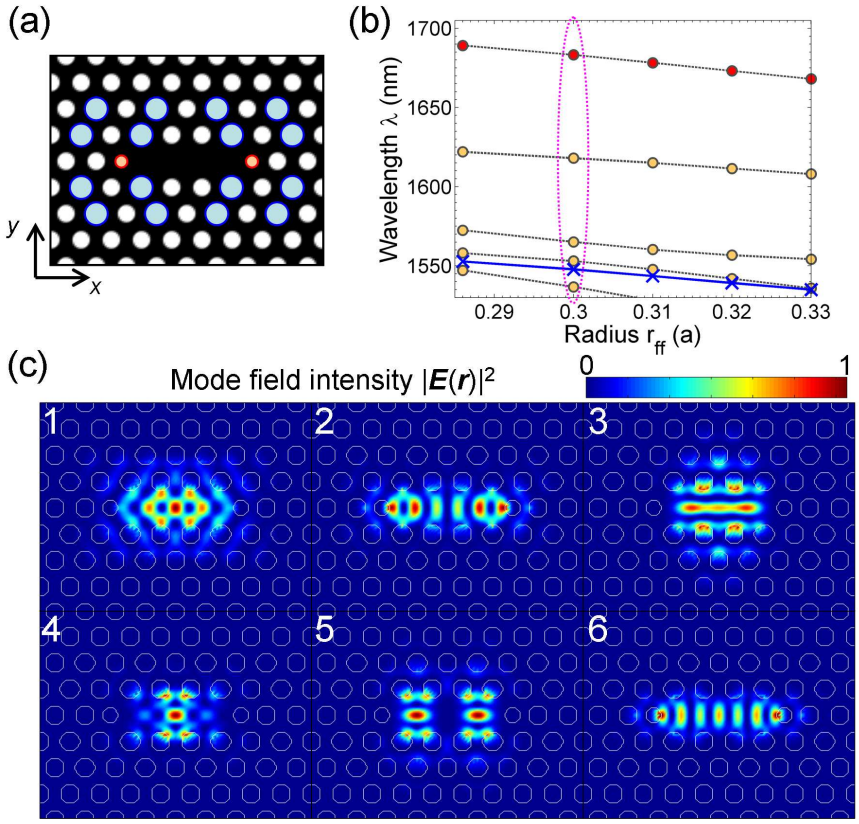

Figure 6.8: (a) Schematic of an L3 cavity designed following gentle confinement and far-field optimisation rules: holes 0 (red) have reduced hole radius $r_0 < r$ and are shifted outwards by $s_{0,x}$; the second periodicity needed for far-field optimisation is achieved by enlarging the blue holes to $r_{ff} > r$. (b) Cavity resonant wavelengths as a function of r_{ff}; other parameters are $a = 460$ nm, $r = 0.286a$, $r_0 = 0.226a$ and $s_{0,x} = -74$ nm. We choose $r_{ff} = 0.3a$ (marked by the dotted curve) as the starting point for optimisation of the modes spacing. (c) Intensity profiles of the six cavity modes for $r_{ff} = 0.3a$; with respect to the basic L3 cavity of Fig. 6.1, modes 2 and 3 are interchanged, and modes 3 and 4 are anti-crossing with each other, which is directly reflected in their shape.

aperture NA (and thus given light cone), different modes will have different coupling efficiency, depending on their radiation pattern. According to the numerical and experimental analysis performed by Fujita *et al.* [253], for an NA 0.4 objective lens the coupling efficiency of modes 3 and 5 of Fig. 6.1c is one order of magnitude smaller than for the other modes. As such, it is not advisable to work with these two modes.

In addition, as the holes 0 are those used to increase the L3 cavity Q-factor through gentle confinement, it is best not to use these holes for fine-tuning of the modes spacing. We will rely on shifting holes A and B instead (see Fig. 6.6a).

In fact, our starting point is not the basic L3 cavity from Fig. 6.1, because we need to include also gentle confinement to improve the Q-factor and far-field optimisation to improve coupling. Following the design rules adopted in Refs. [259, 260], we shifted holes 0 outwards by $s_{0,x} = -0.16a$ and reduced their size to $r_0 = r - 0.06a$; we therefore expect modes 2 and 3 to be interchanged (Fig. 6.6b). We also introduce a second periodicity for far-field optimisation by increasing the size of specific holes to $r_{ff} > r$, as schematically represented in Fig. 6.8a. This increase in r_{ff} also causes a shift in the modes wavelengths, as indicated in Fig. 6.8b, where the calculations are performed for period $a = 460$ nm. We choose $r_{ff} = 0.3a$, so that the increase in radius for this period is only ~6 nm, consistently with the optimum reported in Ref. [259]. Our new zero condition is marked in Fig. 6.8b by a dotted curve. The corresponding intensity profiles are shown in Fig. 6.8c; note, in particular, the alteration of the shape of modes 3 and 4 (corresponding to modes 2 and 4 from Fig. 6.1c), due to anti-crossing.

Starting from this basic design, Fig. 6.9 shows that, simply by shifting holes A (Figs. 6.9a-c) or B (Figs. 6.9d-f) by few nanometres, it is already possible to achieve at least six matches — which from now on we will indicate as A–F — between modes 1 and 4 (Fig. 6.9 displays only the wavelengths of modes 3–5 and the projection of the fundamental mode). In a real sample, the hole size will be another degree of freedom to be used for additional fine-adjustment of the modes spacing. Figure 6.10 shows the Raman volume V_R calculated for each of these designs, with cavities oriented along the [100] direction of crystalline silicon: the values of V_R are all very close to each other.

Note that for each cavity mode one electric field component is spatially symmetric and the other component is anti-symmetric. The polarisation is determined by the symmetric component, as explained in Ref. [253]. The fundamental mode is polarised normally to the cavity axis, whereas mode 4 is polarised along the cavity axis.

Table 6.1 summarises design parameters for all the cavities A–F, including the modes Q-factor calculated with RSoft FullWAVE, a software implementing a finite-difference time-domain (FDTD) numerical method to simulate arbitrary optical structures. The fundamental mode Q-factor of our

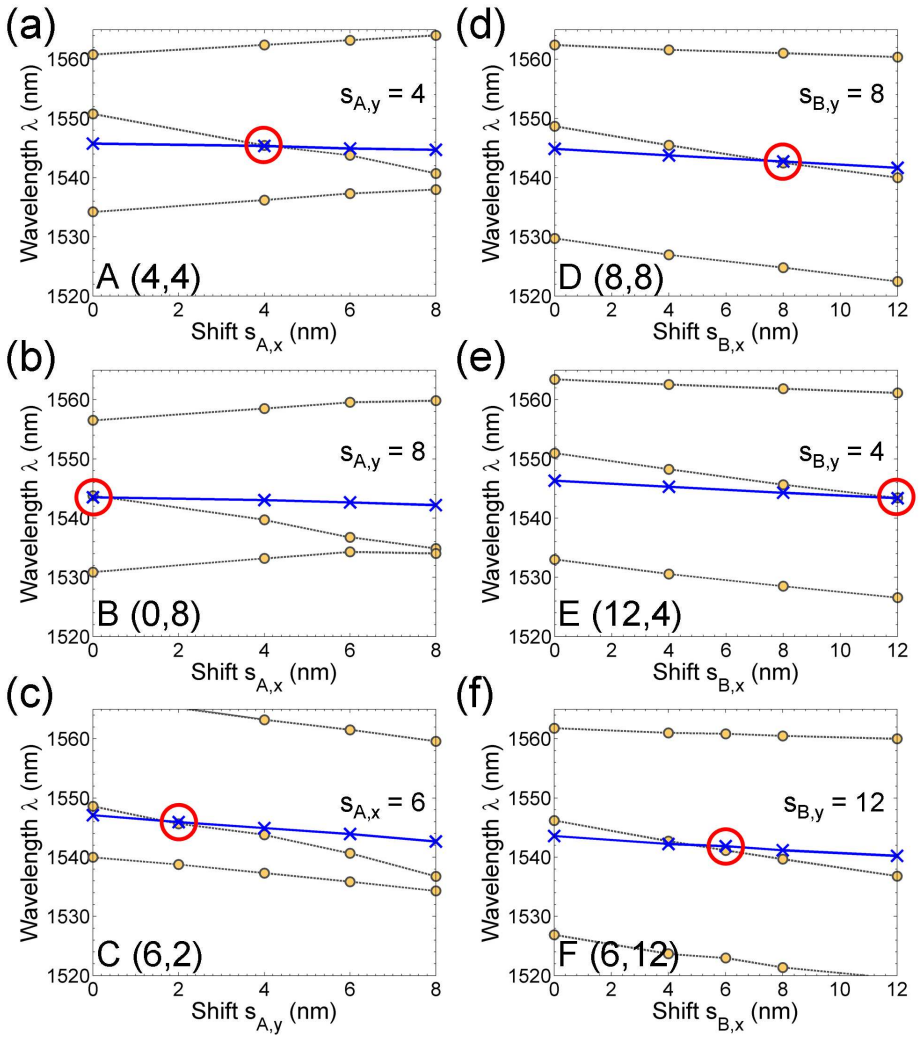

Figure 6.9: Six possible combinations of the shifts of holes A ($s_{A,x}, s_{A,y}$) (a-c) and B ($s_{B,x}, s_{B,y}$) (d-f) in the x- and y- directions, resulting in matching of the fundamental mode with mode 4 from Fig. 6.8c. Here only the resonant wavelengths of modes 3–5 are plotted.

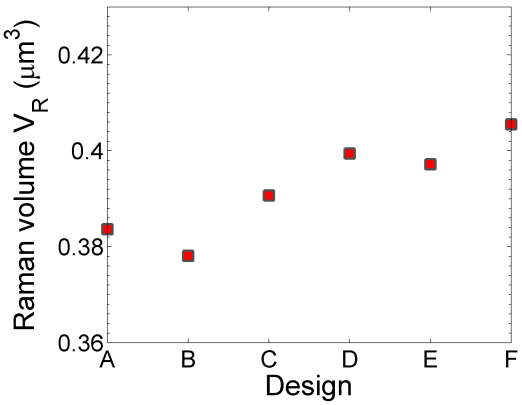

Figure 6.10: Raman volume V_R calculated for each of the six designs A–F using Eq. (6.2) and orientation along the [100] direction of crystalline silicon.

Table 6.1: Summary of the parameters characterising the cavity designs A–F. All designs have period $a = 460$ nm, hole sizes $r = 0.286a$, $r_0 = 0.226a$, $r_{ff} = 0.3a$ and $s_{0,x} = -74$ nm. All shifts and wavelengths are expressed in nanometres.

Design	$s_{A,x}$	$s_{A,y}$	$s_{B,x}$	$s_{B,y}$	λ_p	λ_s	Q_p	Q_s
A	4	4	0	0	1545.5	1680.4	302	42,864
B	0	8	0	0	1543.8	1678.2	239	33,910
C	6	2	0	0	1546.0	1681.2	326	47,534
D	0	0	8	8	1542.5	1677.3	431	13,636
E	0	0	12	4	1543.4	1678.0	496	12,487
F	0	0	6	12	1541.2	1676.2	377	13,208

starting point cavity from Fig. 6.8 was calculated as $Q_s = 31,910$. Shifting holes A inwards as in designs A–C seems to be improving the initial value for Q, whereas acting on holes B as in designs D–F has the opposite effect.

In the experiment we will test, for each set of three, those designs that correspond to the highest Q_pQ_s product, i.e. designs C and E.

6.4 Characterisation of the mode spacing

The photonic crystal cavities were realised on SOI substrate following the fabrication steps described in section 2.3; the fabrication of cavities is actually simplified by the absence of access ridge waveguide, as it does not require definition of an HF mask nor cleaving. The cavities were aligned along the [100] direction of crystalline silicon, and SEM images of the fabricated structures are shown in Fig. 6.11.

The cavities resonances were characterised using a resonant scattering

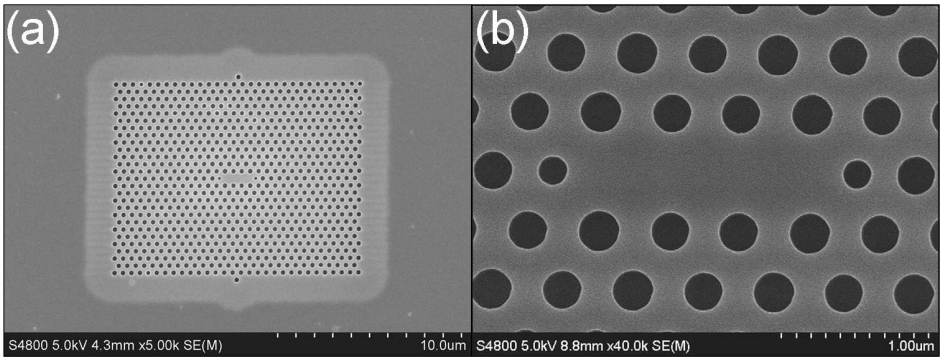

Figure 6.11: SEM images of the fabricated photonic crystal cavities. Structures are aligned along the [100] direction of crystalline silicon.

setup (Fig. 6.12a) similar to that developed in Ref. [263] (Fig. 6.12b), and described in detail in Ref. [262]. Briefly, light from a broadband near-infrared (NIR) source (green path in Fig. 6.12a) is polarised along the x-direction (see notation of Figs. 6.12b-c) by a polariser and vertically coupled to the photonic crystal cavity through a high-NA objective lens. The sample is mounted on a xyz-precision stage, with the cavities oriented at 45° in the xy-plane normal to the propagation direction of the incoming light beam, as schematically shown in Fig. 6.12c. Since the cavity modes are polarised either along or normally to the cavity axis (as indicated by the red arrows in Fig. 6.12c), part of the incident x-polarised light can couple to the cavity, if on-resonance, thus creating a y-polarised component in the re-emitted beam (red path in Fig. 6.12a). This y-polarised component is separated by a polarisation beam splitter (the exchangeable element in the setup of Fig. 6.12a), passes through an analyser and is finally collected by an InGaAs array detector connected to a monochromator. Light that is not on-resonance with a cavity mode is simply backscattered from the sample with no polarisation rotation, and is then filtered out by the polarisation beam splitter and the analyser, generating a small background in the detected signal. The background, however, may increase due to polarisation intermixing generated by the high-NA objective lens [262].

Figure 6.13 shows examples of measured cavity spectra for designs C (Figs. 6.13a-b) and E (Figs. 6.13c-d). Five modes are clearly distinguishable in the wavelength range tested, and they are labelled consistently with the notation of Fig. 6.8c. The high-order modes 2 and 5 (corresponding to modes 3 and 5 from Fig. 6.1c) have lower coupling efficiency than modes 3 and 4, in agreement with the results by Fujita *et al.* [253]. The broad features around 1600 and 1630 nm are part of the background.

In the broad spectra of Figs. 6.13a,c, the fundamental mode of cavity E is more clearly visible than that of cavity C, indicating that the Q-factor

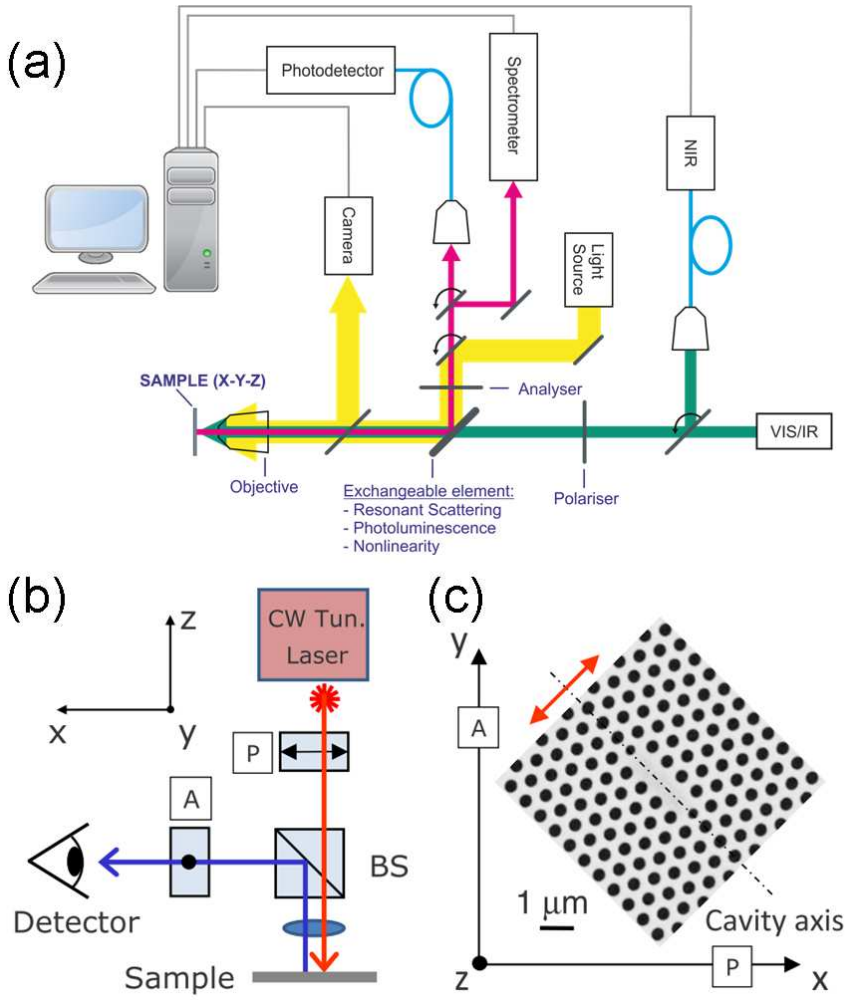

Figure 6.12: (a) Top-view of the multi-functional setup used to characterise the cavity resonances (from [262]). In the case of resonant scattering measurements, the exchangeable element is a polarisation beam splitter. The NIR source consists of an ASE source covering the 1520–1620 nm wavelength range, and a filtered white light source covering longer wavelengths, up to $\lambda > 1700$ nm; the two sources are combined through fibre splitters (not shown here). The light beam from the input arm (green) is focussed on the sample by means of a high-NA objective lens. Alignment to the cavity centre is achieved by illuminating with a visible light source (yellow) and visualising the cavity with a camera. Light re-emitted from the cavity (red) can be directed to a free-space spectrometer (monochromator + InGaAs array detector) or to a fibre-coupled detector. (b-c) Schematic of the sample orientation relative to the polarising optics in the resonant scattering setup (from [263]. Reprinted with permission from [263]. Copyright 2009, American Institute of Physics): the polariser (P) sets the input beam polarisation along the x-axis and the polarisation beam splitter (BS) reflects y-polarised light; the analyser (A) improves the signal-to-noise ratio of y- to x-polarised light in the output arm.

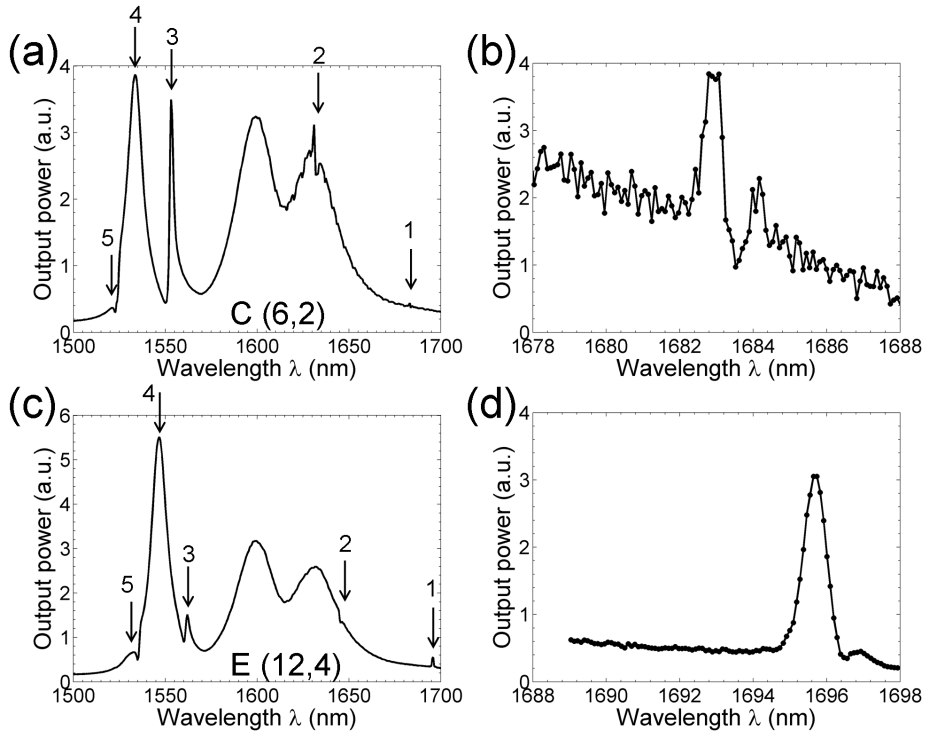

Figure 6.13: Measured cavity spectra for cavities of type C (a-b) and E (c-d). The first five cavity modes are recognisable and indicated with arrows. (b,d) show finer scans of the fundamental modes from (a,c), respectively. The holes shifts $(s_{A,x}, s_{A,y})$ for C (a) and $(s_{B,x}, s_{B,y})$ for E (c) are indicated in brackets, expressed in nanometres. The lattice period is $a = 456$ nm for design C and $a = 464$ nm for design E.

for type E is lower, as expected from the simulations (see Table 6.1). These spectra, however, were recorded with a resolution of ∼1 nm, and therefore cannot be used directly to infer the exact value of the Q-factors: a better resolution for scans with broadband source was not possible with this setup configuration, due to limitations of the monochromator. In general, if a mode lies in the range 1520–1630 nm, we can determine its Q-factor by replacing the broadband NIR source with a narrow line-width tunable laser and scanning its wavelength; light is then detected with a photodetector (Fig. 6.12a). Since the fundamental mode of our cavity designs falls outside the laser wavelength range, however, accurate determination of its Q-factor was not possible with this setup.

Figure 6.14 demonstrates the tuning principle applied to cavities of type C, where the original design $(s_{A,x}, s_{A,y}) = (6, 2)$ nm (black curve) is modified by fine-tuning the position of holes A in steps of 2 nm along the x-direction (Fig. 6.14a) and the y-direction (Figs. 6.14b-d). Due to fabrication tolerances, in the original design the spacing between the fundamental mode and

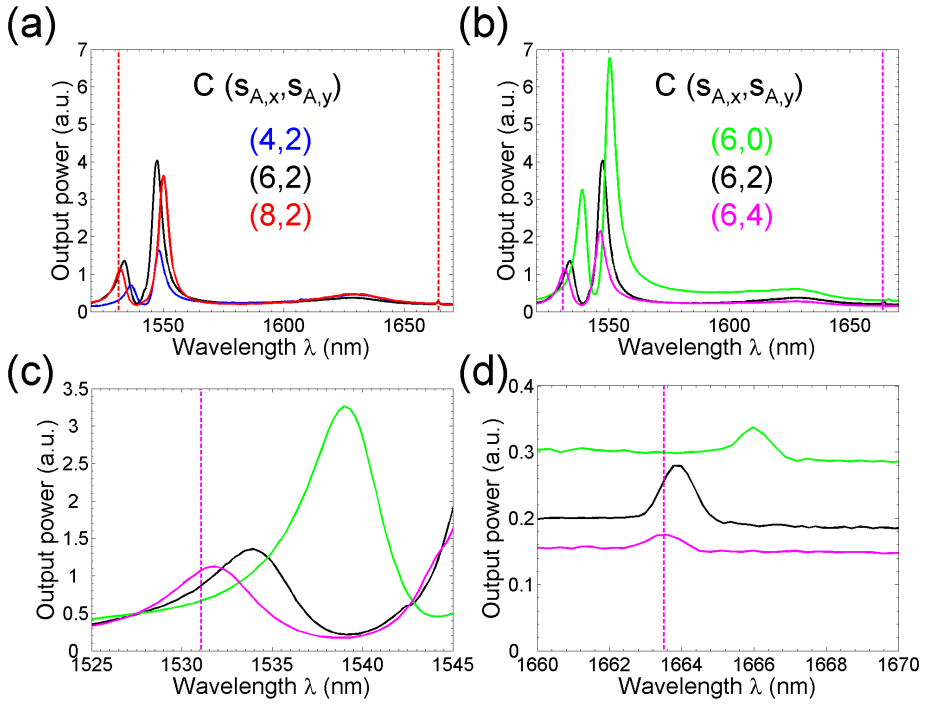

Figure 6.14: Demonstration of the tuning principle applied to design C (period $a = 456$ nm): (a-b) the shifts of holes A from the original design (black curves, see also Table 6.1) are tuned in steps of 2 nm in (a) the x-direction and (b) the y-direction; in each plot, the right dotted line indicates the position of the fundamental mode, whereas the left dotted line indicates where the corresponding pump mode should lie. A good match is obtained in (a) for the red curve and in (b) for the purple curve. (c-d) Enlargements from (b) of mode 4 and the fundamental mode, respectively.

mode 4 is a few nanometres offset with respect to the Raman shift of silicon (indicated by vertical dotted lines); nevertheless, good matches may be obtained with small alterations of the holes position, as for the red spectrum in Fig. 6.14a and the purple spectrum in Fig. 6.14b. This can be seen more clearly in the enlargements of Figs. 6.14c-d. Note that the fundamental mode is less sensitive than mode 4 to variations of the holes positions, and the resonant wavelengths of the two modes are altered at different rates, consistently with the calculations from Fig. 6.9.

Figure 6.15a shows similar spectra for other cavities of type C on the same sample: this set of cavities differs from that of Figs. 6.14a-b only by the radius of the holes, as the two sets were exposed with slightly different doses in the e-beam machine. Comparison between the spectra of Figs. 6.15a and 6.14a-b therefore shows that the radius can also be used as a degree of freedom to fine-tune the modes distance. In the case of Fig. 6.15a a perfect match is obtained for the original design shifts (black curve), as

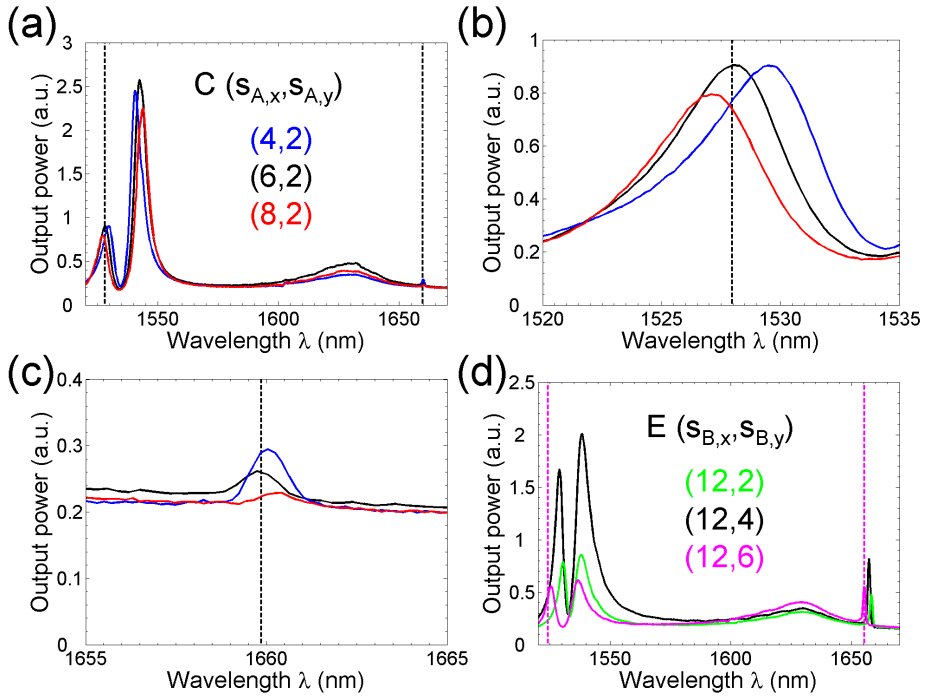

Figure 6.15: (a) Tuning principle applied to design C (period $a = 456$ nm): the shifts of holes A from the original design (black curve, see also Table 6.1) are tuned in steps of 2 nm in the x-direction; the right dotted line indicates the position of the fundamental mode, whereas the left dotted line indicates where the corresponding pump mode should lie. In this case, the original design parameters (black) provide a very good match. (b-c) Enlargements from (a) of mode 4 and the fundamental mode, respectively. (d) Tuning principle applied to design E (period $a = 456$ nm): the shifts of holes B from the original design (black) are tuned in steps of 2 nm in the y-direction; a good match is achieved for the purple curve.

also shown in the enlargements of Figs. 6.15b-c. Finally, Fig. 6.15d shows the tuning principle applied to cavities of type E: although for this radius the distance between modes 1 and 4 does not match the Raman shift of silicon for the original design parameters (black curve), a good match is achieved by adjusting the position of holes B by 2 nm along the y-direction (purple curve).

In conclusion, our work demonstrates the feasibility of using a fine tuning of the holes position to adjust the spacing between resonant photonic crystal cavity modes, in order to match the 15.6 THz Raman shift of silicon. While due to the setup we could not measure accurately the Q-factors for the fundamental mode hosting the Stokes signal, the results from the numerical simulations show that Q-factors of almost 50,000 should still be achievable with these modifications, thus fulfilling the basic requirements from the numerical study by Yang and Wong [43].

6.5 Summary

While in Chapter 5 we have investigated the possibility of enhancing stimulated Raman scattering in photonic crystal slow light waveguides, photonic crystal cavities have also been explored as another type of structures in which enhancement of the optical field intensity over very small volumes may allow useful application of nonlinear effects. In particular, according to the work by Yang and Wong [43], stimulated Raman scattering in photonic crystal cavities holds promise for a miniaturised low threshold silicon laser, provided that both the pump and the Stokes signal are supported by resonant modes. In the first and only experimental demonstration of Raman scattering from photonic crystal cavities by Takahashi *et al.* [261], however, only the pump wavelength corresponds to a cavity mode, whereas the spontaneous Stokes signal is generated at a wavelength outside the photonic band gap. Precise matching of the modes spacing to the Raman shift of silicon in the 1550 nm wavelength range is a challenge in real devices, due to fabrication tolerances.

We have shown how the positions of the holes adjacent a simple L3 cavity can be used as tuning parameters of the spacing between the fundamental mode and the higher-order modes. The key point is that as different cavity modes have different spatial intensity profiles, they sense differently the alterations to specific holes and therefore are shifted in wavelength at different rates. We have shown that while achievement of a low effective Raman volume V_R is not a limiting factor in the choice of the type of pump mode, some restrictions are instead posed by the in- and out-coupling efficiency dictated by the mode's radiation pattern. Starting from an L3 photonic crystal cavity, modified for gentle confinement and far-field optimisation, we have found that small modifications of the holes position are sufficient to obtain good matches between the fundamental mode and the fourth mode of the cavity. We have then applied this tuning principle experimentally and successfully demonstrated photonic crystal cavities with modes matching the 15.6 THz Raman shift of silicon, while retaining a good Q-factor.

These results are only a first step that opens the door to further improvement, and additional modifications of the cavity design may be explored, for example to increase also the Q-factor of the high-order pump mode.

From the experimental point of view, we did attempt to observe spontaneous Raman scattering from some of the measured structures using the tunable laser as a pump (NIR source in Fig. 6.12a). The only other alteration to the setup described in section 6.4 was orienting the cavity axis parallel to the x-direction (Figs. 6.12b-c): this way, the in-coupling is optimised for the x-polarised high-order pump mode, whereas the out-coupling is optimised for the y-polarised fundamental signal mode. Unfortunately, we were unable to detect any spontaneous Raman signal, most likely due to the noise level of the InGaAs array detector, which was too high for our

purpose. The measurement will be repeated in the near future with a more sensitive setup.

In light of the experimental results on stimulated Raman scattering in slow light waveguides discussed in section 5.5, suggesting that the Raman gain coefficient may be lower than 20 cm/GW, we do not expect to be able to go as far as lasing with this particular sample: in the numerical study from Ref. [43], while taking into account all possible detrimental effects, Yang and Wong assume a Raman gain coefficient $g_R = 29$ cm/GW, and predict that photonic crystal cavities can lase for Q_s on the order of a few tens of thousands and $Q_p < 1000$ *in pulsed operation*. In continuous-wave operation, unless the free-carrier lifetime is reduced to values on the order of 100 ps, lasing of the Stokes signal is suppressed due to the combined effect of free-carrier loss and thermal drifts causing shifts of the resonant wavelengths.

The first step in our experimental study is to observe spontaneous scattering, which would be already an important achievement. Noda and coworkers observed ~150 fW of Stokes signal by pumping in the fundamental mode of an L3 cavity, and with the Stokes wavelength corresponding to a leaky mode of the crystal (see Fig. 6.5), so that the 3×10^4 Q-factor of the pump was the only mechanism of enhancement. In our case, even with lower signal Q in the range 1.2–3×10^4 and with pump Q of around 300 (Table 6.1), the product $Q_p Q_s$ is on the order 3.6–9×10^6. We therefore believe it is reasonable to expect spontaneous signals on the order of hundreds of femtowatts, if not few picowatts, in agreement also with the experimental results by Checoury *et al.* [241], who studied a 50 μm-long photonic crystal Fabry-Perot cavity and collected a Raman spontaneous signal of a few tens of picowatts for few milliwatts pump power.

Chapter 7

Summary and conclusions

The prospect for on-chip all-optical signal processing, with silicon photonics complementing silicon electronics, has inspired in the past decade great research effort into achieving active functionalities in silicon miniaturised devices. In this study we have investigated two such active functionalities based on the platform of slow light photonic crystals: tunable delay and optical gain for loss compensation, through stimulated Raman scattering.

In Chapter 3 we have proposed a new device for tunable delay based on a "poor man's" version of a photonic indirect transition, consisting of a two-step change of the frequency and wavevector of light, and equivalent to a wavelength conversion and group velocity dispersion approach (Fig. 7.1a). The wavelength conversion process is based on the relatively new effect of optical adiabatic tuning, and we rely on the precise control of a waveguide's dispersion offered by the techniques used for slow light engineering. We have then applied this scheme to demonstrate continuously tunable delays of optical pulses in a 300 μm-long photonic crystal waveguide (Figs. 7.1b-c) and we have succeeded in the manipulation of a single pulse in a train of two pulses (Fig. 7.1d).

We have already discussed in section 3.6 the advantages of our approach in terms of flexibility and scalability, as well as constraints posed on the bit rate by the free-carrier lifetime, and on the maximum achievable delay by the dispersion and propagation loss. For future development, steps could be taken in the reduction of the carrier lifetime, for instance by integrating a reverse-biased p-i-n junction to sweep out carriers from the short tuning region [56] or by treating the silicon slab with ion implantation and annealing [174]. The distortion of optical pulses due to group velocity dispersion in the slow light regime may be addressed by introducing one or more additional waveguide sections with opposite dispersion, designed to operate dispersion compensation while maintaining the delay.

Propagation loss in slow light waveguides is a more immediate issue, and the prospect for its compensation by introducing optical gain through stim-

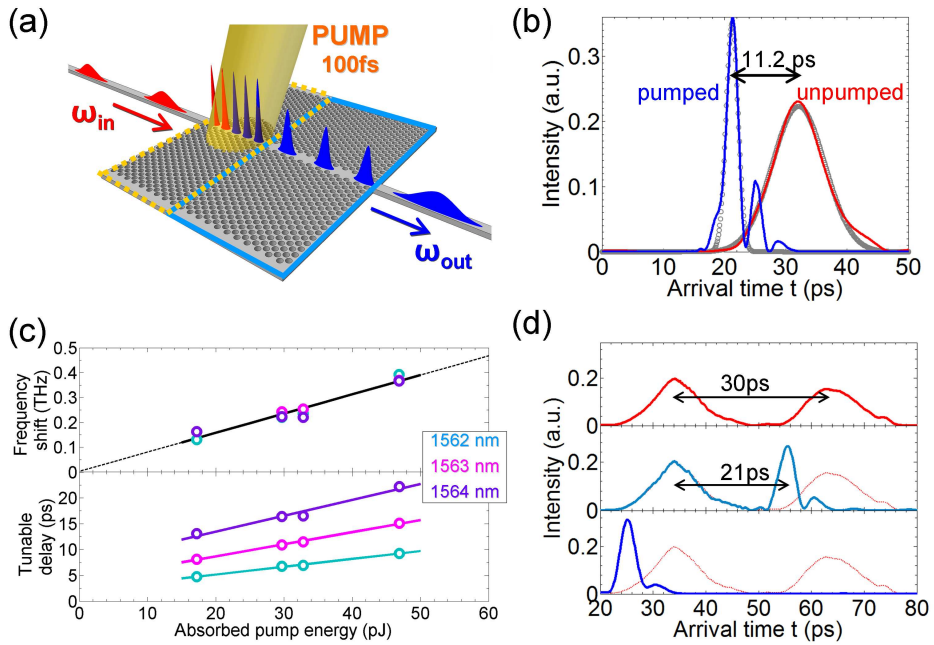

Figure 7.1: Selection of obtained results for tunable delay. (a) Schematic of the working principle of the device based on a two-step indirect photonic transition. (b) Experimental demonstration of tunable delay. (c) Demonstration of continuous tunability. (d) Operation on a single pulse in a train of two bits.

ulated Raman scattering has motivated the work described in Chapter 5. Here we have started by understanding the roles played by the slowdown factors of both the pump and the Stokes beams in determining the total gain of a signal propagating through a silicon waveguide, and we have described a simplified model taking into account the interplay of linear and nonlinear losses. We have then extended this model by including the effect of realistic intrinsic propagation loss and mode profiles; we have also included the anisotropy of Raman scattering in silicon, and shown that with our pump-signal polarisation scheme, orienting the waveguide along the [100] direction of crystalline silicon corresponds to a more efficient Raman scattering than along the standard [110] direction. Our analysis has brought us to the conclusion that, if the design properties are carefully evaluated, stimulated Raman scattering in slow light photonic crystals may surpass the performance of ridge waveguides, in terms of operating power and required propagation length (Fig. 7.2a).

We have then presented the experimental work conducted for the observation of Raman scattering from silicon slow light waveguides. The fabrication process of inverse taper spot-size converters was described in detail in Chapter 4, together with the solution to the challenges arising from the

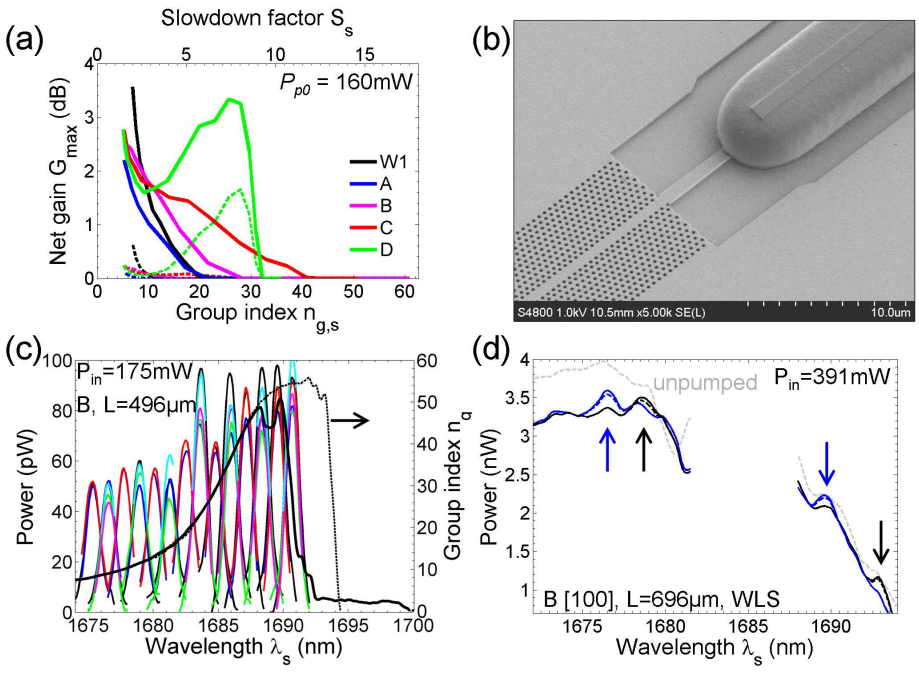

Figure 7.2: Selection of obtained results on Raman scattering in slow light waveguides. (a) Calculated peak net signal gain as a function of the signal slowdown factor for waveguide oriented along the standard [110] direction of crystalline silicon (dashed) and along the [100] direction (solid). (b) SEM image of a fabricated photonic crystal waveguide with integrated inverse taper spot-size converters. (c) Increase of the spontaneous Raman signal (coloured curves) with group index (black) in a waveguide aligned along the standard [110] direction. (d) Spectra measured in a pump-probe experiment, demonstrating stimulated Raman scattering for a waveguide oriented in the [100] direction.

integration of these features with photonic crystals (Fig. 7.2b). Experimental results on the slow light enhancement of spontaneous Raman scattering were discussed in the second part of Chapter 5 (Fig. 7.2c), including some preliminary results on the possible increase of the spontaneous Raman signal for waveguides aligned in the [100] direction. In addition, we have reported on the observation of stimulated Raman scattering in a pump-probe experiment (Fig. 7.2d); while we could not draw conclusions about the influence of slow light on this effect, it is to our knowledge the first time that this type of experiment is reported from silicon photonic crystal waveguides.

It appears from our data, however, that the loss caused by TPA and FCA is higher than estimated, possibly due to the free-carrier lifetime, and that the Raman gain was not always sufficient to compensate for such pump-induced loss, at least for waveguides aligned along the standard [110] direction of crystalline silicon. The measurement results appear to hint to a more efficient loss compensation for waveguides oriented in the [100] direction, but

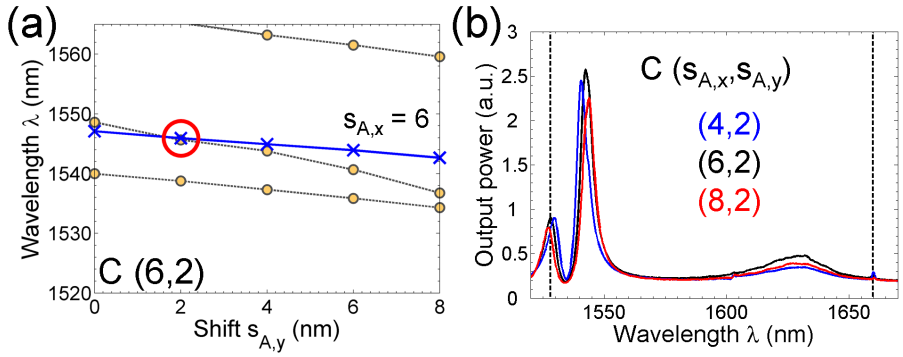

Figure 7.3: Selection of obtained results on tuning of the modes spacing in silicon L3 photonic crystal cavities by shifting the position of the holes, with a match to the Raman shift of silicon obtained for shifts $(s_{A,x}, s_{A,y}) = (6, 2)$ nm (design C). (a) Simulation and (b) experiment.

further investigation and optimisation would be required, especially in the fabrication of waveguides tilted at 45° to the cleaving plane and of the inverse tapers couplers, which could not sustain higher pump powers than those reported in section 5.4.4. The side-wall roughness of the inverse tapers could be improved through surface treatment in an hydrogen plasma or by surface oxidation. It would be also worth exploring the use of grating couplers instead, which would require two different designs for the pump and the signal, due to their very different wavelengths and their opposite polarisations.

If the photonic crystal design would allow at least for a moderate reduction of the group velocity of the pump, the power thresholds needed would be reduced, lowering also the load on the inverse tapers. Some design work towards this end was already initiated, as discussed in section 5.5, but it would require further development before being applied in a real device.

Finally, driven by the prediction that also photonic crystal cavities have the ability to enhance the Raman nonlinearity [43, 261], we have shown in Chapter 6 that it is possible to control the spacing between resonant modes in order to match the Raman shift of silicon, simply by introducing a small shift in the position of the holes adjacent the cavity (Fig. 7.3a). We have also demonstrated such match experimentally for the first time (Fig. 7.3b). We attempted to observe Raman scattering from the fabricated structures, but could not detect any spontaneously generated signal due to limitations of the measurement setup. The measurement will be repeated in the near future with a more sensitive setup, but further improvement is still possible also from the design point of view, and steps can be taken to increase the Q-factor of the pump mode. While lasing is probably still very far, we believe it should be possible to observe at least spontaneous scattering.

During this project I have gained deeper understanding about the properties of silicon photonic crystals slow light waveguides and cavities, and about nonlinear optical effects such as stimulated Raman scattering. I had the opportunity of working on such structures at first hand and acquire experience in all aspects of their development: from the theory and simulations to fabrication and characterisation.

The two most important outcomes of the project are the proposal of a novel concept for tuning the delay of pulses in photonic crystals, with its successful application in a proof of principle experiment, and the understanding of the dynamics of Raman scattering in real-life slow light waveguides, taking into account effects which have been neglected in other studies. In addition, the experimental results on Raman scattering in both slow light waveguides and nanocavities represent a first important step to answer questions that are still open to further investigation.

If time would have permitted, I believe it would have been interesting to further explore the option of the waveguides oriented at 45° to the cleaving plane, *if* the transmission issues could be improved easily. Any major revision to the structure would not be worth unless we could scale the device for operation for a 1550 nm signal with a 1435 nm pump, so that the smaller size would allow for tighter optical confinement and thus enhance the nonlinearity.

The choice of operating at ~1686 nm was determined by the availability of a very good and narrow line-width tunable pump source — the laser amplified by the EDFA — at 1550 nm. In fact, even without considering the theoretical reduction of the nonlinearity due to scaling to higher period, this operating wavelength has proven to be one of the most problematic aspects from a practical point of view. We have already discussed in section 4.2.4 how the absorption of polymers at this wavelength complicates the fabrication process with the addition of a grinding step, which is needed to be able to cleave very short samples. Additional complications arise also in the measurement phase, and are mainly related to the detection system, as such wavelength falls at the edge of the detection capabilities of both the infrared photodetectors used during the alignment and the OSA used for acquiring the spectra. Particular care must also be paid to the available source at this wavelength: the white light source used for both the passive characterisation and the stimulated scattering measurements is pulsed, with peak powers high enough to burn not only tapers, but also photonic crystals and even facets of blank silicon waveguides if not properly filtered and attenuated.

The second major practical issue turned out to be the use of a pump in TM polarisation: the adopted configuration for Raman amplification (section 5.3.3) was a good starting point from the point of view of the photonic crystal properties, because the TM mode is the only available at the pump wavelength of standard or engineered waveguides, and it is very broadband,

meaning that we could focus only on engineering of the slow signal mode with no need to match the modes distance in the design phase and no issues from fabrication tolerances in the experiment. The real issue is to bring such TM-polarised power to the photonic crystal: the TM mode is less confined than the TE mode, meaning that bend radii and nanowire widths must be wider, the spot-size converters are less efficient and any misalignment of the SU-8 waveguides with the inverse tapers can result in very strong wavelength-dependent TM-to-TE conversion. In addition, the TM mode senses defects differently than the TE mode, and we cannot exclude the possibility that the inverse tapers could have sustained higher pump powers in TE polarisation.

Nevertheless, from calculations, the adopted configuration (~1686 nm signal and TM pump) was promising and thus worth trying. I think it is still a promising possibility for some loss compensation at 45° if a suitable pump source would allow for scaling back to 1550 nm, resolving at least the wavelength issue. Switching to a TE pump would require more work on the design to find a matching (and possibly slow) pump mode. My preliminary work towards this goal indicates that such designs are not impossible to find, but they require optimisation to experimentally achieve good coupling with the input access waveguides.

Regarding photonic crystals in general, I believe they are fascinating structures and extremely versatile: they offer many degrees of freedom, and small modifications in their geometry are sufficient to study new effects and achieve new functionalities which are not otherwise possible in standard waveguides. It is true, however, that such positive characteristic becomes one of the major drawbacks of photonic crystals in terms of sensitivity to fabrication fluctuations, as small changes/disorder in any geometrical parameter can cause the spectral properties to change, making these devices still not suitable for real-world practical applications.

As for slow light as a tool to enhance nonlinear optical effects, it is now clear that the enhancement of the interaction by (a power of) the slowdown factor S is only part of the full story: even without considering the competing TPA and FCA processes (which are important at least in silicon), the compression of a signal pulse along the propagation direction is partially counteracted by an enlargement of the mode in the transversal direction, and the linear propagation loss remains a major issue. This does not mean that slow light and dispersion engineering cannot find also very useful applications, as demonstrated by the realisation of compact switches [264] and by the recent employment in the generation of correlated photon pairs [249] with real and promising prospects for on-chip quantum computing. Slow light is perhaps more suited for some applications than for others, but it remains a fascinating feature of photonic crystals, opening the possibility of exploring new and exciting concepts.

As for the use of slow light to improve Raman amplifiers, the achievement

of gains similar to those predicted by the calculations would represent a very important result, but pump powers of the order of 100–200 mW are still too high to be of practical use in real integrated devices. The real limitations of slow light were not clearly understood in the initial studies and it was worth conducting a new analysis on the feasibility of such a device. If net positive gains are out of reach, the question whether Raman scattering can operate loss compensation for a 1550 nm signal is still open to further analysis.

References

[1] M. Hilbert and P. Lopez, "The worlds technological capacity to store, communicate, and compute information," *Science*, vol. 332, no. 6025, pp. 60–65, 2011.

[2] R. S. Tucker, G. Eisenstein, and S. K. Korotky, "Optical time-division multiplexing for very high bit-rate transmission," *J. Lightw. Technol.*, vol. 6, no. 11, pp. 1737–1749, 1988.

[3] R. S. Tucker, "Energy consumption and energy density in optical and electronic signal processing," *IEEE Photon. J.*, vol. 3, no. 5, pp. 821–833, 2011.

[4] K. Hinton, G. Raskutti, P. M. Farrell, and R. S. Tucker, "Switching energy and device size limits on digital photonic signal processing technologies," *IEEE J. Sel. Topics Quantum Electron.*, vol. 14, no. 3, pp. 938–945, 2008.

[5] A. Shacham, K. Bergman, and L. Carloni, "Photonic networks-on-chip for future generations of chip multiprocessors," *IEEE Trans. Comput.*, vol. 57, no. 9, pp. 1246–1260, 2008.

[6] D. A. B. Miller, "Are optical transistors the logical next step?," *Nature Photon.*, vol. 4, pp. 3–5, 2010.

[7] G. T. Reed, "The optical age of silicon," *Nature*, vol. 427, pp. 595–596, 2004.

[8] M. Lipson, "Silicon photonics: the optical spice rack," *Electron. Lett.*, vol. 45, no. 12, pp. 576–578, 2009.

[9] M. Asghari and A. V. Krishnamoorthy, "Silicon photonics: Energy-efficient communication," *Nature Photon.*, vol. 5, pp. 268–270, 2011.

[10] B. Jalali and S. Fathpour, "Silicon photonics," *J. Lightw. Technol.*, vol. 24, no. 12, pp. 4600–4615, 2006.

[11] A. E. Willner, B. Zhang, and L. Zhang, *Slow Light, Science and Applications*, ch. 16, pp. 321–346. CRC Press, 2009.

[12] R. S. Tucker, P.-C. Ku, and C. J. Chang-Hasnain, "Slow-light optical buffers: capabilities and fundamental limitations," *J. Lightw. Technol.*, vol. 23, no. 12, pp. 4046–4066, 2005.

[13] J. Spring and R. S. Tucker, "Photonic 2*2 packet switch with input buffers," *Electron. Lett.*, vol. 29, no. 3, pp. 284–285, 1993.

[14] J. Khurgin, "Adiabatically tunable optical delay lines and their performance limitations," *Opt. Lett.*, vol. 30, no. 20, pp. 2778–2780, 2005.

[15] S. Fan, M. F. Yanik, Z. Wang, S. Sandhu, and M. L. Povinelli, "Advances in theory of photonic crystals," *J. Lightw. Technol.*, vol. 24, no. 12, pp. 4493–4501, 2006.

[16] A. Melloni, F. Morichetti, C. Ferrari, and M. Martinelli, "Continuously tunable 1 byte delay in coupled-resonator optical waveguides," *Opt. Lett.*, vol. 33, no. 20, pp. 2389–2391, 2008.

[17] Q. Xu, P. Dong, and M. Lipson, "Breaking the delay-bandwidth limit in a photonic structure," *Nature Phys.*, vol. 3, pp. 406–410, 2007.

[18] J. Upham, Y. Tanaka, T. Asano, and S. Noda, "Dynamic increase and decrease of photonic crystal nanocavity Q factors for optical pulse control," *Opt. Express*, vol. 16, no. 26, pp. 21721–21730, 2008.

[19] J. Sharping, Y. Okawachi, J. van Howe, C. Xu, Y. Wang, A. Willner, and A. Gaeta, "All-optical, wavelength and bandwidth preserving, pulse delay based on parametric wavelength conversion and dispersion," *Opt. Express*, vol. 13, no. 20, pp. 7872–7877, 2005.

[20] Y. Okawachi, M. A. Foster, X. Chen, A. C. Turner-Foster, R. Salem, M. Lipson, C. Xu, and A. L. Gaeta, "Large tunable delays using parametric mixing and phase conjugation in Si nanowaveguides," *Opt. Express*, vol. 16, no. 14, pp. 10349–10357, 2008.

[21] Y. Dai, X. Chen, Y. Okawachi, A. C. Turner-Foster, M. A. Foster, M. Lipson, A. L. Gaeta, and C. Xu, "1 μs tunable delay using parametric mixing and optical phase conjugation in Si waveguides," *Opt. Express*, vol. 17, no. 9, pp. 7004–7010, 2009.

[22] M. A. Foster, A. C. Turner, J. E. Sharping, B. S. Schmidt, M. Lipson, and A. L. Gaeta, "Broad-band optical parametric gain on a silicon photonic chip," *Nature*, vol. 441, pp. 960–963, 2006.

[23] M. Notomi and S. Mitsugi, "Wavelength conversion via dynamic refractive index tuning of a cavity," *Phys. Rev. A*, vol. 73, no. 5, p. 051803, 2006.

[24] T. F. Krauss, "Why do we need slow light?," *Nature Photon.*, vol. 2, pp. 448–450, 2008.

[25] T. Baba, "Slow light in photonic crystals," *Nature Photon.*, vol. 2, pp. 465–473, 2008.

[26] B.-S. Song, S. Noda, T. Asano, and Y. Akahane, "Ultra-high-Q photonic double-heterostructure nanocavity," *Nature Mater.*, vol. 4, pp. 207–210, 2005.

[27] E. Kuramochi, M. Notomi, S. Mitsugi, A. Shinya, T. Tanabe, and T. Watanabe, "Ultrahigh-Q photonic crystal nanocavities realized by the local width modulation of a line defect," *Appl. Phys. Lett.*, vol. 88, no. 4, p. 041112, 2006.

[28] J. Li, T. P. White, L. O'Faolain, A. Gomez-Iglesias, and T. F. Krauss, "Systematic design of flat band slow light in photonic crystal waveguides," *Opt. Express*, vol. 16, no. 9, pp. 6227–6232, 2008.

[29] S. Hughes, L. Ramunno, J. F. Young, and J. E. Sipe, "Extrinsic optical scattering loss in photonic crystal waveguides: role of fabrication disorder and photon group velocity," *Phys. Rev. Lett.*, vol. 94, no. 3, p. 033903, 2005.

[30] L. O'Faolain, T. P. White, D. O'Brien, X. Yuan, M. D. Settle, and T. F. Krauss, "Dependence of extrinsic loss on group velocity in photonic crystal waveguides," *Opt. Express*, vol. 15, no. 20, pp. 13129–13138, 2007.

[31] L. O'Faolain, S. A. Schulz, D. M. Beggs, T. P. White, M. Spasenović, L. Kuipers, F. Morichetti, A. Melloni, S. Mazoyer, J. P. Hugonin, P. Lalanne, and T. F. Krauss, "Loss engineered slow light waveguides," *Opt. Express*, vol. 18, no. 26, pp. 27627–27638, 2010. The code written by D. M. Beggs and S. A. Schulz for prediction of linear propagation loss is freely available at: www.st-andrews.ac.uk/microphotonics.

[32] R. W. Boyd, *Nonlinear optics*. Academic Press, 2nd ed., 2003.

[33] A. Liu, H. Rong, M. Paniccia, O. Cohen, and D. Hak, "Net optical gain in a low loss silicon-on-insulator waveguide by stimulated Raman scattering," *Opt. Express*, vol. 12, no. 18, pp. 4261–4268, 2004.

[34] T.-K. Liang and H.-K. Tsang, "Efficient Raman amplification in silicon-on-insulator waveguides," *Appl. Phys. Lett.*, vol. 85, no. 16, p. 3343, 2004.

[35] O. Boyraz and B. Jalali, "Demonstration of 11 dB fiber-to-fiber gain in a silicon Raman amplifier," *IEICE Electron. Express*, vol. 1, no. 14, pp. 429–434, 2004.

[36] O. Boyraz and B. Jalali, "Demonstration of a silicon Raman laser," *Opt. Express*, vol. 12, no. 21, pp. 5269–5273, 2004.

[37] H. Rong, A. Liu, R. Jones, O. Cohen, D. Hak, R. Nicolaescu, A. Fang, and M. Paniccia, "An all-silicon Raman laser," *Nature*, vol. 433, pp. 292–294, 2005.

[38] H. Rong, R. Jones, A. Liu, O. Cohen, D. Hak, A. Fang, and M. Paniccia, "A continuous-wave Raman silicon laser," *Nature*, vol. 433, pp. 725–728, 2005.

[39] T. F. Krauss, "Slow light in photonic crystal waveguides," *J. Phys. D Appl. Phys.*, vol. 40, no. 9, pp. 2666–2670, 2007.

[40] J. F. McMillan, X. Yang, N. C. Panoiu, R. M. Osgood, and C. W. Wong, "Enhanced stimulated Raman scattering in slow-light photonic crystal waveguides," *Opt. Lett.*, vol. 31, no. 9, pp. 1235–1237, 2006.

[41] J. F. McMillan, M. Yu, D.-L. Kwong, and C. W. Wong, "Observation of spontaneous Raman scattering in silicon slow-light photonic crystal waveguides," *Appl. Phys. Lett.*, vol. 93, no. 25, p. 251105, 2008.

[42] X. Checoury, M. El Kurdi, Z. Han, and P. Boucaud, "Enhanced spontaneous Raman scattering in silicon photonic crystal waveguides on insulator," *Opt. Express*, vol. 17, no. 5, pp. 3500–3507, 2009.

[43] X. Yang and C. W. Wong, "Coupled-mode theory for stimulated Raman scattering in high-Q/Vm silicon photonic band gap defect cavity lasers," *Opt. Express*, vol. 15, no. 8, pp. 4763–4780, 2007.

[44] X. Yang and C. W. Wong, "Design of photonic band gap nanocavities for stimulated Raman amplification and lasing in monolithic silicon," *Opt. Express*, vol. 13, no. 12, pp. 4723–4730, 2005.

[45] J. D. Joannopoulos, S. G. Johnson, J. N. Winn, and R. D. Meade, *Photonic Crystals, Molding the Flow of Light*. Princeton University Press, 2nd ed., 2008.

[46] R. M. Camacho and J. C. Howell, *Slow Light, Science and Applications*, ch. 1, pp. 3–12. CRC Press, 2009.

[47] S.-W. Chang and S. L. Chuang, *Slow Light, Science and Applications*, ch. 2, pp. 13–35. CRC Press, 2009.

[48] J. Mørk, P. Lunnemann, W. Xue, Y. Chen, P. Kaer, and T. R. Nielsen, "Slow and fast light in semiconductor waveguides," *Semicond. Sci. Technol.*, vol. 25, no. 8, p. 083002, 2010.

[49] S. G. Johnson and J. D. Joannopoulos, "Block-iterative frequency-domain methods for Maxwell's equations in a planewave basis," *Opt. Express*, vol. 8, no. 3, pp. 173–190, 2001.

[50] T. F. Krauss, *Slow Light, Science and Applications*, ch. 4, pp. 59–75. CRC Press, 2009.

[51] M. Soljačić, S. G. Johnson, S. Fan, M. Ibanescu, E. Ippen, and J. D. Joannopoulos, "Photonic-crystal slow-light enhancement of nonlinear phase sensitivity," *J. Opt. Soc. Am. B*, vol. 19, no. 9, pp. 2052–2059, 2002.

[52] D. M. Beggs, T. P. White, L. O'Faolain, and T. F. Krauss, "Ultracompact and low-power optical switch based on silicon photonic crystals," *Opt. Lett.*, vol. 33, no. 2, pp. 147–149, 2008.

[53] Y. A. Vlasov, M. O'Boyle, H. F. Hamann, and S. J. McNab, "Active control of slow light on a chip with photonic crystal waveguides," *Nature*, vol. 438, pp. 65–69, 2005.

[54] L. Gu, W. Jiang, X. Chen, L. Wang, and R. T. Chen, "High speed silicon photonic crystal waveguide modulator for low voltage operation," *Appl. Phys. Lett.*, vol. 90, no. 7, p. 071105, 2007.

[55] L. O'Faolain, D. M. Beggs, T. P. White, T. Kampfrath, K. Kuipers, and T. F. Krauss, "Compact optical switches and modulators based on dispersion engineered photonic crystals," *IEEE Photon. J.*, vol. 2, no. 3, pp. 404–414, 2010.

[56] H. C. Nguyen, Y. Sakai, M. Shinkawa, N. Ishikura, and T. Baba, "10 Gb/s operation of photonic crystal silicon optical modulators," *Opt. Express*, vol. 19, no. 14, pp. 13000–13007, 2011.

[57] T. Kampfrath, D. M. Beggs, T. P. White, A. Melloni, T. F. Krauss, and L. Kuipers, "Ultrafast adiabatic manipulation of slow light in a photonic crystal," *Phys. Rev. A*, vol. 81, no. 4, p. 043837, 2010.

[58] M. Soljačić and J. D. Joannopoulos, "Enhancement of nonlinear effects using photonic crystals," *Nature Mater.*, vol. 3, pp. 211–219, 2004.

[59] B. Corcoran, C. Monat, C. Grillet, D. J. Moss, B. J. Eggleton, T. P. White, L. O'Faolain, and T. F. Krauss, "Green light emission in silicon through slow-light enhanced third-harmonic generation in photonic-crystal waveguides," *Nature Photon.*, vol. 3, pp. 206–210, 2009.

[60] C. Monat, M. Ebnali-Heidari, C. Grillet, B. Corcoran, B. J. Eggleton, T. P. White, L. O'Faolain, J. Li, and T. F. Krauss, "Four-wave mixing in slow light engineered silicon photonic crystal waveguides," *Opt. Express*, vol. 18, no. 22, pp. 22915–22927, 2010.

[61] J. Li, L. O'Faolain, I. H. Rey, and T. F. Krauss, "Four-wave mixing in photonic crystal waveguides: slow light enhancement and limitations," *Opt. Express*, vol. 19, no. 5, pp. 4458–4463, 2011.

[62] C. Monat, B. Corcoran, M. Ebnali-Heidari, C. Grillet, B. J. Eggleton, T. P. White, L. O'Faolain, and T. F. Krauss, "Slow light enhancement of nonlinear effects in silicon engineered photonic crystal waveguides," *Opt. Express*, vol. 17, no. 4, pp. 2944–2953, 2009.

[63] N. Le Thomas, Z. Diao, H. Zhang, and R. Houdré, "Statistical analysis of subnanometer residual disorder in photonic crystal waveguides: Correlation between slow light properties and structural properties," *J. Vac. Sci. Technol. B*, vol. 29, no. 5, p. 051601, 2011.

[64] O. Khayam and H. Benisty, "General recipe for flatbands in photonic crystal waveguides," *Opt. Express*, vol. 17, no. 17, pp. 14634–14648, 2009.

[65] L. H. Frandsen, A. V. Lavrinenko, J. Fage-Pedersen, and P. I. Borel, "Photonic crystal waveguides with semi-slow light and tailored dispersion properties," *Opt. Express*, vol. 14, no. 20, pp. 9444–9450, 2006.

[66] Y. Hamachi, S. Kubo, and T. Baba, "Slow light with low dispersion and nonlinear enhancement in a lattice-shifted photonic crystal waveguide," *Opt. Lett.*, vol. 34, no. 7, pp. 1072–1074, 2009.

[67] T. Kawasaki, D. Mori, and T. Baba, "Experimental observation of slow light in photonic crystal coupled waveguides," *Opt. Express*, vol. 15, no. 16, pp. 10274–10281, 2007.

[68] M. Notomi, K. Yamada, A. Shinya, J. Takahashi, C. Takahashi, and I. Yokohama, "Extremely large group-velocity dispersion of line-defect waveguides in photonic crystal slabs," *Phys. Rev. Lett.*, vol. 87, no. 25, p. 253902, 2001.

[69] A. Y. Petrov and M. Eich, "Zero dispersion at small group velocities in photonic crystal waveguides," *Appl. Phys. Lett.*, vol. 85, no. 21, pp. 4866–4868, 2004.

[70] M. F. Yanik and S. Fan, "Stopping light all optically," *Phys. Rev. Lett.*, vol. 92, no. 8, p. 083901, 2004.

[71] N. Le Thomas, H. Zhang, J. Jágerská, V. Zabelin, R. Houdré, I. Sagnes, and A. Talneau, "Light transport regimes in slow light photonic crystal waveguides," *Phys. Rev. B*, vol. 80, no. 12, p. 125332, 2009.

[72] L. C. Andreani and D. Gerace, "Light-matter interaction in photonic crystal slabs," *Phys. Status Solidi B*, vol. 244, no. 10, pp. 3528–3539, 2007.

[73] V. Savona, "Electromagnetic modes of a disordered photonic crystal," *Phys. Rev. B*, vol. 83, no. 8, p. 085301, 2011.

[74] E. Kuramochi, M. Notomi, S. Hughes, A. Shinya, T. Watanabe, and L. Ramunno, "Disorder-induced scattering loss of line-defect waveguides in photonic crystal slabs," *Phys. Rev. B*, vol. 72, no. 16, p. 161318, 2005.

[75] M. Notomi, T. Tanabe, A. Shinya, E. Kuramochi, H. Taniyama, S. Mitsugi, and M. Morita, "Nonlinear and adiabatic control of high-Q photonic crystal nanocavities," *Opt. Express*, vol. 15, no. 26, pp. 17458–17481, 2007.

[76] J. Li, L. O'Faolain, S. A. Schulz, and T. F. Krauss, "Low loss propagation in slow light photonic crystal waveguides at group indices up to 60," *Phot. Nano. Fund. Appl.*, 2012. in press.

[77] J. P. Hugonin, P. Lalanne, T. P. White, and T. F. Krauss, "Coupling into slow-mode photonic crystal waveguides," *Opt. Lett.*, vol. 32, no. 18, pp. 2638–2640, 2007.

[78] Y. A. Vlasov and S. J. McNab, "Coupling into the slow light mode in slab-type photonic crystal waveguides," *Opt. Lett.*, vol. 31, no. 1, pp. 50–52, 2006.

[79] S. G. Johnson, P. Bienstman, M. A. Skorobogatiy, M. Ibanescu, E. Lidorikis, and J. D. Joannopoulos, "Adiabatic theorem and continuous coupled-mode theory for efficient taper transitions in photonic crystals," *Phys. Rev. E*, vol. 66, no. 6, p. 066608, 2002.

[80] P. Pottier, M. Gnan, and R. M. De La Rue, "Efficient coupling into slow-light photonic crystal channel guides using photonic crystal tapers," *Opt. Express*, vol. 15, no. 11, pp. 6569–6575, 2007.

[81] C. Martijn de Sterke, K. B. Dossou, T. P. White, L. C. Botten, and R. C. McPhedran, "Efficient coupling into slow light photonic crystal waveguide without transition region: role of evanescent modes," *Opt. Express*, vol. 17, no. 20, pp. 17338–17343, 2009.

[82] N. Ozaki, Y. Kitagawa, Y. Takata, N. Ikeda, Y. Watanabe, A. Mizutani, Y. Sugimoto, and K. Asakawa, "High transmission recovery of slow light in a photonic crystal waveguide using a hetero group velocity waveguide," *Opt. Express*, vol. 15, no. 13, pp. 7974–7983, 2007.

[83] A. Melloni, A. Canciamilla, C. Ferrari, F. Morichetti, L. O'Faolain, T. F. Krauss, R. De La Rue, A. Samarelli, and M. Sorel, "Tunable delay lines in silicon photonics: coupled resonators and photonic crystals, a comparison," *IEEE Photon. J.*, vol. 2, no. 2, pp. 181–194, 2010.

[84] A. Sakai, I. Katoh, D. Mori, T. Baba, and Y. Takiguchi, "Anomalous low group velocity and low dispersion in simple line defect photonic crystal waveguides," in *Lasers and Electro-Optics Society, 2004. LEOS 2004. The 17th Annual Meeting of the IEEE*, vol. 2, pp. 884–885 Vol.2, nov. 2004.

[85] J. Adachi, N. Ishikura, H. Sasaki, and T. Baba, "Wide range tuning of slow light pulse in SOI photonic crystal coupled waveguide via folded chirping," *IEEE J. Sel. Topics Quantum Electron.*, vol. 16, no. 1, pp. 192–199, 2010.

[86] D. Mori and T. Baba, "Wideband and low dispersion slow light by chirped photonic crystal coupled waveguide," *Opt. Express*, vol. 13, no. 23, pp. 9398–9408, 2005.

[87] S. Kubo, D. Mori, and T. Baba, "Low-group-velocity and low-dispersion slow light in photonic crystal waveguides," *Opt. Lett.*, vol. 32, no. 20, pp. 2981–2983, 2007.

[88] R. Hao, E. Cassan, H. Kurt, X. Le Roux, D. Marris-Morini, L. Vivien, H. Wu, Z. Zhou, and X. Zhang, "Novel slow light waveguide with controllable delay-bandwidth product and utra-low dispersion," *Opt. Express*, vol. 18, no. 6, pp. 5942–5950, 2010.

[89] S. A. Schulz, L. O'Faolain, D. M. Beggs, T. P. White, A. Melloni, and T. F. Krauss, "Dispersion engineered slow light in photonic crystals: a comparison," *J. Opt.*, vol. 12, no. 10, p. 104004, 2010.

[90] D. M. Beggs, L. OFaolain, and T. F. Krauss, "Accurate determination of the functional hole size in photonic crystal slabs using optical methods," *Phot. Nano. Fund. Appl.*, vol. 6, no. 34, pp. 213–218, 2008.

[91] J. Hou, D. Gao, H. Wu, R. Hao, and Z. Zhou, "Flat band slow light in symmetric line defect photonic crystal waveguides," *IEEE Photon. Technol. Lett.*, vol. 21, no. 20, pp. 1571–1573, 2009.

[92] M. Ebnali-Heidari, C. Grillet, C. Monat, and B. J. Eggleton, "Dispersion engineering of slow light photonic crystal waveguides using microfluidic infiltration," *Opt. Express*, vol. 17, no. 3, pp. 1628–1635, 2009.

[93] L. Dai, T. Li, and C. Jiang, "Wideband ultralow high-order-dispersion photonic crystal slow-light waveguide," *J. Opt. Soc. Am. B*, vol. 28, no. 7, pp. 1622–1626, 2011.

[94] J. Ma and C. Jiang, "Flatband slow light in asymmetric line-defect photonic crystal waveguide featuring low group velocity and dispersion," *IEEE J. Quantum Electron.*, vol. 44, no. 8, pp. 763–769, 2008.

[95] P. Colman, S. Combrié, G. Lehoucq, and A. De Rossi, "Control of dispersion in photonic crystal waveguides using group symmetry theory," *Opt. Express*, vol. 20, no. 12, pp. 13108–13114, 2012.

[96] S. Rawal, R. Sinha, and R. M. De La Rue, "Slow light miniature devices with ultra-flattened dispersion in silicon-on-insulator photonic crystal," *Opt. Express*, vol. 17, no. 16, pp. 13315–13325, 2009.

[97] A. Säynätjoki, M. Mulot, J. Ahopelto, and H. Lipsanen, "Dispersion engineering of photonic crystal waveguides with ring-shaped holes," *Opt. Express*, vol. 15, no. 13, pp. 8323–8328, 2007.

[98] F.-C. Leng, W.-Y. Liang, B. Liu, T.-B. Wang, and H.-Z. Wang, "Wideband slow light and dispersion control in oblique lattice photonic crystal waveguides," *Opt. Express*, vol. 18, no. 6, pp. 5707–5712, 2010.

[99] T. Baba, D. Mori, K. Inoshita, and Y. Kuroki, "Light localizations in photonic crystal line defect waveguides," *IEEE J. Sel. Topics Quantum Electron.*, vol. 10, no. 3, pp. 484–491, 2004.

[100] R. J. P. Engelen, D. Mori, T. Baba, and L. Kuipers, "Two regimes of slow-light losses revealed by adiabatic reduction of group velocity," *Phys. Rev. Lett.*, vol. 101, p. 103901, 2008.

[101] K. Yamada, M. Notomi, A. Shinya, C. Takahashi, J. Takahashi, and H. Morita, "Single-mode lightwave transmission in SOI-type photonic-crystal line-defect waveguides with phase-shifted holes," *Electron. Lett.*, vol. 38, no. 2, pp. 74–75, 2002.

[102] D. Mori and T. Baba, "Dispersion-controlled optical group delay device by chirped photonic crystal waveguides," *Appl. Phys. Lett.*, vol. 85, no. 7, p. 1101, 2004.

[103] D. Mori, S. Kubo, H. Sasaki, and T. Baba, "Experimental demonstration of wideband dispersion-compensated slow light by a chirped photonic crystal directional coupler," *Opt. Express*, vol. 15, no. 9, pp. 5264–5270, 2007.

[104] S.-C. Huang, M. Kato, E. Kuramochi, C.-P. Lee, and M. Notomi, "Time-domain and spectral-domain investigation of inflection-point slow-light modes in photonic crystal coupled waveguides," *Opt. Express*, vol. 15, no. 6, pp. 3543–3549, 2007.

[105] T. Baba, T. Kawasaki, H. Sasaki, J. Adachi, and D. Mori, "Large delay-bandwidth product and tuning of slow light pulse in photonic crystal coupled waveguide," *Opt. Express*, vol. 16, no. 12, pp. 9245–9253, 2008.

[106] N. Ishikura, T. Baba, E. Kuramochi, and M. Notomi, "Large tunable fractional delay of slow light pulse and its application to fast optical correlator," *Opt. Express*, vol. 19, no. 24, pp. 24102–24108, 2011.

[107] M. Notomi, E. Kuramochi, and T. Tanabe, "Large-scale arrays of ultrahigh-Q coupled nanocavities," *Nature Photon.*, vol. 2, pp. 741–747, 2008.

[108] A. Yariv, Y. Xu, R. K. Lee, and A. Scherer, "Coupled-resonator optical waveguide: a proposal and analysis," *Opt. Lett.*, vol. 24, no. 11, pp. 711–713, 1999.

[109] A. Martínez, A. García, P. Sanchis, and J. Martí, "Group velocity and dispersion model of coupled-cavity waveguides in photonic crystals," *J. Opt. Soc. Am. A*, vol. 20, no. 1, pp. 147–150, 2003.

[110] W. J. Kim, W. Kuang, and J. O'Brien, "Dispersion characteristics of photonic crystal coupled resonator optical waveguides," *Opt. Express*, vol. 11, no. 25, pp. 3431–3437, 2003.

[111] A. Melloni, F. Morichetti, and M. Martinelli, "Linear and nonlinear pulse propagation in coupled resonator slow-wave optical structures," *Opt. Quantum Electron.*, vol. 35, no. 4-5, pp. 365–379, 2003.

[112] K. Üstün and H. Kurt, "Ultra slow light achievement in photonic crystals by merging coupled cavities with waveguides," *Opt. Express*, vol. 18, no. 20, pp. 21155–21161, 2010.

[113] S. Olivier, C. Smith, M. Rattier, H. Benisty, C. Weisbuch, T. Krauss, R. Houdré, and U. Oesterlé, "Miniband transmission in a photonic crystal coupled-resonator optical waveguide," *Opt. Lett.*, vol. 26, no. 13, pp. 1019–1021, 2001.

[114] E. Ozbay, M. Bayindir, I. Bulu, and E. Cubukcu, "Investigation of localized coupled-cavity modes in two-dimensional photonic bandgap structures," *IEEE J. Quantum Electron.*, vol. 38, no. 7, pp. 837–843, 2002.

[115] E. Kuramochi, T. Tanabe, H. Taniyama, M. Kato, and M. Notomi, "Observation of heavy photon state in ultrahigh-Q photonic crystal coupled resonator chain," in *Conference on Lasers and Electro-Optics/Quantum Electronics and Laser Science Conference and Photonic Applications Systems Technologies*, p. QMG2, Optical Society of America, 2007.

[116] J. Jágerská, N. Le Thomas, V. Zabelin, R. Houdré, W. Bogaerts, P. Dumon, and R. Baets, "Experimental observation of slow mode dispersion in photonic crystal coupled-cavity waveguides," *Opt. Lett.*, vol. 34, no. 3, pp. 359–361, 2009.

[117] N. Matsuda, T. Kato, K.-i. Harada, H. Takesue, E. Kuramochi, H. Taniyama, and M. Notomi, "Slow light enhanced optical nonlinearity in a silicon photonic crystal coupled-resonator optical waveguide," *Opt. Express*, vol. 19, no. 21, pp. 19861–19874, 2011.

[118] N. Le Thomas and R. Houdré, "Group velocity and energy transport velocity near the band edge of a disordered coupled cavity waveguide: an analytical approach," *J. Opt. Soc. Am. B*, vol. 27, no. 10, pp. 2095–2101, 2010.

[119] J. B. Khurgin, "Expanding the bandwidth of slow-light photonic devices based on coupled resonators," *Opt. Lett.*, vol. 30, no. 5, pp. 513–515, 2005.

[120] J. K. Poon, L. Zhu, G. A. DeRose, and A. Yariv, "Transmission and group delay of microring coupled-resonator optical waveguides," *Opt. Lett.*, vol. 31, no. 4, pp. 456–458, 2006.

[121] F. Xia, L. Sekaric, M. O'Boyle, and Y. A. Vlasov, "Coupled resonator optical waveguides based on silicon-on-insulator photonic wires," *Appl. Phys. Lett.*, vol. 89, no. 4, p. 041122, 2006.

[122] F. Xia, L. Sekaric, and Y. A. Vlasov, "Ultracompact optical buffers on a silicon chip," *Nature Photon.*, vol. 1, pp. 65–71, 2007.

[123] F. Morichetti, A. Melloni, C. Ferrari, and M. Martinelli, "Error-free continuously-tunable delay at 10 Gbit/s in a reconfigurable on-chip delay-line," *Opt. Express*, vol. 16, no. 12, pp. 8395–8405, 2008.

[124] L. O' Faolain, X. Yuan, D. McIntyre, S. Thoms, H. Chong, R. M. De La Rue, and T. F. Krauss, "Low-loss propagation in photonic crystal waveguides," *Electron. Lett.*, vol. 42, no. 25, pp. 1454–1455, 2006.

[125] A. Gomez-Iglesias, D. O'Brien, L. O'Faolain, A. Miller, and T. F. Krauss, "Direct measurement of the group index of photonic crystal waveguides via Fourier transform spectral interferometry," *Appl. Phys. Lett.*, vol. 90, no. 26, p. 261107, 2007.

[126] L.-S. Yan, L. Lin, A. Belisle, S. Wey, and X. S. Yao, "Programmable optical delay generator with uniform output and double-delay capability," *J. Opt. Netw.*, vol. 6, no. 1, pp. 13–18, 2007.

[127] C. Liu, Z. Dutton, C. H. Behroozi, and L. V. Hau, "Observation of coherent optical information storage in an atomic medium using halted light pulses," *Nature*, vol. 409, pp. 490–493, 2001.

[128] D. Dahan and G. Eisenstein, "Tunable all optical delay via slow and fast light propagation in a Raman assisted fiber optical parametric amplifier: a route to all optical buffering," *Opt. Express*, vol. 13, no. 16, pp. 6234–6249, 2005.

[129] K. Y. Song, M. Herráez, and L. Thévenaz, "Observation of pulse delaying and advancement in optical fibers using stimulated Brillouin scattering," *Opt. Express*, vol. 13, no. 1, pp. 82–88, 2005.

[130] Y. Okawachi, M. S. Bigelow, J. E. Sharping, Z. Zhu, A. Schweinsberg, D. J. Gauthier, R. W. Boyd, and A. L. Gaeta, "Tunable all-optical delays via Brillouin slow light in an optical fiber," *Phys. Rev. Lett.*, vol. 94, no. 15, p. 153902, 2005.

[131] J. Sharping, Y. Okawachi, and A. Gaeta, "Wide bandwidth slow light using a Raman fiber amplifier," *Opt. Express*, vol. 13, no. 16, pp. 6092–6098, 2005.

[132] Y. Okawachi, M. Foster, J. Sharping, A. Gaeta, Q. Xu, and M. Lipson, "All-optical slow-light on a photonic chip," *Opt. Express*, vol. 14, no. 6, pp. 2317–2322, 2006.

[133] F. G. Sedgwick, B. Pesala, J.-Y. Lin, W. S. Ko, X. Zhao, and C. J. Chang-Hasnain, "THz-bandwidth tunable slow light in semiconductor optical amplifiers," *Opt. Express*, vol. 15, no. 2, pp. 747–753, 2007.

[134] B. Pesala, F. Sedgwick, A. Uskov, and C. Chang-Hasnain, "Ultrahigh-bandwidth electrically tunable fast and slow light in semiconductor optical amplifiers," *J. Opt. Soc. Am. B*, vol. 25, no. 12, pp. C46–C54, 2008.

[135] F. Shinobu, N. Ishikura, Y. Arita, T. Tamanuki, and T. Baba, "Continuously tunable slow-light device consisting of heater-controlled silicon microring array," *Opt. Express*, vol. 19, no. 14, pp. 13557–13564, 2011.

[136] J. B. Khurgin and P. A. Morton, "Tunable wideband optical delay line based on balanced coupled resonator structures," *Opt. Lett.*, vol. 34, no. 17, pp. 2655–2657, 2009.

[137] J. Cardenas, M. A. Foster, N. Sherwood-Droz, C. B. Poitras, H. L. R. Lira, B. Zhang, A. L. Gaeta, J. B. Khurgin, P. Morton, and M. Lipson, "Wide-bandwidth continuously tunable optical delay line using silicon microring resonators," *Opt. Express*, vol. 18, no. 25, pp. 26525–26534, 2010.

[138] M. S. Rasras, C. K. Madsen, M. A. Cappuzzo, E. Chen, L. T. Gomez, E. J. Laskowski, A. Griffin, A. Wong-Foy, A. Gasparyan, A. Kasper, J. Le Grange, and S. S. Patel, "Integrated resonance-enhanced variable optical delay lines," *IEEE Photon. Technol. Lett.*, vol. 17, no. 4, pp. 834–836, 2005.

[139] J. Yang, N. K. Fontaine, Z. Pan, A. O. Karalar, S. S. Djordjevic, C. Yang, W. Chen, S. Chu, B. E. Little, and S. J. B. Yoo, "Continuously tunable, wavelength-selective buffering in optical packet switching networks," *IEEE Photon. Technol. Lett.*, vol. 20, no. 12, pp. 1030–1032, 2008.

[140] N. K. Fontaine, J. Yang, Z. Pan, S. Chu, W. Chen, B. E. Little, and S. J. B. Yoo, "Continuously tunable optical buffering at 40 Gb/s for optical packet switching networks," *J. Lightw. Technol.*, vol. 26, no. 23, pp. 3776–3783, 2008.

[141] T. Wang, F. Liu, J. Wang, Y. Tian, Z. Zhang, T. Ye, M. Qiu, and Y. Su, "Pulse delay and advancement in SOI microring resonators with mutual mode coupling," *J. Lightw. Technol.*, vol. 27, no. 21, pp. 4734–4743, 2009.

[142] X. Luo, H. Chen, and A. W. Poon, "Electro-optical tunable time delay and advance in silicon microring resonators," *Opt. Lett.*, vol. 35, no. 17, pp. 2940–2942, 2010.

[143] M. Lipson, "Guiding, modulating, and emitting light on silicon — challenges and opportunities," *J. Lightw. Technol.*, vol. 23, no. 12, pp. 4222–4238, 2005.

[144] R. Soref and B. Bennett, "Electrooptical effects in silicon," *IEEE J. Quantum Electron.*, vol. 23, no. 1, pp. 123–129, 1987.

[145] G. Lenz, B. J. Eggleton, C. K. Madsen, and R. E. Slusher, "Optical delay lines based on optical filters," *IEEE J. Quantum Electron.*, vol. 37, no. 4, pp. 525–532, 2001.

[146] S. Fan and M. L. Povinelli, *Slow Light, Science and Applications*, ch. 14, pp. 277–289. CRC Press, 2009.

[147] M. F. Yanik, W. Suh, Z. Wang, and S. Fan, "Stopping light in a waveguide with an all-optical analog of electromagnetically induced transparency," *Phys. Rev. Lett.*, vol. 93, no. 23, p. 233903, 2004.

[148] M. F. Yanik and S. Fan, "Stopping and storing light coherently," *Phys. Rev. A*, vol. 71, no. 1, p. 013803, 2005.

[149] D. M. Beggs, T. F. Krauss, L. Kuipers, and T. Kampfrath, "Ultrafast tilting of the dispersion of a photonic crystal and adiabatic spectral compression of light pulses," *Phys. Rev. Lett.*, vol. 108, no. 3, p. 033902, 2012.

[150] A. W. Elshaari, A. Aboketaf, and S. F. Preble, "Controlled storage of light in silicon cavities," *Opt. Express*, vol. 18, no. 3, pp. 3014–3022, 2010.

[151] J. Upham, Y. Tanaka, Y. Kawamoto, Y. Sato, T. Nakamura, B. S. Song, T. Asano, and S. Noda, "Time-resolved catch and release of an optical pulse from a dynamic photonic crystal nanocavity," *Opt. Express*, vol. 19, no. 23, pp. 23377–23385, 2011.

[152] Z. Zhu, D. J. Gauthier, A. L. Gaeta, and R. W. Boyd, *Slow Light, Science and Applications*, ch. 3, pp. 37–57. CRC Press, 2009.

[153] Y. Dai, Y. Okawachi, A. C. Turner-Foster, M. Lipson, A. L. Gaeta, and C. Xu, "Ultralong continuously tunable parametric delays via a cascading discrete stage," *Opt. Express*, vol. 18, no. 1, pp. 333–339, 2010.

[154] N. Alic, E. Myslivets, S. Moro, B. P.-P. Kuo, R. M. Jopson, C. J. McKinstrie, and S. Radic, "Microsecond parametric optical delays," *J. Lightw. Technol.*, vol. 28, no. 4, pp. 448–455, 2010.

[155] S. R. Nuccio, O. F. Yilmaz, X. Wang, H. Huang, J. Wang, X. Wu, and A. E. Willner, "Higher-order dispersion compensation to enable a 3.6 μs wavelength-maintaining delay of a 100 Gb/s DQPSK signal," *Opt. Lett.*, vol. 35, no. 17, pp. 2985–2987, 2010.

[156] H. Fukuda, K. Yamada, T. Shoji, M. Takahashi, T. Tsuchizawa, T. Watanabe, J.-i. Takahashi, and S.-i. Itabashi, "Four-wave mixing in silicon wire waveguides," *Opt. Express*, vol. 13, no. 12, pp. 4629–4637, 2005.

[157] R. Salem, M. A. Foster, A. C. Turner, D. F. Geraghty, M. Lipson, and A. L. Gaeta, "Signal regeneration using low-power four-wave mixing on silicon chip," *Nature Photon.*, vol. 2, pp. 35–38, 2008.

[158] J. N. Winn, S. Fan, J. D. Joannopoulos, and E. P. Ippen, "Interband transitions in photonic crystals," *Phys. Rev. B*, vol. 59, no. 3, pp. 1551–1554, 1999.

[159] Z. Yu and S. Fan, "Complete optical isolation created by indirect interband photonic transitions," *Nature Photon.*, vol. 3, pp. 91–94, 2009.

[160] H. Lira, Z. Yu, S. Fan, and M. Lipson, "Electrically driven nonreciprocity induced by interband photonic transition on a silicon chip," *Phys. Rev. Lett.*, vol. 109, no. 3, p. 033901, 2012.

[161] S. F. Preble, Q. Xu, and M. Lipson, "Changing the colour of light in a silicon resonator," *Nature Photon.*, vol. 1, pp. 293–296, 2007.

[162] T. Tanabe, M. Notomi, H. Taniyama, and E. Kuramochi, "Dynamic release of trapped light from an ultrahigh-q nanocavity via adiabatic frequency tuning," *Phys. Rev. Lett.*, vol. 102, no. 4, p. 043907, 2009.

[163] J. Upham, Y. Tanaka, T. Asano, and S. Noda, "On-the-fly wavelength conversion of photons by dynamic control of photonic waveguides," *Appl. Phys. Express*, vol. 3, no. 6, p. 062001, 2010.

[164] G. P. Agrawal, *Nonlinear fiber optics*. Academic Press, 3rd ed., 2001.

[165] D. K. Schroder, "Carrier lifetimes in silicon," *IEEE Trans. Electron Devices*, vol. 44, no. 1, pp. 160–170, 1997.

[166] Z. G. Ling, P. K. Ajmera, M. Anselment, and L. F. DiMauro, "Lifetime measurements in semiconductors by infrared absorption due to pulsed optical excitation," *Appl. Phys. Lett.*, vol. 51, no. 18, p. 1445, 1987.

[167] O. Palais and A. Arcari, "Contactless measurement of bulk lifetime and surface recombination velocity in silicon wafers," *J. Appl. Phys.*, vol. 93, no. 8, p. 4686, 2003.

[168] V. R. Almeida, C. A. Barrios, R. R. Panepucci, and M. Lipson, "All-optical control of light on a silicon chip," *Nature*, vol. 431, pp. 1081–1084, 2004.

[169] T. Tanabe, M. Notomi, S. Mitsugi, A. Shinya, and E. Kuramochi, "All-optical switches on a silicon chip realized using photonic crystal nanocavities," *Appl. Phys. Lett.*, vol. 87, no. 15, p. 151112, 2005.

[170] M. Notomi, T. Tanabe, A. Shinya, E. Kuramochi, and H. Taniyama, "On-chip all-optical switching and memory by silicon photonic crystal nanocavities," *Adv. Opt. Technol.*, vol. 2008, p. 568936, 2008.

[171] T. Tanabe, H. Taniyama, and M. Notomi, "Carrier diffusion and recombination in photonic crystal nanocavity optical switches," *J. Lightw. Technol.*, vol. 26, no. 11, pp. 1396–1403, 2008.

[172] T. Kampfrath, D. M. Beggs, T. F. Krauss, and L. K. Kuipers, "Complete response characterization of ultrafast linear photonic devices," *Opt. Lett.*, vol. 34, no. 21, pp. 3418–3420, 2009.

[173] L. Lepetit, G. Chériaux, and M. Joffre, "Linear techniques of phase measurement by femtosecond spectral interferometry for applications in spectroscopy," *J. Opt. Soc. Am. B*, vol. 12, no. 12, pp. 2467–2474, 1995.

[174] T. Tanabe, K. Nishiguchi, A. Shinya, E. Kuramochi, H. Inokawa, M. Notomi, K. Yamada, T. Tsuchizawa, T. Watanabe, H. Fukuda, H. Shinojima, and S. Itabashi, "Fast all-optical switching using ion-implanted silicon photonic crystal nanocavities," *Appl. Phys. Lett.*, vol. 90, no. 3, p. 031115, 2007.

[175] A. Chin, K. Y. Lee, B. C. Lin, and S. Horng, "Picosecond photoresponse of carriers in Si ion-implanted Si," *Appl. Phys. Lett.*, vol. 69, no. 5, p. 653, 1996.

[176] T. Tanabe, E. Kuramochi, H. Taniyama, and M. Notomi, "Electro-optic adiabatic wavelength shifting and Q switching demonstrated using a p-i-n integrated photonic crystal nanocavity," *Opt. Lett.*, vol. 35, no. 23, pp. 3895–3897, 2010.

[177] T. Shoji, T. Tsuchizawa, T. Watanabe, K. Yamada, and H. Morita, "Low loss mode size converter from 0.3 μm square Si wire waveguides to singlemode fibres," *Electron. Lett.*, vol. 38, no. 25, pp. 1669–1670, 2002.

[178] T. Tsuchizawa, K. Yamada, H. Fukuda, T. Watanabe, J.-i. Takahashi, M. Takahashi, T. Shoji, E. Tamechika, S. Itabashi, and H. Morita, "Microphotonics devices based on silicon microfabrication technology," *IEEE J. Sel. Topics Quantum Electron.*, vol. 11, no. 1, pp. 232–240, 2005.

[179] V. R. Almeida, R. R. Panepucci, and M. Lipson, "Nanotaper for compact mode conversion," *Opt. Lett.*, vol. 28, no. 15, pp. 1302–1304, 2003.

[180] S. McNab, N. Moll, and Y. Vlasov, "Ultra-low loss photonic integrated circuit with membrane-type photonic crystal waveguides," *Opt. Express*, vol. 11, no. 22, pp. 2927–2939, 2003.

[181] M. Pu, L. Liu, H. Ou, K. Yvind, and J. M. Hvam, "Ultra-low-loss inverted taper coupler for silicon-on-insulator ridge waveguide," *Opt. Commun.*, vol. 283, no. 19, pp. 3678–3682, 2010.

[182] F. Van Laere, G. Roelkens, M. Ayre, J. Schrauwen, D. Taillaert, D. Van Thourhout, T. F. Krauss, and R. Baets, "Compact and highly efficient grating couplers between optical fiber and nanophotonic waveguides," *J. Lightw. Technol.*, vol. 25, no. 1, pp. 151–156, 2007.

[183] D. Vermeulen, S. Selvaraja, P. Verheyen, G. Lepage, W. Bogaerts, P. Absil, D. Van Thourhout, and G. Roelkens, "High-efficiency fiber-to-chip grating couplers realized using an advanced CMOS-compatible Silicon-On-Insulator platform," *Opt. Express*, vol. 18, no. 17, pp. 18278–18283, 2010.

[184] A. Borreman, S. Musa, A. A. M. Kok, M. B. J. Diemeer, and A. Driessen, "Fabrication of polymeric multimode waveguides and devices in SU-8 photoresist using selective polymerization.," in *IEEE/LEOS Benelux Chapter 2002 Annual Symposium* (T. D. Visser, D. Lenstra, and H. F. Schouten, eds.), pp. 83–86, Amsterdam: Vrije Universiteit Amsterdam, 2002.

[185] M. Grande, L. O'Faolain, T. P. White, M. Spurny, A. D'Orazio, and T. F. Krauss, "Optical filter with very large stopband (≈300 nm) based on a photonic-crystal vertical-directional coupler," *Opt. Lett.*, vol. 34, no. 21, pp. 3292–3294, 2009.

[186] T. Wahlbrink, W. S. Tsai, M. Waldow, M. Först, J. Bolten, T. Mollenhauer, and H. Kurz, "Fabrication of high efficiency SOI taper structures," *Microelectron. Eng.*, vol. 86, no. 46, pp. 1117–1119, 2009.

[187] D. L. Snavely and J. Dubsky, "Near-IR spectra of polyethylene, polyethylene glycol, and polyvinylethyl ether," *J. Polym. Sci., Part A: Polym. Chem.*, vol. 34, pp. 2575–2579, 1996.

[188] A. Tuchbreiter, B. Kappler, R. Stockmann, R. Mülhaupt, and J. Honerkamp, "Near infrared reflection spectroscopy: A versatile tool for rapid characterization of olefin copolymers and high-throughput experiments," *Macromol. Mater. Eng.*, vol. 288, no. 1, pp. 29–34, 2003.

[189] M. Spurny, L. O'Faolain, D. A. P. Bulla, B. Luther-Davies, and T. F. Krauss, "Fabrication of low loss dispersion engineered chalcogenide photonic crystals," *Opt. Express*, vol. 19, no. 3, pp. 1991–1996, 2011.

[190] J. Leuthold, C. Koos, and W. Freude, "Nonlinear silicon photonics," *Nature Photon.*, vol. 4, pp. 535–544, 2010.

[191] B. Jalali, R. Claps, D. Dimitropoulos, and V. Raghunathan, "Light generation, amplification, and wavelength conversion via stimulated Raman scattering in silicon microstructures," in *Silicon Photonics* (D. J. Lockwood and L. Pavesi, eds.), vol. 94 of *Topics in Applied Physics*, pp. 199–239, Springer Berlin / Heidelberg, 2004.

[192] D. Dimitropoulos, R. Claps, Y. Han, and B. Jalali, "Nonlinear optics in silicon waveguides: stimulated Raman scattering and two-photon absorption," in *Integrated Optics: Devices, Materials, and Technologies VII* (Y. S. S. A. Tervonen, ed.), vol. 4987, pp. 140–148, 2003.

[193] Q. Xu, V. Almeida, and M. Lipson, "Time-resolved study of Raman gain in highly confined silicon-on-insulator waveguides," *Opt. Express*, vol. 12, no. 19, pp. 4437–4442, 2004.

[194] Q. Xu, V. R. Almeida, and M. Lipson, "Demonstration of high Raman gain in a submicrometer-size silicon-on-insulator waveguide," *Opt. Lett.*, vol. 30, no. 1, pp. 35–37, 2005.

[195] R. Claps, D. Dimitropoulos, Y. Han, and B. Jalali, "Observation of Raman emission in silicon waveguides at 1.54 μm," *Opt. Express*, vol. 10, no. 22, pp. 1305–1313, 2002.

[196] R. Claps, D. Dimitropoulos, V. Raghunathan, Y. Han, and B. Jalali, "Observation of stimulated Raman amplification in silicon waveguides," *Opt. Express*, vol. 11, no. 15, pp. 1731–1739, 2003.

[197] H. Rong, A. Liu, R. Nicolaescu, M. Paniccia, O. Cohen, and D. Hak, "Raman gain and nonlinear optical absorption measurements in a low-loss silicon waveguide," *Appl. Phys. Lett.*, vol. 85, no. 12, p. 2196, 2004.

[198] J. I. Dadap, R. L. Espinola, R. M. Osgood, S. J. McNab, and Y. A. Vlasov, "Spontaneous Raman scattering in ultrasmall silicon waveguides," *Opt. Lett.*, vol. 29, no. 23, pp. 2755–2757, 2004.

[199] R. L. Espinola, J. I. Dadap, R. M. Osgood, S. J. McNab, and Y. A. Vlasov, "Raman amplification in ultrasmall silicon-on-insulator wire waveguides," *Opt. Express*, vol. 12, no. 16, pp. 3713–3718, 2004.

[200] R. Claps, D. Dimitropoulos, and B. Jalali, "Stimulated Raman scattering in silicon waveguides," *Electron. Lett.*, vol. 38, no. 22, pp. 1352–1354, 2002.

[201] R. Claps, V. Raghunathan, D. Dimitropoulos, and B. Jalali, "Influence of nonlinear absorption on Raman amplification in silicon waveguides," *Opt. Express*, vol. 12, no. 12, pp. 2774–2780, 2004.

[202] T.-K. Liang and H.-K. Tsang, "Role of free carriers from two-photon absorption in Raman amplification in silicon-on-insulator waveguides," *Appl. Phys. Lett.*, vol. 84, no. 15, p. 2745, 2004.

[203] T.-K. Liang and H.-K. Tsang, "Nonlinear absorption and Raman scattering in silicon-on-insulator optical waveguides," *IEEE J. Sel. Topics Quantum Electron.*, vol. 10, no. 5, pp. 1149–1153, 2004.

[204] D. Dimitropoulos, R. Jhaveri, R. Claps, J. C. S. Woo, and B. Jalali, "Lifetime of photogenerated carriers in silicon-on-insulator rib waveguides," *Appl. Phys. Lett.*, vol. 86, no. 7, p. 071115, 2005.

[205] D. Dimitropoulos, S. Fathpour, and B. Jalali, "Limitations of active carrier removal in silicon Raman amplifiers and lasers," *Appl. Phys. Lett.*, vol. 87, no. 26, p. 261108, 2005.

[206] I. D. Rukhlenko, M. Premaratne, C. Dissanayake, and G. P. Agrawal, "Continuous-wave Raman amplification in silicon waveguides: beyond the undepleted pump approximation," *Opt. Lett.*, vol. 34, no. 4, pp. 536–538, 2009.

[207] S. Roy, S. K. Bhadra, and G. P. Agrawal, "Raman amplification of optical pulses in silicon waveguides: effects of finite gain bandwidth, pulse width, and chirp," *J. Opt. Soc. Am. B*, vol. 26, no. 1, pp. 17–25, 2009.

[208] R. Jones, H. Rong, A. Liu, A. Fang, M. Paniccia, D. Hak, and O. Cohen, "Net continuous wave optical gain in a low loss silicon-on-insulator waveguide by stimulated Raman scattering," *Opt. Express*, vol. 13, no. 2, pp. 519–525, 2005.

[209] H. Rong, Y.-H. Kuo, S. Xu, A. Liu, R. Jones, M. Paniccia, O. Cohen, and O. Raday, "Monolithic integrated Raman silicon laser," *Opt. Express*, vol. 14, no. 15, pp. 6705–6712, 2006.

[210] H. Rong, S. Xu, Y.-H. Kuo, V. Sih, O. Cohen, O. Raday, and M. Paniccia, "Low-threshold continuous-wave Raman silicon laser," *Nature Photon.*, vol. 1, pp. 232–237, 2007.

[211] H. Rong, S. Xu, O. Cohen, O. Raday, M. Lee, V. Sih, and M. Paniccia, "A cascaded silicon Raman laser," *Nature Photon.*, vol. 2, pp. 170–174, 2008.

[212] V. Sih, S. Xu, Y.-H. Kuo, H. Rong, M. Paniccia, O. Cohen, and O. Raday, "Raman amplification of 40 Gb/s data in low-loss silicon waveguides," *Opt. Express*, vol. 15, no. 2, pp. 357–362, 2007.

[213] V. Raghunathan, D. Borlaug, R. R. Rice, and B. Jalali, "Demonstration of a mid-infrared silicon Raman amplifier," *Opt. Express*, vol. 15, no. 22, pp. 14355–14362, 2007.

[214] M. Krause, H. Renner, and E. Brinkmeyer, "Analysis of Raman lasing characteristics in silicon-on-insulator waveguides," *Opt. Express*, vol. 12, no. 23, pp. 5703–5710, 2004.

[215] M. Krause, H. Renner, and E. Brinkmeyer, "Raman lasers in silicon photonic wires: unidirectional ring lasing versus Fabry-Perot lasing," *Electron. Lett.*, vol. 45, no. 1, pp. 42–43, 2009.

[216] X. Chen, N. C. Panoiu, and J. Osgood, R. M., "Theory of Raman-mediated pulsed amplification in silicon-wire waveguides," *IEEE J. Quantum Electron.*, vol. 42, no. 2, pp. 160–170, 2006.

[217] C. M. Dissanayake, I. D. Rukhlenko, M. Premaratne, and G. P. Agrawal, "Raman-mediated nonlinear interactions in silicon waveguides: Copropagating and counterpropagating pulses," *IEEE Photon. Technol. Lett.*, vol. 21, no. 19, pp. 1372–1374, 2009.

[218] I. D. Rukhlenko, M. Premaratne, I. L. Garanovich, A. A. Sukhorukov, and G. P. Agrawal, "Analytical study of pulse amplification in silicon Raman amplifiers," *Opt. Express*, vol. 18, no. 17, pp. 18324–18338, 2010.

[219] M. Krause, H. Renner, S. Fathpour, B. Jalali, and E. Brinkmeyer, "Gain enhancement in cladding-pumped silicon Raman amplifiers," *IEEE J. Quantum Electron.*, vol. 44, no. 7, pp. 692–704, 2008.

[220] I. D. Rukhlenko, C. Dissanayake, M. Premaratne, and G. P. Agrawal, "Maximization of net optical gain in silicon-waveguide Raman amplifiers," *Opt. Express*, vol. 17, no. 7, pp. 5807–5814, 2009.

[221] I. D. Rukhlenko, C. Dissanayake, M. Premaratne, and G. P. Agrawal, "Optimization of Raman amplification in silicon waveguides with finite facet reflectivities," *IEEE J. Sel. Topics Quantum Electron.*, vol. 16, no. 1, pp. 226–233, 2010.

[222] T. Y. L. Ang and N. Q. Ngo, "Enhanced tunable Raman amplification/attenuation and parametric Raman wavelength conversion in submicrometer silicon waveguides via temperature variation," *J. Opt. Soc. Am. B*, vol. 28, no. 6, pp. 1556–1565, 2011.

[223] Q. Lin, O. J. Painter, and G. P. Agrawal, "Nonlinear optical phenomena in silicon waveguides: modeling and applications," *Opt. Express*, vol. 15, no. 25, pp. 16604–16644, 2007.

[224] M. Krause, H. Rentier, and E. Brinkmeyer, "Optical isolation in silicon waveguides based on nonreciprocal Raman amplification," *Electron. Lett.*, vol. 44, no. 11, pp. 691–693, 2008.

[225] M. Krause, H. Renner, and E. Brinkmeyer, "Strong enhancement of Raman-induced nonreciprocity in silicon waveguides by alignment with the crystallographic axes," *Appl. Phys. Lett.*, vol. 95, no. 26, p. 261111, 2009.

[226] J. Müller, M. Krause, H. Renner, and E. Brinkmeyer, "Measurement of nonreciprocal spontaneous Raman scattering in silicon photonic wires," *Opt. Express*, vol. 18, no. 19, pp. 19532–19540, 2010.

[227] M. Krause, "Nonreciprocal light transmission in silicon by Raman-induced asymmetry of the permittivity tensor," *J. Appl. Phys.*, vol. 111, no. 9, p. 093107, 2012.

[228] Z. Luo, W. Ye, X. Yuan, Z. Zhu, C. Zeng, and J. Ji, "Complete optical on-chip isolation based on asymmetric stimulated Raman gain/loss," *Opt. Lett.*, vol. 37, no. 4, pp. 674–676, 2012.

[229] I. D. Rukhlenko and M. Premaratne, "Spectral compression and group delay of optical pulses in silicon Raman amplifiers," *Opt. Lett.*, vol. 35, no. 18, pp. 3138–3140, 2010.

[230] M. Khorasaninejad and S. S. Saini, "All-optical logic gates using nonlinear effects in silicon-on-insulator waveguides," *Appl. Opt.*, vol. 48, no. 25, pp. F31–F36, 2009.

[231] S. Abaslou and V. Ahmadi, "Compact all-optical switch for WDM networks based on Raman effect in silicon nanowavegide," *Opt. Lett.*, vol. 37, no. 1, pp. 40–42, 2012.

[232] V. V. Kozlov and S. Wabnitz, "Silicon Raman polarizer," *Opt. Lett.*, vol. 37, no. 4, pp. 737–739, 2012.

[233] A. Liu, H. Rong, R. Jones, O. Cohen, D. Hak, and M. Paniccia, "Optical amplification and lasing by stimulated Raman scattering in silicon waveguides," *J. Lightw. Technol.*, vol. 24, no. 3, pp. 1440–1455, 2006.

[234] M. Woldeyohannes, S. John, and V. I. Rupasov, "Resonance Raman scattering in photonic band-gap materials," *Phys. Rev. A*, vol. 63, no. 1, p. 013814, 2000.

[235] S. V. Gaponenko, "Effects of photon density of states on Raman scattering in mesoscopic structures," *Phys. Rev. B*, vol. 65, no. 14, p. 140303, 2002.

[236] R. G. Zaporozhchenko, S. Y. Kilin, and A. G. Smirnov, "Stimulated Raman scattering of light in a photonic crystal," *Quantum Electron.*, vol. 30, no. 11, pp. 997–1001, 2000.

[237] L. Florescu and X. Zhang, "Semiclassical model of stimulated Raman scattering in photonic crystals," *Phys. Rev. E*, vol. 72, no. 1, p. 016611, 2005.

[238] M. Krause, A. Petrov, and E. Brinkmeyer, "Total gain of silicon Raman amplifiers: Scaling with group velocity in slow-light waveguides," in *6th IEEE International Conference on Group IV Photonics (GFP '09)*, pp. 125–127, sept. 2009.

[239] I. D. Rukhlenko, M. Premaratne, and G. P. Agrawal, "Maximization of gain in slow-light silicon Raman amplifiers," *Int. J. Opt.*, vol. 2011, p. 581810, 2011.

[240] K. Inoue, H. Oda, A. Yamanaka, N. Ikeda, H. Kawashima, Y. Sugimoto, and K. Asakawa, "Dramatic density-of-state enhancement of Raman scattering at the band edge in a one-dimensional photonic-crystal waveguide," *Phys. Rev. A*, vol. 78, no. 1, p. 011805, 2008.

[241] X. Checoury, Z. Han, and P. Boucaud, "Stimulated Raman scattering in silicon photonic crystal waveguides under continuous excitation," *Phys. Rev. B*, vol. 82, no. 4, p. 041308, 2010.

[242] M. Dinu, F. Quochi, and H. Garcia, "Third-order nonlinearities in silicon at telecom wavelengths," *Appl. Phys. Lett.*, vol. 82, no. 18, p. 2954, 2003.

[243] A. D. Bristow, N. Rotenberg, and H. M. van Driel, "Two-photon absorption and Kerr coefficients of silicon for 850-2200 nm," *Appl. Phys. Lett.*, vol. 90, no. 19, p. 191104, 2007.

[244] D. M. Beggs, I. H. Rey, T. Kampfrath, N. Rotenberg, L. Kuipers, and T. F. Krauss, "Ultrafast tunable optical delay line based on indirect photonic transitions," *Phys. Rev. Lett.*, vol. 108, no. 21, p. 213901, 2012.

[245] N. C. Panoiu, J. F. McMillan, and C. W. Wong, "Theoretical analysis of pulse dynamics in silicon photonic crystal wire waveguides," *IEEE J. Sel. Topics Quantum Electron.*, vol. 16, no. 1, pp. 257–266, 2010.

[246] B. Jalali, V. Raghunathan, D. Dimitropoulos, and O. Boyraz, "Raman-based silicon photonics," *IEEE J. Sel. Topics Quantum Electron.*, vol. 12, no. 3, pp. 412–421, 2006.

[247] P. N. Butcher and D. Cotter, *The elements of nonlinear optics*, vol. 9 of *Cambridge Studies in Modern Optics*. Cambridge University Press, 1990.

[248] R. M. Osgood, N. C. Panoiu, J. I. Dadap, X. Liu, X. Chen, I.-W. Hsieh, E. Dulkeith, W. M. Green, and Y. A. Vlasov, "Engineering nonlinearities in nanoscale optical systems: physics and applications in dispersion-engineered silicon nanophotonic wires," *Adv. Opt. Photon.*, vol. 1, no. 1, pp. 162–235, 2009.

[249] C. Xiong, C. Monat, A. S. Clark, C. Grillet, G. D. Marshall, M. J. Steel, J. Li, L. O'Faolain, T. F. Krauss, J. G. Rarity, and B. J. Eggleton, "Slow-light enhanced correlated photon pair generation in a silicon photonic crystal waveguide," *Opt. Lett.*, vol. 36, no. 17, pp. 3413–3415, 2011.

[250] D. Englund, I. Fushman, and J. Vuckovic, "General recipe for designing photonic crystal cavities," *Opt. Express*, vol. 13, no. 16, pp. 5961–5975, 2005.

[251] Y. Akahane, T. Asano, B.-S. Song, and S. Noda, "High-Q photonic nanocavity in a two-dimensional photonic crystal," *Nature*, vol. 425, pp. 944–947, 2003.

[252] J. Vuckovic, M. Loncar, H. Mabuchi, and A. Scherer, "Optimization of the Q factor in photonic crystal microcavities," *IEEE J. Quantum Electron.*, vol. 38, no. 7, pp. 850–856, 2002.

[253] M. Fujita, Y. Tanaka, and S. Noda, "Light emission from silicon in photonic crystal nanocavity," *IEEE J. Sel. Topics Quantum Electron.*, vol. 14, no. 4, pp. 1090–1097, 2008.

[254] Y. Takahashi, H. Hagino, Y. Tanaka, B.-S. Song, T. Asano, and S. Noda, "High-Q nanocavity with a 2-ns photon lifetime," *Opt. Express*, vol. 15, no. 25, pp. 17206–17213, 2007.

[255] Y. Taguchi, Y. Takahashi, Y. Sato, T. Asano, and S. Noda, "Statistical studies of photonic heterostructure nanocavities with an average Q factor of three million," *Opt. Express*, vol. 19, no. 12, pp. 11916–11921, 2011.

[256] Y. Akahane, T. Asano, B.-S. Song, and S. Noda, "Fine-tuned high-Q photonic-crystal nanocavity," *Opt. Express*, vol. 13, no. 4, pp. 1202–1214, 2005.

[257] D. Gerace and L. C. Andreani, "Effects of disorder on propagation losses and cavity Q-factors in photonic crystal slabs," *Phot. Nano. Fund. Appl.*, vol. 3, no. 23, pp. 120–128, 2005.

[258] N.-V.-Q. Tran, S. Combrié, and A. De Rossi, "Directive emission from high-Q photonic crystal cavities through band folding," *Phys. Rev. B*, vol. 79, no. 4, p. 041101, 2009.

[259] S. L. Portalupi, M. Galli, C. Reardon, T. F. Krauss, L. O'Faolain, L. C. Andreani, and D. Gerace, "Planar photonic crystal cavities with far-field optimization for high coupling efficiency and quality factor," *Opt. Express*, vol. 18, no. 15, pp. 16064–16073, 2010.

[260] M. Galli, D. Gerace, K. Welna, T. F. Krauss, L. O'Faolain, G. Guizzetti, and L. C. Andreani, "Low-power continuous-wave generation of visible harmonics in silicon photonic crystal nanocavities," *Opt. Express*, vol. 18, no. 25, pp. 26613–26624, 2010.

[261] Y. Takahashi, R. Terawaki, M. Chihara, T. Asano, and S. Noda, "First observation of Raman scattering emission from silicon high-Q photonic crystal nanocavities," in *Quantum Electronics and Laser Science Conference*, p. QWC3, Optical Society of America, 2011.

[262] C. P. Reardon, I. H. Rey, K. Welna, L. O'Faolain, and T. F. Krauss, "Fabrication and characterisation of photonic crystal slow light waveguides and cavities," *J. Vis. Exp.*, 2012. in press.

[263] M. Galli, S. L. Portalupi, M. Belotti, L. C. Andreani, L. O'Faolain, and T. F. Krauss, "Light scattering and Fano resonances in high-Q photonic crystal nanocavities," *Appl. Phys. Lett.*, vol. 94, no. 7, p. 071101, 2009.

[264] D. M. Beggs, T. P. White, L. Cairns, L. O'Faolain, and T. F. Krauss, "Ultrashort photonic crystal optical switch actuated by a microheater," *IEEE Photon. Technol. Lett.*, vol. 21, no. 1, pp. 24–26, 2009.